SSCL-Preprint-173

CONF-920966--27

SSCL-Preprint-173
September 1992
Distribution Category: 400

L. Cormell

Coping with Distributed Computing

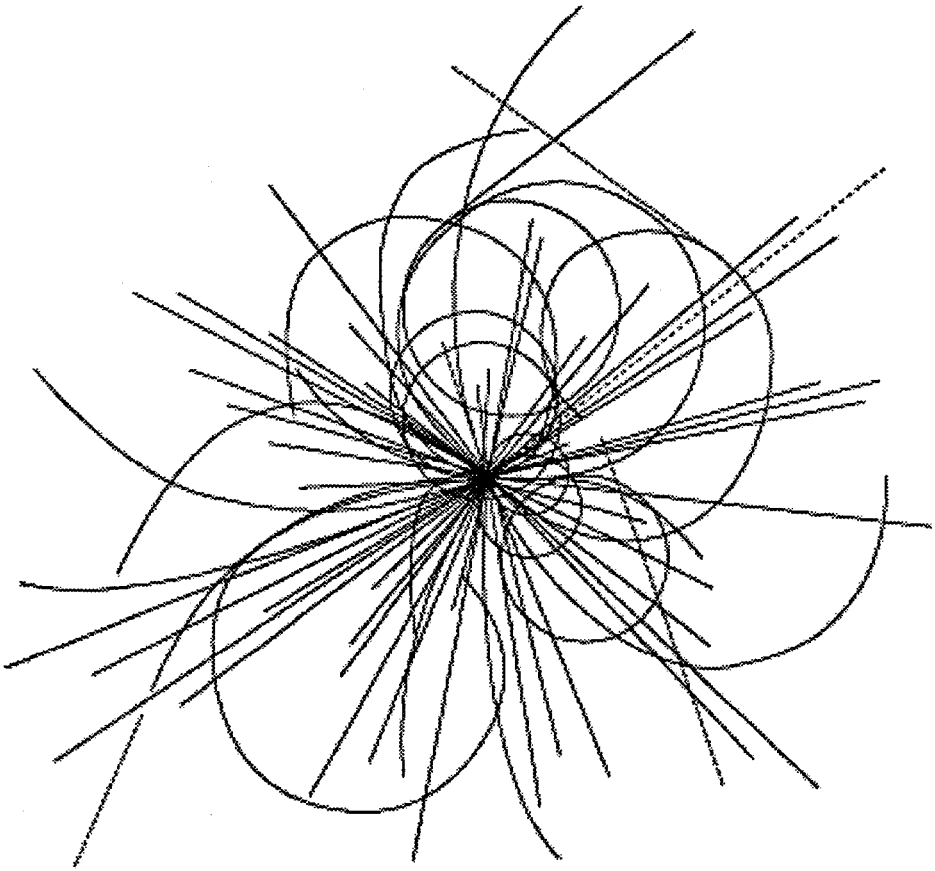

Superconducting Super Collider Laboratory

APPROVED FOR RELEASE OR
PUBLICATION - O.R. PATENT GROUP
BY.....DATE.4/2/93

Disclaimer Notice

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government or any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Superconducting Super Collider Laboratory is an equal opportunity employer.

SSCL-Preprint-173

Coping with Distributed Computing*

L. Cormell

Superconducting Super Collider Laboratory†
2550 Beckleymeade Ave.
Dallas, TX 75237

September 1992

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

---

*Presented at the 1992 CHEP Conference, Annecy, France, September 20–25, 1992.

†Operated by the Universities Research Association, Inc., for the U.S. Department of Energy under Contract No. DE-AC35-89ER40486.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Coping with Distributed Computing[§]

Laird R. Cormell
Superconducting Super Collider Laboratory[†]
Dallas, TX 75237-3946

The rapid increase in the availability of high performance, cost-effective RISC/UNIX workstations has been both a blessing and a curse. The blessing of having extremely powerful computing engines available on the desk top is well-known to many users. The user has tremendous freedom, flexibility, and control of his environment. That freedom can, however, become the curse of distributed computing. The user must become a system manager to some extent, he must worry about backups, maintenance, upgrades, etc. Traditionally these activities have been the responsibility of a central computing group. The central computing group, however, may find that it can no longer provide all of the traditional services. With the plethora of workstations now found on so many desktops throughout the entire campus or lab, the central computing group may be swamped by support requests. This talk will address several of these computer support and management issues by providing some examples of the approaches taken at various HEP institutions. In addition, a brief review of commercial directions or products for distributed computing and management will be given.

