CERN-RRB-2011-019

Computing Resources Scrutiny Group

C.Bozzi (Italy), T. Cass (CERN), G.Lamanna (France), D.Espriu (Spain, *Chairman*), J.Flynn (UK), M.Gasthuber (Germany), D.Groep (The Netherlands), T.Schalk (USA), W.Trischuk (Canada), B.Vinter (Nordic Grid), H.Meinhard (CERN/IT, *Scientific Secretary*)

INTRODUCTION

The Computing Resources Scrutiny Group (CRSG) is an independent committee established by the WLCG Memorandum of Understanding whose members are selected by the Funding Agencies represented in the C-RRB. The purpose of the CRSG is to inform the decisions of the Computing Resources Review Board (C-RRB) for the LHC experiments.

According to the WLCG MoU, every year the CRSG shall scrutinize

- The resource accounting figures for the preceding year
- The use the experiments made of these resources
- The overall request for resources for every experiment for the following year and forecasts for the subsequent two years

The CRSG shall also examine the match between the refereed requests and the pledges from the institutions and make recommendations concerning apparent under-funding.

This report summarizes the deliberations of the CRSG regarding the usage of the computing resources by the four LHC experiments during the whole year 2010. We have also examined the resource accounting figures for the whole year 2010. A partial summary covering the first nine months was presented in the October 2010 C-RRB.

We have scrutinized the experiments' requests for the year 2012 and 2013. At the Chamonix meeting it was agreed to postpone the shutdown of the LHC to 2013. This means that the combined 2010 and 2011 run has in practice been extended to 2012 implying a substantial modification of the planning of resources.

In 2010 the LHC has increased its luminosity by five orders of magnitude The increase in the LHC delivery has been very marked in the last weeks of running reaching and exceeding the targeted 10^32 luminosity. The switch to HI running was equally successful, taking only a few days to record the first Pb-Pb interactions and fulfilling expectations. As a consequence the experiments' computing models have been put to a real test and the Tier 1 and Tier 2 centers have dealt with very substantial amounts of data.

The resources made available to the experiments in 2010 have been sufficient for the experiments to take, record, distribute and analyze their results, proving the overall validity of the computing models. The experiments have continuously assessed the adequacy of their models and have adapted the data distribution to the real needs, often in the direction that was suggested to them by the CRSG. A full account will be provided in this document.

In 2011 the expected LHC schedule will not substantially change with the respect to the one adopted in the scrutiny one year ago. However, the characteristics of the LHC have forced a decrease in the number of bunches in the machine making them more squeezed and longer. As a consequence the experiments are facing increased pile-up implying further reassessment of the data distribution policies. The experiments' perception is that these readjustments as well as the constant optimization of other parameters in their models will allow them to complete their physics objectives in 2011 within the available resources. As these conditions will continue up to the shutdown, they impact the 2012 and 2013 resource request to some extent, in spite of the optimization of the models.

In the previous C-RRB the experiments were asked in to report on ways to mitigate the growth in resources as luminosity increases in order to guarantee the sustainability of the WLCG effort. We have devoted a section in this report to review the steps taken in order to limit the growth in the resources so that they can be accommodated within a reasonable expenditure profile.

The LHC running conditions

The expected schedule of the LHC is the most essential ingredient of the scrutiny.

After the Chamonix meeting early in 2011, it was agreed that, barring unforeseen circumstances, the LHC should run for the best part of 2011 and 2012 (the estimate is 8 months per year), with only a relatively short break at the end of 2011. The shutdown has been postponed by one year The energy has been fixed at 7 TeV (3.5 GeV + 3.5 GeV) as in 2010. About 10% of the time is expected to be dedicated to heavy ion (HI) physics. After this long run a long shutdown will follow to enable the machine to reach the design energy.

For the scrutiny the most relevant quantity is the total number of seconds when the beam is declared to be stable and good for physics. After receiving input from the CERN management the following scheme has been adopted:

Live time: 30 days/month = **720** hours
Folding in efficiencies 720 x 0.7 x 0.4 = 201.6 effective hours/month = **725760** sec/month

RRB year	RRB year start	RRB year end	Months (max) Data taking	Total live time (in Ms)	pp	HI
2010	**June '10**	**March '11**	**8**	**5.8**	**5.1 (~ 3)**	**0.7**
2011	**April '11**	**March '12**	**8**	**5.9**	**5.2**	**0.7**
2012	**April '12**	**March '13**	**8**	**5.9**	**5.2**	**0.7**

The above scheme is an optimal one, assuming no unscheduled interruptions and continuous smooth running of the accelerator. During 2010 LHC has provided in reality approximately 3Ms of usable data, namely 60% of the total expected time delivery. Even if this falls short of expectations it has to be considered an excellent record for an accelerator in its first full year of physics running. In view of the excellent performance of the machine in the last part of 2010 it seems natural to assume that the fraction of time delivered for physics in 2011 and 2012 will be closer to expectations with less room for the experiments to compensate computing shortages.

On average the experiments have recorded about 30% fewer pp events than originally envisaged. Nonetheless the computing models and the WLCG distributed computing have been put to a complete test that on the whole has been passed brilliantly. The ICHEP conference in summer 2010 represented a particular moment of stress in the system when sizeable volumes of data just taken were reprocessed and analyzed in a matter of days. At the time of writing the experiments have completed their analysis for the winter conferences with quite satisfactory results too. At the end of 2010 the first HI run has brought the new challenge of dealing with very large event sizes. ALICE, ATLAS and CMS have managed within the available resources to reach their physics objectives in this area too.

The running conditions have changed in one important aspect. Proton bunches will be injected with expected minimal separation of 50ns but not below that in order to guarantee beam stability. To compensate they will contain more protons, squeezing them as much as possible to sustain a substantial luminosity which is expected to be increased in the coming run by another order of magnitude to ~ 10^33. The consequence is the appearance of events with many interactions at an earlier stage in the planning. The experiments have estimated that this may reach up to ~20 interactions per crossing towards the end of the run and indeed crossings with ~10 interactions have already been observed since the new run started early in March 2011. This has a substantial impact in reconstruction times and on the size of the data sets, typically doubling them.

Interactions with the LHCC

In carrying out the scrutiny of the experiments' requests the scope of this group is largely limited to the implementation of the respective computing models which are periodically reviewed by the LHCC. Obviously there is a gray area where the respective competences of the CRSG and the LHCC overlap and it is not so clear what would represent a change of the computing model or what would just be a natural adaptation of it to the changing physics circumstances.

In practice, after confronting the computing models with real data, a number of changes have been made, sometimes representing limitations in the original model or assumptions, some of which were already mentioned in previous reports, and others forced by the modification of the running conditions.

The CRSG was invited to a LHCC computing review on March 21st at which the experiments presented their requests and the revisions of their respective computing models, which in some cases are of some substance. The LHCC recommendations following this review have been taken into account in this report.

Since the last scrutiny no issues appeared for which we thought it was necessary to refer to the LHCC.

Interactions with the experiments

The recommendations of the previous C-RRB report in October 2010 urged the experimental collaborations to submit a detailed account of resource usage by March 1st. Three of the experiments complied with this deadline while one of the large experimental collaborations provided their formal request only three weeks after the deadline. Once more we have to kindly ask the experimental collaborations to adhere to the deadlines as the short time available for this scrutiny group's deliberations makes our task more difficult and is the ultimate justification for this report being submitted to the C-RRB with such short notice, for which we have to apologize.

Following the reception of the reports, referees were assigned to the different experiments and questions and answers exchanged in a frank and constructive atmosphere. As agreed with the ATLAS and CMS management in 2009 the scrutiny procedure for these two experiments is done by a common team of referees, using a coherent set of principles.

At the October 2010 C-RRB the experiments were asked to provide the CRSG with a list of steps taken to prevent an unsustainable growth of the computing resources as higher luminosities are reached. The experiments have tackled this issue in face of the realistic running conditions that they face now including the larger than planned pile up.

Experiments were also requested to provide the CRSG with enough information so that the gap between our recommendations and their requests can be narrowed. We find that the level of information provided to us is improving and generally sufficient. We nevertheless still find room for improvement in the information provided to us by the ALICE collaboration as the initial report was very short. We emphasize that the core of the ALICE collaboration computing model lies in Heavy Ion (HI) collisions. The initial HI run took place at the end of 2010 so the CRSG feels that following up the developments in this area is quite relevant now.

Generally speaking the interactions with the experiments are quite fluid and we thank the respective managements for their openness and collaboration. Thanks are due in particular to Ian Fisk who compiled and summarized the Tier 1 and Tier 2 usage for 2010 greatly facilitating our task.

Mitigation of the resource growth

Experiments have been urged to use the experience gained in the first year of running to modify the implementation of the computing models in ways that would make them sustainable in the long run and mitigate the growth in resources, In addition the emergence of larger-than-planned pile up has meant a real need for optimization to temper the need for additional resources. In our scrutiny we have seen substantial revisions and adaptations of some computing models

- Experiments have made an effort to reduce the raw event size (and the size of all subsequent derived formats) and event processing time. These efforts have mitigated the serious challenge that pile-up represents.
- The experiments have made compromises in their data distribution policies, reducing the number of copies stored in Tier 1 or Tier 2.
- Continuous optimization of the Monte Carlo simulation generators and of the data reprocessing times.
- Some experiments are introducing dynamical data placing policies. We applaud these initiatives and encourage all collaborations to follow suit.
- The potential proliferation of different data formats serving the same purposes has been checked and the role of different formats revised.
- The Tier 1 resources were underused in 2010. Experiments have been active in redistributing tasks among Tier 1 and Tier2.
- The experimental collaborations have implemented aggressive data cleaning policies.

Recommendations

We list below some recommendations on general aspects of the WLCG functioning, the usage made by the experiments during 2010 and other general aspects of the scrutiny. Specific questions related to the experiments' requests are deferred to the separate scrutinies of the different experiments.

- The WLCG accounting of Tier 2 resources is improving steadily but it is still insufficient: the Installed CPU compared to the pledged and installed disk capacity at the Tier 2 centres is not centrally accounted so far. It would be useful to disentangle the efficiency of organized/chaotic activities.

- Care should be taken that the worldwide LCG resources are used as much as possible as there may be a tendency by collaborations to place heavier demands on CERN resources or suggest that a larger than originally planned part of their analysis should be done at CERN. In 2010 Tier 1 resources were still underused.

- The CRSG encourages close collaboration of the different Tier centers with the experiments to continue the implementation of intelligent storage management policies to allow efficient and cost-effective access to data.

- The implications for best-use of resources of the interplay between improvements in network bandwidth and dynamical data placement policies should be evaluated.

- Experiments should carefully quantify the impact of pile-up and incorporate these values uniformly and realistically. We expect a careful evaluation of this impact by the October C-RRB meeting.

- We recommend the ALICE collaboration to submit more detailed reports to the CRSG. The CRSG would also like to see a better time granularity in the requests of the ATLAS collaboration. .

- On efficiency grounds the CRSG recommends sharing CERN resources when allocations are not fully used, perhaps using low priority queues.

- The CRSG recommends revising the assumed Tier 2 efficiency up to 2/3 (currently is 60%). This represents a saving of 10% in Tier 2 CPU to funding agencies.

- The CRSG does not see how the substantial increases of the data taking rate can be accommodated with the existing computing resources. To do so would require further, substantial modifications to the computing models we were presented with

On the scrutiny process

The CRSG is generally satisfied with the quality and quantity of the information provided by the experimental collaborations. The experiments' reports are well documented.

For future reviews we insist that the following good practices be maintained:

- All changes to the models compared to the previous review should be documented

- We request specific information from the experiments on the evolution of their computing models (see the recommendations section). This is to be provided ahead of the October 2011 C-RRB

- All documents should be provided sufficiently early to allow time for the review, a deadline for the revised requirements should be agreed upon well ahead of the final report deadline. For the upcoming October 2011 C-RRB meeting this deadline is **September 1st 2011**.

The CRSG is committed to making recommendations to the C-RRB with the aim of optimizing the resources already invested and adjusting the future provision to the computing needs.

The CRSG wishes to state that the recommendations contained in this scrutiny are to the best of our knowledge rigorous. They correspond to the real needs of the experiments for a given LHC live time assuming the validity of their computing model implementation. There is no contingency for late delivery or failure to meet the pledges included in our estimates or for less than 100% availability of these resources.

Any shortage of CERN resources implies potential disruption of data taking; therefore, we advocate full support of the recommended CERN resources. Resources at Tier1s and Tier2s are crucial for physics output; any shortage of non-CERN (non-custodial) resources does not impact data samples, but slows down the physics productivity of the experiments

On the CRSG membership

Since the last C-RRB Terry Schalk has replaced Albert Lazzarini as US representative and Giovanni Lamanna has replaced Cristinel Diaconu as the representative of the French funding agencies. We thank Diaconu and Lazzarini for their task within the CRSG. Another member D. Groep (The Netherlands) was reappointed by his funding agency for a further period of three years.

Harry Renshall has been replaced by Helge Meinhard from CERN as scientific secretary. It is my pleasure to thank Harry Renshall for his help and unfailing support to the activities of the CRSG.

After the new appointments and reappointments to the CRSG the mandate of the WLCG MoU concerning the renewal of the scrutiny group has been fulfilled. The process to renew or reappoint the remaining members should commence in about one year from now.

PART A

Scrutiny of the WLCG resources utilization in 2010

This report refers, unless otherwise stated, to the calendar year 2010, from January 1st to December 31st . The experiments are asked to report by **March 1st** on the usage made during the previous calendar year

This report has used the following sources:

1.- Cumulative accounting for Tier 1s and CERN

http://lcg.web.cern.ch/LCG/accounting/Tier1/2010/december-10/Master_accounting_summaries_December2010.pdf ,

2.- Month-by-month accounting of the CPU delivered by the Tier 2s

http://lcg.web.cern.ch/LCG/accounting/Tier-2/2010/

3.- The EGI accounting portal at CESGA

http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php

4.- WLCG accounting reports for non-GRID CPU

5.- 2010 pledges as presented to the C-RRB

6.- The Tier-1 and Tier-2 Usage Reports by Ian Fisk, presented to the LHCC review. The Tier-2 Usage Report is attached to this document.

General comments

In 2010 there was a substantial change in the WLCG usage in both qualitative and quantitative terms. While the 2009 run offered a limited amount of data used primarily for calibration and commissioning processes, 2010 has brought about sizeable amounts of real data valid for physics analysis. As a consequence the experiments have made extensive and intensive use of the WLCG resources.

Generally speaking the experiments' computing models and the WLCG have demonstrated their capability to record, distribute and analyze the rapidly increasing amounts of data. The luminosity increased in 2010 by five orders of magnitude. This rapid increase will be followed by yet another order of magnitude in the coming run. The main problem faced by experiments has not been the sheer amount of data but the different policy concerning the number and separation of bunches that had to be adopted. Protons had to be grouped in fewer but more intense bunches leading to considerable pile-up; a difficulty that will increase in the coming run.

The computing models have also demonstrated their adaptability to the rapidly evolving running conditions. Whenever necessary the experiments have modified some of the computing model parameters to adapt to the required needs. The CRSG generally speaking welcomes these changes that go in the direction of making the models more sustainable in the medium and long run, even if they represent some challenges for the computing management and have implications for other resources such as bandwidth.

These changes to the computing models are scrutinized in the appropriate section of this document.

The performance of the WLCG throughout the year has been generally regular and without any noticeable difficulties even if the rate of events delivered has increased by several orders of magnitude The Grid fabric works well; data distribution and network performance are excellent, much better than could perhaps have been envisaged some time ago. A similar comment applies to the current status of the middleware.

Early in 2010, resources exceeded the experiments' needs and the experimental collaborations had substantial headroom that they employed to increase simulation production, making reprocessing passes more often and making more copies in Tier 1 and Tier 2 to increase accessibility and to ease the access of enthusiastic physicists to the long awaited real data. In fact, the high usage of Tier 2s is partly due to increasing demand by individuals or groups. By comparison, Tier 1s are still somewhat underused at this stage but the experiments have taken steps to redistribute tasks and this situation is evolving rather quickly. Towards the end of the run and in the first weeks of 2011 the situation changed and the already approved resources for 2011 will match the needs only after readjustment of the computing model parameters. These readjustments have been made necessary by the increase of pile-up.

Both ATLAS and CMS make a commensurate usage of CERN and Tier 1 CPU resources. While ATLAS uses a lot more disk than CMS, the trend is reversed for tape. CERN CPU usage is now above 30 % of the nominal capacity (note that since sizeable parts of this capacity are used 100% during data acquisition, a 100% usage on average is impossible).

ALICE usage of CERN resources compared to external resources has increased substantially with respect to the previous report. LHCb relies disproportionately on CERN resources too.

We noted in the October 2010 scrutiny that the Tier 1 CPU resources were underused. ATLAS and CMS are reversing this by doing more Monte Carlo simulation at the Tier 1 and the percentage of use has gone from 50% to 62 %. There is no fundamental reason why this percentage could not be much higher. In view of the investment made by the Funding Agencies we urge the collaborations to move a larger fraction of their activities to the Tier 1.

The following table describes the degree of usage of the different resources.

April 2011

Resource	Site(s)	Used/Pledged Period average	Used/Pledged End of period
CPU	CERN	**32 %**	**---**
	T1	**62 %**	**---**
	T2	**122 %**	**---**
Disk	CERN	**75 %**	**110 %**
	T1	**89 %**	**110 %**
	T2	Not available	Not available
Tape	CERN	**52 %**	**69 %**
	T1	**52 %**	**60 %**

For comparison we reproduce the analogous table presented in the October 2010 C-RRB.

October 2010

Resource	Site(s)	Used/Available
CPU	CERN	**22 %**
	T1	**50%**
	T2	**92 %**
Disk	CERN	**71 %**
	T1	**92%**
	T2	Not available
Tape	CERN	**52 %**
	T1	**49 %**

The CPU figures correspond to a time average over the year 2010, obtained from averaging the monthly figures; those for disk or tape are usage relative to the installed capacity at the end of the accounting period.

The percentage usage of CPU at the Tier 2 is very high. This number refers to the HS06*hours used/pledged rather than to HS06*hours used/installed. In 2010 CPU usage at the Tier 2s had actually amply surpassed the pledge thanks to opportunistic use of other resources and/or more resources than pledged being installed. The efficiency is also much higher than expected (see below)

The time profile of the installed disk capacity at the Tier 2s is not centrally reported and it is not available. The experimental collaborations have tools to know how much disk is available at a given time in a Tier 2, but no central statistics are kept. Generally speaking the accounting tools for the Tier 2s are still not satisfactory and clearly more work on this is recommended.

Efficiencies

The computing TDR estimates the efficiency to be 85% for CPU and 70% for disk in the case of organized (group driven) analysis, reducing to 60% in the case of chaotic (user-driven analysis).

For simplicity it is assumed in the calculation of resources that organized activities are carried out only at CERN and Tier1's while chaotic analysis is carried out only at Tier2's. (Actually CERN is a combination of T0/T1/T2/T3 varying in relative percentage from experiment to experiment). We note that the current efficiency at Tier1s is 80%, close to the reference value, with little spread among the different experiments.

However the assumed efficiency is manifestly incorrect for the Tier 2. Even though a large fraction of the Tier 2 is already user/chaotic analysis (an exception is LHCb where only Monte Carlo production is carried out at the Tier 2), the statistics yield efficiencies that are much higher than 60%; the average is 78%.

In view of these figures we strongly support a revision of the official figure of 60% for Tier 2 CPU efficiency to 66% (2/3)

Efficiency of the utilization of the CPU at Tier 2s per experiment in the whole of 2010 (left column) compared to October 2010 (right column)

ALICE	**73%**	74%
ATLAS	**85%**	84%
CMS	**66%**	62%
LHCb	**88%**	92%

Disk usage

While the interpretation of CPU usage is straightforward, disk usage is more subtle to analyse. A metric based exclusively on disk occupancy does not account for how frequency of access or how efficiently disks are managed.

We are aware of the technical difficulties involved and of the timescale needed to define and implement such a metric. In the April 2010 C-RRB we required the experimental collaborations to provide the disk utilization in terms of the various data types involved (e.g. RAW, RECO, AOD, derived data, group data, user data and so on) and how frequently they were changed/replaced on disk in order to be able to assess the efficiency of disk usage. So far only ATLAS is providing partial information on this.

