

W L C G

Minutes of the 3rd COLLABORATION BOARD Meeting(Held at CERN on 24th April 2008)**Present:**

CERN IT Dept. Head	F. Hemmer (for W. Von Rueden)
CERN PH Deputy Dept. Head	M. Doser
Scientific Secretary	D. Jacobs
LCG Project Leader	I. Bird
<i>LHC Experiment Spokespersons</i>	
ALICE	Y. Schutz (for J. Schukraft)
ATLAS	P. Jenni
CMS	R. Cousins (for T. Virdee)
<i>International Membership</i>	
Austria, Austrian Tier-2 Federation	C.E. Wulz
Belgium, Belgian Tier-2 Federation	G. Bruno (<i>phone</i>)
Canada, TRIUMF	M. Vetterli
Canada, Eastern Tier-2 Federation	P. Savard (<i>phone</i>)
Czech Rep., FZU AS, Prague	M. Lokajicek
France, CC-IN2P3	F. Hernandez (for F. Malek)
France, CC-IN2P3 AF	
France, GRIF, Paris	J.P. Meyer
Germany, ATLAS Federation, Munich	S. Bethke
Germany, DESY, Hamburg	V. Gülzow
Germany, CMS Federation	
Germany, GSI, Darmstadt	P. Malzacher (<i>phone</i>)
Germany, FZK-GridKa	K.-P. Mickel
India, TIFR, Mumbai	A. Gurtu
Israel, ICHEP	L. Levinson
Italy, INFN ATLAS Federation	G. Carlino
Japan, ICEPP, Tokyo	H. Sakamoto
Netherlands, LHC/Tier1	J. Templon (<i>phone</i>)
Nordic Data Grid Facility (NDGF)	L. Fischer
Poland, Polish Tier-2 Federation	R. Gokieli
Portugal, LIP Tier-2 Federation	M. David (for J. Gomes)
Romania, Romanian Tier-2 Federation	M. Dulea
Russian Fed., Russian Data-Intensive GRID (RDIG)	V. Ilyin (<i>phone</i>)
Spain, ATLAS Federation	J. Salt
Spain, CMS Federation	F. Matorras (<i>phone</i>)
Switzerland, CHIP	C. Grab
UK, London Tier-2	D. Colling
UK, NorthGrid	R. Jones
UK, RAL	N. Geddes
USA, BNL	M. Ernst
USA, Northeast ATLAS T2 Federation	J. Shank
USA, FNAL	P. McBride (for V. White)
USA, Nebraska CMS T2	K. Bloom (<i>phone</i>)

Speakers

ALICE Computing Co-ordinator	F. Carminati
ATLAS Computing Co-ordinator	D. Barberis
CMS Computing Co-ordinator	M. Kasemann
CERN-IT Grid Support Group Leader	J. Shiers

Excused

CERN CSO	J. Engelen
LHCb Spokesperson	A. Golutvin
LHCb Computing Co-ordinator	P. Charpentier

Absent

Australia, Univ. of Melbourne	G. Moloney
China, IHEP, Beijing	G. Chen
Estonia, NICPB	M. Kadastik
Finland, NDGF/HIP Tier-2	D.O. Riska
France, LAPP, Annecy	S. Jézéquel
France, LPC, Clermont-Ferrand	D. Pallin
France, SUBATECH, Nantes	L. Aphecetche
Germany, ATLAS Federation, FR/W	J.E. Sundermann
Hungary, HGCC Federation	G. Vesztergombi
India, VECC/SINP, Kolkata	Y.P. Viyogi
Italy, CNAF	M. Mazzucato
Italy, INFN ALICE Federation	M. Masera
Italy, INFN CMS Federation	M. Paganoni
Italy, INFN LHCb Federation	U. Marconi
Republic of Korea, KISTI, Daejeon	S. Hwang
Norway, UNINETT SIGMA Tier-2	J. Koster
Pakistan, Pakistan Tier-2 Federation	H. Hoorani
Slovenia, SIGNET	B. Kersevan
Spain, PIC	M. Delfino
Spain, LHCb Federation	R. Graziani Diaz
Sweden, SNIC Tier-2	S. Holmgren
Taipei, ASGC	S. Lin
Taipei, Taiwan AF Federation	
Turkey, Turkish Tier-2 Federation	
UK, Scotgrid	N. Glover
UK, SouthGrid	J. Tseng
Ukraine, Ukrainian Tier-2 Federation	G. Zinovjev
USA, Great Lakes ATLAS T2 Federation	S. McKee
USA, Midwest ATLAS T2 Federation	R. Gardner
USA, Southwest ATLAS T2 Federation	K. De
USA, SLAC ATLAS T2	W. Yang
USA, Caltech CMS T2	H. Newman
USA, Florida CMS T2	P. Avery
USA, MIT CMS T2	C. Paus
USA, Purdue CMS T2	N. Neumeister
USA, UC San Diego CMS T2	F. Wuerthwein
USA, U. Wisconsin CMS T2	S. Dasu

