# Linux migration of the Ganil control system

*L. David & the Ganil control group*
*Grand Accelerateur National d'Ions Lourds,*
*BP 55027 14076 CAEN Cedex 5 FRANCE*

## ABSTRACT

After three years of studies and evaluations, the control system of the Ganil facility has been upgraded from VMS to Linux during the last winter shutdown. All the beam tuning programs written in Ada have been migrated, keeping the Motif graphical user interface for the operator interfaces and the on-line Ingres databases as the main data repository.

To provide high availability capabilities on the server, a failover cluster has been configured based on the Hewlett-Packard Service Guard package beside of the Red Hat Linux toolkit.

This paper presents the different milestones of the project which allowed to reach the goal of restarting the accelerators within this new environment, also adding two new low energy beam lines (one for exotic ions and one for stable ions).

Off-line development tools and the database environment are therefore described. Then the presentation summarizes the main tasks performed by central services on the cluster such as alarm handling process, front-end configuration servers, archiving, data collectors … It also shows the few changes in the software on the operators stations and the renewal of the real time compiler for the Vme front-end crates within the VxWorks environment.

Lastly, the upgrades that could be useful for the new Spiral2 project are listed.

## THE LEGACY CONTROL SYSTEM

The Ganil facility located in Caen (France) consists of cyclotrons in cascade thus allowing to accelerate heavy ion beams for nuclear physics, atomic physics, radiobiology and material irradiation. Since December 1997, the Spiral extension has been producing radioactive ion beams using the so-called ISOL technique. A new project named Spiral2 has been approved in May 2005 and aims to produce rare ion beams using a Uranium carbide fission process : a Linac accelerator will accelerate deuteron beams impinging on a carbon converter to produce the neutrons to get the fission.

When it was defined in the early nineties, the control system was based on Vax workstations and VMS servers, programmed in Ada. The front-end Camac crates were using RtVax microprocessors.

When Dec decided to stop the Vax family, we began to replace the Camac crates by Vme with PowerPc processors, and we moved from VaxEln system to VxWorks, using Ada95 (object oriented) instead of Ada83.

At this time the network, which was a 10B2 daisy chain, was replaced by a 10BT structure, with bandwidth switching at 10 Mb/s speed linked to the external national research network.

### *Software*

The control system is built upon several levels :

- The graphical user interfaces, used for the beam tuning programs are Motif based with XRT commercial widgets and the algorithms are written in Ada,
- The relational data base Ingres, used by all the programs accessing to the parameters of the beams and to the characteristics of the 4500 devices involved into the accelerator control,
- The network communications using TCP/IP sockets, between the operator consoles and the real time crates, with a client mode on the workstations and a server mode on the crates,
- The drivers for the devices on the Vxworks crates, also programmed in Ada.

### *Development tools*

The Ada programs were developed on a Vms server, with an X terminal for each programmer.

Concerning the real time tools, the Aonix Ada cross compiler was running on Windows PCs.

The offline programs for beam parameters calculations were written in Fortran.

## EVALUATIONS

The aim of this upgrade was to install a solution with new technologies, that can easily evolve in future and to have cheaper operator consoles than the commercial workstations, which means to choose PCs.

Concerning the operating system, the choice was between Linux and Windows.

Linux was chosen for the accelerators controls because of reliability reasons and its wide use among the accelerator control community. It was also easier to move the control programs, keeping the same Motif graphical interfaces. Another reason was the compatibility with Epics, which is a possibility for the Spiral2 project.

Windows was reserved for the supervision interfaces of the Programmable Logic Controllers, with Panorama, a commercial software widely used in Ganil and only available on Windows.

## HARDWARE UPGRADES

### *Client server architecture*

The new control system is built over a cluster server for common services, and PCs clients are used for the beam tuning in the main or auxiliary control rooms.

Several links have been established with the office desktops and the Intranet, so that all the measurements from the accelerator devices, which are archived on the server, can also be read by Windows clients and displayed on the Intranet with all the accelerators informations.

The development software is made on other PCs, Linux or Windows.

The Vme crates were kept as before, but the Aonix cross compiler was changed for the Gnat compiler and the upgrade of the processors towards a new generation will be evaluated later.