From the available information it is claimed that the 70% efficiency assumed in the computing models is a realistic estimate but the CRSG has been unable to verify this claim.

Sharing of the WLCG resources

The following tables give an idea of the use by the different experiments of the disk and CPU made available to them through the WLCG. The percentages refer to the fraction of the total mass storage, disk and CPU used per experiment (therefore all columns add up to 100% up to rounding errors). On the first (CERN+Tier 1) table the last column indicates which fraction of the total CPU that a given collaboration has used has been at CERN rather than using the T1's (and, consequently, does not add up to 100%). For comparison the percentages of October 2010 are shown in a separate table.

Percentage of use of the resources by experiment in 2010 (CERN+Tier 1s)

Collaboration	% of tape in T1+CERN used at end of period	% of disk in T1+CERN used at end of period	% of CPU in T1+CERN used	% of which at CERN
ALICE	8 %	6 %	16 %	33 %
ATLAS	35 %	57 %	59 %	14 %
CMS	52 %	29 %	17 %	20 %
LHCb	5 %	9 %	9 %	46%

Percentage of use of the resources by October 2010 (CERN+Tier 1s)

Collaboration	% of tape in T1+CERN used at end of period	% of disk in T1+CERN used at end of period	% of CPU in T1+CERN used	% of which at CERN
ALICE	5 %	10 %	15 %	26 %
ATLAS	35 %	52 %	59 %	10 %
CMS	55 %	30 %	17 %	22 %
LHCb	5 %	7 %	9 %	42%

In spite of the many-fold increase in luminosity the percentages of usage have stayed relatively constant over 2010. The differences between the different placing and analysis strategies that were visible in the October 2010 report (including data from January to August) have consolidated after the luminosity ramp-up. Note the large use of CPU and disk at CERN and T1s by ATLAS in 2010 (close to 60% in both cases). In contrast, ATLAS has reduced their percentage of CPU consumption at CERN with respect to the total as did CMS and both seem to converge to a 15%-20% bracket.

LHCb has done 46% of their computing at CERN, increasing the already large fraction of reported in October. The CRSG believes that this increasingly large fraction pre-empts to some extent the investment made by the Funding Agencies to secure Grid resources for LHCb. We urge the collaboration to reverse this trend.

The following table contains analogous information, but referred to the CPU usage in the Tier 2s. We see that ATLAS accounts for more than half of the CPU time used. The figures appear to be completely stable.

Percentage of use of the resources by experiment in 2010 (Tier 2s)

Collaboration	% of CPU in T2 used (All 2010)	% of CPU in T2 used (Oct 2010)
ALICE	**7%**	7 %
ATLAS	**59%**	59 %
CMS	**30%**	29 %
LHCb	**4%**	4 %

Delivered versus pledged

The overall level of fulfilment of the pledges can be seen from the following table.

Resource	Site(s)	Available / pledged
CPU	CERN	**129 %**
	T1	**109 %**
	T2	**122 %***
Disk	CERN	**112 %**
	T1	**100 %**
	T2	Not available
Tape	CERN	**100 %**
	T1	**101 %**

The figures refer in all cases to the end of the reporting period. These percentages are very satisfactory. However there are two local anomalies that have not had any real impact on physics:

- NL-LHC/T1 delivered only 79% of the pledged disk
- ASGC delivered 71% of the pledged disk and 82% of the pledged tape only..

The large turnout in CPU at the Tier 2 indicates that the percentage installed is actually above 100%, however automated accounting of this is still not in place. The figure indicated (*) is the delivered versus pledged.

Detailed graphics on the Tier 2 usage are attached at the end of this report.

PART B

Usage by the individual experimental collaborations

Below we provide some tables, ordered alphabetically by experimental collaboration, that give a rather accurate picture of the experiments usage of the resources. Every table is followed by some comments. The data has been provided by the collaborations themselves and cross-checked whenever possible.

ALICE

The following table shows ALICE resource usage for January to December 2010. Pledge values are from the April 2010 RRB. Data for CERN and Tier 1 is from the WLCG Master Accounting Summary for December 2010. CPU data for Tier 2 is from the EGI Accounting Portal (www3.egee.cesga.es). CPU values are total normalized wall time in kHS06 days, divided by the number of days in the year, with numbers in parentheses being the minimum and maximum monthly usage. We do not have straightforward access to Tier 2 disk usage data.

Resource	**Site(s)**	**Pledge**	**Usage**	**Used/Pledge %**
CPU/kHS06	T0+CAF	46.8	25.8 (6.9, 55.2)	55%
	T1	45.6	37.5 (9.1, 73.3)	82%
	T2	52.6	32.3 (5.9, 63.6)	61%
Disk/TB	T0+CAF	5500	2582	47%
	T1	6122	1958	32%
	T2	4326	Not available	--
Tape/TB	T0+CAF	6300	3014	48%
	T1	8485	1229	14%

Comments on the ALICE usage report

The full-year average CPU capacities are below the pledges, but the variability is large and the maxima exceed the pledges. If ALICE can successfully move jobs between tiers, we encourage the experiment to try to balance their computing load across the year and reduce the variability and hence their requirements. Full-year CPU efficiencies (measured from the ratio of CPU time to wall time) are low at the T0+CAF but better at the T1s and T2s, though still behind the best-performing experiments

As noted in the main report, computing resources generally exceeded the experiments' requirements early in 2010. While noting ALICE's low 2010 usage, it is worth considering that CPU usage over the first two months of 2011 is 68.1 kHS06 for T0+CAF, 52.2 kHS06 for T1 and 49.4 kHS06 for T2. Likewise, storage was below pledges in 2010, but is now growing faster. Disk and tape storage at T0+CAF had both grown to 3.1 PB at the end of February 2011, while disk storage at T1s had grown to 2.3 PB. Of course these increases in resource usage in early 2011 are a clear reflection of the presence of the 2010 heavy ion data and so are likely to be a better reflection of ALICE's usage patterns than those early in 2010.

ATLAS

Resource	**Site(s)**	**Pledged [1]**	**Used [2]**	**Used/ Pledged**	**Average CPU efficiency**
CPU (kHS06)	T0+CAF	67	36	54% (36%)	
	T1	211	171	81% (55%)	81% (83%)
	T2	215	183	85% (81%)	84% (84%)
Disk (TB)	T0+CAF	3900	3900	100% (81%)	-
	T1	22018	22000	100% (82%)	-
	T2	21238	12000	56% (57%)	-
Tape (TB)	T0+CAF	8900	9500	107% (72%)	-
	T1	15372	6800	44% (51%)	-

The figures for CPU usage correspond to a time average over the period January-December 2010, those referring to disk or tape reflect the amount of the resource installed at the end of 2010. Quantities in brackets refer to the ones reported in the October 2010 C-RRB.

[1] pledged resources from April 2010 RRB (http://lcg.web.cern.ch/lcg/Resources/WLCGResources-2009-2010_12APR10.pdf)

[2] storage information from WLCG accounting summaries (http://lcg.web.cern.ch/lcg/accounting/Tier1/2010/december-10/Master_accounting_summaries_December2010.pdf) and ATLAS. There is a slight (<10%) discrepancy between the two sets of numbers: ATLAS reports less occupancy than LCG. CPU usage from the EGI accounting portal.

Comments on the ATLAS usage report

ATLAS has accumulated 3.6 Ms corresponding to 1 B events. Analysis activity in the first half of 2010 focused more on detector commissioning, changing rapidly to more physics oriented analysis in the second half of the year. The performance of the ATLAS computing system was generally smooth throughout the whole period; more intensive use is observed, corresponding to simulation production and reprocessing campaigns and to user activities just before conferences.

The reconstruction time per event was more than 50% larger than in the request. This was due to using a lower pt threshold for tracking and other modifications in the commissioning process.

The collaboration simulated 1.4 B events, more than foreseen in the approved request. This was possible thanks to sufficient CPU resources being available due to the smaller simulation time for 2010 events and because the need for large CPU resources arrived later in the year. There was an improvement of about 15% in the simulation time.

A reduction of 20% in the ESD size was achieved. The 2010 RAW event size was also below what was foreseen in the resource request.

At CERN ATLAS used less time than required due in part to a lower requirement for calibration and alignment thanks to the machine and detector stability achieved. Part of the available time was used as Grid resource. The disk and tape are in line with the requests with a slight overuse of tape.

At Tier 1s disk was fully used. This arises from the opportunistic use for temporary placement of additional copies and from the fact that the volume of simulated data is much larger. Group and user data volumes are less than expected on the whole. Because at present the resources at the CAF at CERN are sufficient Tier 1s have not performed these tasks and the corresponding disk allocations have not been used at Tier 1s. When combined with the less group usage, this has reflected in a slight underuse of CPU at Tier 1s. The tape resources were also clearly overestimated.

Concerning Tier 2 usage the main relevant fact is the substantial overcommitment of disk.

CMS

Resource	**Site(s)**	**Pledged [1]**	**Used [2]**	**Used/ Pledged**	**Average CPU efficiency**
CPU (kHS06)	T0+CAF	96.6	21	22% (17%)	
	T1	103.5	46	44% (32%)	77% (71%)
	T2	196	152	78% (76%)	64% (62%)
Disk (TB)	T0+CAF	4100	4396	107% (60%)	-
	T1	12183	9223	76% (69%)	-
	T2	13627	10000	74% (55%)	-
Tape (TB)	T0+CAF	14600	9728	67% (49%)	-
	T1	23677	17676	75% (65%)	-

The figures for CPU usage correspond to a time average over the period January-December 2010, those referring to disk or tape reflect the amount of the resource installed at the end of 2010.

Quantities in brackets refer to the resource usage reported in the October 2010 C-RRB.

[1] pledged resources from April 2010 RRB (http://lcg.web.cern.ch/lcg/Resources/WLCGResources-2009-2010_12APR10.pdf)

[2] storage information from WLCG accounting summaries (http://lcg.web.cern.ch/lcg/accounting/Tier1/2010/december-10/Master_accounting_summaries_December2010.pdf) and CMS. CPU usage from the EGI accounting portal http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php

Comments on the CMS usage report

The performance of the CMS computing system was generally smooth throughout the whole period. The number of events collected was lower than expected due to the LHC turnout. With the expected overlap in the assignment into primary datasets CMS expected to collect 2.2 B events. At the end, in 2010 CMS collected 1.5 B collision events.

CMS observed that in 2010 the rate of Monte Carlo production was much larger than expected in 2010. CMS produced more than 3 B simulated events in 2010.

Concerning CERN usage the CPU utilization was low on average. As the luminosity increased in the second half of 2010 the peak load on the Tier 0 consumed all the available resources but this does not show on averages over multiple days. CPU utilization is largest over the HI running. Disk at the Tier 0 + CAF is fully subscribed.

At the Tier 1 the CPU utilisation peaks at 60% at the end of the year. Even in this case the Tier 1 usage is below expectations and CMS is changing the operational model to push for more Monte Carlo production at the Tier 1. An analysis of CPU efficiency as a function of time shows a dip in summer 2010, which was related to an IO problem identified in the software.

The volume of data on disk at the Tier 1 uses more than 75% of the total pledge. A similar fraction of tape with respect to the pledges is used.

CMS reached 100% usage of the Tier 2 CPU early in the 2010 run and it remained high throughout the year with a mixture of simulation and analysis. There was an increase in the simulation in the second half of the year as the physics needs for specific analysis are better understood. CMS is satisfied with the use of the Tier 2 for analysis during 2010; with some seasonal dips the subscription is very high. The total number of individual submitters per month is larger than 800.

The transition to AOD for analysis was expected to take place in 2010 but it has not progressed as quickly as CMS hoped as RECO was needed for detector commissioning and there was no resource scarcity incentive to move to other datasets. The total number of jobs submitted to AOD is 10% but growing.

LHCb

The usage of the computing resources by LHCb displays a rather healthy situation. The experiment is successful in using the resources and adapting to the unexpected and somewhat difficult situation. It should be noted that the recent physics output of the LHCb collaboration is also an indication that the computing is in good shape and is able so far to process and analyze efficiently the collected data.

Site	CPU [kHS06] April (Oct)	CPU [kHS06] used	Disk (TB) April (Oct)	Disk (TB) used	Tape (TB) (Oct)	Tape (TB) used
CERN	21 (23)	8.4 (12)	1220 (1220)	922	1500 (1800)	844
Tier-1	41 (44)	21.4 (27)	2870 (3220)	2096	2800 (2400)	903
Tier-2	36 (38)	26.5 (34)	20 (20)	n/a	0	0

Comparison between the requests made for April 2010 (approved by the RRB) and October 2010 generated in 2009 and the yearly average CPU Power (from WLCG accounting). In brackets (CPU usage) the variation throughout the year 2010 in monthly averages. CPU numbers from WLCG accounting (April 10 to Jan 11). Disk usage from SLS - tape usage from LFC.

* 70 Tb front-end (buffer) disk space is subtracted from each T1 disk request in order to compare the usage to the effective storage area.

Comments on the LHCb usage report

The table shows the quite extrem peak load for CERN and T1 resources - the T2 resource usage in contrast is expected to be more constant compared to the observation. LHCb verifies the usage numbers by recalculating the same numbers based on known resource usage for specific computing tasks and the CPU-hours gathered by the DIRAC system. The resulting normalization factor (HS06 per CPU/core) is in line with the numbers measured by various centers for their average factor applied to their CPU resources.

The detector is in good shape and the data collection efficiency is now around 90% throughout the active data taking period. The RAW data distribution from T0 to T1s proceeds according to the plan: 155 Tb of data were shared among the 6 T1s outside CERN.

The LHCb report now present numbers for failed jobs - consuming CPU power, usually without valuable results. We appreciate this information as it also demonstrate the efforts spent to further increase the resource utilization.

The running conditions at LHC with high intensity per bunch lead to an average collisions per bunch up to 1.5-2 (with peaks up to 2.7) while the nominal conditions assume 0.4. As a result, significant increases in the event size and in the CPU needs to reprocess each event were observed. The event size increased from 80 kB to 140kB. The reconstruction algorithms were tuned to cope with the high occupancy. The CPU time scales linearly with the average number of collisions per bunch crossing μ, doubling when μ increases from 0.4 (nominal) to 2 as reached during 2010. In addition, looser cuts during the stripping were applied as expected for the first studies using real data.

These effects (event size, CPU, loose stripping) lead to high disk occupancy and to a need to decrease the number of copies in T1s in order to save resources. The average replication factor was in fact around 6 instead of 7. The replication will be further reduced to 3 and some older reprocessing (Reco4) will be removed in order to free resources for the incoming data. Finally the replication factor for the T1 data ends at two.

A campaign of processing of data and Monte Carlo production with best parameter sets deduced from data starts at the end of 2010 to be ready for spring 2011 conferences.

The CPU resources usage is not continuous and the installed power is used on (frequent) processing bursts. The data flow is obviously dependent on machine delivery. This corresponds to a variety of Monte Carlo production which are expected to be available in a short time. The data is used to tune the simulation and this process logically lead to successive processings of Monte Carlo samples In addition, a modulation of the analysis flux is also naturally expected. As a result, the difference between the average CPU usage and the peak usage is a factor of ~3 in T1s. It is likely that rearrangements will occur as the experiment progress through data taking, but the CPU peak power is expected to continue be an important parameter for LHCb.

The usage of CPU is distributed according to pledges for the regular productions (reconstruction, stripping) and reaches about 45% for CERN. The analysis jobs are in turn distributed according to the site availability and reaches about 45% at CERN as well. According to LHCb experts, this is a sign of the availability of the CERN resources when compared with other T1s.

In conclusion, the LHCb computing model is robust and was able to incorporate parameter adjustments beyond the anticipated variations. The usage of the resources is satisfactory, given the running conditions. The used CPU power is different in average from the pledges but incorporates the necessary contingency for burst of data and MC production. As already noted by CRSG, LHCb would benefit from a mutualisation of resources at CERN.The usage of tapes is significantly lower than expected mainly due to less running time in 2010. The disk occupancy is relatively high and is a matter of concern. LHCb experts took correct and timely actions (clean-up older sets, reduction in replication, stripping optimization) to cope with a possible limitation are are confident that the data taking, the MC production and the reprocessing foreseen by the end of the year will be covered within the available resources. Fragmentation of disk resources has been identified as one main reason while not all resources can be used. In addition, the two methods to generate these numbers (SLS, LFC) resulting in discrepancies being under investigations.

Since the data collection rate is continuously increasing throughout 2010 and the needs for 2011 and beyond are finalized, it is likely that the disk space in the T1s and the corresponding data distribution strategy will require some modifications.

PART C

Scrutiny of the requests for 2012 and 2013 (preliminary)

ALICE

ALICE provided us with a spreadsheet which calculated the numbers used in their requirements document. This was extremely helpful. The requirements were calculated based on a month-by-month plan for LHC running. The referees interacted with ALICE by email.

The collaboration's requirements are based on collecting 10^9 pp events and 7.4×10^7 heavy-ion events in each of (calendar years) 2011 and 2012. ALICE has now tested, for the first time, its computing model with real heavy-ion data and this has led to increases in CPU and storage requirements albeit with correspondingly reduced uncertainties. Of course, ALICE has had this new data only for 3 to 4 months, compared to the 15 to 18 months' experience with pp data.

Data sizes and storage

The event data sizes used in the requirements calculations are averages based on data recorded and processed in 2010, but increased to account for the addition of new elements to the detector during the 2010-11 shutdown.

- The raw event size for Pb-Pb Monte Carlo is reduced by a factor 5 from the 300MB used last year. This is because ESDs only are stored for most MC events with raw MC data for "exceptional requirements" only.
- The Pb-Pb ESD for Monte Carlo has increased in size from 6.05 MB to 50.0 MB. This is because most MC production is now for the highest-multiplicity central collisions with various embedded signals, increasing the size compared to minimum bias events that had been used up to now.
- Last year ALICE recorded extra information in the ESDs from real data, but said they would not need to do so in the longer term. Now their real data ESDs are at, or below, 10% of the raw event sizes, which was their computing model target.

The capacity of the CASTOR buffer at T0 is assumed to double to 4.3 PB from 2012 onwards.

CPU requirements

CPU requirements have increased both because of additional tasks and extra processing per event:

- There is an additional reconstruction pass on a fraction of the raw data for calibration (mainly central tracking); Pass 0 processes 20% of the data prior to Pass 1 reconstruction
- Additional post-reconstruction calibration is performed at the beginning of the analysis trains
- CPU power per event is increased by:
    - A new secondary vertex finder algorithm
    - (for MC only) additional details in digitization of data from drift detectors (does not scale with number of tracks)
    - New detector elements (EMCAL and TRD modules)
    - Enrichment by the trigger mix of data with high track-multiplicity events (central collisions for Pb-Pb and high multiplicity for pp)

The CPU requirements for event processing have been updated based on experience with 2010 data, the addition of new detector elements and an increase in collision energy. For Pb-Pb only minimum bias events were recorded in 2010, so the requirement for ion reconstruction has been further increased in anticipation of a data sample enriched in central collisions. The increases are significant: CPU for pp event reconstruction has increased by 65% from the 2010 requirement and has *doubled* for Pb-Pb reconstruction. These increases, especially that for ion reconstruction, are reflected in substantial increases in the T0 CPU request compared to October 2010.

Processing requirements for MC production have increased over the 2010 values by a factor of 5 for pp events and by 39% for Pb-Pb events. The pp requirement per event is less than one percent of that for heavy ions, so that ion MC still dominates despite the larger number of pp events.

The CAF CPU requirement for 2012 and 2013 is not calculated directly from processing tasks. ALICE notes that their current CAF resource is saturated and they double this to obtain the quoted requirement of 31.1 kHS06.

Comments and recommendations

- Reducing heavy-ion Monte Carlo needs. For 2011 T2 CPU usage, ALICE account for doing Pb-Pb MC simulations of only 1/3 of the number of heavy-ion data events, but for 2012 and 2013 their requirement returns to 1 MC event per real event. We expect that they will in practice have to scale back in 2012 and 2013 and encourage the collaboration to contemplate this, appreciating that this will delay physics output. For illustration, reducing the Pb-Pb MC simulation fraction to 80% in 2012 and 2013 would reduce the T2 CPU requirement by about 24 kHS06 and would decrease disk usage at each of T0, T1 and T2 by around 1.2 PB bringing them much closer to the pledges they already have on the books for each of these years. We note further that in the ALICE requirements, the 8-fold increase in the size of Pb-Pb MC ESDs leads to a more than 4-fold increase in the disk space needed per event, which could be mitigated if fewer copies of ESDs on disk were assumed; we encourage the collaboration to minimize the number of copies if possible.
-
- The ALICE computing model assumes 3 reconstruction passes. Reducing this to 2 would make substantial savings in CPU and storage requirements and an obvious recommendation is to encourage all efforts to reduce the (effective) number of reconstructions.

