

Linux for the CLEO-c Online system

H. Schwarthoff, T. Wilksen, Cornell University, Ithaca, NY 14853, USA

Abstract

The CLEO collaboration at the Cornell Electron Positron Storage Ring CESR has completed its transition to the CLEO-c experiment. This new program contains a wide array of physics studies of e^+e^- collisions at center of mass energies between 3 GeV and 5 GeV.

New challenges await the CLEO-c online computing system, as the trigger rates are expected to rise from less than 100 Hz to around 300 Hz at the J/Ψ production threshold, with a moderate increase in data throughput requirements. While the current Solaris and VxWorks based readout system will perform adequately under those conditions, there is a desire to improve the performance of the central components to extend monitoring capabilities and provide larger safety margins. The solution, as in most modern particle detector systems, is to deploy Linux on Intel architecture computers for the performance critical applications. For reasons of hardware and software availability, the existing CLEO online and offline computing environment has been ported to the Linux platform. This development allows the described challenge to be met.

In this presentation, we will report on our experiences adapting the CLEO online computing system for operation under Linux. Issues regarding third party software and code portability will be addressed. Performance measurements will be presented.

THE CLEO EXPERIMENT

For over 20 years, the 10.6 GeV e^+e^- collider CESR at Cornell University in Ithaca, New York, has served as an outstanding particle physics research facility. Its unique detector, constructed and operated by the CLEO collaboration, has been able to provide a wealth of results from its focus on B meson studies and a variety of other topics such as charm and tau physics ([1]). The storage ring has recently been converted to a charm factory (CESR-c,[2]), adapted for center of mass energies between 3 GeV and 5 GeV. It is now serving as a new research tool for precision charm quark measurements, searches for new physics, and tau lepton studies. It is also expected to provide an array of Quantum Chromodynamics (QCD) results that permit comparison to the most recent generation of lattice QCD calculations with unprecedented precision.

In parallel, the CLEO detector, now called CLEO-c, has been modified to allow operation at those lower energies.

CLEO ONLINE COMPUTING

The current system

The current CLEO online computing system was developed for CLEO III, which ended in 2003. It is based on hardware readout through 31 VxWorks PowerPCs, and a central system comprised of three Sun Solaris 8 servers (see figure 1).

All computers are networked using standard 100 MBit/s Ethernet over copper cable connections; the path from the central network switch to the event builder workstation is a 1 GBit/s link. 5 Windows 2000 PCs serve as operator consoles, and 3 additional Windows PCs provide hardware controller capabilities for the CLEO-c gas and cooling systems.

The CLEO-c event size is well below 20 kBytes, with a readout rate of up to 100 Hz (currently).

CLEO-c online software

All slow control and central event building programs are C++ applications, while the graphical user interfaces (GUIs) have been implemented in Java ([3, 4]). The connecting fabric is homogeneously carried out using CORBA middleware ([6, 7]). The event builder stores all recorded event data on five spooling hard disks. From there, a dedicated manager process using the Veritas Media Manager software orchestrates the storage of two independent copies on AIT tape cartridges, located in a tape robot system.

A remote shifting setup has been developed, which allows CLEO collaboration members who reside outside the area to conduct CLEO shifts ([8]).

Requirements for future CLEO-c operation

At a later stage of CLEO-c, the detector will be run at a center of mass energy of 3.1 GeV, the J/Ψ production threshold. In this domain, we expect a production cross section of approximately $1 \mu b$, which determines the maximum expected trigger rates to be approximately 300 Hz ([2]).

While the current, Solaris 8 based, configuration will likely be sufficient to fulfill these requirements, a higher level of performance is desirable:

- The current setup will reach its limits, especially since the central event building node, an Ultra Sparc III workstation, has a clock frequency of only 1.015 GHz.

