MIGRATING THE CERN PS CONTROL SYSTEM TO IBM WORKSTATIONS

Nicolas de METZ-NOBLAT
PS Division, CERN, CH-1211 Geneva 23, Switzerland

ABSTRACT

The workstations used within the control system of the CERN PS accelerator complex are not produced any more. We had therefore to review the software primary used as user interface and we made a port to IBM workstations.

We are also preparing the maintenance of this code for the next ten years with minimal staff. This implies a clear separation between general computing facilities, control system developments, and operation.

In order to share our experience, we will try to summarize various aspects of this migration:

- system installation principles used to speed-up error recovery time and reduce long-term maintenance costs,
- problems resulting from the coexistence of two different platforms during migration,
- software problems due to the platform and operating system changes,
- hidden dependencies on a specific manufacturer.

1. INTRODUCTION

During the last five years, the PS Division has carried out a large project to renew the control system of the CERN PS particle accelerator complex [Ref. 1 & 2].

This resulted in the introduction of a rather large number of UNIX workstations and X-terminals mainly used for the daily operation and the exploitation of the CERN PS complex.

Our original systemused DEC workstations based on 32 bit MIPS microprocessors. The new generation of DEC workstations is based on a new 64 bit chip and a new operating system. The amount of work required to port our control software to this new generation was almost the same as that necessary to port it to any other UNIX workstation. After detailed cost analysis, we decided on IBM RS/6000 workstations based on PowerPC chips and running the AIX operating system.

The introduction of a new generation of workstations is also a good time to review organizational details in order to minimize the staff required for the long-term maintenance of such computing resources.

2. NETWORK TOPOLOGY [Fig. 1]

We have been introducing a large number of workstations (>100) and diskless front-end computers (>100), initially to develop a new control system and then to operate and maintain it. The problem was to keep this maintainable.

We split the equipment into two different networks isolated by a router:

- An office network: This covers development and test equipment
- A control network: This concentrates the equipment needed for normal operation of the complex

Then we split this control network into various subnets, at present isolated by bridges:

- a central control subnet on which we keep some critical services like central timing generator or databases,
- one subnet for each accelerator (or group of accelerators):

- one for the proton and the ions LINACs,
- one for the LEP Pre-Injector complex (LIL and EPA),
- one for the PS Booster (PSB),
- one for the CERN Proton-Synchrotron itself and its transfer lines (CPS).

We organized each subnet to still allow a minimum of operation from a local control room, even when disconnected from the central control subnet. Diskless computers rely on a single file server located on the same subnet to recover after a power-cut.

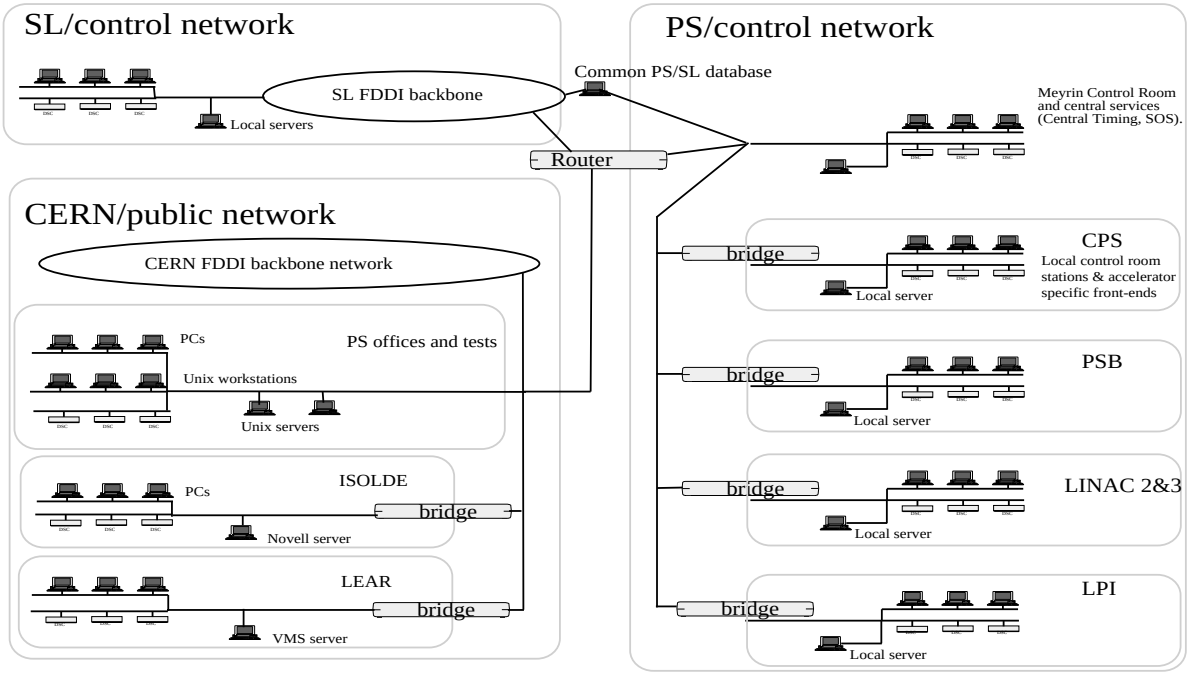

Fig 1: The network topology

3. FILE SYSTEMS ORGANIZATION

We had to decide early on a well-structured file organization that could stay stable during the whole project.