- ALICE's usage of CPU resources in 2010, as measured by average capacity used in each month, was highly variable. We encourage the collaboration to balance computing load across the year where possible to reduce this variability and thence their peak CPU capacity requirements.

ATLAS

The ATLAS requests were shown in a LHCC Referees' meeting held at CERN on March 21st. The written documents were provided to the CRSG some days afterwards with a delay of more than three weeks with respect to the March 1st deadline. The requests were scrutinized in parallel by the combined team of ATLAS and CMS referees. A set of questions were asked and subsequent clarifications were provided by email. The level of detail provided by ATLAS has been very good; the detailed spreadsheets were clear enough so that the referees did not feel necessary to cross-check them with independent calculations.

The major driver in the ATLAS computing requests for 2012 is the increased pile-up foreseen in the operation of the LHC. The resulting, non-negligible impact on basic parameters of the ATLAS computing model and computing requests has been mitigated by a redefinition of the ATLAS data placement strategy.

In the past year, ATLAS has addressed some of the criticalities of their computing model which had a substantial impact on resources and long term sustainability. Progress has been made in the areas of data size, processing times, data format proliferation and distribution. We encourage ATLAS to continue in this way.

However, there is a non-negligible impact on the requests due to the data taking scenario and to the number of pileup-events foreseen during 2012, which was not thoroughly addressed. In summary, the ATLAS requests for 2012 are of the same order of magnitude as the requests for 2013, shown last year in the hypothesis of an LHC shutdown during 2012.
If the conditions assumed by ATLAS are verified, then their requests are justified.

On general terms, we encourage ATLAS

- to look at more centralized ways to perform skimming and slimming activities, in order to limit their requirements on disk space and CPU power, and
- to keep their CERN CPU request in 2013 equal to 2012, in order to alleviate the resource increase at external Tier1s.

We also note that the disk usage at Tier2s is somewhat lower than the disk space allocated to ATLAS. In other terms, the ATLAS Tier2 disk request has not been supported by actual usage so far.

The ATLAS Computing Model

The model ATLAS is using today and its input parameters differ from the one scrutinized in 2010 in a number of ways. Due to pile-up, higher values for events sizes and reconstruction times than in past years are assumed, as shown in the following table.

	Previous years	End-2010	2011-2012
RAW (MB/event)	1.6	1.4	2.8
ESD (MB/event)	0.8	1.6	2.2
AOD (MB/event)	0.15	0.18	0.36
Sim-RAW (MB)	2	2	2
Sim-ESD (MB)	1.1	1.5	3.0
Sim-AOD (MB)	0.18	0.21	0.42
Reco time (data) (HS06-s/event)	80	100	200
Reco time (MC) (HS06-s/event)	135	170	340
Simulation time (MC) (HS06-s/event)	6000	4100	5100

In order to accommodate larger events event sizes arising from pileup, a new, reduced data distribution plan has been adopted for 2011 and subsequent years. Emphasis is given moving analyses to further derived data samples and a dynamic data placement mechanism, already implemented during 2010, which profits of the available bandwidth between computing centers in order to reduce the number of replicas at Tier2 sites and utilize the available CPU resources more efficiently.

Dynamic data placement has been demonstrated and reduces the total amount of expensive disk resources at the Tier-2. While we welcome this development, we nevertheless require that the impact on network bandwidth in future years be thoroughly assessed and the resulting Tier2 CPU utilization be reported accordingly.

ATLAS made progress in order to reduce some of the input parameters to their computing model; the required CPU power to reconstruct an event is now reduced by 20%, the power needed for a full Geant4 simulation by 15% and the ESD event size is reduced by 30%. However, the impact on the requests due to these improvements is “neutralized” by the subsequent increase assumed, due to pileup.

Comments and recommendations

The tables below show the resource estimations for 2011.

ATLAS provided a detailed spreadsheet of their computing needs which allowed us to reconstruct closely enough the rationale behind their requests, the detailed breakdown in terms of activities and data formats. Their time granularity of one year means that their required disk resources are the ones needed at the end of the 2012 RRB period (i.e. March 2013), while their disk resources should be taken as peak values. We encourage ATLAS to increase the granularity of their computing needs, in order to efficiently optimize them. It’s a tool for them in order to optimize their resources.

We asked ATLAS about the actual values of pile-up assumed and the dependence of data size and processing times with respect to pile-up. We quote the answer given by ATLAS below:

“Pile up effects was an estimated average: RAW, AOD, ESD and dESD all double in size. Reconstruction takes twice as much CPU. The numbers are the ones that appear in Table 1 of the resource request document. We have done estimates of pileup effects that justify these numbers, but we are actively working on improving them and are thus reluctant to publish the numbers at this time. We think the average factor of 2 is a conservative estimate.

Our initial estimates, before any mitigation, corresponded to approximately a factor 2.3 increase in resource requirements, both for disk and event size, for mu=10. After the mitigation work done to date, we have now achieved a reduction to only a factor two increase with mu=15. The variation between mu=10 and mu=15 is estimated to be roughly an increase in CPU and size of a factor 1.5. There is no a priori reason why CPU and event size should scale in the same way, but we find that they do so approximately in this range of mu values. We do not expect further substantial improvements in resource requirements, but we are still working on further reductions which may be at the 10-20% level, but are difficult to predict. During the first couple of weeks of data-taking at 7 TeV this year the average pile-up has been of ~ 6 interactions per crossings and up to 13 per crossing.”

We further asked ATLAS whether their input parameters were consistent with what is being seen in the current LHC operations, we got a detailed explanation by ATLAS which can be summarized as follows:

“From the 2011 data so far and after the recent extensive optimization campaign, it appears that we can just meet the reconstruction time and AOD event size estimates of the request, provided the pileup does not regularly exceed mu=15. It is likely that there will be some reduction in the RAW event size per event, although in terms of overall resources the impact of the latter will be small.”

Generally speaking, we feel that the resources are well justified. We require that the impact due to pile-up be quantitatively assessed, based on real data taking in 2011.

The CERN resources are reasonable and should be supported; as noted above, any shortage of resources at CERN would reflect in potential disruption of data, whereas shortage at Tier1s and Tier2s slows down the physics productivity of the experiments.

The recommended resources are shown in the following table. A detailed breakdown is given in the following paragraphs.

CPU [kHS06]	**2012**	**2013**
CERN	***73***	***73***
Tier-1	***259***	***273***
Tier-2	***295***	***315***
Disk [PB]		
CERN	***9***	***10***
Tier-1	***27***	***30***
Tier-2	***47***	***53***
Tape [PB]		
CERN	***18***	***18***
Tier-1	***36***	***40***

Tape resources

As the computing model advocates a removal of the ESD from disk, we asked ATLAS what impact on tape drive occupancy this decision would generate. We quote ATLAS's answer below:

"We have conducted a very extensive review of the usage of ESD, spanning detector and trigger systems, performance groups, and physics analysis over the last months. There is a consensus that full bulk ESD samples are not needed in 2011/12. Some ESD will be available: highly selected events in dESD samples, which will be used mainly for performance and analysis work, and "small stream" ESD samples which will be used primarily for detector and performance work. In addition there is the rolling buffer of ESD, also for unforeseen uses which arise soon after the data are taken. A lot of work has gone into careful definition of these samples, and this will continue in future. If limited ESD are needed for small special events, it is always possible to regenerate it. We will retire the 2010 ESD only when we need the space later in this year, so users can still use the legacy ESD.

Tape usage and throughput go down because of our ESD policy, but the larger event size drive up the requirements."

Tier-0 Tape (PB)	***2012***	***2013***
Total	***18***	***18***

Tier-1 Tape (PB)	***2012***	***2013***
LHC RAW data	9	9
Real ESD+AOD+DPD data	7	8
Simulated data	14	17
Group + User	2	2
Cosmics and other data	4	4
Total	***36***	***40***

Disk resources

CERN disk has fixed buffers from 2010 to 2013. These numbers, based on ATLAS's running experience, are kept fixed, even though event sizes have doubled. They think there is some uncertainty in what the optimal value for each component is, but they should all fit in the total, fixed buffer space of just under 1PB There is a surprisingly large disk request for the ATLAS Tier 2, driven by a large increase in real data from 13PB in 2011 to 21PB in 2012. ATLAS comments as follows:

"The large part of the increase is simply due to the increased data volumes, primarily AOD. A smaller increment comes from the more gradually increasing volumes of simulated data. The 8 copies of the AOD distributed over the Tier-2scomes to a total of 14 PB (12 PB latest reconstruction version and 3 PB previous, rounding to the nearest PB), and the DESD samples (4 copies, latest processing only) correspond to 6 PB.The increment between 2010 and 2011 is smaller essentially because we have reduced the numbers of replicas between the two years, as we explain in the request document."

We think it would be helpful for them and for future scrutiny to understand the performance of the new dynamic data placement. We therefore encourage ATLAS to implement a disk utilization metric, which can be used to investigate possible optimization and mitigation strategies.

Tier-0 Disk (PB)	***2012***	***2013***
Buffer for RAW and processed data	0.65	0.65
Buffers for merging	0.10	0.10
Tape buffer	0.10	0.10
Total	***0.85***	***0.85***

CAF Disk (PB)	***2012***	***2013***
Calibration and alignment	0.7	0.7
Derived detector data	4.1	4.1
Derived simulated data	2.3	3.0
Group data (grid aware)	1.0	1.0
Group data (non grid aware)		
User data (grid aware)	0.4	0.4
User data (non grid aware)		
CAF Total	***8.5***	***9.2***

Tier-1 Disk (PB)	***2012***	***2013***
Current RAW data	4.3	4.3
Real ESD+AOD+DPD data	5.4	5.4
Sim. RAW+ESD+AOD+DPD data	5.6	6.8
Calibration and alignment outputs	0.4	0.4
Group data	6.5	7.6
User data (scratch)	0.6	0.6
Cosmics	0.2	0.2
Processing and I/O buffers	4.3	4.4
Total	***27***	***30***

Tier-2 Disk (PB)	***2012***	***2013***
Current RAW data	0.0	0.0
Real AOD+DPD data	21	21
Simulated RAW+ESD+AOD+DPD data	16	21
Calibration and alignment output	0.3	0.3
Group data	7	8
User data (scratch)	2	3
Processing buffers	1	1
Total	***47***	***53***

The Tier1 disk requirement is dominated by the storage of real data, as well as simulated data. In both cases, most of the disk space is needed for RAW, AOD and derived data, contrary to current usage, dominated by ESD.

With the current ATLAS analysis model, activities performed at the level of physics/detector groups result in large requests for space to support skimming and slimming. Indeed, the disk space needed to store data skimmed and slimmed by physics/detector groups will give the highest single contribution to Tier1 disk in 2012 and 2013. We encourage ATLAS to study the possibility of implementing these activities in a more centralized way, and to use real experience with data to minimize the disk space reserved for group data by fully exploiting derived physics data formats.

Non-negligible disk space is needed for processing and I/O buffers (about 0.3PB per Tier1)

The disk space at Tier2 is also dominated by real and simulated AOD+DPD data. The same comments made about group space at Tier1 apply to Tier2 as well.

We welcome that, with respect to previous years, ATLAS now takes into account a dynamic data placement mechanism in order to increase the CPU utilization and reduce the number of pre-placed copies at Tier2s by a factor 2, thus alleviating disk requests. Unfortunately, this reduction is fully eaten up by the increased data size due to running conditions. We recommend ATLAS to fully exploit their networking capabilities in order to mitigate their disk requests at Tier2s.

CPU resources

The Tier0 requests are firm and should be granted. The CPU request for servers might be reduced by profiting of virtualization. A somewhat high CPU request at CAF for 2013, where activities at CERN should be reduced due to the LHC shutdown, was explained by ATLAS as follows:

“During 2013 we will largely enable the CERN CPU resources for Grid jobs (this is a capability that is already in place for times when they are idle).We thus expect a mix of uses as typical on the Grid: reprocessing, group and user analysis, and simulation. The normal roles of T0 and CAF are not strictly needed in a shut down year, but since these resources will not disappear, we will use them for critical physics analysis needs not covered by normal user T2 activity during the year.”

We encourage ATLAS to use also the Tier0 for Tier1-like activities in 2013, for instance simulation production, thus alleviating the resource needed at the other tiers. The associated, modest, CPU requests are indicated as “other activities” in the tables below. We correspondingly increase the ATLAS CPU requests at CERN and decrease the resources required for simulation production at Tier1s and Tier2s.

Tier-0 CPU (kHS06)	***2012***	***2013***
Full reconstruction	28	28
Partial processing and validation	2	2
Merging and monitoring	2	2
Automatic calibration	5	0
Other activities	0	5
Servers	1	1
Total	***38***	***38***

CAF CPU (kHS06)	***2012***	***2013***
Partial, reconstruction, debugging and monitoring	4	0
Non-automatic calibrations	4	0
Group activities	12	12
User activities	4	4
Other activities	0	8
Servers	12	12
Total	***35***	***27***

Tier-1 CPU (kHS06)	***2012***	***2013***
Re-processing	71	74
Simulation production	77	70
Simulation reconstruction	39	52
Group (+user) activities	72	76
Total	***259***	***273***

Tier-2 CPU (kHS06)	***2012***	***2013***
Simulation production	56	50
Group activities	58	58
User activities	180	207
Total	***295***	***315***

The Tier1 CPU requests for 2012 and 2013 are dominated by Monte Carlo production and user/group activities. In 2010 and 2011, the contribution due to group and user activities corresponded to reasonable assumptions on the number of physics groups, on how often they pass on the entire datasets and on the CPU power required to process a single event.

Last year we noted that, contrary to CMS, data skimming and slimming are not done centrally and in parallel to data reprocessings, but at the group level and more frequently. The corresponding resources are therefore higher by a factor of about 10. We encourage ATLAS to explore the possibility of and eventually implement skimming and slimming in a centralized way.

Since resources needed to reprocess real data and, to a lesser extent, slimming and skimming by physics groups, are the major element defining the maximum CPU needs at Tier1s, we asked ATLAS how the impact on requests could be mitigated by relaxing the assumed time window. ATLAS answers as follows:

"The reprocessing represents 20% of our total Tier-1 CPU request in 2011(41/202) and 27% in 2012. Unlike for the other experiments, the CPU requirement for reprocessing is thus not the overriding factor for the Tier-1 CPU total. In fact we model that the reprocessing is more spread rather than concentrated in specific months. We model that the final reprocessing of the year, in the last quarter, takes place over 80 of the 90 days."

We note that the average CPU time required to fully simulate an event with GEANT4 in 2010 has been 4100HS06 seconds, whereas the value used in the requests is 5100. ATLAS comments that

"The 4100 figure was the measured CPU requirement for the G4 simulation step for the event samples generated in 2010 - it is a mean over all the different types of events generated. Especially in the early part of the year 2010, we generated less complex events (minimum bias...) which had a substantially lower CPU requirement. In2011 and 2012 we expect that the typical events simulated are more active(higher-pt) and thus the CPU requirement increases. We note that in both cases the CPU requirement is lower than in our 2010 request (6000 units per event) due to the continuing optimisation programme for our full simulation.

The pileup is overlaid in the digitisation and reconstruction step for simulated events: this is accounted for in the "Sim recon" entry in table 1. (This is the main reason why this figure is higher for simulated events than for data)."

The Tier2 CPU requirements are mostly due to user analysis activities and, to a lesser extent, to group activities and simulation production. The group and user activities have been computed by using reasonable assumptions on processing times, number of passes and time required for a single pass on the datasets.

CMS

The CMS requests were submitted within the March 1st deadline, and also shown in the LHCC Referees' meeting held at CERN on March 21st. CMS presented two sets of requests, referring respectively to trigger rates of 300Hz and 400Hz. The latter was motivated by the increased resulting physics performance of CMS. We understand that the LHCC has discussed this matter.

"The LHCC encourages CMS and ATLAS to explore ways to record the data needed to maintain high sensitivity to new physics at low thresholds as the luminosity increases. Such an initiative is supported to fulfil the physics goals of the 2011-12 run. In case the available computing resources are restricted, the committee recommends delaying the processing, analysis or simulation, as necessary"

Therefore we examined the requests corresponding to the 300Hz trigger rate

We also share the following LHCC recommendation, which also applies to the other LHC experiments.

"In the longer term, all experiments are encouraged to develop more sustainable computing models, with the aim of limiting the further increase in resources required at CERN and at Tier-1 and Tier-2 sites, in light of the restricted funding available."

The requests were scrutinized in parallel by the combined team of ATLAS and CMS referees. Two sets of questions were asked and subsequent clarifications were provided by email. The level of detail provided by CMS has been very good; the detailed spreadsheets were clear enough so that the referees did not feel necessary to cross-check them with independent calculations. The required computing resources are calculated with the time granularity of one month. It is therefore possible to follow closely the time evolution and easily determine possible optimizations in their computing model.

The major driver in the CMS computing requests for 2012 is the increased pile-up foreseen in the operation of the LHC. The resulting, non-negligible impact on basic parameters of the CMS computing model and computing requests has been mitigated by a redefinition of the CMS data placement strategy.

The majority of the requests is firm and justified. However, when computing their requests, CMS used an LHC running time of 5.7×10^6 s instead the "standard" 5.2×10^6 s. We asked CMS to use the latter and rescale their requests accordingly. CMS provided an update resource request document and the relevant spreadsheets on April 5th.

The CMS computing model

The model CMS is using today and its input parameters differs from the one scrutinized in 2010 in a number of ways. Due to pile-up, higher values for events sizes and reconstruction times than the ones seen at the end of 2012 are assumed, as shown in the following table.

	Previous years	End-2010	2012
Pile-up scenario		4	18
RAW (MB/event)	1-1.5	0.31	0.82
RECO (MB/event)	0.5	0.39	0.94
AOD (MB/event)	0.2	0.17	0.34
Sim-RAW (MB)	1.5	1.5	1.5
Sim-RECO (MB)	0.6	0.49	1.14
Sim-AOD (MB)	0.25	0.22	0.39
Reco time (data) (HS06-s/event)	100-125	28	140
Reco time (MC) (HS06-s/event)	100-125	65	212
Simulation time (MC) (HS06-s/event)	700-900	500	500

CMS assumes explicitly an LHC running scenario resulting in mean values of interactions per collisions ranging from 4 at the end of 2010 to 18 in 2012, according the the month of data taking. For 2012, these values are always set to 18, which is about twice as what they observe in the current running conditions. CMS has performed detailed studies on the impact due to pileup on their computing model. In 2012, they expect basically to double their event sizes and increase even further their processing times with respect to the end of 2010.

The sensitivity of the computing requests to pileup was commented by CMS as follows:

"We are reasonably sensitive to pile-up. The reco size and time track essentially linearly with pile-up. If the pile-up at the end of the year was 20% higher we will struggle with Tier-0 resources and have longer re-reconstruction passes. We have attempted to choose a realistic but conservative pile-up prediction for 2011."

CMS plans to alleviate the increased need for computing resources due to pileup by reducing the number of preplaced AOD copies on Tier1 disks (from 7 to 2) and by implementing a reasonably aggressive deletion scheme for derived data. The latter implies that analysts stay reasonably current with respect to reconstruction version. This might be a risk for analyses which, due to a number of reasons, might use a previous version of data processing.

Generally speaking, a number of optimizations in the CMS computing model were already performed in the past, so their space for optimizations is tight. Their comment on possible mitigations to their requests was

"A number of mitigation strategies were deployed before the data distribution model changes. A series of campaigns to reduce the AOD size were conducted. Substantial improvements were achieved from the first size estimates. RECO has also been examined with smaller improvements. A series of studies was also done for reconstruction time. CMS examined increasing the tracking Pt cut as well as other early cuts. These yield speed improvements, but also cut the sensitivity."