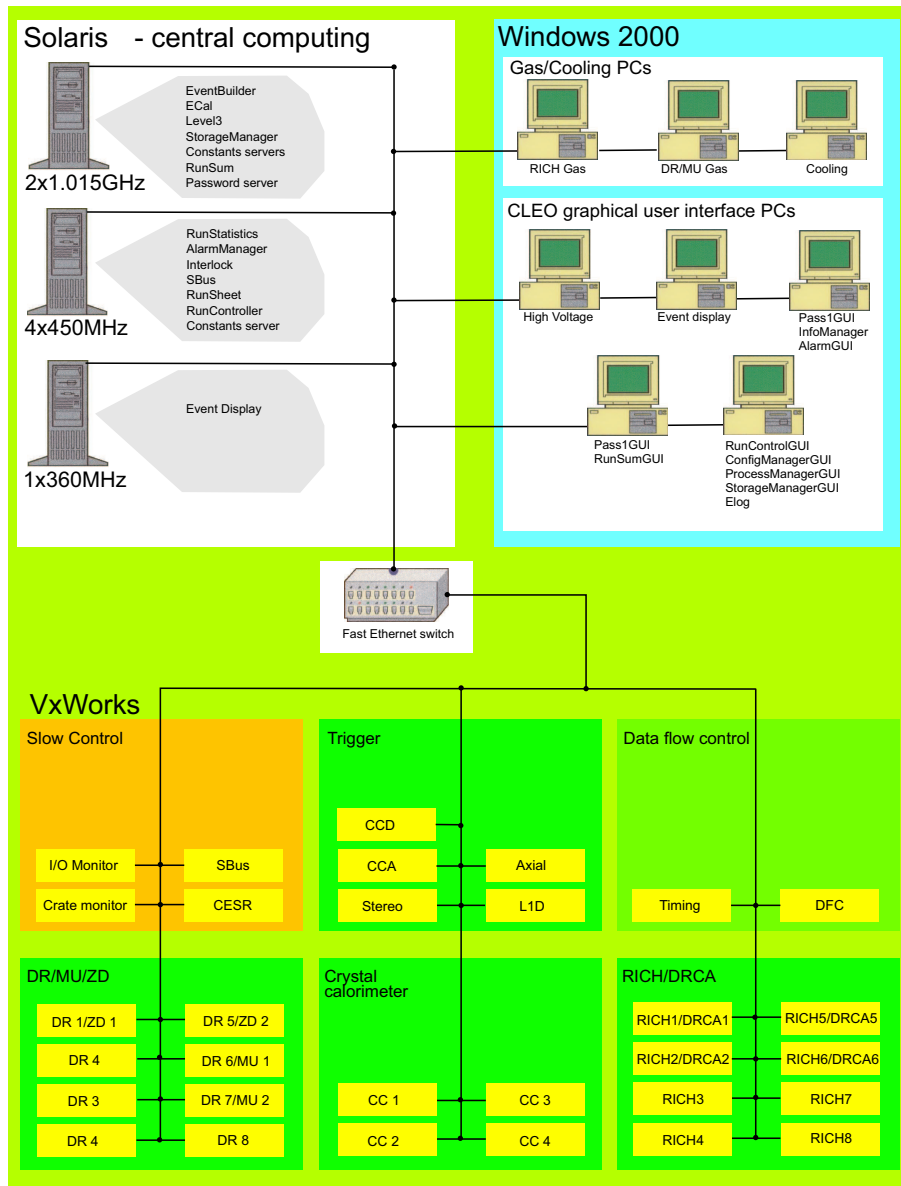


Figure 1: The CLEO-c online computing system

- A powerful monitoring tool is the *Pass1* application. This program performs basic data checks and simple track reconstruction to allow a high degree of real-time monitoring of detector functionality. Due to the complexity of this task, and because it is not capable of parallel processing, *Pass1* is always CPU performance limited and is only able to process a fraction of all events in real time. At the same time, its monitoring quality rises with the event processing fraction.

THE LINUX ALTERNATIVE

While the CLEO software system was historically OSF/Alpha based, the CLEO III development became strongly centered around Solaris, due to the organization around the Objectivity event database. Following the emer-

gence of the Linux on Intel architecture platform, the high energy physics community has adopted Linux as a viable alternative to other Unix platforms, as it provided an excellent price-performance ratio and wide reaching support from the open source community.

The CLEO offline software Linux port

To benefit from this development, the CLEO generic and offline software has been ported to Red Hat Linux 9 ([9]). Even though a number of obstacles had to be overcome, the collaboration is now able to enjoy the advantages: The full CLEO-c event reconstruction and analysis software can now be run on Linux.

The following issues were relevant during the porting effort:

- The main effort was to attain code compatibility with the Linux GNU compiler. Since a large number of packages had been developed by various physicists on OSF and Solaris, using the respective native compilers, they frequently did not comply with the stricter GNU compiler rules.
- A new, high performance, event store system ([10, 11]) was developed. It does not depend on commercial software and is highly flexible with input/output formats.
- CLEO-c is in the process of adding more and more Linux nodes to its analysis farms, as replacement for the existing Solaris hardware.
- At the time of this conference, none of the tape robot controllers used for CLEO-c data analysis have been moved to Linux.

