GG

CERN - CN 93-01

8x 3305 9

CERN - Computing & Networks Division
CN/93/01
D.O. Williams
January 1993

CERN LIBRARIES, GENEVA

CM-P00065730

**DISTRIBUTED COMPUTING AT THE EUROPEAN LABORATORY
FOR PARTICLE PHYSICS**

(Presented at the 1992 IEEE Region 10 International Conference on
Computer, Communications and Automation towards the 21st Century (TENCON '92),
Melbourne, Australia, 11-13 November 1992)

DISTRIBUTED COMPUTING AT THE EUROPEAN LABORATORY FOR PARTICLE PHYSICS

David O. Williams
Computing and Networks Division, CERN
1211 Geneva 23, Switzerland
<davidw@cernvm.cern.ch>

Abstract

CERN operates experimental facilities generating Terabytes per year of data for a world-wide community of 6000 particle physicists. The paper reviews CERN's history of involvement in local and wide area distributed computing and discusses how it has responded to the increasing impact of networked RISC systems in order to provide high quality services. These services include a program library supported on many different platforms, a distributed file and tape management system, centrally operated RISC services, data dissemination by satellite, and a distributed hypertext information system. In conclusion the likely impact of an emerging integrated approach to distributed data analysis for particle physics on the future use of mainframe computers in the discipline is discussed.

Introduction

This paper starts with a brief description of CERN, the European Laboratory for Particle Physics, and of particle physics in general. Following a description of the type of experiment performed at CERN we review some of the background to CERN computing. A description of very early approaches taken at CERN towards distributed computing leads into a series of case studies of some present services. In conclusion, some speculation on the likely future evolution of these services is proffered.

CERN and Particle Physics

CERN, the European Laboratory for Particle Physics, occupies two major and several smaller disjoint sites in an area roughly 15 kilometers across spanning the French-Swiss border near Geneva. It is an international organisation with 18 member states¹ who contribute to its funding in relation to their GNP. The Laboratory operates several particle accelerators and colliders, of which the latest is LEP (the Large Electron-Positron collider) which is installed in a tunnel of some 4 meters diameter and 27 kilometers circumference about 100 meters below the local countryside. Inside the vacuum pipe of LEP, beams of electrons and positrons collide head on at an energy of 200 Giga-electronvolts, producing a very high energy density, similar to conditions that existed shortly after the Big Bang and allowing for the study of a range of more or less short-lived elementary particles.

The aim of particle physics is to study the smallest components of matter (the elementary particles) and the forces that hold them together, in order to try to understand the nature of matter. In the late nineteenth century the forces of electricity and magnetism, initially thought of as separate, were unified into a description of a single "electromagnetic" force. Following work by theoretical physicists in the 1970s, the unification of the electromagnetic force and the

weak nuclear force was confirmed by two key experiments performed at CERN, for one of which Professor Carlo Rubbia, presently CERN Director-General, was awarded the Nobel Prize for Physics in 1984.

CERN itself employs a small number of research physicists, but most of the roughly 3000 staff are scientists, engineers, technicians and mechanics providing development and operational support for the accelerators and other infrastructural services. The CERN facilities attract physicists from all over Europe and the rest of the world who are eager to participate in the experimental collaborations studying phenomena which can be observed only at CERN and a handful of comparable laboratories. About 6000 scientists are registered as CERN users, of whom about 2500 are on the site on an average day, the remainder being at their home institutes, which are typically university physics departments. There are some 200 such institutes throughout Europe, and many others throughout the rest of the world. CERN is the largest centre in the world for particle physics experiments, and has become a principal hub of the particle physics community and its distributed computing environment. Some 15 Australian particle physicists, mainly from the Universities of Adelaide, Melbourne and Sydney, are registered as CERN users.