In case of resource shortages

"The response depends largely on the type of shortage. A Tier-0 shortage could be responded with a cut in the trigger rate, or migration to the Tier-1s for a fraction of the prompt reconstruction or a time delay of a portion of the prompt reco. This only works if there are sufficient tape and Tier-1 resources. A shortage of Tier-1 CPU would result in an extension of the processing times, and hard choices about what data and MC could be reprocessed. Reductions in TIer-1 and TIer-2 would result in a reduced simulation and analysis capacity, which slows down analysis. CMS already expects in 2011 to have to enforce priority choices in the resources allocated for analysis. "

We finally asked CMS whether their input parameters were consistent with what is being seen in the current LHC operations. CMS gave us a detailed explanation which can be summarized as follows:

"We have been monitoring the RAW and RECO event sizes and the prompt reconstruction time. RAW and RECO sizes are lower than our estimates by 10-20%, which is reasonably close. We will make additional measurements after the scrubbing runs.

For CPU reconstruction time, we looked at the time averaged over each primary dataset individually. The CPU efficiency is good, indicating that the infrastructure is working, but the time is about 40% higher than our estimates. Our model estimates are based on averaging between reconstruction time for complicated events like top and simple events like minimum bias. The fraction of simple events in the 2011 dataset is dropping, which is one of the effects. […] There is no evidence the Tier-0 or Tier-1 processing request is going to be over estimated."

Due to a material error, CMS presented their requests by using 5.7×10^6 s LHC live time in 2012. The CRSG asked them to use the agreed 5.2×10^6 s and rescale their requests accordingly. Resources depending mostly on trigger rates, e.g. Tier0 CPU and disk, do not change. Other resources are reduced by about 10%. We note the following comment by CMS:

"Tier-1 Disk should also drop by about 10%, but the Tier-1 disk has a model assumption that we are less comfortable with. If you look at the disk resources, in the charts they exceed the pledges by the end of each year. We had to assume the aggressive deletions before the new reprocessed data was available. This 10% would put the Tier-1 request at just in time in April, but we wouldn't exceed it."

Requests and recommendations

The tables below show the resource estimations for 2012 and 2013.

Generally speaking, we feel that the resources are well justified. The CERN resources are reasonable and should be supported; as noted above, any shortage of resources at CERN would reflect in potential disruption of data, whereas shortage at Tier1s and Tier2s slows down the physics productivity of the experiments.

The recommended resources are shown in the following table. A detailed breakdown is given in the following paragraphs. All tables show the resources after rescaling the LHC live time from 5.7 to 5.2 million seconds.

CPU [kHS06]	2012	2013
CERN	*121*	*119*
Tier-1	*145*	*145*
Tier-2	*350*	*350*
Disk [PB]		
CERN	*5*	*4*
Tier-1	*22*	*27*
Tier-2	*26*	*26*
Tape [PB]		
CERN (including HI)	*23*	*23*
Tier-1	*51*	*59*

Tape requests

We find that the tape requests at CERN and Tier1s are generally justified.

CERN Tape (PB)	*2012*	*2013*
Tier-0 RAW pp data	5.8	5.8
Tier-0 RECO pp data	9.5	9.5
Tier-0 alignment and calibration pp data	0.6	0.6
Tier0 HI data	3.2	3.2
CAF tape	4.6	4.6
Total	***23***	***23***

Tier-1 Tape (PB)	*2012*	*2013*
LHC RAW data	3.5	3.5
Real RECO+AOD data	11.4	12.9
Simulated data	33.7	40.3
Skimmed data	2.4	2.5
Total	***51***	***59***

Disk requests

The CMS disk request at Tier0 is reasonable and justified by the Tier0 being on the critical path.

For Tier1s, the disk space for raw MC data increases from 2011 to 2012 by 46%, whereas it increases by 95% for reco MC data. CMS explains this as follows:

"In the model 10% of the MC RAW and RECO is on disk with heavy reliance on staging. The RAW is only produced once so the increase in 2012 is the new MC from 2012 and some fraction from 2011 that is still active (not yet deleted). The RECO is generated more than once during the year. While the old copy is eventually cleaned up it also stays active for some time and 10% of it resides on disk. The faster increase for RECO is caused by more RECO events."

We recommend CMS to fully exploit their networking capabilities in order to mitigate their disk requests at Tier2s.

Tier-0 Disk (PB)	***2012***	***2013***
Buffer for RAW and processed data	0.3	0
Buffers for merging	0.2	0
Tape buffer	0.5	0
Total	***1***	***0***

CAF Disk (PB)	***2012***	***2013***
Express data	1.7	0
Prompt reco data	0.6	0
Simulation	0.3	0
Validation	0.2	0
CAF Tier2	0.6	0.9
Stager pool	0.6	0.6
Analysis space		2.5
CAF total	***4.0***	***4.0***

Tier-1 Disk (PB)	2012	2013
RAW data	2.1	2.1
Real RECO+AOD data	10.5	10.0
Simulated RAW+RECO+AOD data	6.0	9.1
Group data (skimming)	2.3	2.5
Processing and I/O buffers	2.2	2.7
Total	***23***	***26***

Tier-2 Dis (PB)	2012	2013
Real RECO data	1.0	1.0
Real AOD data	8.0	7.7
Simulated RECO+AOD data	12.4	13.0
User data	3.6	3.6
Processing buffers	1.0	1.0
Total	***26***	***26***

CPU requests

The Tier0 CPU resources needed for full reconstruction are straightforward to compute starting from the trigger rate end event processing time. CMS plans to increase this to 225 Hz, i. e. 75% of their trigger rate in 2010 and 2011

Since these are on the critical path, the resources required for these tasks should be warranted.

We asked CMS for a clarification on the CAF usage during 2013:

"We assumed that the CAF and Tier-0 processing and disk were available for Tier-2 analysis functionality. The CERN grid interface would allow the resources here to appear as any other TIer-2. By the end of 2011 we will have transitioned the CAF storage to a disk-only solution largely decoupled from Castor. When calculating the overall Tier-2 resources needed we used the existing analysis model and included contributions from the distributed Tier-2s and CERN. In the absence of CERN resources, the distributed Tier-2 request would have to be increased to compensate."

Differently to past years, CMS will attempt to use the Tier1 centers for simulated event production first, which will overflow into the Tier2s when the Tier1 resources are full. This will allow Tier2s to be heavily utilized with analysis.

A request for possible mitigation strategies for the CPU required at Tier1s by extending the time window allowed for a reprocessing step was answered as follows

"The time window is constrained on both ends. The winter conferences define the completion time and the start time is when we can realistically have the release and conditions for the winter conferences. In the model our goal is to process the data and simulation (updated conditions and realistic MC) for the winter conferences ,which was defined as being finished by January 15th for this year. We can reprocess the MC as soon as we have a release and understand the realistic conditions we are simulating for. In 2011 we used a start at the beginning of October for simulation work with data reprocessing concentrated later. In 2012 we used the beginning of September. As we stretch farther into the summer we are less likely to have completed all the preparation for a reprocessing. The other possible mitigation would be to further prioritize the samples to be reprocessed and rely on prompt reconstruction or less realistic simulation for some analysis. "

The Tier2 CPU requests are reasonable. Assumptions similar to the ones made for ATLAS lead to similar results.

Tier-0 CPU (kHS06)	***2012***	***2013***
Full reconstruction	61	0
Partial processing and validation	6	0
Merging and monitoring	3	0
Automatic calibration	1	0
Servers	11	0
Analysis	0	60
MC production	0	20
Total	***82***	***80***

CAF CPU (kHS06)	***2012***	***2013***
Processing activities	38	38
Interactive activities	1	1
Total	***39***	***39***

Tier-1 CPU (kHS06)	***2012***	***2013***
Processing	145	145
Total	***145***	***145***

Tier-2 CPU (kHS06)	***2012***	***2013***
Simulation production	100	100
Group and user activities	250	250
Total	***350***	***350***

LHCb

The LHCb computing resource estimates for 2012 based on an evolved model, the 2010 experience, the changed LHC parameters and running periods. The requested resource numbers are now based on a new model implementation with dedicated parameter files and a python implementation[1]. This allows for more detailed profiles, i.e. time dependent resource profile.

LHCb prepared one document - LHCb-PUB-2011-00 - (plus presentation for the LHCC review meeting) and one update to the main document.

The CRSG referees had a fruitful email exchange and would like to thank the LHCb colleagues for their collaborative attitude, timely and detailed responses.

Computing models

The simplified model (implemented as an excel spread-sheet, used in the previous scrutinises) together with the new model and its implementation (see 1) has been used to compare and evaluate the requested resources. The referees welcome the effort to create a new model implementation allowing for finer tuning and larger versatility. The old model implementation is only used to verify the model at the overlapping years (2010 and 2011). The model reproduces the global features of the LHCb computing model and is in agreement with the resources request. Applying the new base numbers for 2011 to the old model implementation still results in only 5% deviation to what the new model produced.

[1] The new simulation tool is a set of python modules steered by configuration files that describe the different activities included in the simulation. The code and configuration files used can be obtained at **https://coma.ecm.ub.es/svn/lhcbcompsim/tags/CRSG-2011-03-04**

LHCb running scenario

The assumed running time is $5.0x10^6$ seconds for 2011 and 2012 assuming an overall LHC duty cycle of 30%. According to the new LHC schedule, data taking is planned for 2011 and 2012, whereas 2013 LHC will produce no new data.

The running time estimates used by LHCb are based on the agreed machine running scenarios. They are summarized in Table 1 , with some adjacent assumption on availability and efficiency. Contrary to the other LHC experiments, LHCb will not take data during the heavy ion runs.

The period for 2011 is defined as the time between the start of LHC in 2011 and April 1st 2012 (when the 2012 pledges have to be made available) while the running time for 2011 is assumed to encompass the period from March 2011 to end October 2011.

There are two main effects producing the differences for resource estimates in the past. One is the luminosity increase achieved by in increase of pp interactions per bunch crossing (μ) and secondly the new 1kHz charm physics data rate. The new 'μ' results in 1.3 - 1.6 times increase of the event size for the different data types and an increase of CPU work (kHS06*s) per processed event of x2 up to x4.5.

LHCb request for 2012 and outlook for 2013

The request of computing resources is summarized in Table 2.

Table 2: LHCb resources requests for 2011 (update), 2012 and 2013 outlook (the number in parents showing the estimates from 2010 for that period)

Date	Site	kHS06 peak-power	Disk (PB)	Tape (PB)
2012 period	CERN	34 (21)	3.5 (1.7)	6.4 (3.0)
	Tier-1	113 (65)	9.5 (3.8)	6.2 (4.4)
	Tier-2	48 (36)	0	0
2013 period	CERN	33	4.0	7.7
	Tier-1	110	11.1	8.0
	Tier-2	48	0	0

The CPU power is the parameter promoted by LHCb to represent the peak CPU needs, as shown in Table 2. The available power is important for the reprocessing steps near the end of the data taking period and also for the MC production campaigns, that may focus in a short period of time resources beyond the normal usage. Nevertheless, an efficient usage of resources is important and in case of large latencies observed during the normal running conditions, changes in the functional model and in the interactions with the computing centers should be proposed. The integrated CPU power shown in Table 3 was also presented as an alternative, in order to compare with other experiments.

Table 3: Integrated CPU requests.

kHS06*year	2012	2013
CERN T0 + T1	24	22
Tier1s	80	74
Tier2s	48	48
Total	151	144

The distribution of the computing power among sites follows the pattern already presented last year. LHCb has already adapted the disk usage (i.e. reduced replicas, less space tokens) to increase the disk usage efficiency. The additional second tape copy for raw data at CERN slightly increase the tape resources at T0, whereas the T1 tape is reduced because of the replica changes.

The CERN resource usage has been reduced and following the new requests will be approaching the 25% level. If there are available resources anywhere, user analysis and simulation will be scheduled, regardless of the pledged resource distribution between CERN and the rest. This scheduling is driven by data and cpu availability at any site.

The strategy to use the on-line farm during the shutdown cannot be applied to 2011/2012 shutdown which is expected to be short. In addition the required effort to make use of the on-line farm resources are underestimated. So far, no human resources have been found to work on that issue. There is still an optimistic goal to use the on-line farm for the shifted one-year shutdown in 2013.

The number of reprocessing passes is assumed to be 2 in 2011 and 2012. These passes consider one partial reprocessing in the first half and one full reprocessing at the end of the year. The data processing will be accompanied by a stripping, however the stripping is in principle decoupled and can be performed as soon as a significant understanding of the data requires a re-made of the selected samples. The numbers of strippings are assumed to be 4 (max) in 2011 and 2012.

Starting 2011 all MC computing will take place at the T2s whereas user activities will be focused on the T1s.

It should be noted that LHCb propose some relaxed timing for the resource provisioning at T0 and T1 which could help sites to practically do all the resource ramp-ups.

Conclusions

The resources requests for 2012 are based on the realistic estimates already thoroughly discussed during the previous scrutiny and the additional charm physics channel.

The referees express their concerns about the usage of CERN based resources, even if available and idle. The Job scheduling (pilot scheduling controlled by LHCb) should always try to use all of the T1 resources before scheduling starts to CERN resources (especially for the user analysis and MC jobs). LHCb should have a 'plan B' if the supplied resources do not allow the peak power requested for certain times (i.e. on-line farm usage, split the full reprocessing, etc.).

Like in 2010 we would like to recommend an intermediate scrutiny by autumn 2011 in order to prepare the resources scrutiny round of 2012.

Experiment Tier-2 CPU Reports:
ALICE Reports

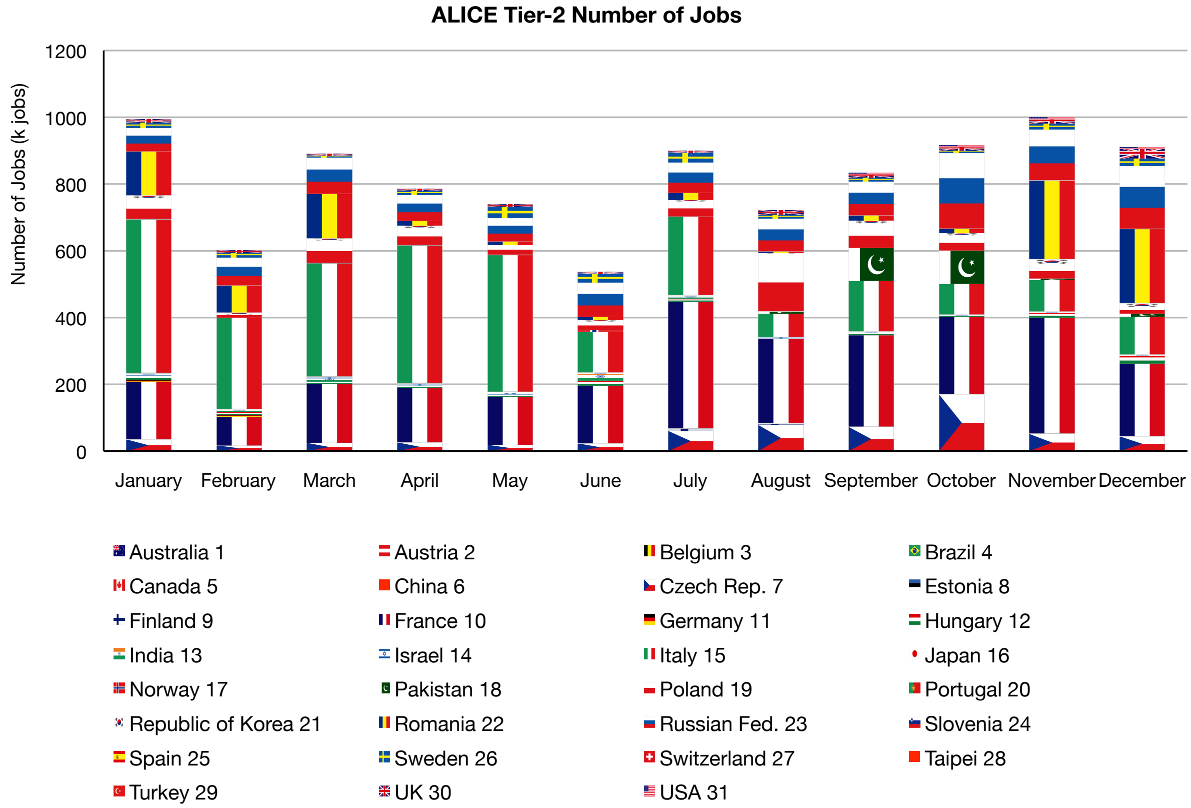

Figure 1: ALICE Jobs Executed on Tier-2 sites during 2010

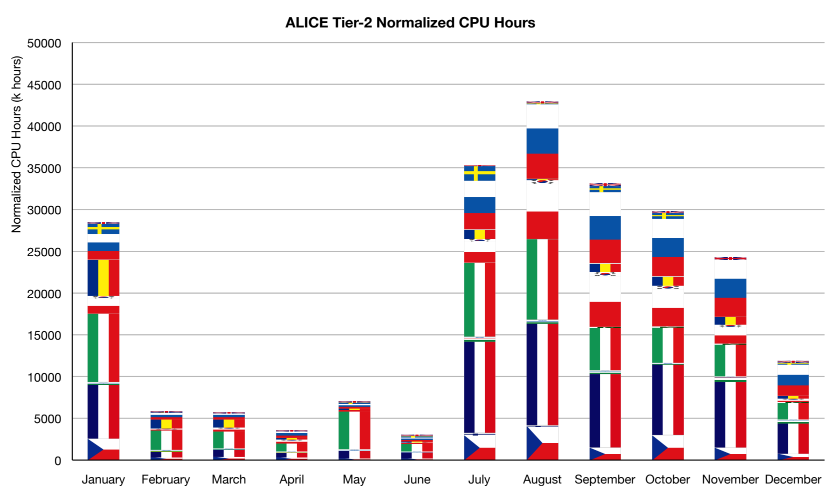

Figure 2: ALICE Normalized CPU Hours at Tier-2 sites in 2010

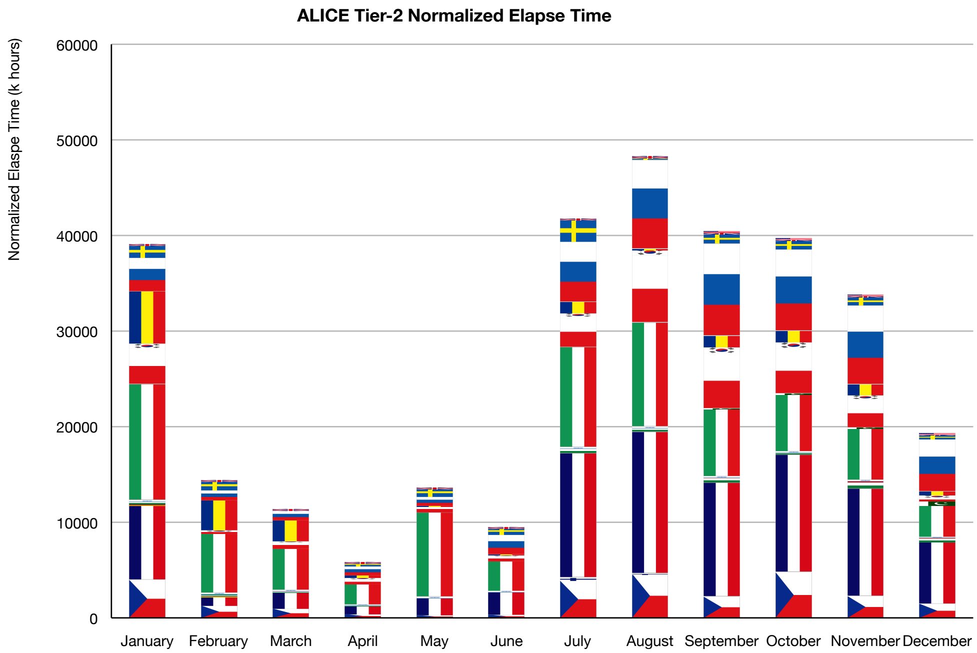

Figure 3: ALICE Normalized Wall Clock Time Used at Tier-2 sites in 2010

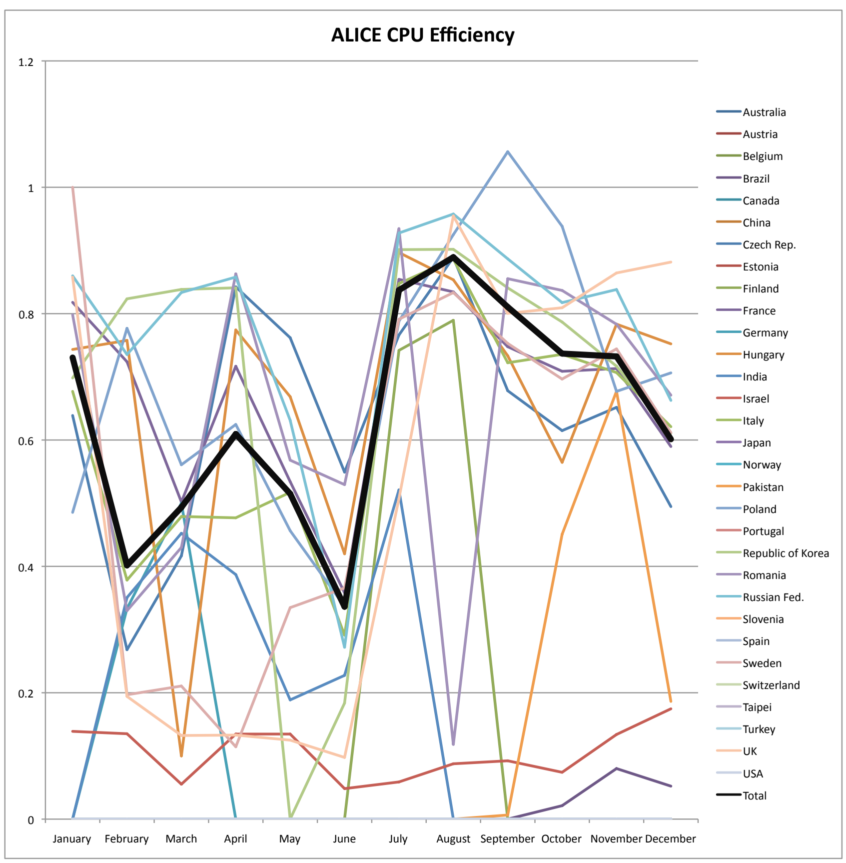

Figure 4: ALICE CPU Efficiency at Tier-2 sites

ATLAS Reports:

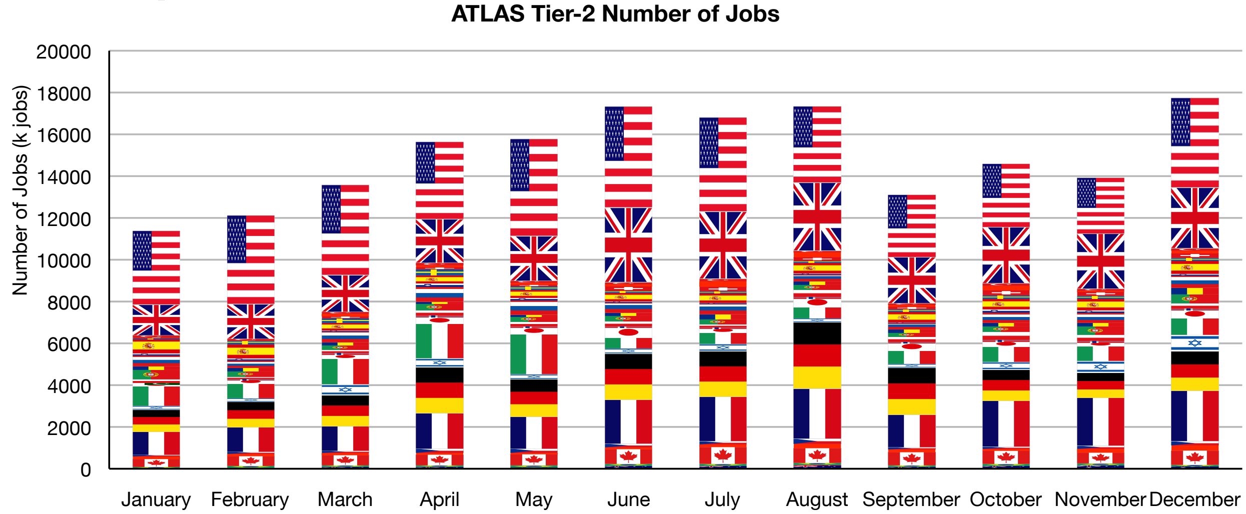

Figure 5: ATLAS Jobs Executed on Tier-2 sites in 2010

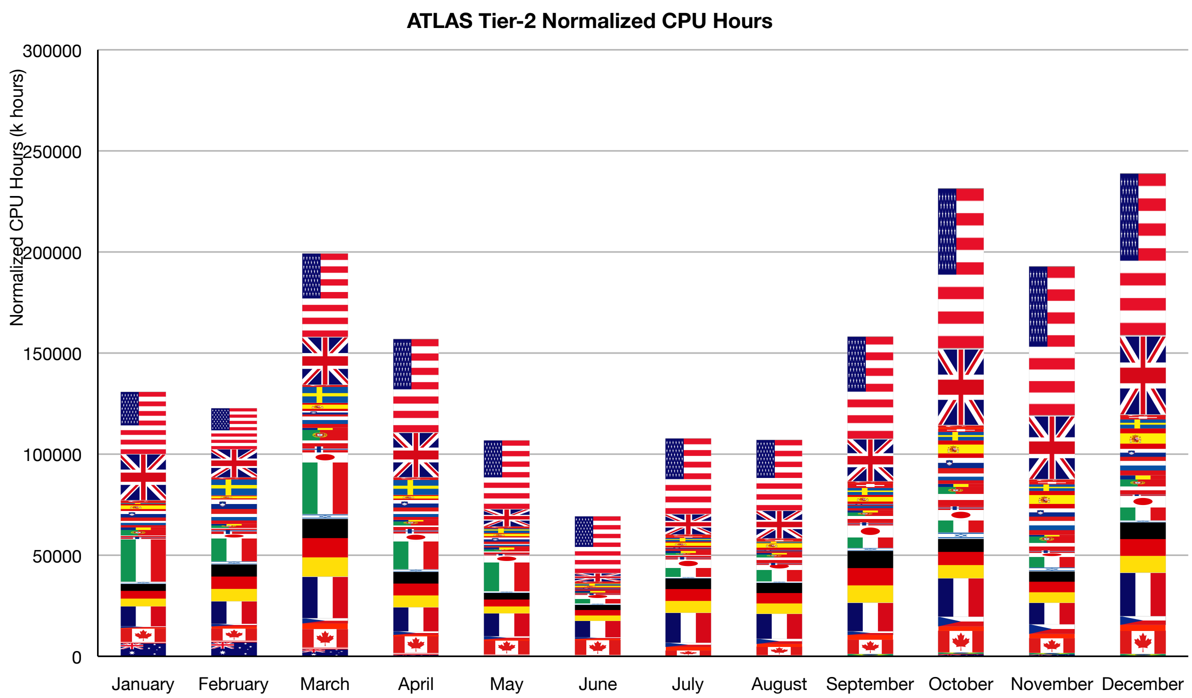

Figure 6: ATLAS CPU Time Used on Tier-2 sites in 2010

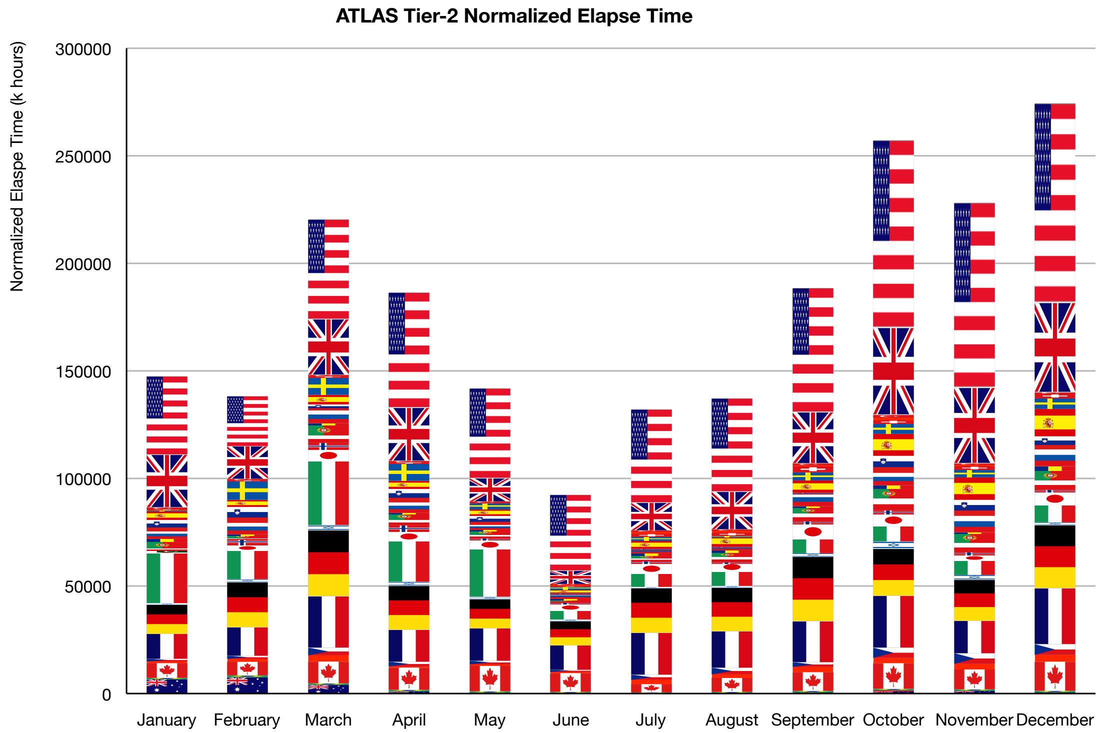

Figure 7: ATLAS Wall Clock Time Used on Tier-2 sites in 2010

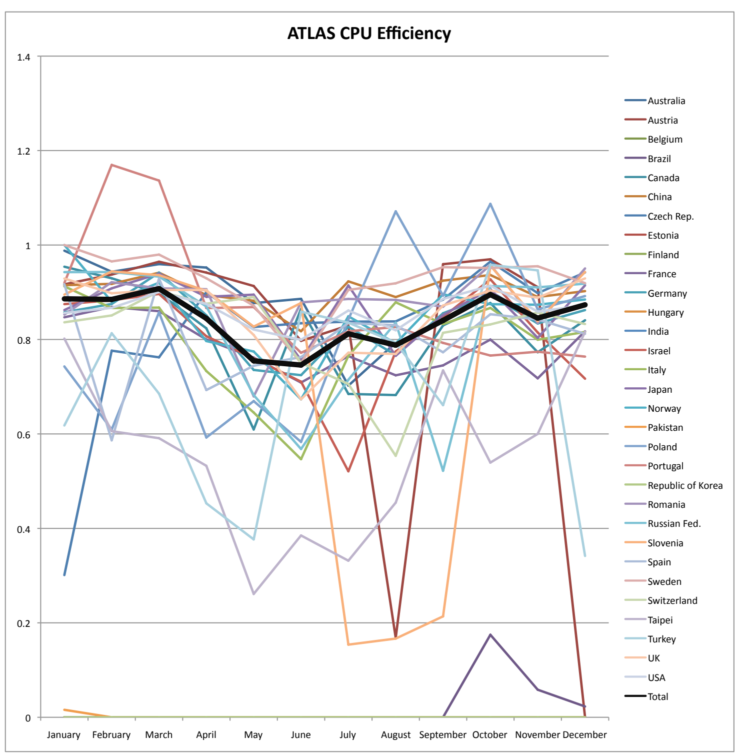

Figure 8: ATLAS CPU Efficiency at Tier-2 Sites in 2010

CMS Reports:

