

The New Generation of PowerPC VMEbus Front End Computers for the CERN SPS and LEP Accelerators Control System

A. Bland, P.Charrue, F. Ghinet, P.Ribeiro, M. Vanden Eynden
SPS and LEP Controls Group
European Laboratory for Particle Physics
CH1211 Geneva 23 - Switzerland

ABSTRACT

The CERN SPS and LEP PowerPC project is aimed at introducing a new generation of PowerPC VMEbus processor modules running the LynxOS real-time operating system. This new generation of front-end computers using the state-of-the-art microprocessor technology will first replace the obsolete XENIX PC based systems (about 140 installations) successfully used since 1988 to control the LEP accelerator. The major issues addressed in the scope of this large scale project are the technical specification for the new PowerPC technology, the re-engineering aspects, the interfaces with other CERN-wide projects, and the set up of a development environment. This project also offers support for other major SPS and LEP projects interested in the PowerPC microprocessor technology.

INTRODUCTION

The mission of the European Laboratory for Particle Physics is pure science - particle physics - probing the innermost constituents of matter to find out how our world and the whole of the Universe works. To do this, it operates a number of particle accelerators, the largest of which is the LEP electron-positron collider. Beams of electrons and their antimatter counter-parts, positrons, are circulating in a 27 kilometer underground ring at nearly the speed of light and collided together. A chain of accelerators, including the Super Proton Synchrotron (SPS), provides beams for LEP. The LEP machine consists of various equipment installed in the 27 km ring tunnel and in 8 surface buildings, their associated pits and underground areas. The SPS and LEP Control System extends from a central control room to all the surface buildings and underground areas in order to provide command and acquisition facilities for the collider and the particle beams.

THE CURRENT SPS AND LEP CONTROL SYSTEM ARCHITECTURE.

The architecture of the SPS and LEP control system is modeled on three layers [1] [2]: the Control Room Layer, the Front End Computing Layer, and the Equipment Control Layer.

The Control Room Layer

The Control Room Layer is composed of modern UNIX workstations, servers and X-Terminals. Human Computer Interfaces are developed using an X-Window and OSF/Motif User Interface Management System. Several other UNIX servers are used for file storage, for public displays, for the management of alarms and for an ORACLE on-line relational database. The current computer installation for this layer is the following :

- 30 Hewlett Packard 700 series workstations
- ~100 X-Terminals
- 7 DEC workstations

The Front End Layer

The Front End Computing Layer consists of Front End process computers (FEs) based on PCs and VMEbus crates. Their main function is to provide a uniform interface to the equipment as seen from the workstations and act as data concentrators for equipment interfaced via various fieldbuses. The task assignment between FEs is made on geographical or functional criteria. The communication between the Human Computer Interfaces (HCI) running in the SPS and LEP control room and the FEs is achieved through Remote Procedure Calls (RPC) or direct TCP/IP socket connections. The current computer installation for this second layer is as follows:

- 150 OS/9 VMEbus Systems
- 110 LynxOS 2.2.1 80486 PC systems
- 45 LynxOS 2.1.0 VMEbus 680x0 systems
- 150 XENIX 2.3.1 80386 PC systems

The Equipment Layer

The Equipment Control Layer consists of Equipment Control Assemblies (ECAs) connected to the FEs via various equipment fieldbuses (1553, GPIB, BITBUS, JBUS) or via RS232/422 links. The ECAs range from G64 6809 systems running FLEX to 3U and 6U VME 68k systems running OS-9. The accelerator equipment is distributed in underground halls and in surface buildings.

The SPS-LEP Network .

Network communication is made by local Ethernet segments bridged to large Token-Rings, one for the LEP general services linking all the surface buildings and others for the accelerators. These Token-Rings have been partially replaced by 100 Mbit FDDI backbones which will eventually cover the entire CERN site. The data distribution is based on the TCP/IP protocol suite and the network management is achieved with SNMP.

CURRENT LEP FRONT ENDS ARCHITECTURE

The LEP front ends were installed in 1988 to control most of the LEP accelerator services (i.e. vacuum system, tunnel cooling and ventilation, electricity, water distribution, magnets, and so forth). They are Olivetti 386-16Mhz PCs running SCO-XENIX v2.3.1. They are disk based, i.e. each has the complete XENIX installation on its 130 Mb disk. They are connected to the network via a Token-Ring card and have a CERN-made 1553 connection to control the 1553 ECAs. The XENIX software drivers for Token Ring and 1553 have been developed at CERN. These computers run on a 24h/day and 365 day/year basis

These front ends run well now since we no longer add or modify software and/or hardware on a large scale. But as we did not follow system updates and we now have TCP/IP and NFS on our major platforms, we definitively need to upgrade these front ends. The Olivetti 386-16Mhz PCs are not available anymore due to their obsolescence and we are organizing the maintenance via a stock of spares. Our main hardware problems were local disk failures, so we renewed many of the disks in 1992. Now we start to experience Power Supply failures. On the software side, the distribution of system files to every hard disk requires complex management. In addition, the lack of standard TCP/IP and NFS now makes us reluctant to develop any further on these platforms.

