

## FOCUS

### FORUM ON COMPUTING: USERS AND SERVICES

#### **MINUTES OF THE 16<sup>th</sup> MEETING OF FOCUS HELD ON THURSDAY 2<sup>nd</sup> DECEMBER 1999**

Present: J.Altaber, T.Cass, M.Cattaneo (Secretary), B.Clare, M.Delfino, F.Etienne, F.Gagliardi, B.Gobbo, R.Gokieli, A.Grant, H.-F.Hoffmann, D.Jacobs, S.Jarp, P.Jeffreys (Chairperson), L.Mapelli, N.McCubbin, E.McIntosh, H.Meinhard, A.Norton, S.O'Neale, M.Pimiä, F.Ranjard, L.Robertson, K.Safarik, A.Silverman, P.Vande Vyvre, W.von Rueden, R.Voss

Invited: L.Cons, J.Shiers\*) (replacing A.Pfeiffer)

Apologies: G.Kellner, A.Sandoval, L.Taylor

Absent: W.Lerche, J.May, M.Mazzucato, H.Renshall, E.Valente, R.Cashmore

\*) part time

#### AGENDA:

- 1.1 Consideration of Agenda
- 1.2 Minutes of last meeting and matters arising
- 1.3 Chairman's comments
  - Status of LHC Computing review (Hans Hoffmann)
2. FOCUS highlights of the year (Chairman)
3. The Linux 2000 project (Lionel Cons)
4. COMPASS computing (Benigno Gobbo)
5. Support of compilers and operating system versions for 2000 (Jamie Shiers, for A.Pfeiffer)
6. Proposed timescales for RISC Unix decommissioning
  - Experiments' reaction (Marco Cattaneo, plus experiments representatives)
7. Update on ongoing IT activities (Manuel Delfino)
  - Overview of past year and coming year
8. Actions outstanding
9. A.O.B.
  - E.U 5<sup>th</sup> Framework Program (Hans Hoffmann)

## 1.1 CONSIDERATION OF AGENDA

Item 3 is an action item from the 14<sup>th</sup> meeting. Item 5, which had originally been foreseen after Item 6 as the two discussions are linked, had to be moved due to illness of the original speaker and the last minute availability of the replacement speaker

Hans Hoffman added an item under A.O.B. concerning the EU 5<sup>th</sup> Framework program.

The proposal for an archive policy, originally foreseen for this meeting, has been postponed to the next meeting at the request of IT Division.

## 1.2 MINUTES OF THE LAST MEETING

The minutes had already been approved by E-mail. There were no further comments.

## 1.3 CHAIRMAN'S COMMENTS

The chairman had no further comments, and asked Hans Hoffman to give an update on the organisation of the LHC computing review.

### LHC COMPUTING REVIEW (Hans Hoffman)

Hans announced the chairmanship of the review:

Review Committee Chair:                    Prof. Dr. Siegfried Bethke  
Professor Bethke has had a career in CMS before being appointed Head of the Max Plank Institute in Munich. By virtue of this he has now joined Atlas. Hans commented that it is appropriate the review should be led by an external user.

Review Committee Secretary:            David Jacobs

Hans has spoken to all experiments and noted a willingness to work together and share resources. The experiments' representatives in the steering group are currently being chosen by the experiments. IT division will be represented by its hierarchy.

The chairs of the technical panels are as follows:

Computing model:                    Denis Langlin  
Head of IN2P3 Computing Centre in Lyon

Software project:                    Matthias Kasemann  
Head of FNAL Computing Division

Resources & planning:            Mario Calvetti  
Head INFN "Gruppo 1" Commission (experimental particle physics), formerly spokesman of NA48

The first meeting between the review Chairman and the panel Chairmen is scheduled for the week of 6<sup>th</sup> December. The Review is scheduled to produce its final report by the end of next spring (21<sup>st</sup> June 2000)

## 2. FOCUS HIGHLIGHTS OF THE YEAR (P.Jeffreys) (see [slides](#))

Paul presented a review of the work of FOCUS in his first year as chairman. He classified the topics covered in three categories: topics where a decision on policy was required (about 25% of items), reviews leading to policy (50 % of topics, equally split between IT- and experiments-led) and purely informational items (25%). Paul asked if this is the right mix for FOCUS.

The main policy theme for 1999 was a review of platforms at CERN. Storage is likely to be an important topic in 2000, as well as the migration to Linux and a re-discussion of the move away from the Fortran CERNLIB.