Besides working to fully exploit the data coming from LEP and the other accelerators, CERN is presently seeking authorisation to construct a new accelerator, the LHC or Large Hadron Collider. The CERN Council has indicated that it believes that construction of this machine, in which protons will collide against protons at 16 Tera-electronvolts, would be a logical next step for CERN, and for science, and has asked for detailed plans and cost estimates to be prepared for a decision at the end of 1993. The LHC would take roughly six years to construct and it would generate data rates far in excess of those at the current experiments. We expect to need all of the formidable computing power and network speed that might be available in 2001.

¹Austria, Belgium, Czech and Slovak Federal Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Spain, Sweden, Switzerland, United Kingdom.

The nature of the experiments and the data

The experiments that are designed to detect and measure the particles of interest are performed by very large teams of physicists coming from many different institutes. For example, there are four experiments at LEP, each with typically 400 authors on their publications, most of whom are Ph.D. research physicists. They come from about 40 different institutes in 10 different countries. For example, the spokesman of the LEP3 experiment, Professor S.C.C. Ting, is based at M.I.T. and the collaboration has strong representation from the U.S.A, China, India, the former U.S.S.R and Europe.

These experiments are typically large, complex and expensive, with many hundreds of thousands of individual detecting elements instrumenting a volume of up to several thousand cubic meters. The particle beams cross or collide at LEP every 22 microseconds, and this is considered "slow" in comparison with many accelerators. The analogue and digital outputs from the detectors must all be sampled and preprocessed at a rate dictated by the beam crossing time and the frequency with which particles in the crossing beams interact. This sampling and pre-processing is typically performed in a parallel manner using electronics and computing facilities conceived, assembled and operated by the experiment, in a multi-stage selection process. The useful data is written onto magnetic storage at a rate which is limited more by storage technology and an assessment of how much data can be subsequently processed, than by the intrinsic desire of the physicists. A typical LEP experiment writes a new 200 Mbyte tape cartridge every few minutes. The experimental environment of the future LHC is normally considered to be about three orders of magnitude more complex than that of LEP.

The particle accelerators and colliders at CERN typically operate on a 168 hours per week basis, with short technical shutdowns programmed every week or ten days, and a major shutdown after Christmas, when the cost of electricity is highest, lasting for two to three months. Many experiments are active at any one time, and the largest experiments generate several terabytes of permanently stored data per year. These data are then available for subsequent analysis by physicists from any of the universities participating in that experiment. Physicists must be able to carry out analysis work both when they are at CERN and when they are back in their home institute anywhere in the world. In that context, where lower transfer speed and reliability are sometimes acceptable, helical-scan video cartridges (Exabytes) can be suitable for the interchange of processed data, but at present the medium for primary data storage, where high transfer speed and very high reliability are critical factors, is almost always the industry-standard IBM 3480 compatible magnetic tape cartridge.

General background on computing at CERN

CERN has to conceive, implement and operate state-of-the-art systems for real-time data acquisition and storage, and a need for well organised analysis of large

quantities of data by a mobile user population with a global spread. To keep CERN's science at the competitive forefront that it reached during the 1980's, it has been necessary to continue a long history of the aggressive use of leading edge computing and network technology.

If there has been one feature that has characterised physics computing during the past two decades it has been the need to handle both very heavy CPU loads, which were and are typical of much scientific computing, *together with* the very heavy movement and manipulation of data, which has historically been more typical of commercial computing. Seismic studies and satellite-based research are now tending towards the same requirements, but they do not yet seem to have the same need as particle physics for multiple people at multiple different sites to collaborate on the analysis of shared data.

The need to provide a very high quality of service continues to be one of our major concerns. CERN's accelerators and experiments run around the clock, and the key computing services have to run on the same basis if they are to be useful. Users must be able to rely at least as much on the availability of very complex distributed computing systems as on the other tools of their everyday working life. So the challenge for computing staff at CERN, as in other similar centres, is to optimise the services made available to our staff and users, so that those services cost as little as possible while still enabling the laboratory to carry out science that is competitive on the world scale. Just to re-emphasise the point - we are not carrying out computer science research - we are in the business of providing leading edge services - and it is a most fascinating challenge.