1. Introduction

Dramatic changes have occurred throughout the computing industry and its user community since the revolution of distributed desktop computing began in the 1980's. The abundance of cheap, powerful computing engines is driving this revolution at a feverish pace. Participants in the revolution are the high speed networks and distributed peripherals such as laser printers, SCSI disks, and 8-mm tapes that make distributed computing a viable option to the more traditional mainframe-oriented, centralized computing environments. Networked desktop computing has changed forever our computing and working environments. In fact, the sociology of how we do our work has also changed. Some of that change can be traced back to the interactive time-sharing environments of the 70's, of course. But in the 90's environment, a typical user is no longer dependent on the central computing organization to meet his needs. The computer has become another tool in the integrated office/work place of the 90's. Around the corner lie many new integrating tools in the area generally referred to as groupware that should further expand the utilization of computing in the office. These products include integrated phone, voice mail and e-mail services, desktop video conferencing, multi-media applications, etc. While a discussion of groupware is beyond the scope of this paper, planning for the management and support of these products must begin soon. There are many management issues to be addressed, but at this time there does not seem to be that many solutions readily available. We will provide a few examples of the methods employed by several HEP labs to address these problems and then discuss some efforts under way in the computer industry.

2. Distributed Computing Management Issues

We've already alluded to the fact that there are both good and bad aspects of distributed computing environments. There are some simple, perhaps obvious, observations that we can make. First the advantages.

2.1 Advantages of Distributed Computing

Some of the advantages of distributed computing include:

- Distributed computing working environment is good. Most people do most of their work in their offices. Efficiency and productivity increase when computing and other resources are available in the office work place.
- Working group clusters are good. Most people work in groups on problems common to that group. A distributed environment that is sub-netted by group can be be very effective in providing common shared resources that are group specific. For example, an engineering group may need CAD and finite element analysis packages, but have no need for the CERN Program Library. Client/server architectures can be utilized to their fullest capability with minimal effect on the network backbone.
- Dedicated resources are good. Placing a cheap RISC work-stations on every desktop can virtually eliminate time sharing. A typical user will have his own dedicated resource.
- Independence is good. Individual users and working groups relish the independence they have when in control of their own environments. They can choose configurations that are tailored to their needs rather than the general purpose environments found in centralized mainframe environments.
- The economics are good. First, RISC/UNIX workstations and personal computers, for that matter, are cheap. And as mentioned above, working groups or individuals buy only the applications and tools that meet their needs. They configure the hardware to meet their needs alone, no more. Shared multiple users licensing can keep the cost of software down. Installation and maintenance fees are typically small.

2.2 Disadvantages of Distributed Computing

There are likewise a number of disadvantages associated with distributed computing environments.

- Distributed computing working environment is bad. The individual user must become his own system manager to some extent. He must become concerned with issues of installation maintenance, updates, backups, etc. If he is not prepared to perform these functions himself, where does he turn, who provides the support?
- Working group clusters are bad. Individual working groups may chose to handle their own support issues. This can lead to duplication of efforts in several service areas. A further complication arises when different groups develop completely different solutions to the same problem.
- Dedicated resources are bad. When dedicated computing is provided for individual users, there are typically many compute cycles and other resources going to waste.
- Independence is bad. Individual users and working groups often don't have the experience or knowledge to configure their environment to meet their needs in an effective efficient manner.
- The economics are bad. While it is true that RISC/UNIX workstations are cheap, there are often hidden costs. There is a need to build up a support staff familiar with several UNIX (or PC) platforms. There are potentially several areas for duplication of efforts. The purchase of software and services may not be coordinated throughout the site and thereby lead to wasted expenditures.

So where does this leave us? While there are certainly some disadvantages associated with the management of distributed computing environments, most of the arguments come down in favor of such environments. The raw computing power, versatility, and freedom offered by distributed computing at an unrivaled price/performance point makes them extremely attractive. The "pros and cons" of centralized vs. distributed computing environments may be debated at length, but there is no escaping from the fact that the era of distributed computing is upon us.