### *The cluster server*

To provide high availability capabilities on the server, a failover cluster architecture has been configured based with the HP Service Guard package beside of the Red Hat Linux toolkit.

So the main server is set up with two similar nodes, working in this cluster mode.

If one of them stops, the services running on it can move to the other server.

To get more reliability, the disks are in an external SAN array, connected through two optical switches, so that both of the servers can use these disks.

To make sure that the redundancy is effective, all the links are always dual, even for the switch and the disk controller. In case of failure, the second link will be automatically activated and the disks will still remain available.

These disks are mirrored into two different racks, so that a disk failure won’t have any consequence on the data integrity. Two other disk-racks are also installed in the same cabinet and are connected to a separate Windows cluster for the office work.

In this SAN array, the Linux disks are only available for the Linux cluster, and the Windows disks are only available for the Windows cluster.
A tape library for archiving facilities is also connected to one of the optic switches, managed by the Windows servers, and using direct access to the SAN array, as for Linux disks as Windows disks.

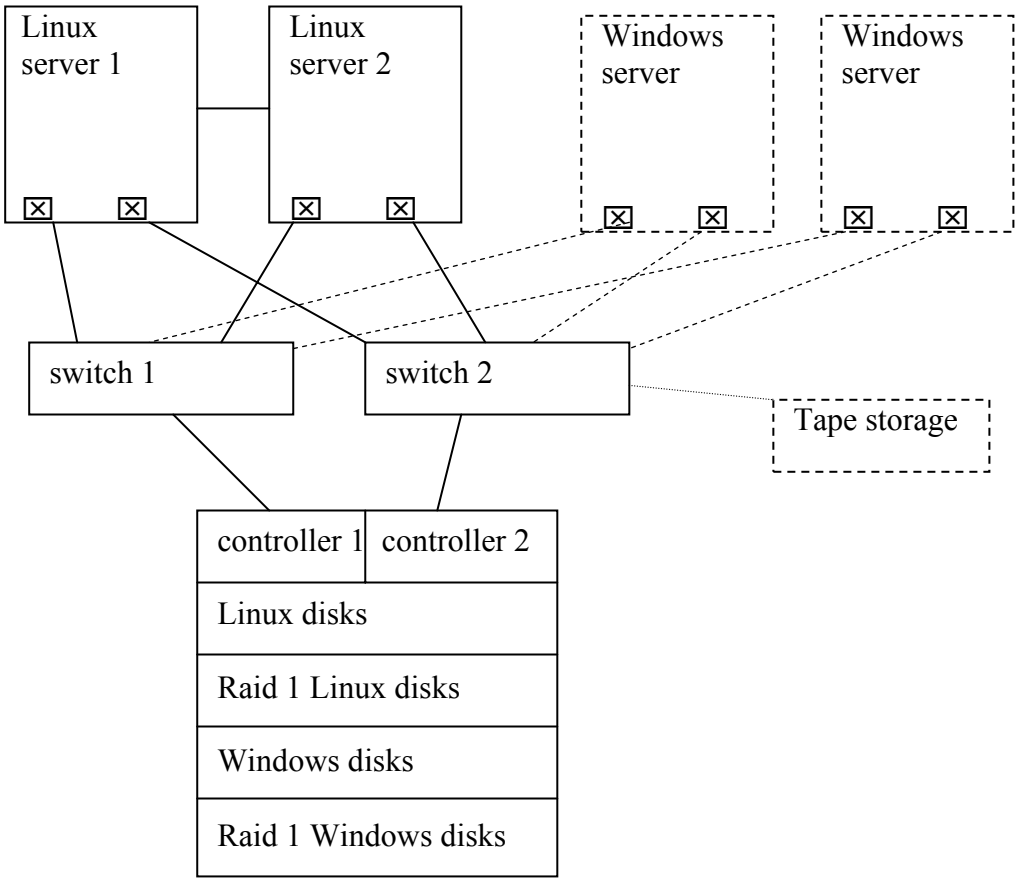


*The operators PCs*

New dial boxes with shaft encoders have been studied and produced for the operator PCs. They have the same function as the old ones, but with an USB output, so they can be connected to laptop PCs or to new PCs without any serial port.