The CLEO online software Linux port

The above-described port made the conversion of the CLEO online software possible.

- As this system had been developed by a small number of programmers, there were few compiler compatibility issues to resolve.
- The main effort proved to be related to the third party software packages used in the online environment. Notably, the deployment of VisiBroker 6 on Red Hat Linux 9 was delayed for several months, simply due to the fact that the vendor had complicated its licensing scheme considerably.
- The fact that the Linux operating system uses little endian byte ordering, as opposed to the other platforms involved in the binary data path, required some code adaptation.

Performance comparison between Solaris and Linux

As benchmark for performance comparisons we chose the CLEO online *Pass1* application, described above. This program is a collection of C++ and Fortran modules, most of which are routinely used in CLEO offline reconstruction and analysis. During regular data taking operation, the performance of online Pass1 is usually limited by the available CPU power. Table 1 shows a performance comparison between our old, Solaris 8 based, configuration and a Linux test setup.

As the comparison shows, the performance gain on Linux is greater than what we would expect by just scaling the CPU clock speed.

Table 1: Performance: Solaris vs. Linux. The processing rates are determined from wall clock times and are not results of rigorous benchmarking.

	Solaris	Linux
Hardware	UltraSparc III	Intel Pentium 4
Clock frequency	1.015 GHz	2.4 GHz
Operating system	Solaris 8	Red Hat Linux 9
Processing rate	5.2 Hz	17.2 Hz

Future plans

A number of Linux tests have been completed successfully:

- All Java GUIs can be operated on Linux without restrictions.
- The central data readout path on Linux is operational.
- All central slow control processes, except for the tape robot managers, have been tested successfully.
- The high voltage control application is being tested at this time during routine running.
- Tape robot controlling has not been deployed on Linux, due to third party licensing issues.

Except for the mentioned mass storage setup, we are now able to upgrade the complete central CLEO-c online computing system to Linux. We expect to complete the migration in the coming months, after resolving any remaining challenges related to the tape robot system.

The fact that third party software licensing has recently become a major nuisance is another blow to the attempt of many high energy physics experiments to deploy commercial packages instead of developing their own solutions. The current circumstances force us to take a renewed look at open source packages and custom developments.

ACKNOWLEDGMENTS

This work has been supported by the US Department of Energy, the National Science Foundation, and the Alexander von Humboldt foundation.

REFERENCES

- [1] *The CLEO III detector, design and physics goals*. CLNS Report 94/1277.
- [2] For a summary, see D.G. Cassel, Cornell University, LEPP, *CLEO-c and CESR-c: A NEW FRONTIER OF WEAK AND STRONG INTERACTIONS*, proceedings of the International Conference On The Structure And Interactions Of The Photon And 14th International Workshop On Photon-Photon Collisions (Photon 2001), 301-308.

- [3] *On the way to maturity - The CLEO III Data Acquisition and Control System*, presentation at Computing in High Energy and Nuclear Physics 2001, Beijing, China.
- [4] *Object Database for Constants: The common CLEO Online and Offline solution*, presentation at Computing in High Energy and Nuclear Physics 2003, La Jolla, CA, USA.
- [5] Objectivity object database, by Objectivity Inc., Version 7. See <http://www.objectivity.com>.
- [6] CORBA (Common Object Request Broker Architecture) is an industry standard for remote method invocation, see <http://www.omg.org>.
- [7] VisiBroker by Borland, Versions 3.2.2 (VxWorks), 5.1 (Solaris and Windows), 6.0 (Linux). References can be found on <http://www.borland.com/visibroker>.
- [8] T. Wilksen, H. Schwarthoff, *Remote Shifting at the CLEO Experiment*, presentation at this conference, No. 313.
- [9] V. Kuznetsov, *Porting CLEO software to Linux*, presentation at this conference, No. 183.
- [10] C. Jones, *Controlling Reconstruction using a Fault-tolerant, Scalable Web Service*, presentation at this conference, No. 198.
- [11] C. Jones, *EventStore: Managing Event Versioning and Data Partitioning using Legacy Data Formats*, presentation at this conference, No. 199.