Many of the computers installed at CERN are not owned by CERN, but by the visiting teams. This effectively precludes most forms of very strong hardware and software standardisation. While it may not seem surprising that teams of visiting scientists arrive with their own PCs and workstations, or even with major minicomputers for data acquisition, there are also two examples of teams installing their own mainframes.

One of the basic tenets of CERN's computing policy is to encourage the users to carry out as much as possible of their computing in their home institutes, so that they can fully participate in the life of that institute, and not need to come to CERN just in order to compute or analyse data. In this context it has proved important to try to provide uniform computing environments at the locations where physics computing is carried out in Europe, and, indeed, throughout the world. One particularly successful example of this has been the HEPVM collaboration, which has "standardised" the environment seen by the CERN user physicists on more than 40 medium and large scale IBM mainframes running the VM operating system. Attention is now being devoted to creating a similarly "standardised" Unix environment.

Twenty five years ago almost all of CERN's computing power was installed in the computer centre. Today, an educated guess is that not more than one-third of the capacity is installed there. It is hard to be more precise, since, as we already mentioned, CERN is far from being the owner of all of the onsite computers. To give the flavour, the computer centre operates an IBM ES9000 Model 900, a Comparex (Hitachi) 8/850, a Cray X-MP/48, two large VAXclusters with a VAX 9000 Model 410 and a total of some ten VAX processors, and a set of services based on RISC processors, which we discuss in more detail below. Attached to these processors we find something like 750 Gbytes of disk space, and more than fifty magnetic cartridge drives, on which we mount an average of more than 2000 cartridges per day. Outside the computer centre we can find at least 200 Vaxes, some of them very powerful systems, being used in data acquisition roles, some 600 VAXstations running VMS, roughly 1300 Unix workstations and servers, more than 2000 IBM and compatible PCs and roughly 2000 Macintosh PCs.

In its relationships with computer vendors CERN has sought to make sensible strategic selection of hardware and software suppliers, to work as closely as possible with the selected vendors in order to improve their products and the applicability of those products to the particle physics environment, and, as far as possible, to maintain at least two competing suppliers in all critical areas. In this context we should mention that, above and beyond the staff needed by the vendors to keep their hardware and software operational, we have small teams of engineers from Digital, HP and IBM onsite, working on joint projects to explore the applicability of their latest hardware and software products to CERN's computing requirements.

Of course there are many different aspects to computing in an organisation the size of CERN, and the Computing and Networks Division is not involved in all of them. Among the topics which we will not really discuss in much detail in this paper, we can mention data acquisition at experiments, the desktop computing environment, computing for accelerator controls, and administrative computing, where CERN has recently introduced a very modern system based on Silicon Graphics multi-processor hardware, Unix, Oracle and Oracle-based applications software.

A distributed computing mentality

As early as 1975, a major CERN particle spectrometer known as Omega used a network of various mini-computers for data acquisition and analysis [1]. By 1980, a 2 Mbit/s campus network known as CERNET [2] linked some 100 mini-computers and mainframes. By 1985, CERNET served as a campus backbone interconnecting several Ethernets. Today, in mid-1992, some 70 Ethernets with 4000 connected devices are interconnected by an FDDI backbone, and several smaller FDDI networks are in use as well as a small Gbit/s network using Ultrahnet™ technology. Over the same timescale, the wide-area network bandwidth connecting CERN to external physics institutes has grown from a few Kilobit/s to about 12 Megabit/s.

Although CERN has always had at least one powerful scientific mainframe computer, from a Ferranti Mercury in 1958 through a CDC 6600 in 1965 to an IBM 9000/900 today, networked computing has been a reality for more than 15 years. It is no surprise, therefore, that the advent of inexpensive RISC computers and of standard and relatively inexpensive networking has had a dramatic effect on CERN's computing in the last few years. Not long ago, decision taking about the evolution of central computing services was related to big, fairly simple issues such as the timing of mainframe upgrades and the wisdom or otherwise of investing in a vector processor. Today, the decisions may concern smaller sums of money at one time, but every decision is linked to every other decision for the simple reason that every computer, and every piece of software, is potentially linked to every other one.