2.3 Management Issues

In the era of distributed computing and heterogeneous environments, diverse systems are networked throughout the HEP community. Unfortunately, the diversity that lets users choose the system that best meets their needs also creates an administrative nightmare: it may require system administrators to develop different management scheme for each hardware platform. It is possible to integrate large multi-vendor systems when standards can be applied. These standards, however, are not all employed in the same way by every vendor. Most users enjoy the independence offered by RISC/UNIX workstations, but many of them still expect (and sometimes demand) the services normally offered by more centralized operations. System managers must find novel ways to deal with these expectations in multi-vendor configurations; they must address several issues including:

- Training. Support personnel must receive training for several different hardware platforms and operating systems. Users also need training. Clearly UNIX managers must provide training for their personnel, but who trains the users?
- Installation and setup. Workstation vendors are typically operating on small profit margins and may not provide free installation. Users often don't want to pay for the additional installation and setup services. But the typical physicist/user may have little or no experience in system management or much knowledge about UNIX environments. How do UNIX managers provide the support?
- Hardware maintenance. UNIX workstations are inexpensive and fairly reliable, but they still must be maintained throughout their life cycles. There are typically product enhancements, firmware fixes, and upgrades offered by the vendors that extend the useful life of the equipment. While system administrators are familiar with this phenomena, it's not clear that the end-user "keeps up" with the latest improvements or that he knows what to do. Who should provide the maintenance and ensure that the system is current?
- Software maintenance. Likewise for software upgrades, only more so. Bug fixes, patches, and new software versions are released on a regular basis. How does a UNIX manager install these on all platforms, how does he maintain configuration control, who tests new versions for compatibility?
- Network management. In many distributed environments the network is an integral part of computing, or as some have stated: "the network is the computer". In working group clusters files and applications are shared in a client/server arrangement; workstations and X-terminals are mounted to hosts; mail, ftp and other communication services are crucial. The HW and SW issues mentioned above apply to NW management as well. In addition there are questions of how to configure the network to provide optimized access to file servers, printers, etc.
- Distributed services. There are a number of services traditionally provided in centralized configurations that must be dealt with in new and innovative ways. How does a system administrator provide uniform printing, backup, licensing, batch, and tape services that users require?
- Accounting. How does an administrator in a distributed environment collect statistics on system utilization, disk usage, print service, and traffic over the network? At some sites the computer group provides resources on a "charge-back" basis. How should these charges be assessed in a distributed environment?
- Documentation. The usefulness of a system is often determined by the quality and accuracy of its documentation. How can the UNIX manager provide current documentation to all the users when some are located not only on the other side of campus but on the other side of the world?
- Host management. In a distributed environment the system administrator is faced with the difficulty of providing user and group accounts and privileges on a number of different hosts. In addition operations on the system such as shutdown/reboot, naming, addressing etc. must be performed.

Obviously there are a number of challenges to be faced in managing distributed systems. So, how do we cope with distributed computing? In the next section we discuss some examples taken from FNAL, CERN, SSCL, and LBL. In the section following that we will identify some new developments from the commercial side that may provide some solutions.

3. Distributed Computing Management Examples

System administrators at every HEP lab have faced these difficulties to one degree or another. The following descriptions are based on information provided by UNIX system managers at their respective laboratories.

3.1 Management of Distributed Computing at FNAL [1]

The first goal at Fermilab is to create a uniform environment. This has been the goal of the Fermi UNIX Environment and products such as UNIX Product Support (ups), UNIX Product Distribution (upd), UNIX Products Pull (upp), and several utilities developed by the Distributed Computing Department (DCD). The Fermilab UNIX Environment has these common aspects [2-6]:

- Provides a single way of obtaining products through subscription to the upd system.
- Automates the maintenance of product updates using the automatic "pull" features of upp.
- Provides a consistent way for users to maintain their environment through access and installation of ups products.
- Provides a consistent login environment, including standard cshrc files, and DCD utilities such as flpr, logdir, less, etc.
- Provides some basic system management facilities such as adding accounts, resetting passwords, etc that are consistent across all UNIX platforms. Management data are recorded in a central uid and gid database.

The second goal at Fermilab is to provide a consistent scheme for systems administration. This goal has been more informal, but the Distributed UNIX Support Group has attempted to create standards for naming conventions, for example, so that systems throughout the site are setup uniformly. They have achieved a reasonable degree of success, thus far, in several key areas including:

- Uniform backup procedure. They have developed a system called Fermi Modular Backup that does remote backups.
- Naming conventions. They are developing a consistent scheme for mount points and for use of the auto-mounter.
- AFS. They are participating in an evaluation of the Andrew File System (AFS).
- New compilers and OS. They have developed formal procedures for evaluating and certifying new releases of the OS and compilers. These procedures include evaluation and testing by UNIX System Support, certification by interested groups with some beta testing and user feedback at selected sites prior to installation.

