SWITCHING THE JLAB ACCELERATOR OPERATIONS ENVIRONMENT FROM AN HP-UX UNIX-BASED TO A PC/LINUX-BASED ENVIRONMENT

T. McGuckin, Jefferson Lab, VA 23606, U.S.A.

Abstract

The Jefferson Lab Accelerator Controls Environment (ACE) was predominantly based on the HP-UX Unix platform from 1987 through the summer of 2004. During this period the Accelerator Machine Control Center (MCC) underwent a major renovation which included introducing Redhat Enterprise Linux machines, first as specialized process servers and then gradually as general login servers. As computer programs and scripts required to run the Accelerator were modified, and inherent problems with the HP-UX platform compounded, more development tools became available for use with Linux and the MCC began to be converted over. In May 2008 the last HP-UX Unix login machine was removed from the MCC, leaving only a few Unix-based remote-login servers still available. This presentation will explore the process of converting an operational Control Room environment from the HP-UX to Linux platform as well as the many hurdles that had to be overcome throughout the transition period. It will conclude with a current assessment of the change-over status as well as an examination of what future steps will complete the project.

HISTORICAL CONTROL ROOM CONFIGURATION

Prior to the upgrade of the MCC Control Room the configuration of the computer systems for running the Accelerator was based around a centralized-model. A single, large server (opsrv) provided virtually all services for the Operations subnet (OPS-subnet), including acting as a boot host, web server, file server compiler and login server. This server was an (aging) HP k370 that had gone through multiple upgrades over the years and had essentially reached the limits of its capabilities. It was also becoming very expensive (in both money and time) to maintain as its aging hardware began to fail more and more.

The workstations in the Control Room that were used by the Operations staff to run the Accelerator were HP B2000 workstations with PA-RISC2 450Mhz processors running HP-UX 11.11. The machines had also undergone several upgrade iterations (including an OS-upgrade from 10.2 to 11.11) and had largely reached their limits of memory and processing capacity.

In Aug. of 2004 a large-scale upgrade of the MCC Control Room took place and it was decided at this time to re-evaluate the computer systems in the Control Room and the servers supporting them to find a better solution.

CHOOSING REDHAT LINUX

Reasons and Goals

In deciding on a new architecture for the MCC controls network and the Control Room, several factors were considered:

- HP-UX had become too expensive and difficult an operating system to maintain, another choice had to be found.
- Linux was a cost-effective, widely-used alternative that several users and developers had already adopted.
- Switching over to Linux would allow for utilization of many inherent capabilities of more modern PC architectures (sound support, multiple monitor displays, more extensive driver support, etc.)
- Broad availability of easily-supported Open Source utilities (such as OpenOffice, Firefox, Thunderbird, etc.)
- Orders of magnitude increase in processing power was possible by going from older HP-machines (450 MHz PA-RISC2 processor to Quad-core 2.4GHz Intel processors).
- Full EPICS-support had become available for Linux.

For these reasons it was decided that Redhat Linux (specifically Redhat Enterprise Linux) would be the chosen architecture for the upgrade. The subsequent goals for the upgrade were identified:

- First and foremost to replace the aging hardware and software on the OPS-subnet to more modern workstations and servers.
- To move from a centralized-model of computing services to a more distributed-model, with several smaller, faster, more cost-effective servers replacing the one large server (opsrv) used previously.

It was also critical that all existing tools used to operate the Accelerator either be supportable under the new architecture, be replaced with updated versions, or be left running on legacy machines.

Finally it was critical that the upgrade have a minimal impact on Accelerator operations and resulted in as close to zero loss of beam-time as possible.

Redhat Enterprise Linux (RHEL)

It was decided to go to RHEL as the chosen architecture over another “flavor” of Linux (such as Ubuntu or Fedora) for several reasons:

- An enterprise solution offered stable versioning control and guaranteed patches, as well as a stable means of receiving and installing new patches via Redhat's Satellite Server.
- Several users across the Accelerator-site (and lab-wide) had already begun using various versions of Redhat Linux, and it had proven to be an excellent platform for them.
- RHEL was supportable under the Lab's existing purchasing agreement with Dell Computers and could be installed on an identical system to what was being purchased for the Accelerator's Windows XP user machines.

Cost Analysis

The final reason for the switch-over from HP-UX to RHEL was a simple cost-analysis that was done both of initial system purchase and maintenance for existing systems (see Table 1).

Table 1: Cost Analysis of HP-UX vs. Linux

System	Cost (HP-UX)	Cost (RHEL)
Server	~$40,000	~$4000 x 6 machines
Workstation configuration	~$8000	~$3000
Maintenance	~$80,000/year	~$4000/year

HURDLES

Switching the entire OPS-subnet over from one computer architecture to another was a major undertaking. As part of the process several categories of hurdles had to be overcome.

Programming Hurdles

Although the scope of the project was very large, much of the ground work for the programmatic changes that would be required had already been setup. As result very few (5 of 100) programs had to be re-written to be compatible with the architecture and only a small portion (24 of 100) had to even be recompiled. The vast majority of programs could be installed on the new architecture with no changes required at all.