Some case studies

In the remainder of this paper several case studies of how CERN has responded to the challenges of supporting a distributed, data-intensive user population are presented. It goes without saying that this presentation cannot be comprehensive. There many interesting developments in computing being carried out at CERN that cannot be covered in this brief sketch.

The CERN Program Library

The analysis of particle physics data was first computerised by 1960 and Fortran, in all of its various incarnations, has been the applications programming language of choice since then. From the earliest days of computer use at CERN there has been a continuous thread of development of algorithms and support routines for data analysis, which led to the creation and evolution of the CERN program library. This library now contains some ten major packages and many independent subroutines. The vast majority of the library is now maintained in Fortran-77, though there is an increasing component of C code, especially in systems related areas, and extensive tests have already been carried out using Fortran-90. There is a total of more than one million lines of code. The library is used at CERN, by all institutes involved in CERN experiments, by many other physics institutes, and to some extent outside the world of particle physics (e.g. in the space science community). It is distributed to about 600 sites in 55 countries, with a new release twice per year.

A major feature of the library is a means of handling input/output of (structured) physics datasets in a computer-independent manner and with maximum efficiency. This is an essential feature in order for data tapes to be shipped from CERN to other sites, running other types of computer. Thus, the library handles issues such as floating point format and byte ordering in a reasonably transparent way. The library [3] also contains modules for dataset management, particle physics simulation, statistical data analysis, minimisation techniques, and data visualisation.

Such a library clearly raises serious problems of version control. Before the advent of RISC workstations, it was sufficient to test only some five different binary versions of each new release in order to cover the main types of computer in use for particle physics. Workstations, with their variety of instruction sets, and with multiple versions of (sometimes immature) operating systems and Fortran compilers in simultaneous use, have complicated life. Today, the number of versions has risen to about 20 and systematic testing has become a major burden, especially as the number of major packages and subroutines continues to grow. There is also a looming need to move from 32-bit to 64-bit floating point numbers as the default arithmetic.

An unexpected side effect of the growth in networked computing compounds the problems of version control. The CERN Program Library is not in the public domain; it is released only under well defined conditions (which are not at all restrictive for *bona fide* research usage). However, we now find that networked servers, which make life very much simpler for our collaborators, make it very difficult to retain control of the dissemination of the library, leading in effect to loss of version control.

File and Tape Management

With terabytes of permanently stored data to manage on an international basis, special tools for distributed dataset management have become essential. At the lowest level, all tapes are labelled both magnetically and visually, in the latter case with a bar code as well as with legible alphanumeric symbols. Since tapes may be exported from CERN to a variety of other computer centres with other users from other disciplines besides particle physics, it was necessary to define a labelling scheme compatible with those used in as many of these other centres as possible.

At the next level, any site dealing with hundreds or thousands of tapes needs a tape management system including an on-line database to keep track of tapes as they are moved between archival storage, racks in the tape vault, active tape drives, and other sites. At CERN, we use a tape management system known prosaically as TMS which originated at the Rutherford Appleton Laboratory in England, and which is now maintained as part of the HEPVM collaboration (see above).

However, physicists working on data analysis are not primarily interested in tapes as such, but in the datasets stored on them. These datasets must be identified and catalogued in a meaningful way, for example such that a dataset written on a given date for a given experiment can be located and retrieved without prior knowledge of the tape on which the information is stored or the exact file name on the tape. Furthermore this must be achieved on a distributed basis, so that members of an experiment can locate a dataset wherever they or it may be, and, if it exists in several incarnations, be directed to the most convenient copy (disk before tape, local before remote, etc.). To achieve these goals, CERN has implemented FATMEN (File And Tape Management

for Experimental Needs) [4]. The heart of FATMEN is a hierarchical catalogue which presents a Unix-like naming structure to the user. The catalogue is in fact stored in a relational database and is accessed (by program or interactively) via a server, normally running on a Unix platform. Multiple FATMEN servers can communicate via a variety of network protocols (TCP/IP, DECnet and RSCS) such that a distributed dataset catalogue can be constructed and kept consistent. FATMEN does not itself contain any of the requested datasets, but it can either return to the user the location of the datasets required, or request a network copy of the datasets to be made.