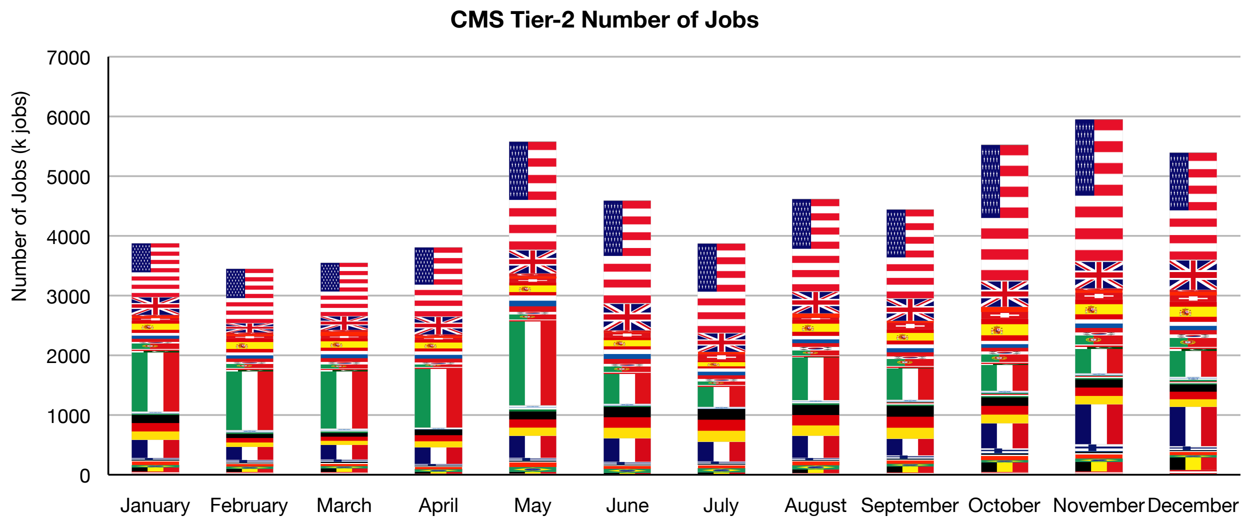

Figure 9: CMS Jobs Executed on Tier-2 sites in 2010

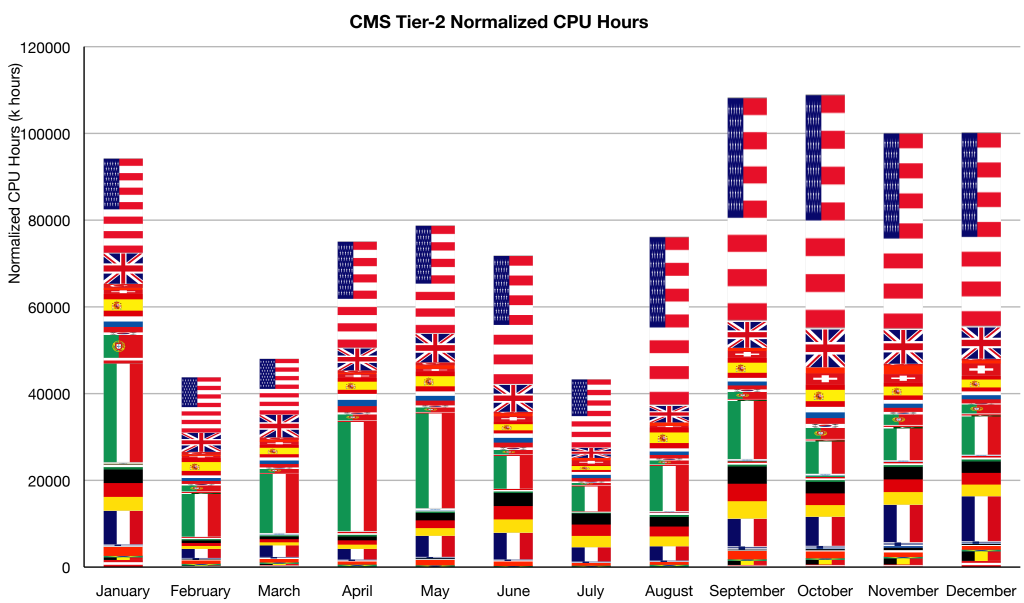

Figure 10: CMS CPU Time Used on Tier-2 sites in 2010

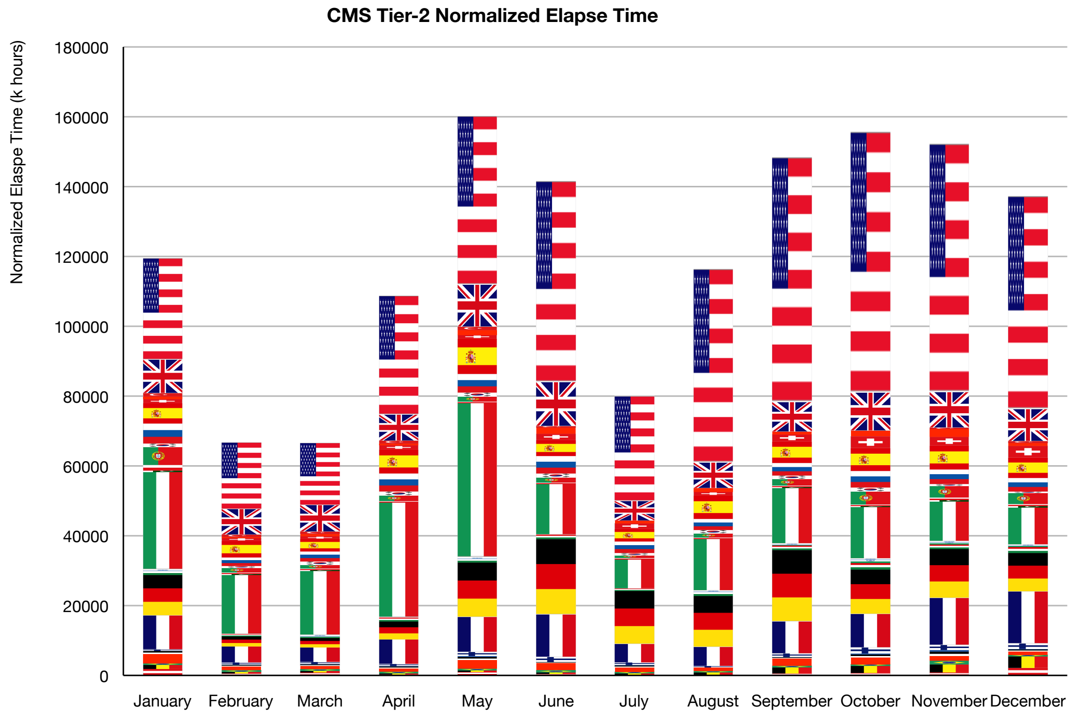

Figure 11: CMS Wall Clock Time Used on Tier-2 sites in 2010

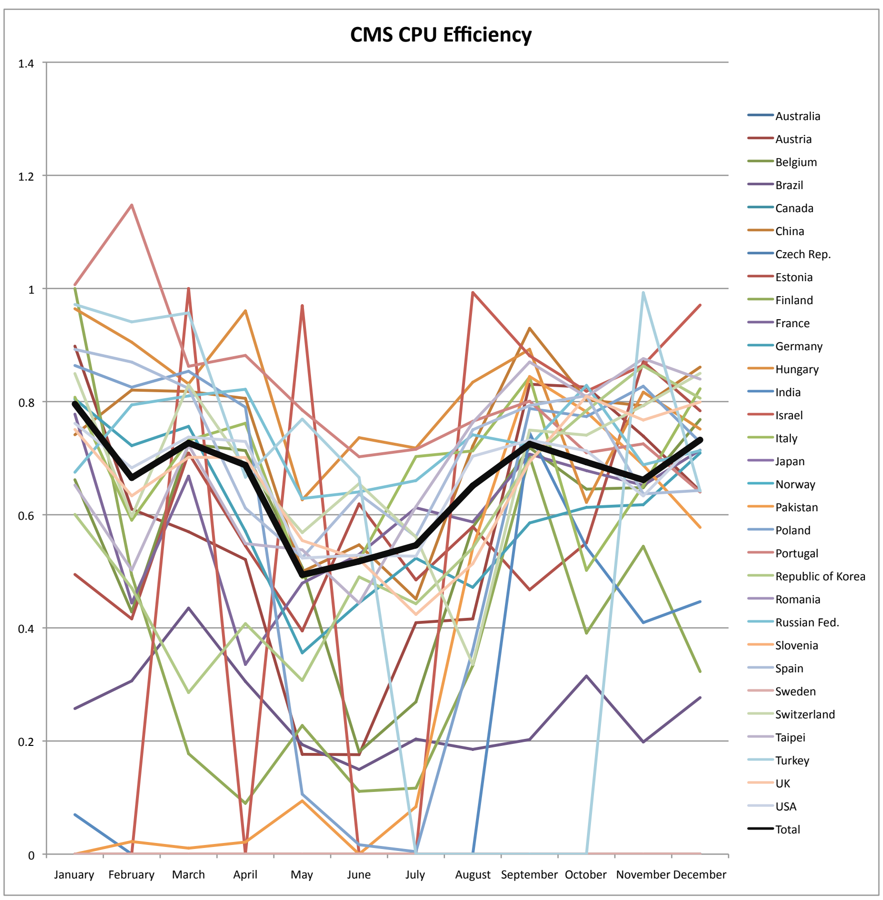

Figure 12: CMS CPU Efficiency at Tier-2 Sites in 2010

LHCb Reports

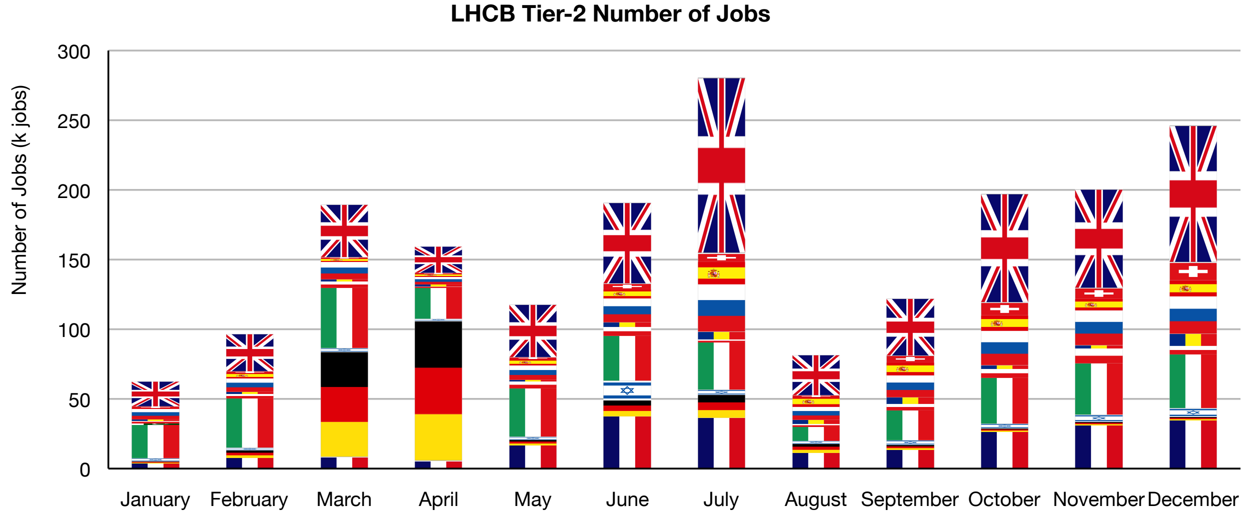

Figure 13: LHCb Jobs Executed on Tier-2 sites in 2010

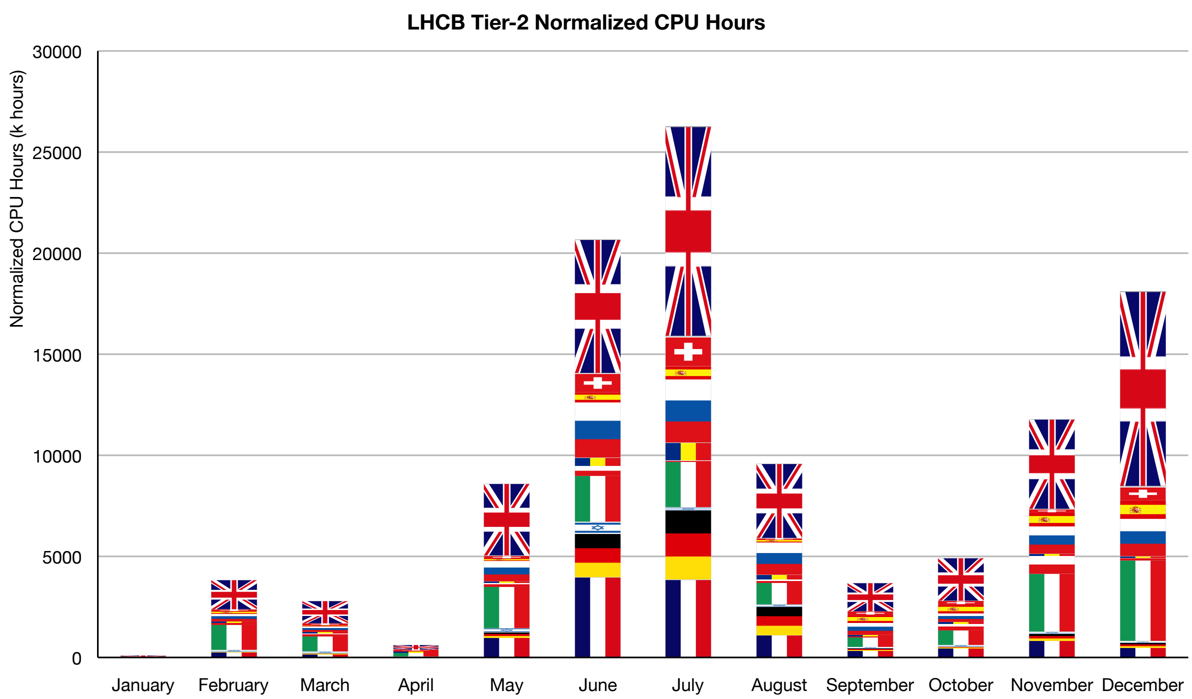

Figure 14: LHCb CPU Time Used on Tier-2 sites in 2010

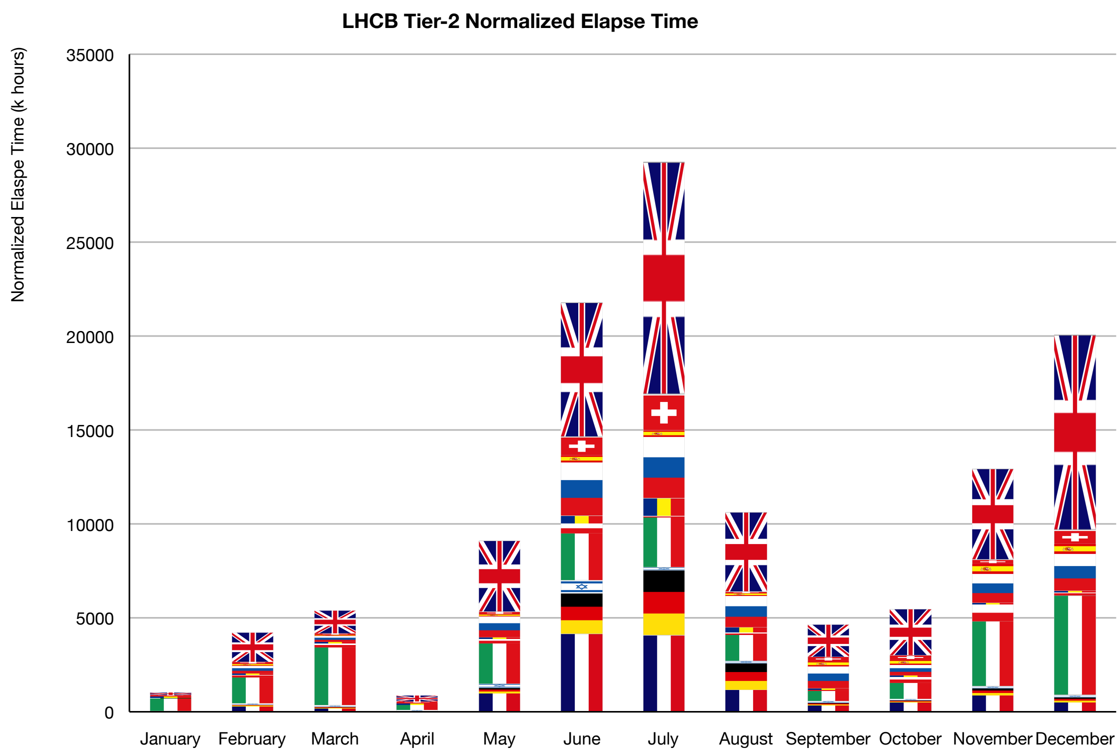

Figure 15: LHCb Wall Clock Time Used on Tier-2 sites in 2010

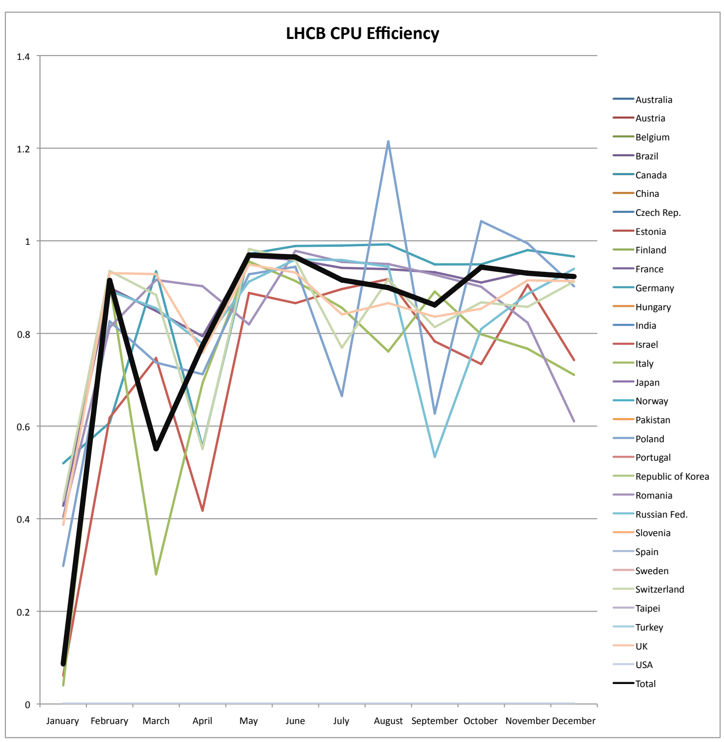

Figure 16: LHCb CPU Efficiency at Tier-2 Sites in 2010

Country Usage Reports for Tier-2s

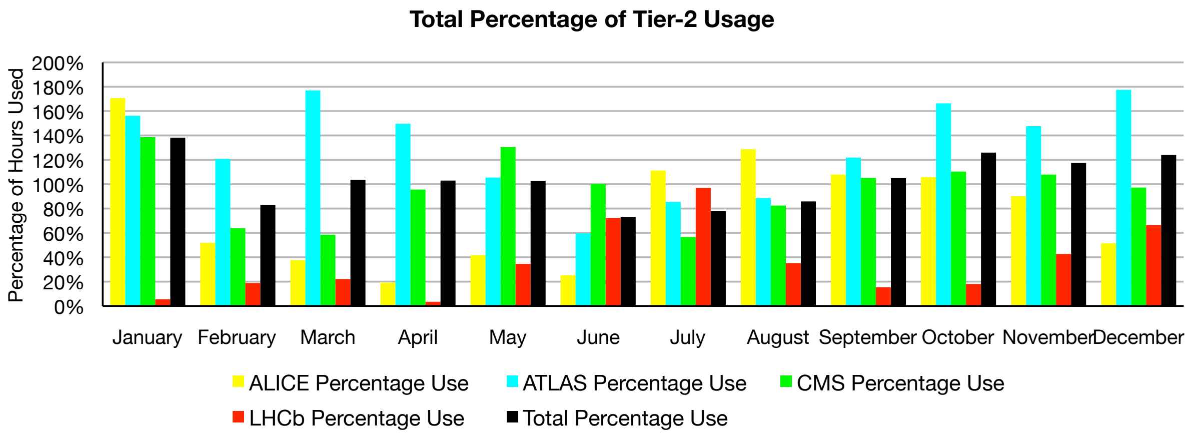

Figure 17: Total Usage of Tier-2 Sites by Experiment

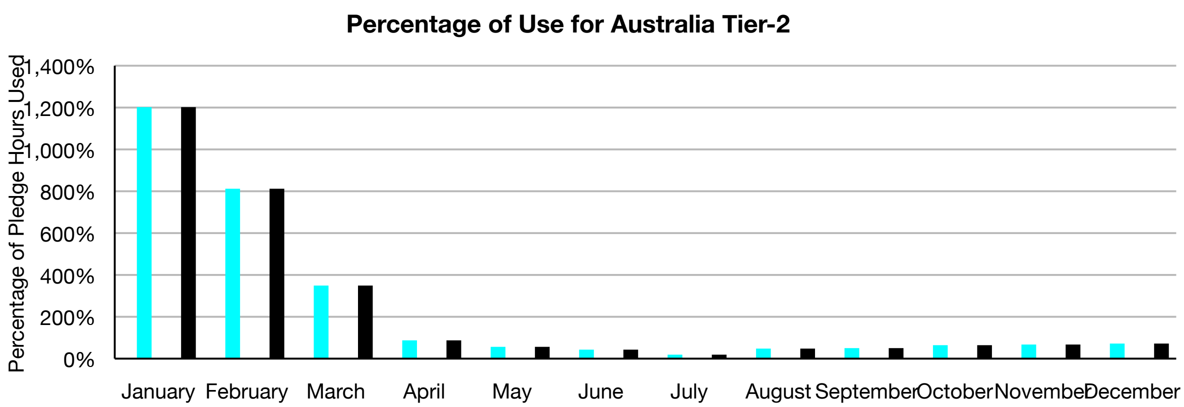

Figure 18: Australia Tier-2 Usage

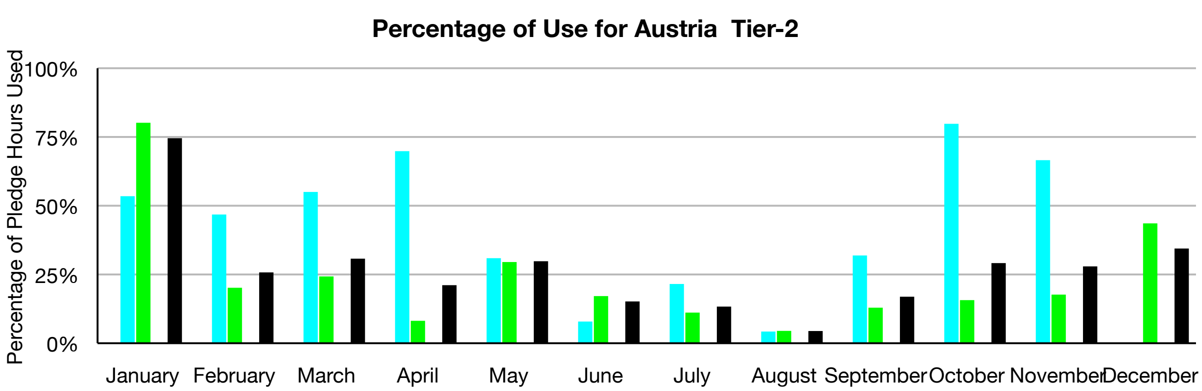

Figure 19: Austria Tier-2 Usage