PROJECT MOTIVATION

The first motivation is that we want to use the VMEbus standard for the LEP Front Ends and to use a VME CPU card based on a powerful processor. The main interests of the VME standard are its big and flexible I/O address space and its good interrupt capabilities. The second important motivation is the desire to use the Network File System (NFS) and the LynxOS operating system in all our FEs to simplify our maintenance effort, since today several generic software packages, like the one used for Remote Procedure Call, have to be maintained on too many different operating systems. But the interest is not only limited to these aspects; several major rejuvenation projects dealing with the SPS power converters (ROCS project) and the LEP Beam Synchronization Timing (BST project) have high performance requirements and will benefit from the new proposed technology.

The **two objectives** of the project are :

- the replacement of the operational PC XENIX systems (70 installations) by VMEbus systems running the LynxOS operating system. The old PC XENIX console functionality will be replaced by already-installed X-Terminals connected to the Hewlett Packard workstations.
- the setup of a development environment for the SPS and LEP VMEbus users. This environment shall provide the necessary tools to migrate the existing PC XENIX operational software (more than 100 software modules) and to perform low level development at the operating system level (i.e. drivers and libraries).

MARKET SURVEY AND TENDERING

The choice of the new VMEbus platform was made in two stages following the CERN purchasing rules. The first step was to make a large market survey to get enough information on what was available in industry to write and submit an official tender for the supply of a VMEbus processor module.

The market survey [3] was designed to detect which was the most appropriate VMEbus processor card running LynxOS with an ethernet connection. The result of this preliminary enquiry was used to write the official invitation to tender for the supply of these VMEbus CPU boards. The market survey was sent to 82 companies in 12 countries in June 1994. We received 22 replies from 5 countries in August 1994 which were proposing 22 different CPU boards. The most popular processor proposed in these commercial offers was the PowerPC family: 2 proposals for the 601, 6 for the 603 and 3 for 604. As we felt that the PowerPC family [4] and in particular the 603/604 processors will be positioned as the high performance successors to the 680x0 family, we decided to focus the invitation to tender on the PowerPC technology.

The invitation to tender [5] was sent to 12 companies from 5 countries in October 1994. It stated that CERN asks for a PowerPC processor implemented in the VMEbus Standard running the LynxOS Real-Time Operating System. The processor module must be a manufacturer's standard product nearest to the specifications. The main requirements that were laid down for the Processor module were:

- PowerPC 603 easily upgradable to 604.
- 32Mbyte of RAM
- 512Kbytes of flash EPROM.
- 32-bit DMA interface to standard Ethernet IEEE 802.3.
- A16-A32/D8-D64 bit wide VMEbus interface.
- One VMEbus slot wide for up to 64Mbytes DRAM memory
- One or two PCI mezzanine slots.
- Standard LynxOS from LynxRTS for execution in diskless and network bootable.
- TCP/IP, NFS plus standard LynxOS drivers (VMEbus, SCSI, Ramdisk, ...).

The tendering exercise resulted in 8 replies. The successful bidder was the CES company [6] of Geneva, Switzerland, and the VME CPUboard proposed was the RIO2-8060 based on a 603-64Mhz with 32MB DRAM, 4MB Flash Eprom, 10baseT ethernet connection and 2 PCI mezzanine slots. The datasheet of this board specified 65 SPECint92 and 60 SPECfp92 for a 603/64Mhz and 192 SPECint92 and 198 SPECfp92 for a 604/96Mhz.

NEW FRONT END ARCHITECTURE

The new generation of FEs will consist of a standard VMEbus crate, a VMEbus processor module based on the PowerPC family, various fieldbus controller modules and equipment dedicated I/O modules. LynxOS is the chosen operating system for the new FEs.

The LynxOS installation will also include :

- host-based TCP/IP networking software for Ethernet;
- Network File System (NFS), client and server;
- NetBoot firmware for bootstrap loading LynxOS systems over the network;
- LynxOS drivers for VMEbus, Ethernet, SCSI-II, and RAMDISK.

The connection of the FEs to the local Ethernet segment will be made directly from the processor module. Connections to the future FDDI network will be made via a separate Data-Link module tightly coupled to the processor via a dedicated PCI interface card, conforming to the PMC specification.