At present, some 150,000 datasets averaging 150 Mbytes each are catalogued in the principal FATMEN server at CERN (known, naturally, as FATMAN). About 15 servers are running at other sites including some in North America. Since the CERN vault contains at least 340,000 cartridges, it is clear that not all of their data is managed by FATMEN. However, the proportion of datasets catalogued is rising fast, and a tool such as FATMEN will be obligatory in order to handle the vastly increased quantities of data expected when the LHC experiments start towards the end of the millenium.

CHEOPS

Despite having several 2 Mbit/s links to other institutes, in general it has proved impossible to transfer physics datasets over wide area networks due to the bulk of data involved. Typically, tape cartridges are sent by road or air transport to remote computer centres. This has obvious disadvantages, not least being the administrative work of keeping track of the datasets, including notifying the FATMEN system when they are no longer available at the sending site, and then, some days, or more likely weeks later, indicating to FATMEN their renewed availability at the destination. During this intermediate period, unless all of the tape cartridges had been copied, the datasets that they contained are, of course, unavailable to the experimenters.

It is hoped that the broadband networks of the future, if and when they become available internationally at affordable tariffs, will allow terabytes per year of experimental data to be disseminated worldwide. Today, the only technology which could allow such dissemination across Europe is a private satellite network. CHEOPS [6] is an experiment using the European Space Agency's Olympus satellite to link CERN to institutes in Finland, Greece and Portugal. (CHEOPS stands for CERN Helsinki Ellas Olympus Portugal System.) At each site a powerful Sun disk server is linked to a private earth station including an 8 Mbit/s modem. Dataset transfers to or from CERN to the other sites are requested at any moment and deferred, using FATMEN facilities (see above), until the next (typically nighttime) slot when Olympus is available. The ability to schedule transfers at off-peak times and to integrate the dataset management with FATMEN are a model for a generalised high-volume data dissemination system in the future. In the short term they would also permit bulk data transfer on normal

high speed land- or sea-based links at off-peak periods, if appropriate tariffs were available.

Centrally Operated RISC Environment

In 1990, 100% of the CPU power for data analysis delivered to physicists by the CERN computer centre was supplied by mainframes. By mid-1992, despite a sizeable mainframe upgrade from an IBM 3090/600EVF to a top-of-the-range ES9000 Model 900, the figure had dropped to 30%, and by the end of 1993 it may be as low as 10%. When CERN's Cray X/MP-48 arrived in 1987 it doubled the available CPU power, even when only taking into account its scalar capacity, but when it is phased out in 1993, its departure will have only a minor impact on the available CPU power. The question of whether CERN will ever need to upgrade its mainframes in the future now has to be seriously debated. What has triggered this fundamental change?

It is easy to see that the cost per equivalent CPU hour delivered is many times lower for a RISC server than for a mainframe, and that the cost per gigabyte for SCSI discs is many times less than for mainframe discs. The question for a large scientific computer centre with thousands of users is whether it is possible to bring services based on RISC and SCSI equipment into operation with quality and manpower costs similar to those for a mainframe service.

