Installing and Managing the Higher Level Network of the ELETTRA Control System

C. Scafuri, F. Radovcic
Sincrotrone Trieste, Padriciano 99, 34012 Trieste, Italy

Abstract

A brief description of the higher level network of the ELETTRA control system[1] is given. This network links the UNIX operator consoles, the main server and the OS-9 Local-Process Computers by means of interconnected fibre-optic backbones. A detailed description of the software architecture follows, which points out the advantages of using de facto standards - such as NFS, NIS, Domain Name Server, IP protocols - for building a powerful, flexible and yet easily manageable distributed multi-platform system. Finally, as a result of our experience gained during the commissioning period, some practical considerations from the system manager point of view are presented.

1. The Higher Level Network

The higher level network is formed by operator consoles, one system server and by local process computers (LPC). Operator consoles and the server are placed in the main control room. One operator console is placed in the linac control room. The LPCs are rack mounted VME systems. These racks are placed in two sites around the ring building and one site in the linac building. Two more sites are set up in the ring building but are not yet used. A local area network links all these computers by means of IP protocols.

1.1 The Local Area Network

The disposition of control rooms and of the various sites maps naturally to a star network topology: from the main control room five fibre-optic backbones reach the four sites in the main ring and the site in the linac building. Each of the active sites is also configured as a star. A hub is connected to the end of each backbone. Each rack has its own thin Ethernet branch departing from the hub.
Fibre-optic provides excellent immunity from electromagnetic interferences and grounding problems. But fibre-optic is not guaranteed to withstand radiation without loosing its characteristics. That is why the section of the backbone to the linac building crossing the transfer line tunnel, where there is the possibility of high radiation doses, runs on conventional Ethernet cable (yellow cable). This section is laid down inside a grounded metal pipe to improve noise immunity.
From the analysis of the typical network communications requirements of our control system, we identified two main flows of data: one between the main control room and the ring LPCs, the other between the linac control room and the linac LPCs. In order to reduce bandwidth occupation and to delay possible congestion problems, we exploited the pattern of the traffic by inserting a bridge between the two data flows: local traffic is not transmitted across the bridge, while communications between the two sections of the network is not affected. A second bridge insulates the high level network from the development system network.

1.2 Operator Consoles and System Server

There are 6 operator consoles, which consist of Hewlett-Packard series 9000 model 720 graphics workstations equipped with 48 Mbyte of RAM and one 420 Mbyte hard disk, 19 inch colour monitor.
The system server is a Hewlett-Packard series 9000 model 827/s with 96 Mbyte of RAM one 420Mbyte and two 1.3 Gbyte hard disks. In order to improve reliability and performances each 1.3 Gbyte disk has its own I/O interface and they are operated in disk mirror mode. Disk mirror is an extension of the operating system which keeps two identical copies of a logical disk on two physical disks. If one of the physical disks fails, operations on the logical disk are not affected and are carried out on the working disk only [2].
All these systems run under HP-UX operating system, the Hewlett-Packard implementation of UNIX System V release 4.

1.3 Local Process Computers

The Local Process Computers are VME board modular computers equipped with Motorola 68030 at 25 MHz processor, 16 Mbyte of RAM and local area network interface. The LPCs run OS-9 real time operating system. There are 15 of them installed in the various sites.

2. System Organization

The computing system is organized in order to distribute computing power, share resources and centralize administration as much as possible. We use a distributed client-server model in which the operator consoles and LPCs use the services available on the network. The services can be divided in two main groups: control services and system services. Control services are available from the LPCs and from the server. These services are all custom developed and based on RPC protocol. System services are available from the system server only.

2.1 System Services

System services are the services related to the administration and configuration of the Higher Level Network, and sharing of its resources. The available system services are: Network Information Service or NIS, Network File System or NFS, remote boot via trivial file transfer protocol or TFTP, Domain Name Service or DNS, printing and time synchronization. Excluding the last one, which is home developed, all the other are consolidated industry standards, readily available an almost every computer platform.
NIS [3] is a distributed database service, which replaces or extends common configuration files with a centralized management facility. In particular we use it to distribute and manage

the password file.
NFS [3][4][5][6] is a distributed file system providing access to remote disks. With NFS client computers can access and share a disk (or better a file system) physically attached to the server machine.
TFTP [7] is a standard protocol used to transfer files between two network connected computers. It is very simple and with a relatively small code size. A typical usage of this protocol is by diskless systems to load the kernel image from a central server.
DNS [8] is a service used to map host names to network addresses.
The printing service allows operators use a fast character printer and a colour printer from every node on the network.
The time synchronization service is used to set the time of all the nodes of the network to the same hour of the server with a one second precision.