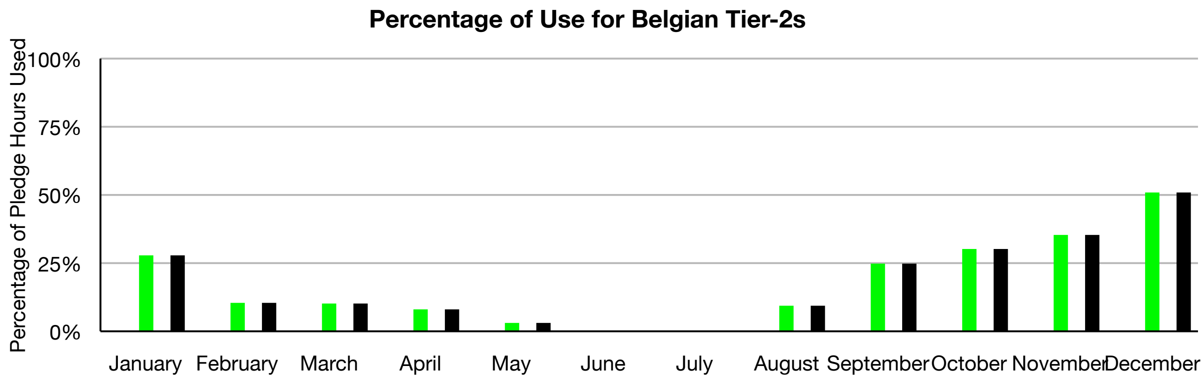

Figure 20: Belgian Tier-2 Usage

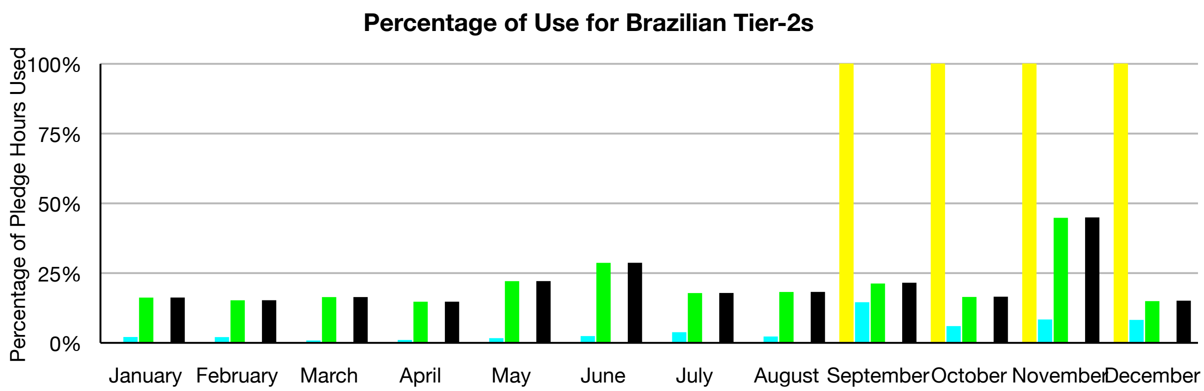

Figure 21: Brazilian Tier-2 Usage

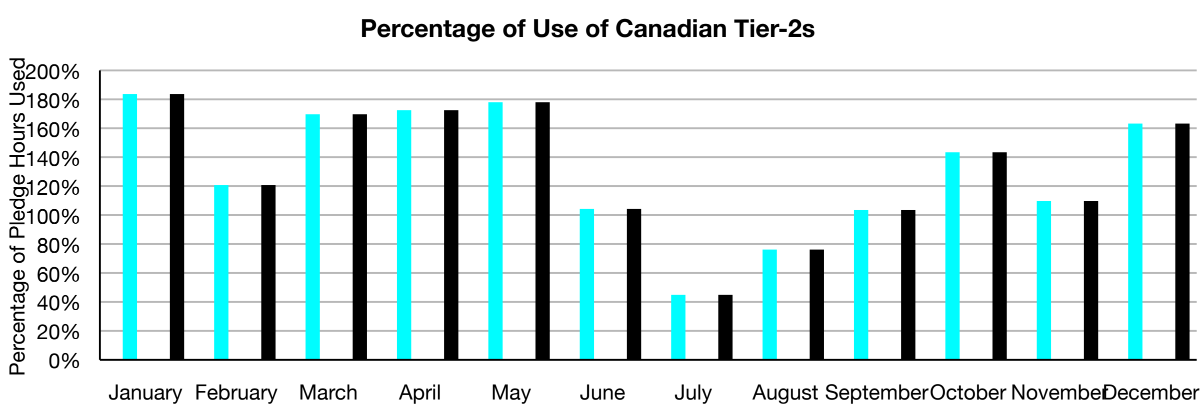

Figure 22: Canadian Tier-2 Usage

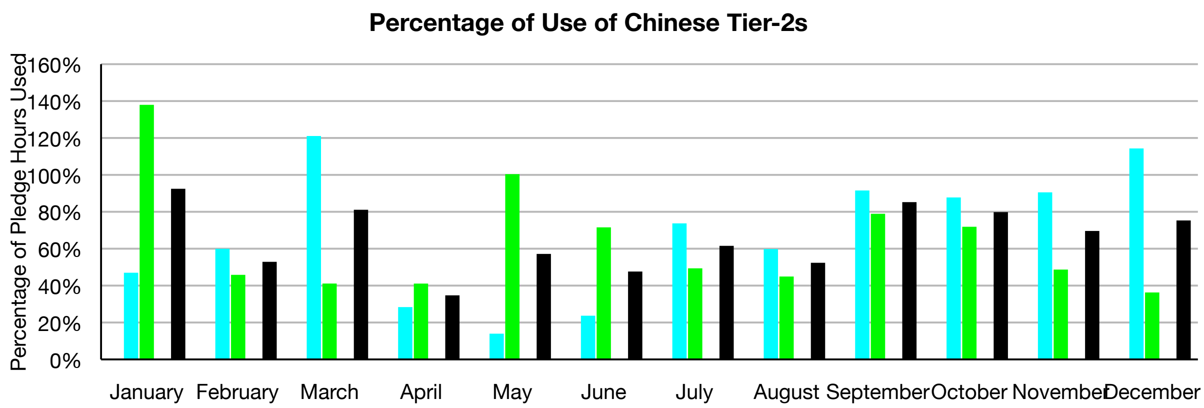

Figure 23: Chinese Tier-2 Usage

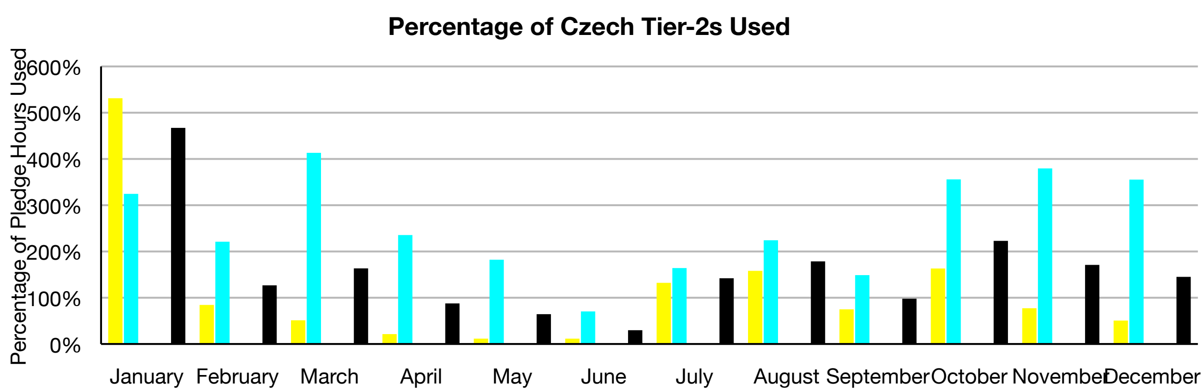

Figure 24: Czech Tier-2 Usage

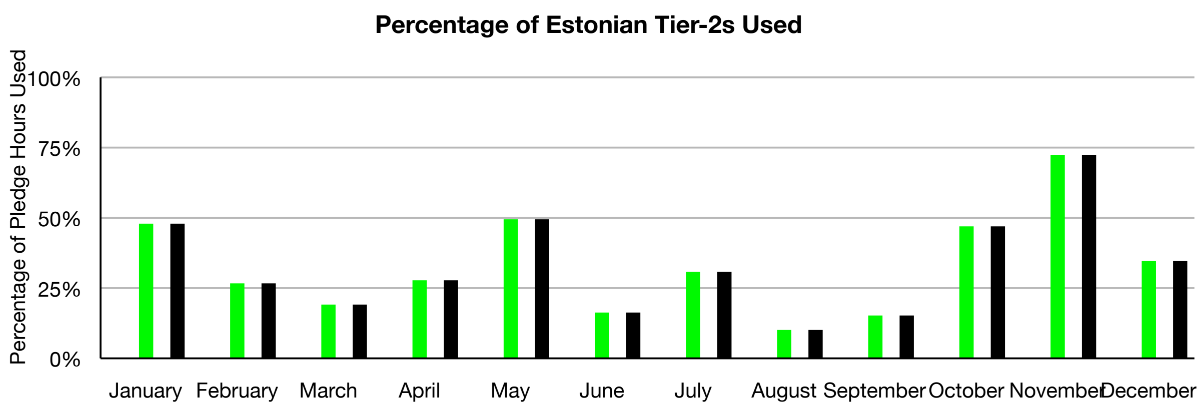

Figure 25: Estonian Tier-2 Usage

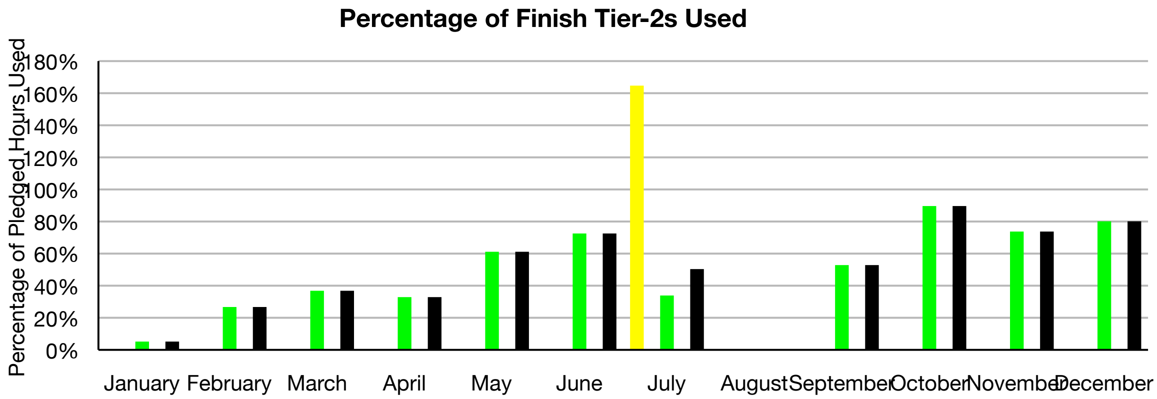

Figure 26: Finnish Tier-2 Usage

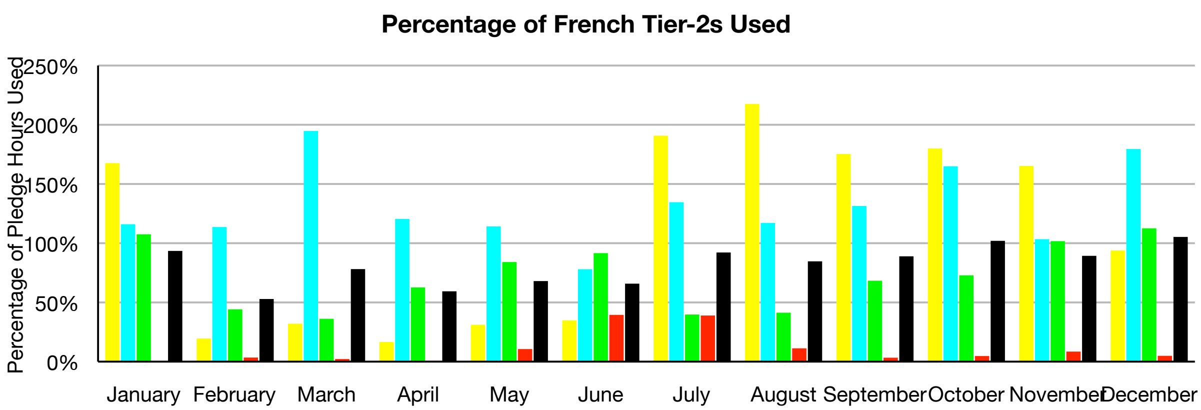

Figure 27: French Tier-2 Usage

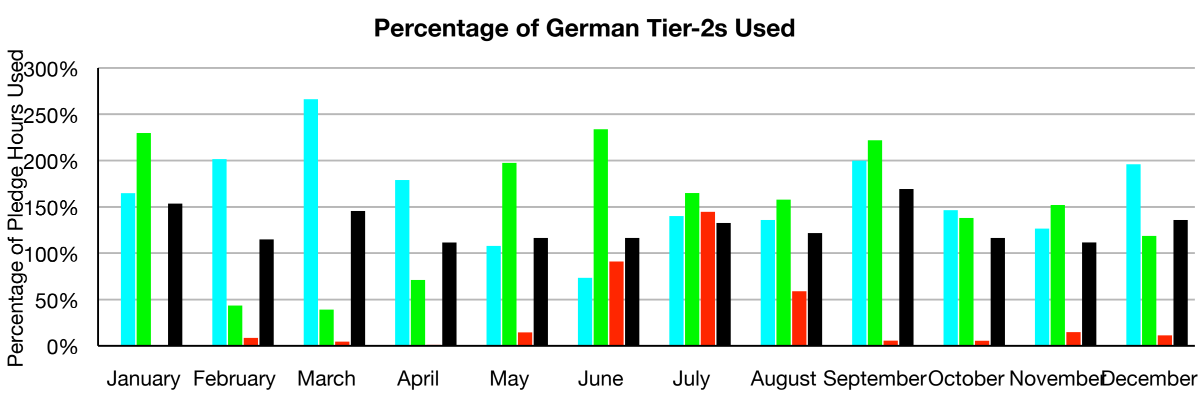

Figure 28: German Tier-2 Usage

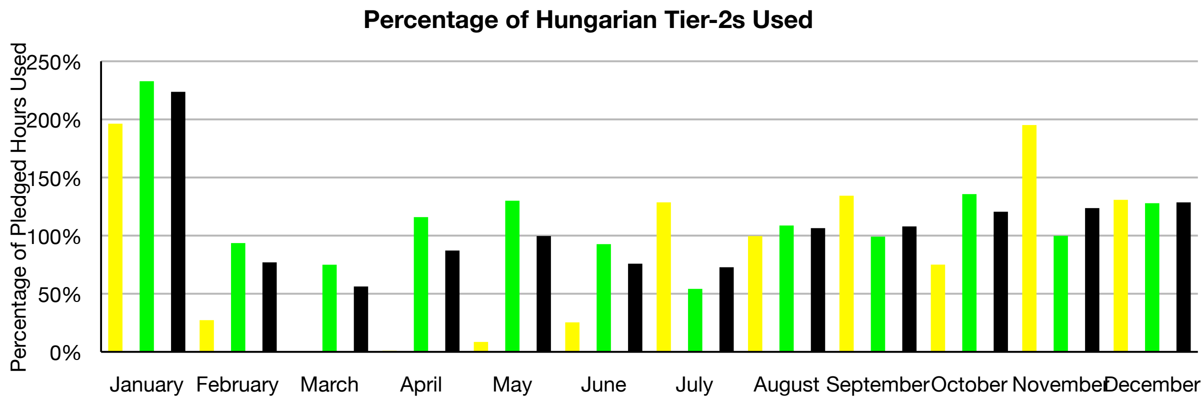

Figure 29: Hungarian Tier-2 Usage

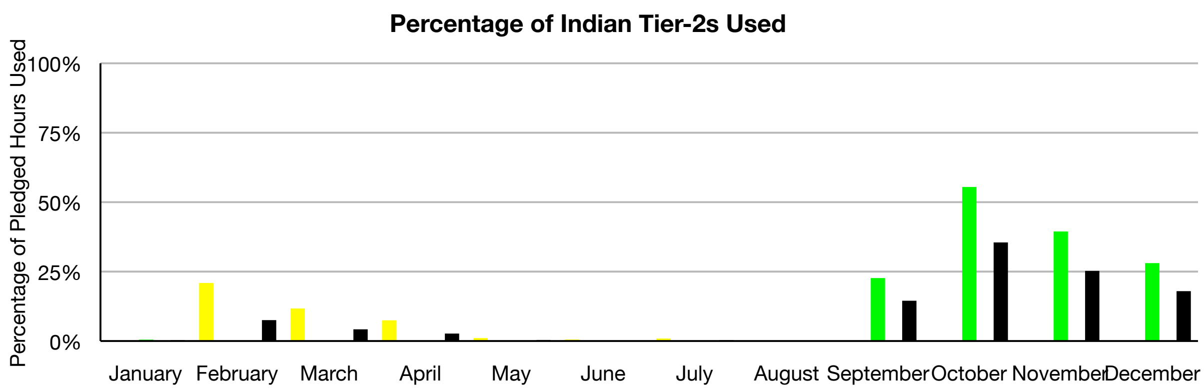

Figure 30: Indian Tier-2 Usage

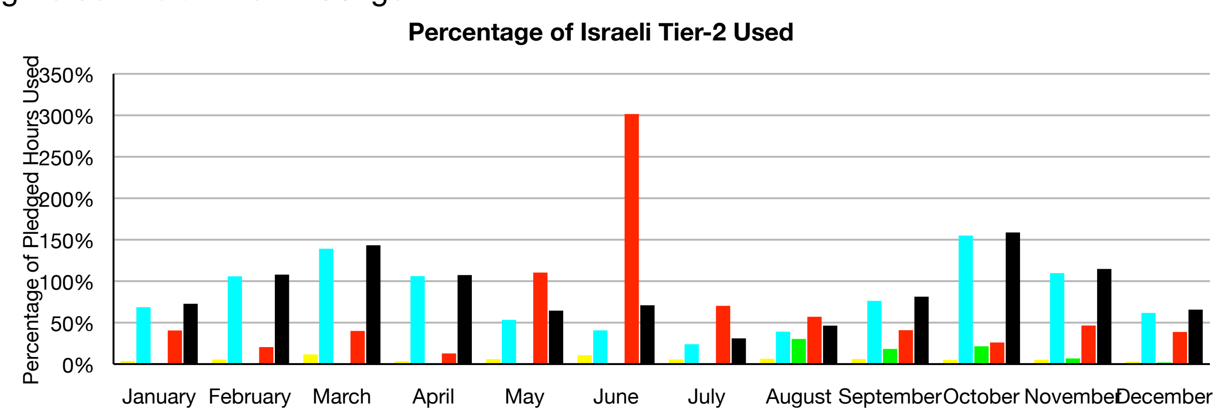

Figure 31: Israeli Tier-2 Usage

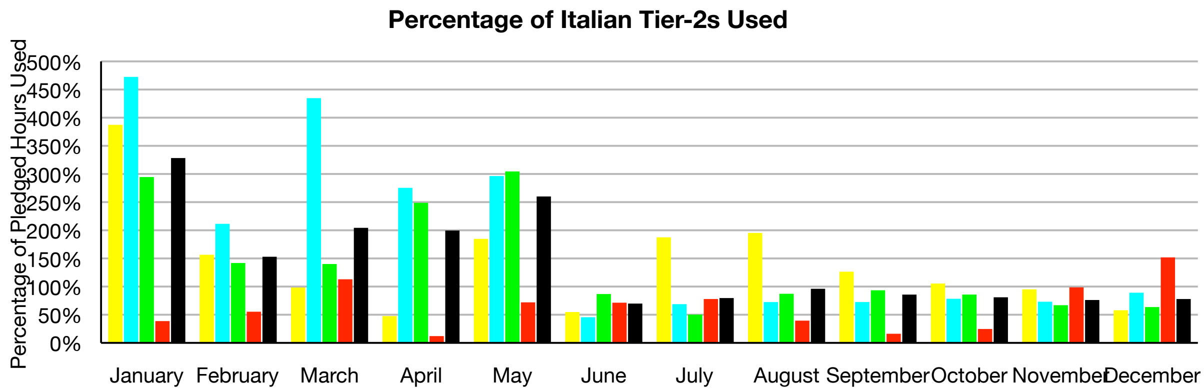

Figure 32: Italian Tier-2 Usage

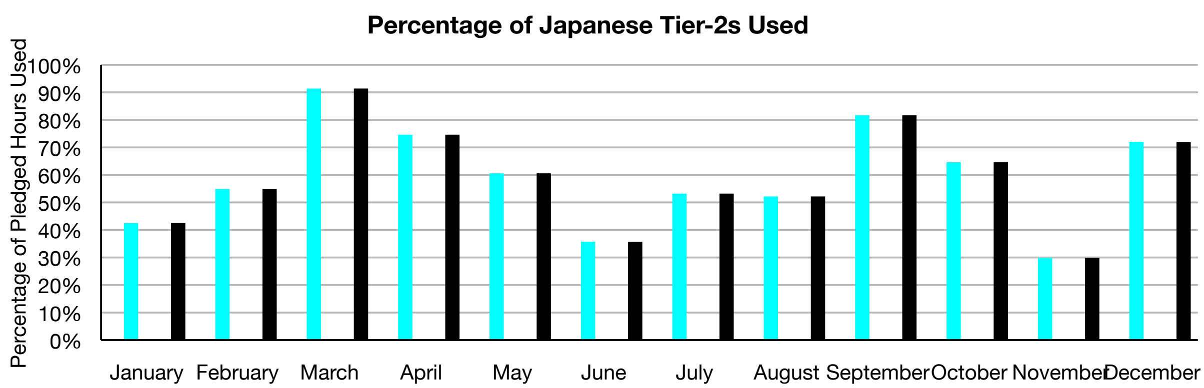

Figure 33: Japanese Tier-2 Usage

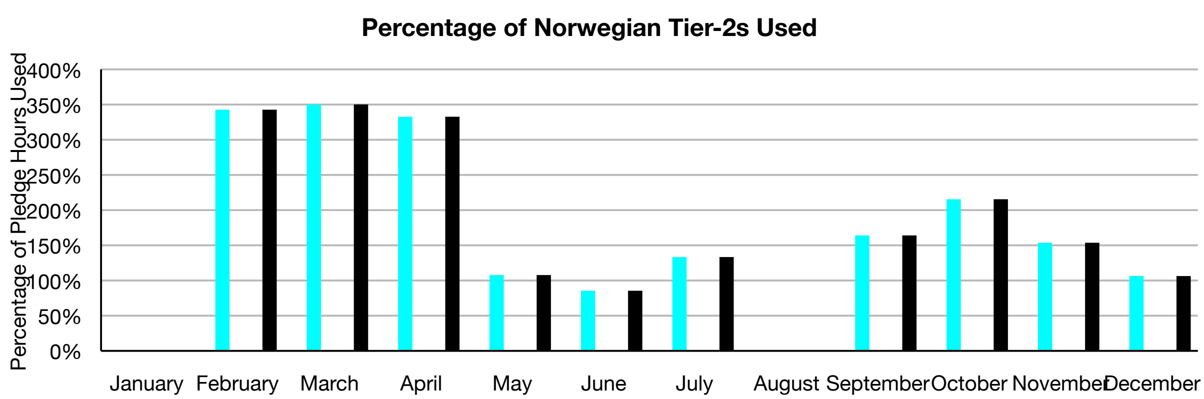