Analysis of mainframe usage at CERN shows that the vast majority of the computing time is consumed by a small fraction of the users, and that the vast majority of login sessions consume trivial amounts of CPU time, mainly for reading electronic mail or for text processing. To be specific, 90% of the CPU time is consumed by about 10% of the users, and more than half of all login sessions consume only a few seconds of CPU time. Further analysis shows that the heavy usage can be classified as belonging to one of a small number of relatively specialised applications, such as extraction of subsets from physics datasets, simulation of particle collisions, interactive physics analysis, or accelerator design simulations. Each of these applications can, in itself, run very efficiently on a modern RISC system with SCSI disks, provided that it has been supplied with the necessary access to I/O devices and an appropriate scientific software environment (good Fortran 77 compiler, linker and debugger), including the CERN program library.

CORE (the Centrally Operated RISC Environment) takes advantage of these observations by providing several specialised distributed RISC systems for these classes of applications. One system is SHIFT (Scaleable Heterogeneous Integrated FaciliTy) [5] in which a set of Sun and Silicon Graphics CPU servers are linked to Silicon Graphics disk servers connected to some 150 Gbytes of disk and a Sun tape server via an Ultranet Gbit/s network. Ultranet is needed not so much for its raw hardware transmission speed as for its efficient protocols which allow network data transfer at up to 7 Mbyte/s from disk to memory and 13 Mbyte/s from memory to memory. The Cray is also connected to Ultranet, and is used as an alternative to the Sun

tape server, and especially to give access to robotically mounted tape cartridges. SHIFT is optimised for the extraction of subsets of data from multiple datasets; it generates about 10% of all tape mounting requests in the computer centre. Jobs in SHIFT are managed using NQS (Network Queuing System) and locally written code to access datasets via the network.

An important component of particle physics computation is the simulation of the events which follow particle collisions, so that the real experimental results can be compared with theory. Since quantum physics has a random nature, these comparisons are statistical, and must be based on large, statistically significant, samples of observed and simulated events. Event simulation, including the tracking of particle motion in magnetic fields and the simulation of the behaviour of the experimental detectors, consumes more CPU time than the analysis of real data. The second component of CORE is therefore CSF (Central Simulation Facility), a cluster of some 25 HP 720 servers dedicated to such simulation tasks. This simulation is very CPU intensive, so the CSF servers need only be linked by Ethernet. Typically each CPU works independently on the simulation of one event at a time, and, even in this simple mode, CSF can achieve a very high degree of parallelism. CSF uses an Ethernet connection to SHIFT to write its output data to tape.

Two other RISC clusters are joining CORE at the time of writing, PARC (Parallel Architecture RISC Cluster) for engineering applications and accelerator simulation and PIAF (Parallel Interactive Analysis Facility) for interactive physics analysis and visualisation. These are based on IBM RISC6000 and HP 760 systems respectively.

Experience so far shows that RISC clusters each dedicated to a narrow class of applications, with networking and I/O optimised accordingly, can deliver a computer centre quality of service at modest manpower cost. Overall, the set of about 65 RISC CPUs in CORE run at satisfactory efficiency (typically 80% and more) and high (>98%) availability. We have real experience that this approach offers a smooth upgrade path through simple scaling-up, for modest increases in capacity. We obviously intend to increase our degree of understanding of the eventual limits to the modular expansion, or scalability, of these systems. We have recently started to explore the use of RAID disk technology in certain critical areas of CORE, to try to see where the extra costs can be justified in terms of increased reliability and ease of operation.

A World-Wide Web of Information

Information must be disseminated in the CERN user community, as well as experimental data. Extensive use is made of traditional network services such as electronic mail and news groups, but unfortunately much vital information is only made available, or, to be more precise, is only kept up-to-date, as raw text stored in a wide variety of formats on a wide variety of computers. One approach to facilitating access to such information is the creation of a distributed hypertext

environment, and CERN has implemented this approach under the name World-Wide Web (W3) [7]. W3 defines a hypertext information server protocol; an information provider must offer access to both raw text information and any hypertext information via a W3 server in order to join the web. A user calls an interactive browser which can follow hypertext links around the web until the required information is located. Several browsers exist, of which the most attractive runs under X windows.