Concerning the membership, COMPASS was added to the represented experiments, and there was new representation for Atlas. HEPCCC has suggested a review of the external membership, e.g. adding a member representing DESY (ACTION).

Paul then presented ideas on how FOCUS could be made more effective: increased user participation, closer links with COCOTIME, looking for "owners" of reviews. He also asked whether it is correct that FOCUS should concentrate mainly on hardware and only specific software support issues - should the time horizon of FOCUS be stretched to cover the startup of LHC, and thus be more closely linked with the LHC computing review?

Steve O'Neale asked that E-mail sent to the FOCUS mailing list be archived and made available (ACTION).

***FOCUS is pleased to record a good year with an interesting programme. Feedback on the presentation, in particular on the specific questions, is encouraged from members of the committee (ACTION), by E-mail to the chairman ([P.W.Jeffreys@rl.ac.uk](mailto:P.W.Jeffreys@rl.ac.uk)).***

## 3. THE LINUX 2000 PROJECT (L.Cons) (see [slides](#))

Lionel's presentation began by pointing out that Linux is already strategic at CERN, making up almost half of the ASIS clients and almost 60% of the CPUs handled by IT/PDP. The goal of the L2K project is to "enable and promote Linux on Intel based hardware as a fully functional computer platform for CERN usage" across all sectors.

L2K is intended to bring together all people working on Linux inside IT. Phase 1, targeted to finish at the end of 2000, is a fact-finding phase, reviewing the services provided by IT and studying integration of batch and desktop, interoperation of Linux and Windows, use on laptops and at home, as well as identifying users and defining their role.

The project was approved on 5<sup>th</sup> November. Current urgent work is on Red Hat 6.1 certification. An early goal is to identify what is missing and launch dedicated sub-projects, to be reviewed throughout next year.

Norman McCubbin asked what is meant by certification and how long it takes. It means making sure that all services work with the operating system version, and that the version works on all hardware. It takes several weeks, due to the diversity of software and hardware, and has external constraints such as the availability of LSF. Note that this procedure is not limited to Linux but applies to all operating systems.

Helge Meinhard asked how the interests of outside users are addressed: how can they maintain interoperability with what goes on at CERN? This is a worry, but the goal is to try to reach consensus, for example through HepiX. However HepiX is not a decision-making body. CERN will encourage e.g. FNAL to set up a similar project so that the two can be linked. Smaller universities usually look to CERN to show the way and to provide quality.

Marco Cattaneo said that a question often asked by external users is why platforms other than Intel are not considered. The answer is that CERN and even all of HEP are too small to drive the market, and these battles are best left to others such as banks and insurance companies. Currently, the savings in buying non-Intel computers would be far outweighed by personnel costs to develop, certify and maintain diverse environments. The choice of Intel should be publicised as much as possible, to avoid misunderstanding.

Helge asked what resources inside IT will be allocated to the project. Manuel Delfino said this is hard to assess now since the evaluation is not yet done. The realistic answer is that as much manpower will be devoted to this as can be released without compromising existing services to the experiments.

***FOCUS recognises that the L2K project is the conclusion of more than one year's discussion between users and IT Division, and warmly endorses the proposal. It supports limiting the duration and scope in phase 1 as suggested and reviewing progress in summer 2000 (ACTION)***

#### 4. COMPASS COMPUTING (B.Gobbo) (see [slides](#))

Benigno began by reviewing the data processing requirements of the COMPASS experiment. An important point is that many of the problems that COMPASS faces (data rate, dataset size, etc.) are similar in scale to the LHC experiments, but for an experiment that will be commissioned in 2Q 2000.

The raw data and the reconstruction will be at CERN. Analysis objects will be generated at CERN and exported to the home-labs. Monte-Carlo production will be exclusively in the home-labs. The computing farm at CERN uses standard Linux PCs; data recording is to 50GB Redwoods because of their (currently) unbeatable cost. The software is based on both FORTRAN (for the simulation) and C++ components (for reconstruction and analysis, both from ROOT and LHC++), the choice of technology being dictated by its readiness for production use.

The online system uses the DATE software from Alice, currently running on Digital Unix but soon on Linux PCs. Software development tools are taken largely from the IT SDT service, a problem is that some of these tools do not run on Linux.

Benigno concluded by thanking IT/PDP for their support, and listed three requirements:

- Periodic frozen releases of the LHC++ software

- A unique CERN software bundle (including ASIS, LHC++, ROOT etc.)

- CERN support for Linux is welcomed, but software support for other Unix dialects should continue for some years.