3.2 Management of Distributed Computing at CERN [7]

The UNIX Support Group in the CN Division at CERN consists of 6 FTEs. They currently support about 300 DECstations, (and a few VAXstations running ULTRIX), 350 SUNs, some 150 HP Series 700, 300 Apollo/Domain systems; and about 20 RS/6000s. They try to support only system administrators for the various groups rather than end-users, but small groups and single users don't have system administrators, of course, so they must deal with end-users as well. The group only considers tools which operate across multiple families (e.g. they recommend Motif and have pre-purchased Motif licenses for SUN from DEC which they resell to users). They have written installation guides for ULTRIX and HP-UX and another is in preparation for SUN. These tailor the vendor's own guides for use specifically at CERN (e.g. how to access CERN name and time servers, how to spool all printing through a central

print server, etc). They feel that this has reduced their load to a point where they are able to cope with the rapid workstation population explosion (up to 50 new stations per month) without adding new staff. The staff is encouraged to spend up to 25% of their time on development projects such as tests of AFS, Kerberos, OSF beta tests, etc.

They provide a number of central workstation servers for the distribution of vendor software kits to CERN users. Distributed shared licensing is a problem. HP and DEC have remote distribution tools such as Remote Installation Service (RIS). It is difficult to prevent further copying of the applications after the initial copy is distributed. The vendors themselves should fix this problem.

Backup is currently under study. They have been evaluating some tools. For HP and Apollo they recommend Omniback and are evaluating DECnsr. They have a joint project with IBM to port their WDSF client to HP-UX and ULTRIX. They use WDSF inside the Division as part of its evaluation and are currently studying how to offer a wider service without flooding the network and the IBM disc catalogues.

3.3 Management of Distributed Computing at LBL [8]

In the multi-purpose lab environment of LBL, the central computing group is largely a self-supporting cost center. They provide support for VMS, UNIX, computer graphics, output services, user services, and workstations (including PC and MAC). They charge about $200 per month per workstation for system management and backup. This cost includes system and layered product software maintenance. They currently manage about 150 workstations, mostly Sun Sparcstations, but also a few SGIs and HPs. This represents roughly one-third of the UNIX/RISC workstations at LBL. Their experience has shown that 1 FTE can provide quality support and management for about 4 file servers and 20 diskless workstations [9]. Recent implementation of Modules has simplified greatly the management of commonly-used software packages. they maintain one central machine for users to mount these files. As this machine is essential to the operation of many workstations, it is on UPS and has a RAID disk system.

Both the VMS and UNIX projects are currently self-supporting recharge operations. Support of a critical mass of capable systems programming personnel is enabled by income from both central machine recharges and workstation management fees. Further eroding of central machine income would force an increase in the system management fee, which at ~$200 per month already meets resistance.

They maintain a distributed printing system which features over 200 Laserwriters site-wide on LocalTalk LANS connected to the LBLnet ethernet subnets. A software queueing system built upon lpr provides uniform access to these printers from varying hosts. Twenty of these printers are in public areas, the remainder are user-owned. There is a small fee per page printed. Programming and operational support for this system has been paid for by overhead funds. They are in the processing of adding color printers and plan to raise the rates over the next two years to cover all support costs.

In a highly-competitive open market, they feel that they are surviving the transition from central cycles to workstation support. The future is uncertain, as the number of workstations increases. Hopefully, system vendors will provide tools to simplify management to the point where it can be done easily.

3.4 Management of Distributed Computing at SSCL [10]

The SSCL is committed to developing a distributed computing environment. Use of a centrally controlled but otherwise decentralized and distributed environment allows maximum functionality and flexibility of access to shared data, print services, electronic mail, backup systems and nameservice while minimizing network overloads across the primary backbone.

Operations support. The Informations Services Operations Group has 25 employees that provide all levels of support from Help Desk assistance to system programming on all central computing resources at the SSCL. They support five operating systems (UNIX, VMS, MS/DOS, Novell, and MAC-OS.) and approximately 3000 user accounts. The central resources include two VAXclusters, the PDSF, QuickMail servers, Appleshare servers, and

several Novell servers. There are approximately 1500 Macintosh systems, 300 PCs and 700 UNIX workstations. Some groups (particularly those in Accelerator Systems Division) provide their own system support for specific functions.

