

## *Plenary Meeting of the LHC Resource Review Boards RRBs, 23<sup>rd</sup> April 2001*

### **1. Welcome** (*R.J. Cashmore, Director for Collider Programmes*)

The Chairman, R.J. Cashmore, welcomed RRB delegates of ALICE, ATLAS, CMS and LHCb. He announced that, having failed to take place at the last meeting, visits to ATLAS and CMS areas are scheduled during the present meeting. He also asked delegates to fill out a questionnaire concerning the use of e-mail and the Web for LHC RRB business.

### **2. Status of the LHC** (*L. Maiani, Director General*)

Professor Maiani started with an artist's impression of the LHC machine *in situ*. Reporting on progress in underground/overground civil engineering, he highlighted the impressive work at the ATLAS and CMS areas. Excavation and lining of the ATLAS vault is nearing completion, after which excavation will begin of the cavern itself, a 6-storey-deep hole below the vault. At CMS, a huge reinforced central pillar is being constructed to support the roof before excavation of the two adjacent caverns can start.

The DG showed a picture and training curve of the first superconducting dipole ready for installation, one of pre-series of 30 dipoles per firm ordered last year. The magnet experienced its first quench just below the nominal field of 8.3 Tesla, reaching the ultimate field of 9 Tesla without a further quench. The DG expressed the wish that the remaining 1235 dipoles should be equally well behaved.

Following discussions between those responsible for the experiments and for the machine, the LHC commissioning schedule has been brought into line with the current status:

April – August 2004	first octant test
March 2005	last dipole delivered
end 2005	ring closed and cold
February 2006	first circulating beams
April 2006	1 month pilot run with collisions
May – July 2006	shutdown
August 2006 – February 2007	7 months of p-p collisions at $L > 2 \times 10^{33}$
April 2007	6 weeks of Pb-Pb collisions

The DG remarked that the revised schedule is not so different from the one drawn up in 1996 when difficulties such as the adverse geology and pillar excavation in the CMS area were as yet unknown. He also noted that not having winter shutdowns in 2005 and 2006 will affect CERN's electricity contracts.

A realistic cost-to-completion estimate for the machine should be possible by the end of the year, after some big contracts for dipole assembly have been adjudicated in September. The DG concluded by pointing out that costs for LHC computing and detector Maintenance and Operation, to be discussed at this and at the October RRB meetings, should also be fairly well known by the end of 2001.

### 3. Maintenance and Operation (*R.J. Cashmore*)

Cashmore opened by reminding delegates of his presentation on M&O at the October 2000 RRB meeting. This had been followed by several bilateral conversations and a paper (*RRB-D 2001-04*) had been distributed in preparation for detailed discussions at the current meeting.

He drew attention to the scale of LHC experiments, which, taken together, are some 4 times bigger than the LEP experiments in capital investment, the number of people involved and M&O costs. During 2000, a CERN Working Group on M&O had collated input from the experiments, from LEP physicists and from CERN experts to prepare preliminary cost estimates for the years up to 2007.

The main principles concerning M&O cost sharing were developed by the CERN Management and presented to the Scientific Policy Committee in March 2001. The SPC expressed its general support for an approach that would: follow the best practices used at LEP, HERA etc.; aim for uniformity and transparency; establish total costs including core manpower; share general costs by established scientific author; leave sub-detector maintenance with the constructors; allow payment in cash and in kind but with a minimum cash contribution; take into account contributions that Member States and some non-Member States have made to building the LHC machine. These basic principles should be discussed during these current RRBs so that a draft MoU for M&O from 2003 onwards can be prepared for October. Interim arrangements for 2002 will have to be agreed in October, as substantial M&O costs are already being incurred.

Cashmore then listed the cost categories; A for common operations, Common Fund items and collaboration-wide services, B for maintenance of items provided by sub-sets of the collaboration, and C for items that are the responsibility of host laboratory. ICFA guidelines, signed MoUs for LHC detector construction and CERN's General Conditions for Experiments had been taken into account when categorising cost items. In essence, CERN must provide safe, bare experimental caverns and ancillary buildings, and the collaborations must cover costs incurred because detectors are installed.

Cashmore emphasised that present cost estimates are not yet reliable for several reasons. All items must be examined for completeness, correctness and categorisation and a clear definition must be given for each one so that costs can be correctly estimated and double counting avoided. The collaborations must bring their cost estimates in line with the 2006 machine start-up schedule and refine them, if necessary, to fit the tighter item definitions. In particular, sub-detector maintenance costs need careful revision. That having been done, scrutiny groups with some members nominated by the RRBs must examine the final estimates to assure the RRBs that they are reasonable.