At present some 25 W3 servers are available at 12 sites in Europe and the USA. W3 has attracted interest throughout the Internet and is already by no means limited to particle physics information. Indeed this is the attraction of the hypertext approach, since in principle hypertext links can lead from anything to anything. One imminent challenge for W3 will surely be the integration of multimedia information. Another will be to learn how the great flexibility of hypertext can best co-exist with the need for a very disciplined information environment required in large collaborations such as we find with big particle physics experiments.

Speculation on the future evolution of services

In the beginning we needed a mainframe since it was the only way to perform serious scientific calculations. At that time it was very expensive and needed a whole team of systems programmers, applications specialists and operations staff to optimise the return on the huge investment.

Then minicomputers arrived, and we had a comfortable division of responsibilities between the minis that recorded the data at our experiments, and the mainframes that provided the subsequent high quality batch processing. The growth of the acceptance of laboratory-wide general-purpose timesharing during the latter half of the 1970s helped to amortise the high costs of mainframes, since they played a dual role, acting as major timesharing providers during the day and offering major batch processing capacity outside prime shift.

More recently workstations have set another level of improved price-performance, and the battle between low-end workstations and high-end PCs seems likely to ensure continuing progress in this area. For a while the strongest supporters of mainframes felt that they would assume an obvious role as data servers. There seem to be two problems with this idea; first of all the arrival of SCSI disks means that the cost of data storage on the mainframes is now noticeably more expensive than on workstations, and secondly the cost of driving network bandwidth from mainframes appears to be very high.

CERN has always paid considerable attention to the quality of its computing services, and it has been making serious experiments over the past two years with centrally organised services based on clusters of workstations. As discussed above, we can report that the results of these experiments are encouraging.

Unless the price-performance ratio of mainframes improves quickly and dramatically we can expect to see a increasing fraction of CERN's investment go towards centrally organised workstation services in the future. Price-performance ratios appear to be moving against the mainframes and in favour of the workstations even for heavy scientific computation with serious I/O requirements, and it is not obvious that massively parallel processors will succeed to reverse this trend.

The manpower costs and reliability of distributed RISC systems are far more favourable than they were for minicomputers a few years ago. Dependency on networks is now universal, even though network bandwidth is still far too low, particularly over long distances. If present technology trends continue, and hopes for wide-area broadband networks are realised, the goal of an integrated global distributed system for particle physics data analysis will be attainable.

Acknowledgements

Much of modern science and technology requires the common efforts of large groups of people. This is especially true of the type of distributed computing environment that has been discussed here. It is a pleasure to acknowledge the input and hard work of colleagues in the Computing and Networks Division at CERN, of colleagues in other CERN Divisions, and of those particle physicists from Europe and elsewhere who never cease to express their requirements with clarity, thereby keeping us up to the mark.

Brian Carpenter and Tony Osborne helped me considerably in the production of this paper.

References

1. Parkman, C.F. and J.G.Lee, OMNET High Speed Data Communications for PDP-11 Computers, CERN 79-11, Geneva, 1979.
2. Gerard, J.M. (ed.), CERNET, a High Speed Packet Switching Network, CERN 81-12, Geneva, 1981.
3. CERN Program Library Manual, available from Program Library Office, CERN, 1211 Geneva 23, Switzerland.
4. Shiers, J.D., Distributed Storage Management in High Energy Physics, Proc. Eleventh IEEE Symposium on Mass Storage Systems, Monterey, October 1991, 109-112, IEEE Computer Society, 1991.
5. Baud, J.P., C.Boissat et al., Mainframe Services from Gigabit-Networked Workstations, Proc. Summer 1992 USENIX Conference, San Antonio, 165-176, USENIX Association, 1992.
6. Altaber, J., B.Carpenter et al., CHEOPS: Really Using a Satellite, Computer Networks and ISDN Systems, 23(1-3)131-142, 1991.
7. Berners-Lee, T., R.Cailliau et al., World-Wide Web: the Information Universe, Electronic Networking 2(1)52-58, 1992.