There were however some applications that simply could not be converted over for various reasons:

- Some hardware (such as the OPS large-format display wall) could not be setup to run under Linux.
- Some services (such as NFS, NIS and DNS) ran better/more robustly on other architectures (in this case Solaris 10).
- Some legacy software only ran on other OSs (some Windows applications, some Sun-only applications) either because they were written in-house for that OS, or the cost to update the software was prohibitive.
- Finally, there hasn't yet been time or resources to convert over some hardware (such as scanners and paging-systems) and these systems currently run on legacy HP-UX machines as result.

Psychological Hurdles

The concept of completely changing architectures/ OSs can be a daunting one to developers and user, and the psychological aspects of pushing a project like this forward can't be discounted.

One of the biggest psychological hurdles to overcome was the nervousness of Control Room staff to use an unfamiliar PC-architecture to perform their Accelerator-critical work. To overcome this trepidation, it was necessary to gradually introduce Linux-based machines into the Control Room environment. It was also necessary to insure as many systems worked on the new machines as possible and to have a fast problem-reporting/resolution turn-around time to reinforce the concept that this was an architecture that would be supported long-term.

In contrast to this, once users (especially in the Control Room) started using the new systems they quickly realized how much faster and more robust these machines were and these machines rapidly became the preferred platform to use.

While this was a very desirable result, it did result in a "More, more, more" PC-mentality where additional functionality beyond what was previously possible quickly moved from anticipated to expected. It became necessary to temper the user's expectations of what the systems were capable of to a reasonable level. It was also necessary to reassure users that more expanded functionality (the support of audio-video options for example) would be possible as the project moved forward, but that decisions on configurations and supportable software-packages had to be made over time.

Financial Hurdles

Finally, as this was a large-scale project, it had to be planned with budgetary constraints in mind. The cost of upgrading the entire OPS-subnet (and other Accelerator subnets along with it) necessitated a multi-year plan. Because the project would be completed over multiple years, this dictated hardware choices that would be robust enough to last for a multi-year cycle and also be available over a multi-year cycle (so that multiple, different configurations wouldn't have to be installed as machines were upgraded across multiple years).

Dell Computer systems were chosen primarily for these two reasons. Because they had proven to be a stable and robust architecture that would last for many years (and that included multi-year support) and because an existing site-purchasing contract with Dell Inc. kept the cost of new computers low.

FUTURE UPGRADES

There are few steps still left to be completed as part of the Accelerator-site upgrade from HP-UX to RHEL.

Several core services still run on older, legacy HP machines such as printing (using HP's lp drivers, which will updated to CUPS on Linux) and scanning (again currently running under HP, and to be converted to running completely under SANE on Linux).

The next big push for upgrades is also taking place in the other two control rooms at the lab:

- The Free-Electron Laser (FEL) Control Room is in the process of switching out their HP-UX machines for Linux.
- The Central Helium Liquefier (CHL) Control Room has not yet had any of its computers upgraded to Linux, but has recently had its servers upgraded off of HP-UX and on to Sun-servers. This will be the first step in migrating the HP-UX workstations out of the CHL Control Room in favor of Linux workstations

Finally, the last few software packages still used that are not supportable under Linux are being examined to either be phased out completely, or upgraded to versions supportable under the new architecture. These include software packages such as FrameMaker, Optim and CapFast.

CONCLUSIONS

Converting the Accelerator OPS-subnet (and expanding this conversion site-wide) allowed for significant upgrades to the Accelerator Computing Environment for multiple reasons:

- Cost – the upgrades that were done could not have been without utilizing lower-cost PC solutions and a Linux architecture. The cost to perform these upgrades under another architecture/OS would have simply been too high to be economically feasible.
- Usability – the expanded usability of the Linux environment and the increased support available through a PC-based architecture allowed for development of improved tools to make operation of the Accelerator easier.
- Processing Power – the inherent boost in processing attained by switching to a PC-architecture allowed for expanded capability of the Control Room machines and also aided in development of more advanced tools and scripts.

However, it had to be accepted that Linux was not a panacea for all situations. In some cases other architectures were utilized (Solaris systems for file services) or legacy HP-UX system had to left in place to run older utilities until another option can be found.

There also had to be an assured level of buy-in to the upgrades from all levels:

- Manager-level buy-in had to be attained by proving the cost-effectiveness of the upgrades.
- Developer-level buy-in had to be attained by proving the long-term savings in time and resources (for some short-term difficulty upfront).
- User-level buy-in had to be attained by proving that the new architecture offered improved functionality and was inherently "better" than the old, and that any issues that came up during testing of the new architecture would be resolved in a timely manner.

Finally, because these upgrades were taking place on an operational Accelerator Control Room environment, upgrades had to be planned around the Accelerator schedule such that they had a minimal impact on beam delivery.

This necessitated a very steady and careful progression of the project that extended its overall lifetime over the course of several years. This however was turned to an advantage in some respects by allowing a gradual introduction of the new Linux architecture to both developers and users over the course of the upgrade.