3rd Meeting of the WLCG Collaboration Board (CB), 24 April 2008

Documents can be found via: [WLCG CB Home Page](#)

1. Introduction

CB Chair N. Geddes

N. Geddes welcomes those present, noting the good attendance, especially the representation of the Tier-2's, which hitherto has been somewhat disappointing. The CB is taking place in the middle of a Collaboration Workshop, which is also very well attended (~240 participants).

He conveys apologies from the CERN CSO J. Engelen, who is unable to attend.

The revised minutes (draft 3 dated 12 February 2007) of the 2nd CB meeting on 25 January 2007 are approved without further comment.

2. Election of the next CB Chair

OB Secretary D. Jacobs

Voting slips carrying the names of the three candidates – Roger Jones, Jose Salt and Michel Vetterli – are distributed to the voting members present with the instruction that they indicate with a cross the candidate of their choice. 31 votes are cast in addition to the 7 email votes already sent to D. Jacobs by those participating on the phone. Both the paper and phone votes are counted and **Michel Vetterli** is declared the winner. He will therefore assume the role of CB Chair from the end of this meeting.

M. Vetterli expresses his thanks for the confidence placed in him.

3. Summary of OB and C-RRB meetings

OB Secretary D. Jacobs

Overview Board (OB)

The OB, comprising management representatives of the experiments, the LCG project and the Tier-1's, has met five times since the last CB meeting. Further meetings are planned in July, October and December. The meetings have treated the progressive (but not always smooth) ramp-up over the last year of software readiness, processing and storage capacity, monitoring and accounting capabilities, and the human aspects of service readiness.

During the period, Ian Bird has taken over from Les Robertson as LCG Project Leader and the task of Resource Manager has passed from Chris Eck to Sue Foffano. It has also been decided (subject to confirmation) that CERN will continue to run its WLCG activity as a project at least until end-2009.

The WLCG MoU has now been signed for all Tier-1 sites and good progress has been made with signature for Tier-2 sites and Federations. Open issues concern Brazil (reluctance of the Funding Agency to sign), along with the open signature timetables of Austria and the Czech Republic. US-ALICE has submitted a capacity request to the DOE.

The perennality of Grid Infrastructure has become a concern. Six of the seven (NDGF absent) European Tier-1's have confirmed that they have fallback plans for the case that EGI (or other EU activity) does not provide the hoped-for Grid Infrastructure effort in 2010 and later. The same applies with respect to OSG for the US and Canadian Tier-1's (Taipei was absent). Nevertheless, a high value is set on a common European Grid Infrastructure and so pressure is being maintained on the EU to provide it.

Computing Resources Review Board (C-RRB), 15 April 2008

The CERN personnel + material funding, which was at the level of ~40 MCHF in 2006 and 2007, was reported as ~48 MCHF in 2008. The amounts planned over 2009 to 2012 decrease slowly from ~38 MCHF to ~36 MCHF (including an infrastructure allowance for increased power and cooling needs).

In view of the imminent LHC operation, pressure is being put on the funding agencies to meet their pledges for 2008 and to give pledge values for 2009-12 as soon as possible. Provided that the 2008 pledges are met, there is unlikely to be a capacity problem this year, although there is still a large (40-50%) shortfall in pledges for ALICE, relative to expressed needs (which were based on there being a Heavy Ion run). In future years there are overall shortfalls, the exact impact of which is hard to determine, since some capacity is not exchangeable amongst experiments. Nevertheless, it is heartening to note that Tier-1 installed cpu capacity actually got ahead of usage since mid-2007, driven by the slope needed for the ramp-up to 2008 levels.