All these PCs are strictly the same, each of them with a dual channel graphic card, one VGA and one DVI, to drive either two flat TFT screens, either both one CRT and one TFT.

On the new flat screens, the pixel definition is linked to the size of the screen, so if the definition is decreased to obtain large windows, the graphical aspect is worse.

21" TFT screens with 1600x1200 pixels will display smaller windows than 19" with 1280x1024 pixels. So the screens in the control room are TFT 19", until the availability of a new generation or changes in our graphical software to display larger windows on any screen.

*Development PCs*

Each member of the Computer group and of the Beam tuning group, working on Windows PCs, also needs to use the Linux server for parameters calculations or Ada compilations.

An X emulator on Windows was tried to make connections to the Linux server with graphical display, but a problem appeared : the XDM server on the cluster prevented to perform the switch of the NFS package from one node to another.

So the PCs with the X emulator Cygwin are now connected to an X server on a front node, with SSH access to the cluster. So there isn't any X server blocking the NFS package as described above.

For simple commands, the so-called Putty utility on Windows allows to open a terminal window with a secure protocol. It could also be used during external interventions through ADSL connections.

### *Network*

The network has been separated in two different TCP/IP subnets : one is public for the office work while the other specific for the control is hidden from the Internet networks.

The control network is based over two main and three remote Gigabit switches (one for the injector area, one for the Spiral area and one for the instrumentation gallery).

In the control room, the Linux servers and PCs are plugged on the main switches, with unshielded twisted pairs (UTP). In the remote areas the Vme crates can be connected with 100 Mb/s UTP links while deported PCs are reached through UTP Gigabit links. Lastly, office and control networks are connected to the other Ganil networks by a Gigabit optical fiber link.

## SOFTWARE UPGRADES

### *Common services*

The cluster is used for two main functions : the control of the accelerator and off line services such as the development of the Ada programs, the Intranet server, the printing server and the file access from the windows clients. So each server runs a separate database system, one for the control and the other one for the development.

Except the databases that have direct access to the disks, the other services perform disks access as clients of a central NFS service on the development server.

On the control server, common services are always running and restart in case of failure :

- Ingres database server for the requests from the control programs, the alarm server …
- Alarms display, sent by the Vme crates or the tuning programs
- Several data collectors, which read a list of devices on the Vme crates and keep the measurements in memory. So the different requests issued by all the clients won't create as many network connections as applications having to reach those devices
- Data collectors dedicated to external clients from the control system, such as the acquisition process for the identification of exotic ions, or the standalone irradiation target control
- Periodical archiving of measures on most of devices of the accelerator, so that they can be displayed on Windows PCs later
- Downloads of the front-end databases with the devices characteristics, needed by the Vme crates
- On-line display of the accelerator status, generated by HTML sheets and broadcasted by a web browser and a video network

On the development server, off-line services are running :

- Ingres development database (preparation of the beams, beam statistics, …)
- Cups printer server
- Apache web server
- Samba server to give file access for the Windows PCs
- Nfs server for remote access from the Linux control PCs
- Gnat compiler, with the GPS graphical tools

All these services can move from one node to the other, if the first node stops or shutdowns.

### *Databases*

The Ingres relational database management system is involved into the whole control system so being the core of the system. It deals with the front-end real-time crates configuration, the pieces of equipment definition and specification, the alarms archiving system, an electronic log-book for the machine operation, the on-line configuration of the machine status, the dynamic beam parameters management, the beam optics definition and the links with the off-line computing programs.

The Ingres software has been configured to be fully integrated into the cluster environment adopted for the servers. Therefore, two different Ingres installations are running at the same time ; a first one is

used in an off-line mode for the beams preparation or the pieces of equipment configuration while the second one is devoted to the on-line control of the accelerators. Each of these two Ingres servers is normally run on each of the two nodes but in case of failure they are able to be executed simultaneously on the same node.

In addition, the Ada / Sql gateway between the beam tuning programs and the Ingres database has been migrated from the Ada 83 Vms compiler to the Gnat Ada 95 environment within the Linux operating system, so being able to keep most of the existing code (more than 200 Ada / Sql procedures have been developed previously).