Steve O'Neale asked what is used for bookkeeping. This has not been addressed explicitly, as current thinking is to store the detector data in an object database that would bring bookkeeping automatically through collections, but the concept of a run remains. Martti Pimiä asked for more details about the computing model at the home-labs. The home-labs are asking for funds for Monte

Carlo farms and also for small Monarc-like farms, but currently it is planned to do only the final analysis in the home-labs.

Sverre Jarp asked whether the heterogeneous nature of the software made maintenance a problem. The software technology is in fact consistent within a given program. The biggest worry is the future of Objectivity - how could data be migrated to a possible alternative database?

Livio Mapelli pointed out that Atlas DAQ has the same problem concerning the availability of SDT on Linux - one cannot get away from Solaris. In the absence of a central Solaris service this forces the experiments to have their own Solaris support. It was pointed out that users could access the tools directly on the SDT service Solaris server (confirmed by E-mail after the meeting). Les Robertson pointed out that this is not inconsistent with having Linux as the production platform, since the development environment is not necessarily the same as the production environment.

Concerning the choice of DATE for the online system, its developers were of the opinion that C is more appropriate than C++ for the real-time environment.

***FOCUS is pleased to have a very interesting presentation from the COMPASS experiment. It is interested to note the hybrid system that COMPASS has been obliged to adopt, and recognises the role that the experiment has as a "test-bed" for the LHC experiments. It notes the request for well-defined periodic releases of LHC++, and COMPASS's thanks to IT/PDP for good support.***

#### 5. PLATFORMS AND COMPILERS FOR 2000 (J.Shiers for A.Pfeiffer) (see [slides](#))

Jamie agreed to give this talk at the last minute, due to illness of Andreas. The first slides show details of the OS and compiler versions currently supported by CERNLIB, LHC++ and GEANT 4.

Concerning new platforms for 2000, Linux RedHat 6.1 with egcs 1.1.2 is proposed. For Solaris the proposal is to use Solaris 7 with SunPro C++ V5.0. In both cases CERNLIB and LHC++ are currently under certification, GEANT 4 will follow. For DEC it is planned to upgrade to 4.0f. HP-UX is frozen at 10.20, making LHC++ support doubtful in 2000, as this version is too old for commercial packages. IBM AIX and SGI IRIX support for CERNLIB will continue, but is there any demand for LHC++ on these platforms?

The discussion concentrated on Fortran. Eric McIntosh said that a commercial compiler on Linux (the Portland Group compiler) is being tested, to address the problem of a missing public domain high performance compiler. Norman McCubbin said that there are currently no plans to move some important event generators away from Fortran. Manuel Delfino concluded that the question of the future of Fortran must be addressed (see also the discussion after point 6).

***FOCUS appreciates the clear statement made on the present status of platforms and compilers for 2000. FOCUS notes that a key issue is the interoperability of components within the libraries. It was agreed that we need to begin to determine a policy for future use and support of FORTRAN (ACTION).***

## 6. RISC UNIX DECOMMISSIONING PROPOSAL: SUMMARY OF THE EXPERIMENTS' REACTIONS (M.Cattaneo) (see [slides](#))

Marco summarised the reactions that he received from the LEP experiments, the LHC experiments, the neutrino experiments, Compass, NA48, NA49 and NA57 (see also the individual responses on the [FOCUS web site](#)).

Marco began by some definitions: "System support" means support for new OS versions and compilers. "Frozen support" means no new hardware and software (both applications and libraries) frozen at a particular OS/compiler combination, but with fixes applied if necessary to keep the system working. "Decommissioning" means getting rid of the hardware and not applying even essential fixes to the software.

The LEP experiments need to keep their shift systems (SGI+DEC) until 2003, although there is some flexibility after 2001 if a viable alternative is proven. OS and software support are needed on the current analysis platforms until 2003. The implications of "frozen support" need to be studied. There is no requirement for an HP-UX based CSF service after 2000. CERN is encouraged by the LEPC to satisfy these requirements.

Linux is well accepted in the LHC experiments, with a strong preference for Solaris as the second platform (WNT in LHCb) - these experiments say it is unacceptable to freeze Solaris support. The current frozen HP-UX support is needed in Atlas test beams (until 2002) and by NA49/NA57. The NA57 DAQ requires AIX support until 2001.

The needs of the other experiments (COMPASS, NA48 and current and future neutrino experiments) are, or soon will be, met by Linux PCs. The conclusions concerning RISC decommissioning are compatible with the requirements of the LHC and LEP experiments.

Many experiments are worried about the implications for their home-labs of decisions taken at CERN; one suggestion is to develop a model of HEP-wide software support for those platforms not maintained at CERN.