D. Jacobs' opinion is that some funding agencies are clearly confused by the requirements/pledges/installed/used relationships, against the background of a shifting LHC schedule. They are no longer content with being exhorted to honour pledges and pledge more. They are eager to see reliable, understandable accounting and anxious to be reassured by the newly created Computing Resources Scrutiny Group (C-RSG) that the requirements of the individual experiments are valid in time and size, and that they have taken sufficient account of efficient use of computing resources.

4. CCRC'08

Grid Support Group Leader J. Shiers

Outcome of the February Challenge and plans for May

J. Shiers recalls that the objective of CCRC'08 is primarily to demonstrate that both the sites and experiments can run together at the 2008-scale of production, exercising all functional blocks of the data-processing schemes. Failure in the first phase, in February, would have inevitably implied starting discussions on scaling back the production service for 2008. Fortunately the results indicate that this is not necessary, although much remains to be done before the summer.

The February phase went reasonably well but the exercise underlined the desirability of some form of automated reporting. It was hoped to have the "service dashboards" ready for May but this is unlikely to be the case. It was clearly very useful to have pre-established the metrics. The need was shown to clarify reporting and problem escalation lines, as well as some form of lightweight "post mortems" when MoU targets are not being met. Continued improvements are needed in the clarity and transparency of reporting, as well as further automation of monitoring, logging and accounting, such as single grid maps per experiment that allow users to home-in on the details. Some people seem to feel that, having "proven" operation with CCRC'08, everyone can relax but this is not so. Daily operations meetings will remain a necessity and it appears desirable to repeat CCRC's in future years before the start of data-taking. It will certainly be an issue how to fit these in with the experiments' reprocessing plans but they will undoubtedly be necessary if some components of WLCG (e.g. networking) change in a major way. Even this year there will probably be some things (e.g. the ATLAS conditions database) that cannot be tested in May.

Several first conclusions can be drawn from the February phase. Communications need improvement, particularly those with the Tier-2's. In this respect, the Asia-Pacific area is clearly at a disadvantage due to the time difference and an attempt is being made to address this issue by holding a specific conference call at 08:30 Geneva time. In addition, rather

simple service guidelines are being established so that it is easy to understand whether services respect them or not. These approaches are also influencing other Grid projects. Work continues on clarifying the reporting and problem-escalation lines, including a new lightweight mechanism by which named experts in the experiments can get to LCG experts (via CERN) out-of-hours (to be expanded in due time to the Tier-1's).

Amongst the achievements of CCRC'08 so far, the preparation itself provided a useful focusing effect. Many middleware bugs were fixed "just in time" and a clearer view was obtained on what to install. Further improvements in this domain are expected for May. SRM v2.2 was deployed to ~160 sites by end-February and there were major deployments of CASTOR etc.. It was also challenging to demonstrate operation in normal "service" mode (as against the unsustainable "hero" mode needed previously). Those involved were fully occupied by the effort but at an acceptable working stress level (except perhaps for the experiment contact people at the sites). The daily operations meetings treated on average ~3 issues per VO.

The May challenge is planned to run 5-30 May. A big effort is being made to improve on February in terms of being ready well in-time. As for February, all the details will be summarised in a few tables on the CCRC'08 wiki.

Databases represent a major service issue, since every top critical service of WLCG involves a database, which clearly must thus itself be treated as a critical item. The intent had been to have the CERN Oracle services moved to new hardware but this has only just been done now (one hardware installation per service). The Oracle version upgrade was postponed but may now take place in June. It is however good that many sites will have been converted to Oracle before data-taking starts, as this fixes a number of bugs. There is also now a fallback solution tested and in-place.

For storage systems the main points concerned versions and associated issues. One such issue is the timescale on which sites can decommission SRM v1.1 services. The experiments seem still to need SRM v1.1 for some aspects of their work and so it is not likely that decommissioning can be completed much before end-08.

For the middleware components, it is worthwhile noting that the recommended versions for May are all based on production releases.

Overall, the February challenge was successful, demonstrating the usefulness of pre-defined metrics that could be monitored. Much remains to be done in and after the May challenge to demonstrate full readiness but there is now an increasing awareness that this is not another prototyping phase but rather the WLCG production service in its initial state.