Figure 34: Norwegian Tier-2 Usage

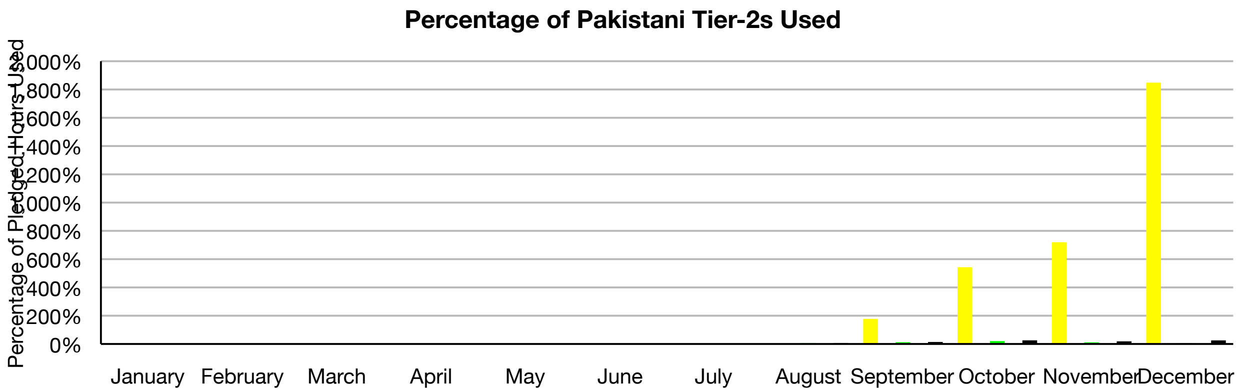

Figure 35: Pakistani Tier-2 Usage

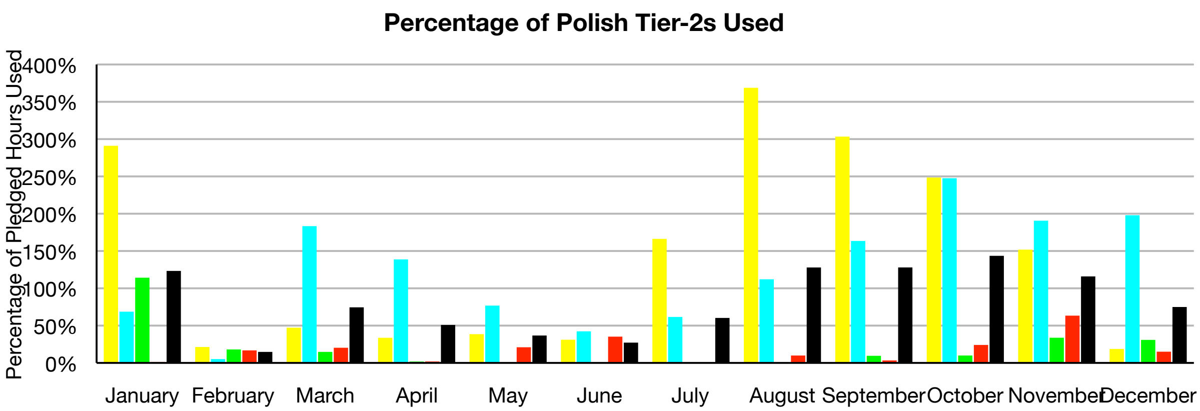

Figure 36: Polish Tier-2 Usage

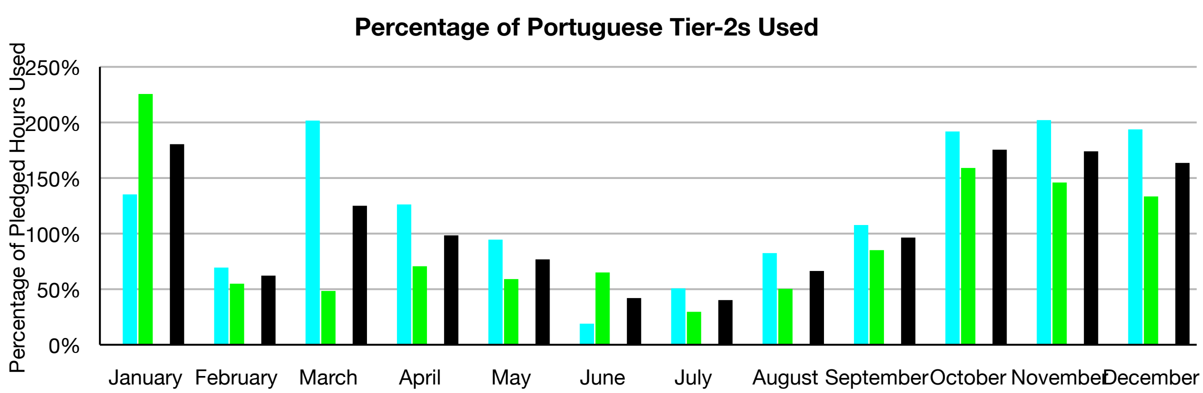

Figure 37: Portuguese Tier-2 Usage

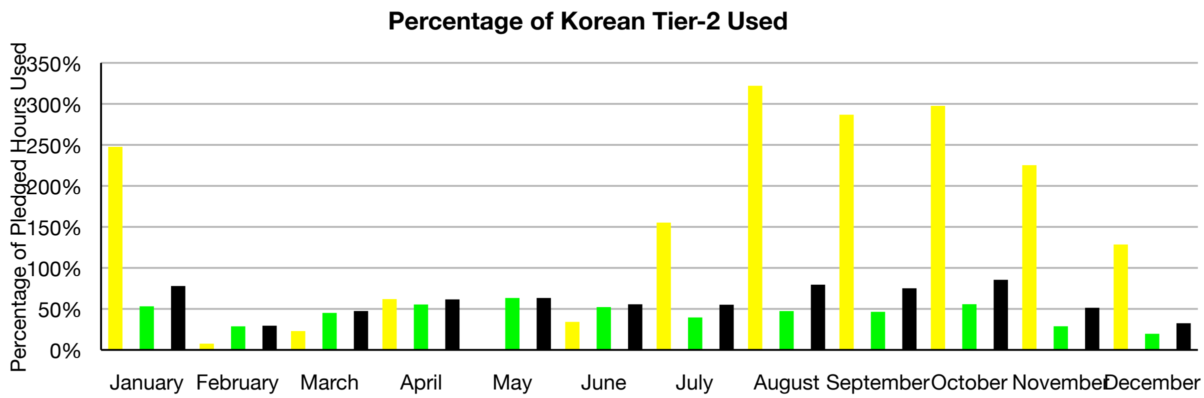

Figure 38: Korean Tier-2 Usage

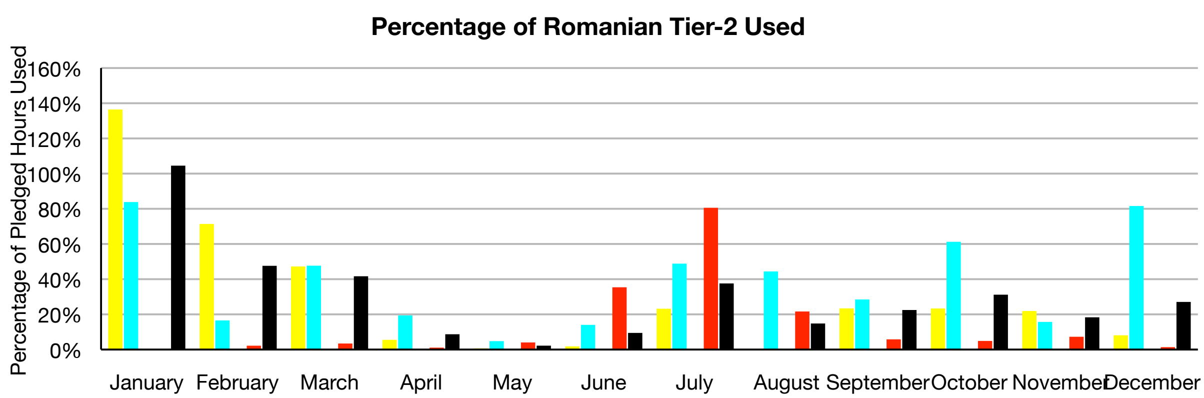

Figure 39: Romanian Tier-2 Usage

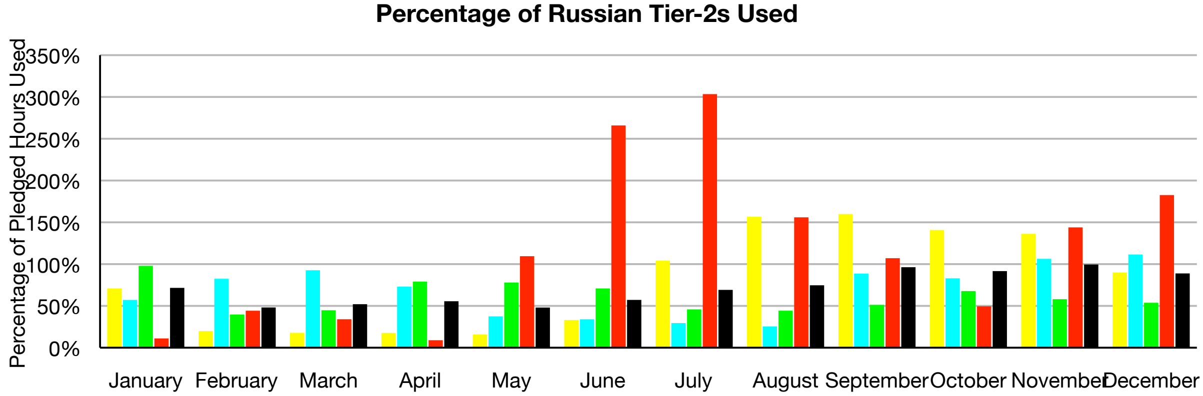

Figure 40: Russian Tier-2 Usage

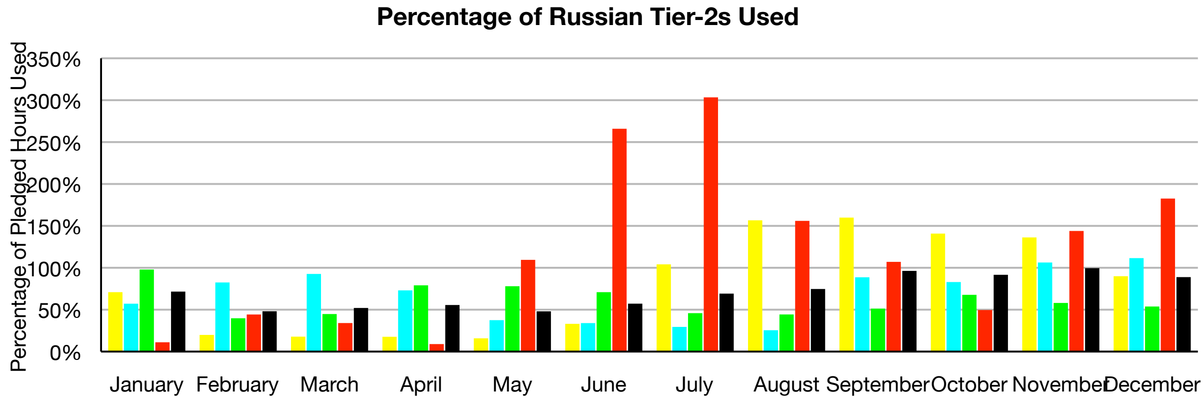

Figure 41: Slovenian Tier-2 Usage

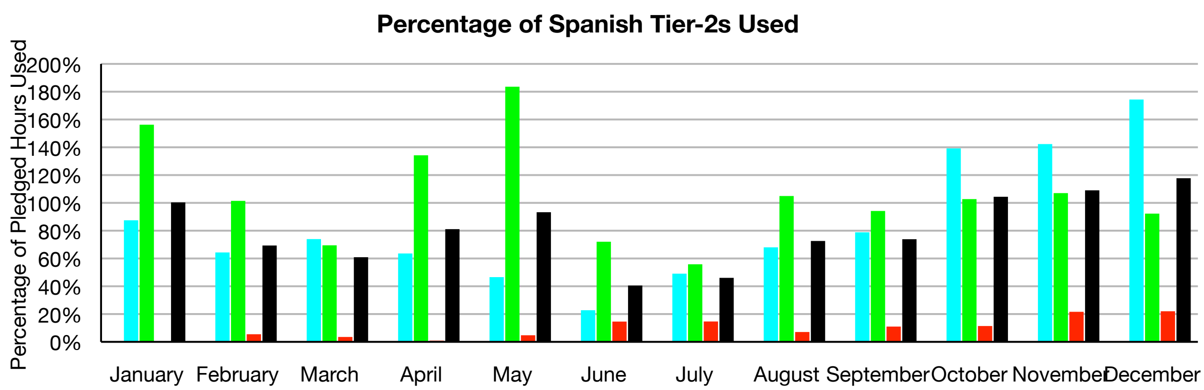

Figure 42: Spanish Tier-2 Usage

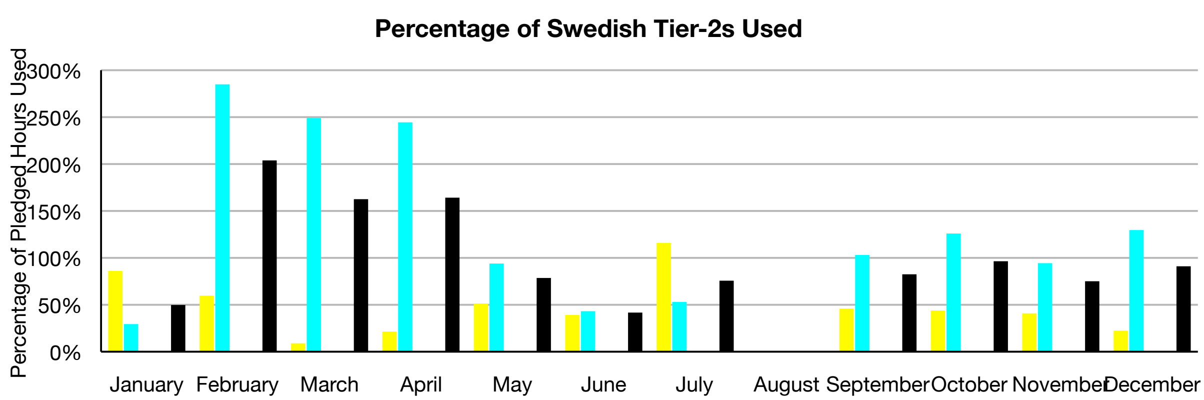

Figure 43: Swedish Tier-2 Usage

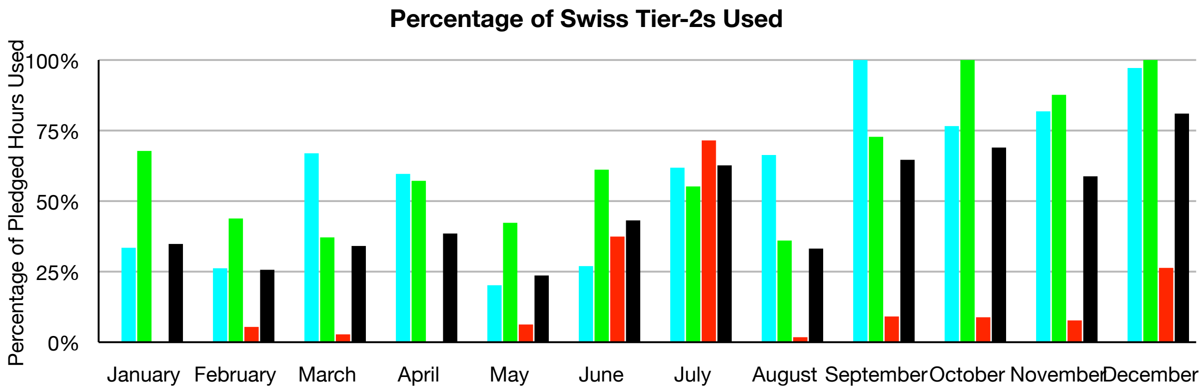

Figure 44: Swiss Tier-2 Usage

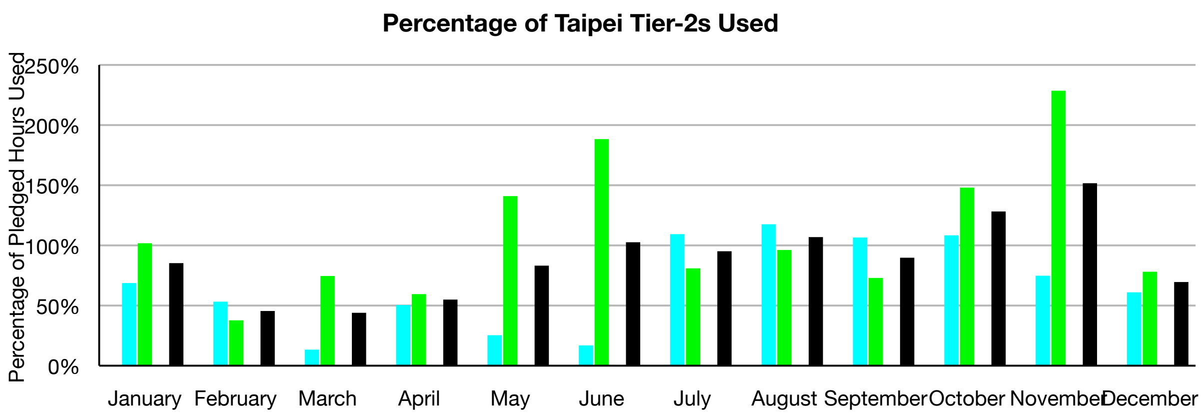

Figure 45: Taipei Tier-2 Usage

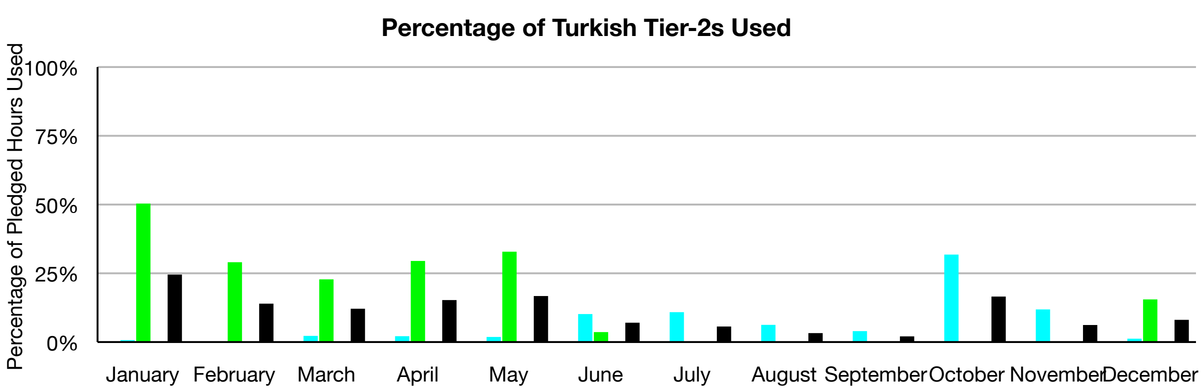

Figure 46: Turkish Tier-2 Usage

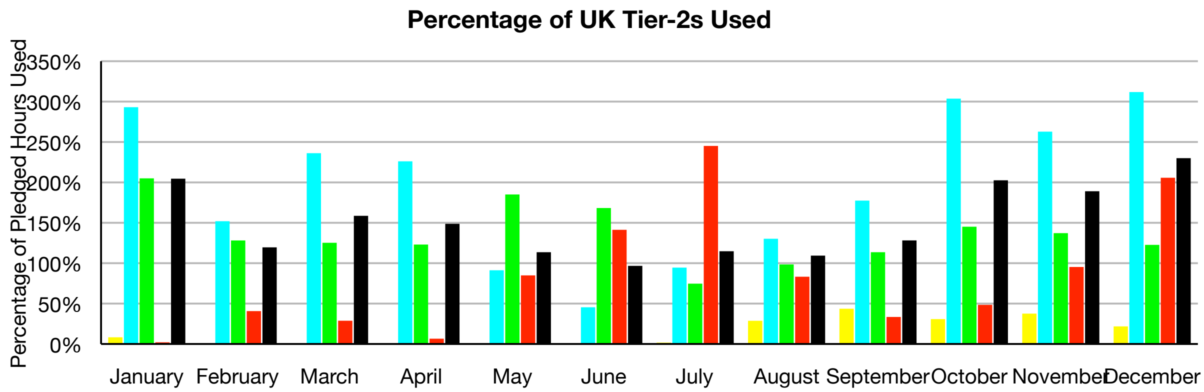

Figure 47: UK Tier-2 Usage

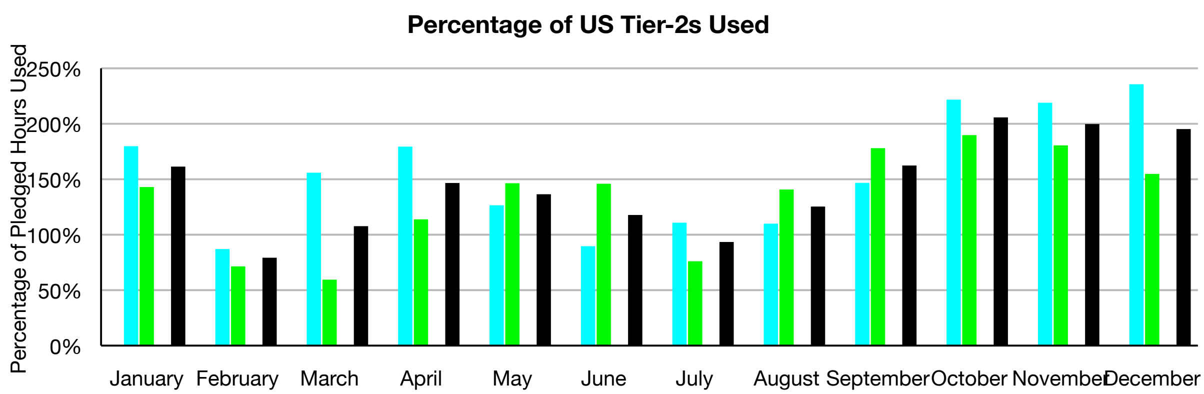

Figure 48: US Tier-2 Usage