2.2 Operator Consoles

The operator consoles are configured to give users a consistent environment. On each of them they can log in with the same user name and password, find exactly the same environment, see the same data files and run the same application programs. The only difference is the name of the host. This is obtained from the combined use of the available services and a careful configuration of each workstation operating system.
The standard user validation system is extended by using NIS. When a user name and password are given during the login procedure, the NIS server is queried and the user/password couple is checked against the stored database. When a password is changed, the central database is automatically updated.
The workstations are configured as NFS clients in order to mount the remote file systems containing application programs and libraries, configuration files and the home directory of each user. On the internal disk there are only the operating system and the swap space.
In this way it is extremely easy and economical to manage the system.
Using a networked file system imposes the penalty of a slower file access, however this fact has negligible consequences for us: most of the disk operations are performed at the startup of programs only. On the other hand we have several advantages:
- reliability: user data is stored on a mirrored disk which is intrinsically highly reliable. This disk is also frequently and regularly backed up on magnetic tapes.
-consistency: since the file system is shared among the workstations, each operator sees always the same set of files from every console.
-ease of management: only one basic file system to care for, update and backup regularly. System disks of the operator consoles are backed up only when the operating system or basic configuration is modified.
-ease of software update administration: we have to update, check version and test new programs on the server only: the new software release is automatically available to all the machines.

2.3 Local Process Computers

The server act as as boot and disk server for the OS-9 Local Process Computers. The LPCs are diskless computers equipped only with the boot program and basic configuration (e.g. network address and related informations) stored in ROM. The kernel, which can be customized for each LPC, is stored on the main server and is loaded via TFTP protocol. Once started, the LPCs make use of NFS to remotely mount their own file systems from the server.
Here again we have all the advantages of centralizing the management of software releases and versions, moreover there are also some economical and reliability benefits to be gained.
In fact we have replaced many local hard disks with a single highly reliable mirrored disk. Considering that the failure of a single LPCs -or more easily of its disk - can actually stop the operations of the whole accelerator, the probability of a disk induced fault is greatly reduced.

2.4 Auxiliary Services

DNS is used to maintain a centralized network addresses database. LPCs and operator consoles refer to network nodes with symbolic names. These symbolic names are automatically mapped into network addresses by queries to the DNS server. In case we need to change the network address of a node, updating the centralized database only is sufficient to make the new address known to the whole system.

2.5 Users, Groups and Access Policy

We have two sets of user accounts, with different environments and capabilities, which are enabled according to the operational status of the machine. During a shift period we activate the operational accounts, during a shutdown period we activate the maintenance accounts instead.
We have three operational accounts: operator, machine physics and guest. The operator account is used to run all the official control programs. The machine physics account is used to test machine physics programs and run some special measurement programs. Usually a machine physics program is transferred to the operator account after it has been thoroughly tested and corrected. The guest account runs general purpose program, temporary applications and is used for early tests of new programs. This account will be phased out at the end of the commissioning period.
Maintenance users are grouped by division, these accounts are used for testing, calibrations and other maintenance duties. By discriminating on user and group identifier it is possible to limit the access to control programs and data. User and group can be also used at LPC level to limit the access to field equipment.
The use of a centralized password database allows us to switch from operational users to maintenance users and vice versa. This operation is performed by some custom scripts based on standard NIS facilities.

3. Our Experience

We have seen that, at least during the commissioning period, the most demanding task for system administrator is to keep updated the releases of custom developed control pro-

grams, configuration files, X11 resource files and environment variables. During this period in fact many programs can be fully tested for the first time, the inevitable bugs and problems are found and corrected and new versions are released very frequently. A uniform and centralized management of application programs is the only way to handle this situation in an orderly and relatively effortless way.

An extremely important system feature that we exploited during the commissioning phase is that of shared libraries[9]. This feature is available on HP-UX and some other versions of UNIX. Like normal libraries, shared libraries are collection of already compiled portions of code, but this code is not appended to the executable object file when a program is linked. The shared libraries code is loaded into memory from disk and linked to the program the first time program makes a call to library code. By using shared libraries it is possible to make corrections or improvements to the our basic software components -like widgets, or data acquisition procedures-and have these improvements automatically available to all our already compiled programs (more than 100 programs) without relinking them. Obviously our shared libraries are stored on NFS shared file system, making the correction of a library fast and easy. This is made on one file only.

We have seen that the use of X11 resources files must be planned and well coordinated in order to obtain a consistent and pleasant graphical user interface as well as to avoid undesired side effects with conflicting resources. Our solution is to use a single common resource file, publicly readable and centrally maintained.

4. References

[1] D. Bulfone, Status and Prospects of the ELETTRA Control Sysem, Proc. Int. Conf. on Accelerator and Large Experimental Physics Control Systems, Berlin, Germany, 1993, to be published on Nucl. Inst. and Meth. for Phys. Res, sectionA.

[2] How HP-UX Works: Concepts for the System Administrator, Hewlett-Packard b2355-90029, 1992.

[3] H. Stern, Managing NFS and NIS, Sebastopol, O'Reilly & Associates, 1992.

[4] Installing and Administering NFS Services, Hewlett- Packard B1013-90009, 1992

[5] Sun Mycrosystems Inc., NFS File System Protocol Specifica tions, RFC 1094, 1989.

[6] L. Garlick, R. Lyon, L. Delzopo, B. Callaghan, "The ONC Envi ronment", SunTechnology, vol. 1, num. 2, pp. 42-46, 1988.

[7] K.R. Sollins, TFTP Protocol, RFC 783, 1981.

[8] P.V. Mockapetris, Domain Names-Concepts and Facilities, RFC 1034, 1987.

[9] Programming on HP-UX, Hewlett-Packard b2355-90026, 1992.