As an invitation for reflection, J. Shiers reviews a number of points that arose when considering the outcome of the February challenge.

Conference calls etc. - it is hard to know what is enough. More awareness is needed of the available tools (e.g. the wiki of the Management Board). The set of communication channels can certainly be improved to remove some duplication (but not for May). Although the number of new problems per day was rather low, the follow-up was not perfect and must be improved in May. RAL was the site most consistently participating in the calls. The calls last <15 minutes and each involves 8-12 participants.

Following the proposal made in January, names have been put forward for Tier-2 co-ordinators (distributed around the Tier-1 sites) and also a database co-ordinator at CERN. DB co-ordination is now considered to be solved, with most of the issues being handled at CERN. For the Tier-2's, it is less obvious if the people involved are aware of their functions in this respect, as the sites still complain of not having enough information. Special efforts (by conference calls etc.) were made to help the Asia-Pacific sites. Collaboration workshops

such as the present one are generally felt to be very useful communication channels, also for promoting personal contacts on issues, but they have their own limitations.

Some communication issues remain to be addressed. For instance many feel that the daily 15:00 (Geneva time) Operations meeting functions like joining a CERN meeting remotely – those outside CERN are disadvantaged. Efforts are under way to improve the agenda and to make people more aware of the daily minutes and weekly summaries. Both of these are relatively readable and it is easy to pick out relevant information.

Regarding service reliability, a set of targets in terms of intervention and problem resolution times has been established for the Tier-0. Tier-1's and Tier-2's are being asked to adopt similar targets but for many the target of 30' operator response to alarms is simply not realistic. More important is to press all sites to have implemented consistent use of all WLCG service standards by end-2008. If this is not achieved there will be a follow-up. Although the targets must be less stringent for the Tier-2's, they should make more effort than at present to respond to alarms. For all sites, the first step towards meeting service reliability targets is obviously to begin measuring their achieved reliability in the various categories. When targets are being missed it is important to understand why. Post-mortems are triggered (only two needed during the February challenge). The Collaboration workshop should be helpful in raising awareness of service reliability issues for May.

J. Shiers concludes that there is now a working WLCG Collaboration and a service – a big improvement with respect to the situation at the time of the last Collaboration meeting. The capability to handle complex situations has been demonstrated and there is confidence that problems arising can be fixed in an acceptable time.

N. Geddes comments that he has gained the impression from the workshop that the February challenge was largely successful and asks what should be the first point of contact for people to enquire about the present LCG status. J. Shiers replies that all is available on the LCG home page (although not yet very well organised). N. Geddes suggests that there should be bi-weekly LCG status summaries.

N. Geddes speculates that it may be useful in the operational phase to co-opt more people onto the Overview Board (as provided for in the WLCG MoU). He leaves it as an exercise to his successor to move to enhance the OB's coordination role.

5. Collaboration-wide Issues

LCG Project Leader I. Bird

I. Bird focuses on the effort needed in the Tier-2's, as the issues at these centres cannot be addressed one site at a time as was the case with the Tier-1's.

In several aspects – reliability, accounting, pledged/installed resource capacity and milestones – the need is felt for a more structured dialogue with the Tier-2's. There are also other issues such as the engagement in EGI, the NGI's etc. for grid infrastructure, the management of procurement, and the best ways to promote information flow to the Tier-2's and to discuss technical matters with them. The Tier-2's need a voice and their participation in the CB is clearly not enough on its own.

Almost full accounting information for sites is now available, although some Tier-2's are not yet reporting (perhaps due to mistaken site or host names). The collection of Tier-2 accounting figures started in mid-2007 and they have been published since September. They have elicited rather little feedback and it is unclear if people are checking them or are simply satisfied. The information is available by VO and by country but it is as yet unclear how best to present it to the C-RRB. The plots also show that it is hard to interpret some data such as pledged vs. used capacities. It is becoming essential to know the actual installed capacities

at the Tier-2's. This should be possible, at least at the Federation or national level. When accounting information is given by Federation, it is vital that the individual sites check and comment on the data before they are made official.

All Tier-1's and ~100 (from ~280) Tier-2's now report reliabilities. The monthly publishing of Tier-1 figures led to some measurement- and site-related problems being uncovered, but the situation is now understood. It will be important to similarly follow-up reliability measurements site-by-site for the Tier-2's (although they are clearly already performing rather well, especially in terms of the fraction of processing power represented).