Software Development Environment

As the new generation of FEs will be diskless, bootstrap load files are prepared on development systems, stored on a boot server and loaded into target systems over the network. Application programs will be compiled, linked and debugged with the GNU software (gcc, g++, gdb). This development activity will be possible either from PowerPC LynxOS native development systems, or by using cross-development facilities running on an IBM AIX workstation. This second option offers to the developers the comfort and performance of a classical workstation development environment. The binaries produced will then be stored on a NFS file server from where they will be loaded onto the target systems.

PROJECT MANAGEMENT

Preliminary Considerations

At the beginning of the project it was decided to use mechanisms to monitor progress, generate milestones and reduce the risk linked to the technology involved in the project. Various aspects such as:

- the project scale (140 XENIX systems)
- the new PowerPC technology
- the different categories of users (application programs developers but also low level operating system specialists)
- the interface with major projects like ROCS and BST
- the re-engineering aspects
- the operational constraints (continuous LEP accelerator operation required)

led us to use software and hardware engineering standards to manage the project. The IEEE software engineering standards [7] and the European Space Agency (ESA) PPS-05 standards [8] were used to address both technical and managerial project issues.

Project lifecycle

The project organizational structure, the organizational boundaries and interfaces and the individual responsibilities are specified in the project management plan (PMP) document. The project lifecycle is divided in six phases : the user requirements phase, the system requirements phase, the architectural design phase, the detailed design phase, the transfer into operation phase, and the operation and maintenance phase.

USER REQUIREMENTS PHASE

This is probably the most crucial phase of the entire project. The objectives of this phase are to identify all entities involved in the project (this activity was mandatory due to the lack of documentation concerning the existing software modules) and capture the requirements for the future operational and development systems. Several technological evolutions (i.e. cross-development environments) as well as the experience gained during the past seven years with the XENIX systems raised new user requirements. After being published, the User Requirements Document (URD) was reviewed and accepted as the baseline document for the next project phases.

PROJECT PROTOTYPING ACTIVITY

According to the technical issues involved in the project, it was decided during the system requirements phase of the project to spend time on prototyping. As shown in figure 1, this prototyping activity is aimed at :

- improving the definition of the user requirements. Building a prototype may lead to new requirements or to the discovery of incomplete or contradictory requirements;
- gaining knowledge about the technical feasibility, the system requirements, and any aspect related to the operation of such a system;
- producing a model for the final development and operational systems.




Figure 1: Project Prototyping Phase.

PROJECT DESIGN PHASES

The architectural and detailed design phases are aimed at defining a collection of software and hardware components and their interfaces and building the final system. The installation of the development environment, the writing of the System User Manual (SUM) and the migration of the XENIX software modules will take place during the detailed design phase.

CONCLUSIONS

At present we are tackling the prototyping phase of the project. The acceptance tests for the first PowerPC 603 processor board have been successful and the bulk delivery will start before the end of 1995 and will continue during the first quarter 1996. The prototyping activities will be developed in two directions: several target systems will be installed in parallel with the existing XENIX PCs to check their behavior under operational conditions and, at the same time, developers will try to migrate some of the XENIX software modules, using either the new cross-development environment or native LynxOS development systems, depending on the nature of that software.

The architectural and detailed design phases will start before December 1995 and the timetable for the transfer into operation has still to be negotiated.

The new FEs are expected to provide a gain by a factor of ten to twenty in computing power compared with the old XENIX FEs.

REFERENCES

- [1] SPS and LEP Controls, Status and Evolution Towards the LHC Era - This conference - R. Lauckner, R. Rausch - CERN
- [2] Interfacing Industrial Equipment to CERN's Accelerators and Services Control System - P.Anderssen, P.Charrue, R.Lauckner, P.Lienard, R.Rausch, M.Tyrrell, M.Vanden Eynden
- [3] Market Survey for the supply of processors in VMEbus standard with LynxOS real-time operating system software - CERN MS-2285/SL
- [4] PowerPC News home page - <http://power.globalnews.com/ppchome.htm>
- [5] Specification for the Supply of PowerPC VMEbus Processor Modules Running LynxOS Real-Time Operating System Software - CERN IT-2285/SL - SL/Tech.Spec.94-05-CO
- [6] Welcome to CES 's WWW Server - http://www.ces.ch/CES_info/Welcome.html
- [7] Software Engineering - IEEE Standards Collection - Institute of Electrical and Electronics Engineers, Inc
- [8] Software Engineering Standards - Prentice Hall - C.Mazza, J.Fairclough, B.Melton, D.De Pablo, A.Scheffer, R.Stevens