The conclusion is that the de-commissioning schedule proposed in July is unacceptable to the experiments. AIX and CSF can be stopped as planned, HP-UX, DUX and SGI must be kept as long as required by the running experiments. Solaris must continue to be fully supported.

Rudiger Voss opened the discussion asking why two platforms are needed, given the momentum of Linux. Maintaining portability of software and the availability of commercial development tools were seen as the main issues. It was argued that portability could be maintained by using a second (commercial?) compiler on Linux, while development tools could be run on any system (Solaris being the preferred platform now, but not necessarily in future). Les Robertson proposed having Linux as the only production platform, with Solaris kept only for validation purposes and not for physics production, so that it could be changed in future without argument. A counter argument was that detailed checks can only be made by processing several thousands of events. Rudiger and Hans Hoffmann closed this discussion requesting that a user group be formed to collect the arguments for the need for a second OS, for presentation at a future FOCUS meeting (ACTION)

Concerning AIX, there were no objections to the suggestion to stop AIX at the end of 2001, though this may cause some problems to some L3 home labs.

The rest of the discussion focussed on exactly what level of support is needed for HP-UX. The sum of all HP-UX desktops is much more powerful for individual L3/Opal physicists than the shift servers. The full analysis chain needs to be supported, it would be unthinkable to move this to the SGI servers which are purely data servers with no interactive environment, not even graphics licenses. A migration from HP to SGI would make no sense, Linux would be more sensible. While there are manpower problems with this (e.g. porting of the L3 event display), the financial implications of replacing hundreds of desktops are more important.

Manuel concluded that the goal of obtaining a better understanding of the experiments' needs has been achieved and that now IT division could take this input and come back in March with a reaction (ACTION). Hans supported Manuel in his intention to simplify; this means leaving things behind, and he strongly supports porting of applications to Linux.

***FOCUS thanks the experiments for their input and is pleased to observe a degree of consistency between the LEP experiments and between the LHC experiments. It notes the LEPC minutes of 7/9/99 where the committee encourages CERN to maintain the current analysis platforms during the period 2001-2003. This position is supported by the experimental returns. There is a wide statement for continuing support for Solaris, and it is proposed to form a user group to advise on requirements (ACTION). FOCUS concludes that RISC Unix will need to be supported for a longer period than originally proposed (except for AIX), and believes that IT now has the information required to formulate a proposal for decommissioning time-scales to be presented at the next FOCUS meeting (ACTION).***

#### 7. IT DIVISION IN 1999 AND PROSPECTS FOR 2000 (M.Delfino) (see [slides](#))

Manuel began by outlining the new environment that has been emerging: the exploding use of Internet applications and the migration to Intel PCs running Linux or Windows.

The main policy themes of 1999 were: streamlining of services for ease of use (e.g. only IP only over Ethernet) and cost reduction (e.g. CDR and virtual elimination of tape mounts), the transition of central services to commodity-based solutions and the minimization of RISC. Tapes remain a problem, media costs for LHC will need to be discussed. There has also been work on data management and storage, and improved understanding and co-operation with the experiments concerning software needs after CERNLIB - this could become a future theme in FOCUS. On the management side, there has been a transition to Activity Based Costing for determining the IT budget.

Strategic directions for 2000 include: the L2K and W2K projects, authentication services, LXPLUS, LXBATCHE and COMPASS commissioning, a computing Farm test-bed, technologies for the desktop, results of the LHC computing review and the next step in physics software. On the management side, a continued effort to understand real costs and prepare for the planned staff reductions, continued use of divisional projects, and a commitment to make the division an attractive service provider, open to collaboration and dialog.

Manuel then showed a pie-chart of the draft IT budget for 2000, where he highlighted the 38% which is purely services for physics (as opposed to laboratory wide) and the 8% to pay off money already spent on the structured cabling project. Physics users benefit as well from many of the general services corresponding to the balance of the budget.

He concluded by giving a preview of the new organisation of IT Division from next year, motivated by changes in the division itself, in the environment surrounding it and, very importantly, in preparation for the 30% staff reduction which must be achieved in the period 2000-2005. He highlighted the new groups: DB grouping together relational and OO databases, WIN for Windows infrastructure and support, SD for all Unix support including desktop (old PDP and DIS groups) as well as data handling, and the appropriately named API for application software (old ASD and IPT groups). He will make the names of the new group leaders available once he has announced them to the Division, but he wished to take the opportunity to acknowledge the role of the previous groups in delivering services to enable the physics activities to be achieved in record time and to thank the leadership of the previous groups. The final version of the IT organizational chart for 2000 is accessible [here](#).