The MB recently set up a group to work on filling the reporting holes by reaching agreement on equivalence of tests in other grid infrastructures, firstly with NDGF (now in production) and then with OSG. There are, however, reporting holes in other places as well. Meaningful comparison between sites is also an issue (e.g. weight them by number of CPU's?) but the greatest need is to publish figures, at least by Federation. It is essential to have named Federation representatives, who can follow up these issues.

Recapitulating the reporting issues, I. Bird asks all site representatives to check carefully the site and host names being used for the accounting and reliability information. Any necessary corrections should be sent to the LCG team at CERN [**action: Tier-2 site representatives**]. He reiterates the need to know the installed resources at the Tier-2's [**action: Tier-2 site representatives**] and to have named contact persons for each Tier-2 Federation [**action: Representatives of the Tier-1's on which the Federations depend**].

There are a number of outstanding resource issues. There was a goal to have all of the pledged Tier-1 cpu capacity for 2008 installed by April but most sites had real problems with this and it will probably be necessary to learn to live with delivery delays. The same is true for disk, flagged as a critical resource. There are various reasons (late budget distribution, procurement process problems, late delivery) but the main point is that the Tier-1's found that, even if they had the money, they could not acquire and install equipment rapidly. It is unclear if the Tier-2's suffer from similar problems. The only solution may be for agencies to advance the availability of funds in the future so that the annual ramp-up process can be started earlier. The point was raised in the C-RRB but met with no reaction.

The management of most parts of the project except the Tier-2's has been aided by the setting of milestones. It is probably now time to extend this scheme to the Tier-2's and milestones will be proposed, probably at the level of overall targets for regions. This will imply the regions having the appropriate technical co-ordination to follow-up on issues.

As already mentioned, better ways need to be found to keep Tier-2's informed of issues. They should happen via the Management Board but it seems they do not pay attention. The GDB is a monthly technical meeting that affects Tier-2's. It was hitherto rather closed but is now open. This is the place where issues such as accounting, pilot jobs etc. come up. Some Tier-2's already join the meetings but there should be many more. Apart from these bodies, the Collaboration Board is presently the only place where the Tier 2's have a voice. I. Bird invites input on how best to structure the communication with the Tier-2's and how to represent the Tier-2 status to reporting bodies.

There are several miscellaneous technical issues on which Tier-2 input is lacking: OS version moves - SL5/SL6; is there sufficient fabric monitoring; are there missing tools - also for security?. For security an attempt has been made to provide codes of "best practice" but there has been little response. All of these points do get discussed in the GDB, in which the Tier-2's should participate, but they should say if something else is needed.

Lastly, I. Bird addresses WLCG's concern with the future of EGI grid infrastructure. Although this is a European matter, the issues could easily be generalised to OSG. A

blueprint for the future of EGI should be available in June – clearly the wrong time for LCG to be disrupted. The NGI's have now recognised that they must be concerned by grid infrastructure but it is not clear if issues such as funding have been understood. Tier-2's should also play their part in making sure that their NGI representatives understand. All Tier-1's have confirmed that they have some fallback solution, should EGI grid infrastructure not be available (albeit with some associated service cut-backs due to staffing limitations). The situation for Tier-2's must be assessed and an overall decision made on how to address the issue.

On Tier-2 communications, J. Salt points out that the Tier-2's could take advantage of the already-defined regions of EGEE projects. This would yield 10-12 contact persons who could collect information and communicate with the Tier-2 Federations in their region. L. Levinson points out, however, that some EGEE regions span very many countries (e.g. in South-East Europe).

Regarding Tier-2 feedback, D. Colling thinks that Tier-2's often get the feeling that they are simply being issued with instructions. He is much in favour of them all joining the GDB. I. Bird responds that it is up to the Tier-2's to propose what they want to do. The fundamental question to them is how they want to be organised.

M. David considers that, rather than leaving it to the Tier-2's to organise themselves, it would be better to ask the Tier-1's and ROC's to act to contact the Tier-2's and resource centres below them in order to drive the process forward. M. Vetterli agrees that the only way to get feedback is to interact downwards – dissemination followed by feedback. P. Malzacher suggests instead that Tier-2's organise themselves along experiment lines, resulting in one Tier-2 representative per experiment. M. David recognises that Tier-2 issues vary according to experiment and thinks that both groupings are valid. A. Gurtu, while feeling that Tier-2's consider their main responsibility to be towards the experiments, is in favour of them interacting via their Tier-1's.