Distributed servers. At present, management functions and critical applications are primarily performed using centralized server systems. The network impact of centralized data has reduced the system's efficiency somewhat. They plan to place a number of servers in strategic laboratory locations. These server locations are being selected on the basis of user population density and networking considerations. The critical applications and services to be distributed include the various file systems (AFS, NFS, Appleshare, and PC volume services), network applications (nameservice, NTP, and USENET), print services (IMPRINT, Alisashare, and Pathworks), mail gateway services, BITNET routing , and software metering for Macintosh and PC software permitting efficient sharing of infrequently used desktop applications.

Distributed Backups. The goal at the Lab is to provide administratively centralized, unattended backups on most SSCL servers by the end of '92. The volume has grown to greater than 250 GB of routinely backed-up data on UNIX, VMS, Macintosh, and Novell systems. In most cases the data is copied to 8-mm tapes in stacker units from remote systems over the FDDI backbone. System backups are currently performed using native operating system utilities. These utilities all fall short in some areas and none provide tape management or simple user archival capabilities.

Site Licensing. The SSCL is site-licensed for the SUN Operating system and is in the process of procuring a site-license for SUN unbundled products. Procedures have been established for the distribution of copies of media and documentation from a centralized source to technical contacts throughout the SSCL. Other software that does not have licensing restrictions is made available through AFS or FTP. Client-server software (particularly that which is run on MACs or PCs) is a difficult problem. The SSCL is currently working on procedures to control and automate distribution and installation of client software through use of tools such as Timbukto for Macintosh.

Distributed Printing. The Operations Group has developed a distributed printing utility, lwprint, that permits access to most of the laser writers on AppleTalk or Ethernet. Print jobs are spooled and queued by a central UNIX print server. Users can select the printer of their choice from an interactive menu. Post script is the default format and several options such as landscape and two pages per sheet are being implemented [11].

4. Distributed Management Environment (DME): Plans and Status

It's a reasonable thing to expect system management solutions from the computer vendors. After all, traditional services such as printing, backups, accounting etc. have always been provided by the mainframe or mini-computer vendors. The bad news is that UNIX vendors have not provided integrated management tools to date, but the good news is that when they do, the systems should be open and interoperable. Some key developments in this area are coming from the Open Software Foundation's (OSF) DME [12,13,14]. The DME project is being developed in a manner similar to OSF's Distributed Computing Environment (DCE). DME will build upon previous efforts by the OSI/Network Management Forum, the Object Management Group (OMG) and its own DCE. DME products are being developed as object-oriented systems where data-like objects or process-like objects are managed throughout the entire network. The products will be integrated into an overall layered architecture that isolates the user or system administrator from the OS and other system peculiarities. The OSF has accepted several currently available products as DME technologies. These products will be refined and evolve into the final suite as additional technologies are accepted. The current plan is to provide DME developer's kits in '93 with delivery of the full environment in early '94.

3.4 DME Architecture

OSF's vision of DME is to unify the management of systems in a heterogeneous environment. Key requirements of this vision include:

- Consistency. DME will provide a common "look and feel" GUI for managing all networks, systems, and applications.
- Interoperability. DME will be based on network standards such as SNMP and will share a common definition of objects across platforms. Gateways to proprietary systems will be possible.
- Scalability. DME will follow a three-tiered hierarchy of management levels that creates a more easily controlled environment. The DME will be scalable from single nodes to large enterprises. It will provide a more centralized management of distributed environments.

The elements of theDME architecture are shown in Figure 1. At the highest level is the management user interface and at the lowest are the object services that interact with an object request broker and the OS.