## 8 ACTIONS FROM LAST MEETING

Minuted/Section	Action	Who	Status
01/07/1999 14/2	Organise discussion of Storage Area Networks at a future meeting	M.Cattaneo, P.Jeffreys	Scheduled for FOCUS 18 (8 <sup>th</sup> June 2000)
01/07/1999 14/2	Organise review of Mass Storage technologies at October 2000 meeting	M.Cattaneo, P.Jeffreys	Pending
01/07/1999 14/3	Make proposal for an archive policy	J.Richards	Postponed to FOCUS 17 (2 <sup>nd</sup> March 2000)
01/07/1999 14/3	Investigate mechanism to keep for some time archived files belonging to a deleted account	IT Division	To be part of J.Richards' proposal at FOCUS 17
01/07/1999 14/3	Organise discussion of backup and archive policies	M.Cattaneo, J.Richards	To be covered by J.Richards' presentation at FOCUS 17
01/07/1999 14/4	Look into new representation on DTF	P.Jeffreys	Ongoing
01/07/1999 14/6.1	Update estimates of maintenance cost of LEP shift systems until end 2002	L.Robertson	See below, Closed
01/07/1999 14/6.1	Organise discussion at December meeting of time-scales for ending RISC Unix support	M.Cattaneo, P.Jeffreys	See these minutes, section 6
01/07/1999 14/6.3	Organise discussion of Linux 2000 project at October meeting	M.Cattaneo, P.Jeffreys	See these minutes, section 6
08/10/1999 15/9.2	Update CERNLIB web pages with compiler/operating system combinations for "2000" release	I.McLaren	Done

Les Robertson informed the committee of the maintenance costs of the LEP experiments' shift facilities:

Aleph	27 kCHF per year
Delphi	10 kCHF per year
L3/Opal	60 kCHF per year (N.B. basic maintenance until 2003 was included in the purchase price. The sum here represents the additional cost of 24h/day, 7days/week maintenance, required until the end of next year. This replaces the 2kCHF charge per callout that existed in previous years but which is no longer available).

## 9. A.O.B.

### 9.1 E.U 5<sup>th</sup> Framework Program (Hans Hoffmann)

Following the participation of a EU official at the last HEPCCC meeting, there has been much enthusiasm at the prospect of obtaining funding under the EU 5<sup>th</sup> Framework program. Hans

stresses that requests for funds must be properly co-ordinated and approved by the directorate. The dates mentioned at HEPCCC for a bid against the 5<sup>th</sup> Framework are not realistic. Any bid or contact with the E.U. must be channelled through David Williams, and also Fabrizio Gagliardi if it is IT related.

### PENDING ACTIONS

Minuted/Section	Action	Who	Status
01/07/1999 14/2	Organise discussion of Storage Area Networks at a future meeting	M.Cattaneo, P.Jeffreys	Scheduled for FOCUS 18 (8 <sup>th</sup> June 2000)
01/07/1999 14/2	Organise review of Mass Storage technologies at October 2000 meeting	M.Cattaneo, P.Jeffreys	Planned for FOCUS 19
01/07/1999 14/3	Make proposal for an archive policy	J.Richards	Scheduled for FOCUS 17
01/07/1999 14/3	Investigate mechanism to keep for some time archived files belonging to a deleted account	IT Division	To be part of J.Richards' proposal at FOCUS 17
01/07/1999 14/3	Organise discussion of backup and archive policies	M.Cattaneo, J.Richards	To be covered by J.Richards' presentation at FOCUS 17
01/07/1999 14/4	Look into new representation on DTF	P.Jeffreys	
02/12/1999 16/2	Review external membership of FOCUS	P.Jeffreys	Planned for FOCUS 18  Review of requirements planned for FOCUS 19  Planned for FOCUS 17
02/12/1999 16/2	Create cern-focus E-mail archive	M.Cattaneo	
02/12/1999 16/2	Provide feedback to Paul on his review of FOCUS in 1999	All FOCUS members	
02/12/1999 16/3	Organise review of progress of L2K progress	M.Cattaneo, P.Jeffreys	
02/12/1999 16/4	Determine policy for future use of FORTRAN CERN libraries	IT Division, FOCUS	
02/12/1999 16/6	Prepare arguments for the need of a second OS apart from Linux, for presentation at a future FOCUS	P.Jeffreys to form user group	
02/12/1999 16/6	Proposal for RISC decommissioning time-scales	IT Division	