J.P. Meyer comments that in France there are fortnightly meetings of Tier-2's but communications are vertical - not between Federations.

R. Jones thinks that more than one dedicated meeting is needed for the Tier-2's, e.g. T2-specific GDB meetings as well, as some issues concern policy and others are technical.

C. Grab reports that in Switzerland most of the Tier-2 contact is through the Tier-1. He understands, however, that some communications do not reach the top and is in favour of T2-dedicated meetings of the GDB.

C. Wulz warns of the risk of creating too many redundant contacts. For her the main communication vehicle is the CMS facilities meeting. M. Vetterli adds that ATLAS also communicates with its Tier-2's via the Tier-1's.

N. Geddes asks the Tier-1's if they feel this contact role to be part of their job. L. Fischer confirms that this is so for NDGF. There is tacit agreement to I. Bird's conclusion that the Tier-1's can thus be asked to follow-up on the Tier-2 communication issue but some additional factors are pointed out: the Tier-1's must name individuals with this Tier-2 communication role (M. Vetterli); the onus is on the experiments to ensure that each of their Tier-2's is associated with a Tier-1 (I. Gaines); some Tier-2's are associated with different Tier-1's according to experiment (R. Jones).

After some further discussion, the meeting accepts N. Geddes' conclusion that the Tier-1's should logically be responsible for downwards and upwards communications with the Tier-2's connected with them, a task for which they need named persons. At the same time, well-working communication channels that already exist should not be killed. It is also

recommended that regional groupings of Tier-2's get themselves organised. D. Barberis registers his dissent, holding that only the experiments are in a position to manage the communication with the Tier-2's and that no new arrangements should be made but N. Geddes points out that the complaints have come from the Tier-2's themselves. V. Gülzow, supported by M. David, strongly proposes that the Tier-2's need an additional channel for communication amongst themselves, since there are also problems not related to the Tier hierarchy.

Summing up, N. Geddes notes the encouragement to the Tier-2's to participate in the GDB and the desirability to include them in all information chains. The CB recommends starting with communication improvement via the Tier-1's, a move that will probably solve most problems (I. Bird comments that the exceptions can surely be handled). It is agreed to be a logical responsibility for the Tier-1's to follow-up on Tier-2 communications.

6. "Express line" issues

Computing Co-ordinators, introduced by N. Geddes

ALICE

Y. Schutz comments for ALICE in the place of F. Carminati. He points out that the idea of an express line is not new, having already featured in the computing TDR. ALICE's goal is to rely fully on the CAF for first analysis, implying collection of the promptly reconstructed data from the Tier-0 - on-line for pp collisions and delayed for HI. The CAF should provide low-latency interactive-mode analysis. ALICE has configured a few machines at CERN with PROOF and these are being regularly used by 10's of end-users (a critical aspect is that the data must reside on the local disk of the worker nodes). Although there have been some problems, the users are globally satisfied.

PROOF is of course new. It requires maintenance and, in the future, further development. ALICE has only one expert and his contract ends in September. In the absence of a fallback solution, it may be necessary to discontinue the PROOF cluster.

N. Geddes notes that there was also concern about increasing RAM requirements but Y. Schutz gives assurance that ALICE is making every effort to stay within 2 GB per core.

ATLAS

D. Barberis shows the ATLAS data flow from the Tier-0 to the CAF. ATLAS regards the CAF as the 2nd part of its computing facility at CERN. He shows the 2008 capacity figures and explains the usage - access to larger samples of RAW data than at Tier-1's; higher priority queues for detector monitoring and calibration/alignment; access to data samples for algorithm development. He shows the Castor data pools concerned and emphasises that the CAF shares by data type are constantly being reviewed and revised. At the start it is likely that mostly raw data will be concerned, with more processed data later.

For ATLAS the primary role of the express stream is fast processing (using the most recent calibration) of data at the Tier-0, so as to produce a new calibration. Its main scope is therefore data quality and calibration monitoring. It is important to note that the express stream is a sub-sample of the bulk - not a separate stream. It is thus expected to have a lifetime measured in days but it will also be distributed to the Tier-1's for validation.