Applications written in alphanumeric mode (such as the uploading procedure for the output data from the off-line computing program or the preparation of the set of beams to be accelerated) have been migrated into this new environment. Specific graphic applications developed within the OpenRoad environment are also available into this new context, allowing people to configure any piece of equipment, to retrieve archived alarms according to various criteria or to update the on-line machine operation logbook. Lastly, Intranet access using a PHP gateway gives users the ability to access the beam parameters database or to manage the list of the specific recommendations or warnings for the daily accelerator operation.

### *Beam tuning programs*

On Vms the fifty beam tuning programs were developed with the Motif graphical user interface, so that they could have been moved to Linux with few changes inside their source code, only in a year instead of several if all of them had to be rewritten completely.

Later, this library could be replaced, because of the lack of evolutions and functions only available in newer standards such as GTK, which will be necessary for new projects.

The Vms system procedures calls which were already embedded in Ada packages have been modified to fit to the Linux run-time.

The Gnat compiler includes a project manager, source oriented, that needed to reorganize the tree structure of all the program files, including binaries or data associated with each of them.

### *Off line parameters*

Some of the beam parameters calculations are made off-line by Fortran programs such as Galop or Param. A reduced release of Galop has been moved to Linux, without the Cern and Nag libraries. This program is running fine, and the next milestone is to install a full size development workbench including the Cern libraries on the development server.

The Param program, based on a lot of different generations and writers procedures, has been moved without any change, keeping data access through flat files instead of migrating them into the Ingres database.

### *Epics binding*

In the Spiral2 project, several Vme crates running Epics are planned to be installed by the Dapnia, for the primary beam control.

Within this context, Epics, which was developed by a community of American laboratories, is built over VME front-end controllers (IOC) with I/O slots embedded into Vxworks crates and graphical interfaces (OPI) on Linux PCs.

The procedures available to get access to the devices (Channel Access) are written in “C”. An Ada binding has been developed to interface these Epics client procedures. So it will be possible, from the current control system, to get access to the Epics Vme crates running an Epics server for specific devices control such as the RFQ or the ion sources. For the exotic beam lines, the Ada control system used for the Cime cyclotron will be extended and the Epics binding described above won’t be necessary.

If more specific tests are needed, the graphical Epics tools (MEDM) will have to be used and learned by the Ganil programmers, although based on Motif. Also, according to the size of the Spiral2 project, a relational database will have to be considered and designed later in this Epics system.

### *Real time*

The real time programs, running on the front end crates, have also been changed, to follow the new beam tuning interfaces and central services (alarms and database downloads).

A new Ada cross compiler for Vxworks, GnatPro, was installed on Windows, with a project oriented manager GPS providing the same environment as the native Linux one. So the communication packages and the interrupt handling have been modified to fit this new system surrounding. Later a new release of Vxworks and associated tools will allow the programmers to develop either from Windows or from Linux.

## CONCLUSION

The basic access to the accelerator devices were tested during the last run of year 2004 with a Linux console operated among the existing Vms environment. So during the winter shutdown period, the servers were set up, and the 30 operators PCs were installed and connected to the network. As the algorithms had not been changed, the beam tuning programs were all tested in the first two weeks of accelerator restart in early 2005.

From the operation point of view, the functionalities have been kept and the operator interface is the same as in the previous control system ; one major feature to be noticed is the response time improvement, mainly for the database access for all the beam parameters management.

The total cost of the project has been kept to, and the complete move to Linux has been achieved one year sooner than expected. This year will be useful for the primary tests of instrumentation over the Spiral2 project, concerning first the R-F amplifiers, and larger tests with the Epics binding if necessary.

With this new Linux control system, it becomes easier to add other control PCs in remote control rooms, dedicated to specific beam lines. Recently two auxiliary control rooms were added for the low energy beams (Irrsud for stable beams and Lirat for exotic beams) and linked to the control system.

The next step will be the control of the Spiral2 project and its integration beside or within the existing control system.

## REFERENCES

[1] L. David & Ganil control group, “Evaluation of a Linux control system”, PCaPAC, Frascati, Italy October 2002 and ICALEPCS, Gyeongiu, Korea, October 2003

[2] E. Lécorché, P. Lermine “Use of the Ingres RDBMS inside the new Ganil Linux based control system”, this conference.