The preliminary total M&O cost estimates are some 18MCHF in 2002 rising to 65MCHF in 2007. These numbers reduce to 13MCHF and 44MCHF respectively after subtracting category B costs and CERN's C and A costs.

Cashmore reported a widespread preference for sharing A costs by qualified scientific authors of publications or Technical Design Reports, that is, based on exploitation of the detector. If B costs are based on retaining responsibility for sub-systems, cost sharing details will be left to the collaboration. However, these sub-systems are highly complex, and responsibilities and commitments will have to be clearly reported to the RRBs. A common recovery plan should be envisaged in case of a major disaster, an issue that is linked to CERN's Insurance policy, presently under review.

Addressing the fact all Member States and some non-Member States have contributed to building LHC, Cashmore mentioned the idea of rebates, whereby CERN pays a fraction of the A costs of Agencies and Institutes belonging to those States. The SPC had been enthusiastic about this suggestion. The level of the payments and the formulae by which they would be apportioned would require endorsement by the CERN Council.

Cashmore then described the proposed procedure for dealing with M&O costs. In April of year N-1, the RRBs would be given preliminary estimates of A and B costs for year N and a forward look for N+1, N+2 and N+3. They would also receive the final accounts of A costs and final reports of B costs for year N-2. In October of year N-1, the RRBs would be given scrutinised estimates for A and B costs for the years N, N+1, N+2 and N+3. They would approve final allocations for A costs and note final arrangements for B costs for year N.

Cashmore expressed his hope that the current RRBs would discuss arrangements for handling M&O in 2002 and agree that draft M&O MoUs be prepared for examination in October. Final versions of the MoUs could then be presented for approval to the RRBs in April 2002 and thereafter circulated for signature.

In conclusion, Cashmore stated that M&O is an important issue for commissioning and exploiting LHC detectors. Serious costs are already being incurred and it is becoming urgent that CERN, the collaborations and the RRBs establish formal arrangements for handling these costs, to guarantee the success of the LHC experiments.

#### *Discussion*

I.F. Corbett (GB) pointed out the considerable overlap between the experiments on issues of both M&O and LHC computing. He asked whether the October meetings could be organised to take account of that and optimise the use of what will be a very limited amount of time. Cashmore agreed, adding that it is important to best use the time before as well as at the October meetings.

#### **4. The LHC Computing Review** *(S. Bethke, Steering Committee Chairman, LHC Computing Review)*

*Computing Review Steering Group Report (RRB-D 2001-03) is available at <http://lhc-computing-review-public.web.cern.ch>*

LHC computing is a vital prerequisite for the success of the experiments and a potential source of spin-offs yet unimagined. Bethke explained that LHC offline computing was not included in detector MoUs because of severe uncertainties in extrapolating more than 5 years ahead in this rapidly developing field. It is now both appropriate and necessary to finalise LHC computing plans, start serious prototyping and secure the resources needed for timely completion.

The LHC Computing Review was conducted by three independent panels reporting to a Steering Committee: a WorldWide Analysis / Computing Model panel (chaired by D. Linglin, CC-IN2P3/CNRS), a Software Project panel (M. Kasemann, FNAL) and a Management & Resources panel (M. Calvetti, INFN Florence). Membership encompassed representatives of the four experiments and CERN IT Division, and experts from around the world.

LHC computing is an unprecedented challenge for the HEP and IT communities, as Bethke illustrated with statistics that included total annual storage requirements of 7 PBs ( $7 \times 10^{15}$ ) for raw data and 3.2 PBs for simulated data. Each year, Tier0+Tier1+Tier2 centres (see later) will need a total tape storage capacity equivalent to 40 million CD-ROMs and disk storage equivalent to 140 thousand 75 GB disks. Their combined CPU capacity will equal that of 360 thousand of today's PCs with a total Tier0 $\leftrightarrow$ Tier1 WAN bandwidth of 5000 Mbps to serve the four experiments. The Review accepted the scale of these resource requirements estimated by the experiments.