ATLAS has an issue with memory size, having asked some time ago that Grid nodes be provided with 2 GB memory per core. Full Geant4 simulation needs just over 1 GB/core. Reconstruction currently needs ~2 GB if the trigger and all reconstruction algorithms are on (which is not the long-term standard running model) but efforts are under way to reduce this before the start of data taking. Analysis needs <1 GB as such but the framework is flexible and users can call-in any part of reconstruction. A task force is looking at a number

of performance issues including use of 64bit processors but tackling the points related to job size has higher priority. ATLAS maintains its request for 2 GB per core.

CMS

M. Kasemann reviews CMS' major use cases for the CAF (more than just the express stream). These comprise: detector & trigger performance monitoring and commissioning; controlled alignment and calibration activities; physics performance monitoring; short latency physics analysis and interactive login facilities for all CMS users (work is still needed to provide high quality disk-space management that scales with user numbers).

The CMS data flow includes 20% (a copy) for express stream and calibration, serving the described use cases. The results are fed back to the Tier-0 for standard reconstruction via a buffer with space for ~24h. They are also distributed to the Tier-1's. The express stream contains about 10% of all event data (based on trigger lists). It is planned also to be available for offsite copies (if quick access is wanted).

Memory usage is a big issue that absorbs much effort. It was intensively studied in 2007 by a taskforce and the work also resulted in a good set of monitoring tools. Single workflows (simulation or reconstruction) presently need ~1 GB (without pile-up). More complex workflows (multiple input or output streams or dealing with pile-up) use 1-1.5 GB (with tails >2 GB). When the limit is exceeded, jobs have to be split - a considerable effort.

CMS has not converted its software to 64bit processors but this would clearly have a very large impact on memory usage. Overall, CMS recommends that new worker nodes are equipped with at least 2 GB of memory per core. For special activities it may be necessary to add yet more (there are already cases up to 4 GB). It is also recommended that the Tier-2's provide 2 GB per core. It is recognised however that, should memory requirements rise above 2 GB, it will be necessary to renegotiate some cpu pledges.

LHCb

No LHCb representative is available. The following statement has been prepared for the minutes by Ph. Charpentier.

"The Computing Model of LHCb foresees that the first pass processing of the raw data will be shared between Tier0 and the 6 Tier1s that serve LHCb. No event streaming is foreseen online, hence no "hot physics" stream is foreseen to be analysed on special resources for physics. However some selected events will be reconstructed online at the pit and possibly stored on the Tier0 SE. This reconstructed data as well as the CERN share of the full reconstructed stream will be used for checking the calibration and alignment of the LHCb detector, and possibly to provide a new calibration. It is foreseen that this takes place partly in the Online system (for monitoring), partly on a Calibration and Alignment Facility (CAF) for computing calibration and alignment constants, located at the CERN Computing Centre. The current estimate is that this task will not require a large amount of CPU, but needs guaranteed response time. Therefore LHCb is investigating with CERN-FIO how to best provide these resources: either with a dedicated interactive cluster (restricted to the calibration and alignment teams) or with a dedicated LSF interactive queue (provides possibly better scheduling of tasks). As a starting point, LHCb would like to exercise this model with an lhcbcaf providing between 12 and 16 current top-level cores. The amount of resources will have to be reviewed in light of experience with first data. Like the LHCb VOBoxes and build servers, it is proposed these resources will be taken out of the LHCb share at CERN."

7. Summary and Future Activities*N. Geddes*

N. Geddes reiterates the request to report the makeup of Tier-2 Federations.

Summarising, he emphasises that improved engagement with the Tier-2's is important - they represent 55% of the currently available cpu capacity.

It is also necessary to know the Tier-2 plans for the post-EGEE era. Someone will be tasked to gather the information for a presentation at the next CB meeting.

Several resource issues were raised. The LCG project should be proactive in this respect towards the experiments and the Tier-2 centres.

7.1. Date of Next Meeting

No date is fixed but it is agreed that there is a good case for >1 meeting in the coming year.

7.2. Any Other Business

M. Vetterli, as incoming Chairman, conveys the CB's appreciation to N. Geddes for his leadership in the first, formative years of the Board's existence (applause).

There being no other business, the Chairman thanks participants once more and closes the meeting.