We decided on the following organization:

- One file system per accelerator: This receives all programs required for the minimum operation and all software specific to this accelerator, including software running in diskless front-ends.
- One file system for central services: This receives programs that are not critical for the minimum operation and used for the control of more than one accelerator..
- One file system for the control software source repository: this receives sources of all the programs developed for the purpose of the accelerator control. Its structure reflects the target file systems in order to ease long-term software maintenance.
- One file system to receive libraries and associated include files for each different target. This file system also receives specific programs required either to produce the applications or to maintain operational data.
- One file system to receive Word-Wide Web documentation used for on-line help.

The CERN central computing service furnishes other resources not used for normal operation:

- Home directories can be located on central servers,
- Public domain software must continue to be maintained centrally,
- Whenever possible, commercial products must also be centrally maintained.

4. MIGRATION STRATEGY

Before deciding on new UNIX workstations, we did a preliminary inventory of all our software and our dependence on specific features. Nearly 30 different programmers had developed about 100 applications. Then, we ported a significant application to various platforms and evaluated the constraints.

The next step was to start the porting with a small number of workstations and servers in order to prepare a validation of almost all applications in a real context co-existing with the current environment [Ref. 3].

We had to minimize the staff required and interference with on-going developments and normal beam production. We prepared an independent environment to receive sources and libraries and a clone of existing operational file systems. This allowed us to prepare for the final software distribution and to execute applications in a realistic environment, keeping the same data source database for software customizations such as program menus and equipment lists.

The next step (March 1996) will reverse the situation between old and new workstations in order to avoid the simultaneous maintenance of two binaries for the same applications.

The final step will be to reduce the platform diversity in order to reduce system management complexity and to ease new application development.

5. PROBLEMS ENCOUNTERED DURING THE MIGRATION

While developing our applications, we tried to avoid the use of manufacturer-specific extensions to standards and the migration was a good occasion to verify this point.

We did not encounter any major problem while porting almost all the software to the new platform.

Some problems originated from librarys updates for such as standard Motif and X11, or from our own libraries. With a new development environment on both platforms, it was possible to isolate and cleanly these problems.

Other problems were due to dependencies on the environment:

- Some programs were using specific fonts.
- Simultaneous use of various UNIX variant compatibility routines (Berkeley Software Distribution, System V and POSIX) does not always mix properly.
- Our front-ends already had a different byte order and most programs were already ready for such a change.
- C compilers and system include files differ, but without any major incompatibility.
- We tried to avoid as most as possible the use of Display Postscript extension, except for two applications.
- One application used a specific external video input extension and multimedia libraries. We will need to adapt it and hope that with current developments in the multimedia area some standards will be available for this field.

The migration itself was an occasion to locate small programming mistakes, or surprising limitations introduced by the programmer. A typical example was the limitation to six characters of a host name!

We also had to prepare the future development environment. One example is that we had to decide on a new User Interface Builder program.

6 OTHERS SYSTEM CONSIDERATIONS

We are trying to prepare long-term exploitation solutions:

- We try to use the same system installation procedures as other CERN central UNIX services with minimum specific customization, in order to share experience and support.
- We prepare all critical computers for remote exploitation. This applies to servers and to front-ends.
- Our network layout is clean, easy to understand and is organized to allow longer intervention delays on non-critical equipment.
- We keep coherent contents within file systems and have automated backups.

7. CONCLUSION

This exercise of porting our control software was very positive. It was an occasion to evaluate in details our dependencies, and to check its completeness. It demonstrated the ability of a UNIX system to maintain a large amount of software written in C. However, it is very important to choose emerging standards and to avoid any specific extensions to these.

During the next year we expect to reach a very clean and stable situation with workstations, servers and control application software. That might open the way to some out-sourcing solutions of our problems as we have already experienced with the system management of the LEAR control system.

ACKNOWLEDGMENTS

I would like to thank Michel Arruat, Franck di Maio, Serguei Moukhine, Orel Oleg and Fabien Perriollat for their great help in achieving this work within the planned timetable. I would like also to thanks all applications developers for their collaboration and the various people who have been testing the resulting environment.

REFERENCES

[1] F.Perriollat, Ch.Serre, “The new CERN PS control system, Overview and status”, ICALEPCS, Berlin, Germany, October 18-23, 1993, Nucl. Instr. And Meth. A352 (1994) 86.

[2] R.Rausch, Ch.Serre (editors), "Common Control System for the CERN Accelerators", Proceedings ICALEPCS, Tsukuba, Japan, 11-15 November 1991, p.54-58
The PS and SL controls groups, "PS/SL Controls Consolidation Project", Technical Report, CERN PS/91-09(CO), CERN SL/91-12(CO), April 91.

[3] M.Boutheon, F.DiMaio, A.Pace, "General Man-Machine Interface used in Accelerator Controls; Some Applications in CERN-PS Control System Rejuvenation", Proceedings ICALEPCS, Tsukuba, Japan, 11-15 November 1991, p.452-455