The World Wide Analysis / Computing Model panel recommended adopting the distributed, hierarchical model developed by MONARC, a collaborative effort on MOdels of Networked Analysis at Regional Centres. This model consists of a Tier0 centre at CERN, to store all raw data and perform some reconstruction, connected to several regional/supranational Tier1 centres for analysis, Monte Carlo generation and the like (one Tier1 centre being at CERN). Each experiment will require the resources of about five Tier1 centres plus a larger number of similar but smaller national/intranational Tier2 centres. The Panel assumed that institutional Tier3 facilities and end-user Tier4 workstations will anyway be available and so did not cost them. GRID technology was recognised as suitable for the transparent and efficient use of these distributed resources, and the need was identified for affordable networking at 1.5 - 3 Gbps per experiment by 2006.

The Software Project Panel called for joint efforts and common projects between the experiments and CERN-IT, with support for widely used packages. A matrix showing who is developing/maintaining/using which software packages illustrated the varying popularity and vulnerability of these products. Data challenges of increasing size and complexity were seen as essential steps in developing production software and CERN was asked to back the transition to OO programming. Areas of concern were the limited maturity of planning and resource estimates, insufficient development and support of simulation packages, and inadequate support for and evolution of analysis tools.

The Management & Resources panel underlined that current cost estimates are based on evolutions forecast by the PASTA committee, the technology tracking team for processors, memory, storage and architectures set up by IT Division and the LHC Computing Board. This foresees logarithmic increases in capacity/performance and decreases in unit cost with time for computer hardware, trends that are currently observed but not guaranteed to continue. The hardware costs of the initial Tier0+Tier1+Tier2 centres are estimated at 240 MCHF, a third of which is for the CERN-based Tier0+Tier1 pair. These numbers will have to be reviewed every couple of years because of uncertainties in the actual performance of the LHC machine, the detectors, triggers, backgrounds and so on. Assuming that LHC starts up in 2006 and reaches design luminosity in 2007, the investment will have to be equally spread through 2005, 2006 and 2007.

A major concern is the chronic understaffing of the teams producing core software (which is everything except physics codes and GRID middleware). Shortfalls range from 28 FTEs already in 2000, reaching 42.5 FTEs in 2002 and slightly falling to 35.5 in 2005, a problem that must be addressed by the collaborating institutes. In the same context, the Review found the planned reduction of CERN-IT staff to be incompatible with providing CERN-based LHC computing and software support.

Estimates for Maintenance and Operation of the LHC computing system are based on rolling replacement within a constant budget. About 30% of the initial investment would be needed each year, namely some 80 MCHF world wide, a sum which would include the steady evolution of capacity. Bethke noted that a similar approach at LEP had given an increase in computing capacity of a factor of 1,000 during the 15 or so years between 1985 and switch-off.

To develop the final LHC computing system, a common prototype must be set-up as joint project between the experiments and CERN-IT with the participation of some major Tier1 and Tier2 centres. By 2003/4, the prototype should reach about half the complexity (not capacity) of one LHC experiment. The estimated cost of the prototype, about 18 MCHF, is not included in the initial investment costs and so an agreement is urgently needed on how to construct and finance it.

The Review strongly recommended setting up an LHC Software and Computing Steering Committee (SC2) composed of the highest level of computing management in the experiments, CERN-IT and regional centres, to steer the development and deployment of the entire system. The SC2 should establish Technical Assessment Groups (TAGs) to launch specific tasks and projects.

Each collaboration must prepare an MoU for LHC computing that defines the overall required funding and agreed responsibilities. As an interim measure, IMoU's or software agreements should be in place by end of 2001 if possible.

The Review has shown the enormity of the LHC computing challenge and Bethke reiterated the non-negligible chance of spin-offs. It has underlined the crucial importance of proper funding, planned development and timely realisation of the entire LHC offline computing system, and its later maintenance and operation. As Bethke had earlier remarked, without adequate and appropriate computing facilities, the LHC machine and its detectors will be of little use.

#### *Discussion*

Asked by S. Bhave (IN) whether cost sharing had been addressed, Bethke replied that the Review had considered costs but not their sharing. Some countries are already setting up Tier1 centres and it will be the RRBs' job to discuss cost distribution as well as how to finance connections to countries with only small national facilities. G. Wormser (FR) remarked that the Review had done a good job defining a common hardware infrastructure. He asked whether, to mitigate the lack of manpower, there is a schedule for choosing between the several software options still open and what role CERN-IT will be able to play in the decision process. Bethke acknowledged the importance of these remarks, adding that software choices had been one of the most difficult issues faced by the Review. An attempt to streamline the use of packages had not had much success because most experiments have already invested considerable development effort based on particular packages. While many packages are used by several experiments, some are used by only one or two and Bethke concluded by saying that common attempts at streamlining had been started and must continue.