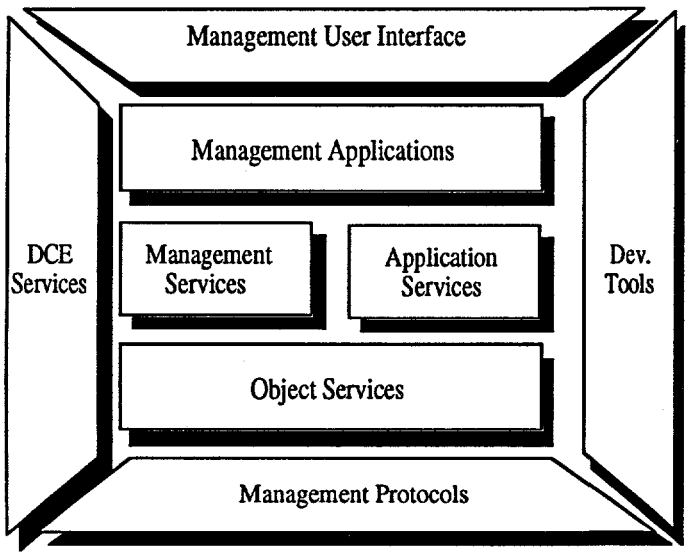

Figure 1. DME Architecture

- Management User Interface includes a Motif-based GUI, command-line interface, and a network mapping and management interface. Selected products are WizDOM from Tivoli, OpenView from HP and OSF Motif.
- Management Applications will provide host management from Tivoli WizDOM Host/User/ Subnet (HUGS) Services and Primary Resource Management (PRM). The services include remote rebooting, workstation configuration, host changes and Kerberos security.
- Management and Application Services include software distribution and installation, networked print services, licensing, event notification and event management. Distribution and installation services are provided by a joint IBM/HP Software Distribution Utilities product. Distributed printing is provided by MIT (Project Athena) Palladium Print System, v2. Licensing services will be provided by HP's NetLS v2.0 from HP and Gradient Technologies. Event management services from Banyan Systems and event notification service from Tivoli PRM will provide the system administrator rapid and automatic notification of system problems throughout the network.
- Object Services maintain DME objects. In DME system resources and applications are encapsulated as object. Management request brokers register, locate, and authorize access to managed objects.
- Management Protocols facilitate communication between managed objects. The DME protocols will include CMIP from HP, SNMP from HP/IBM, and RPC from OSF.

In addition DME will provide interfaces to DCE, development toolkits, and PC integration capabilities.

5. Summary and Conclusions

High energy physicists are resourceful and creative. Given that, it's not surprising to see: first, that physicists, ever eager to obtain more computing power, have exploited RISC/UNIX systems with a passion; and second, that they have attempted, with considerable success, to develop integration and management tools to make effective use of these new systems. Some of the products such as those found in the Fermilab UNIX environment are quite useful.

But we should expect, in fact demand, tools and integrated packages from the workstation vendors. These products should be interoperable, coherent, scalable, and affordable. There

are several promising efforts underway, by groups such as OSF, but how long can we wait and how successful will their products be? As a concluding remark, or question rather, shouldn't issues of this type be taken up by HEPix?

References

§ Presented at the Conference on Computing in High Energy Physics (CHEP92), Annecy, France, Sept., 1992.

† Operated by the Universities Research Association, Inc. for the U.S Department of Energy under Contract No. DE-AC02-89ER40486.

[1] Matt Wicks, FNAL Distributed Computing Group, private communication.

[2] J. Nicholls, Fermilab UNIX Environment, in the Proceedings of the International Conference on Computing in High Energy Physics 91, Tsukuba, Japan, p. 605, Y. Watase and F. Abe, ed., Universal Academic Press, 1991.

[3] Overview of Product Installation and Distribution on UNIX Machines, Fermilab report GU0006, July, 1992.

[4] UPS User's Guide, Fermilab report PN426, v2.4, April, 1992.

[5] W. Bliss, et al., UNIX Product Distribution User Guide, Fermilab report PN435, v2.1, June,1 992.

[6] V. Gurbani, et al., upp User's Guide, Fermilab report PU0135, March, 1992.

[7] Alan Silverman, CN Division, CERN, private communication.

[8] Marv Atchley, Information and Computing Services, LBL, private communication.

[9] C. A. Eades, The Role of the UNIX Central Computing Facility in a Multi-Purpose National Laboratory, in the Proceedings of the International Conference on Computing in High Energy Physics 91, Tsukuba, Japan, p. 583, Y. Watase and F. Abe, ed., Universal Academic Press, 1991.

[10] Brenda Ramsey, Information Services Department, SSCL, private communication.

[11] Dave Somogyi, Information Services Department, SSCL, private communication.

[12] Distributed Management Environment, An Overview, Open Software Foundation, September, 1991.

[13] Distributed Management Environment, Rationale, Open Software Foundation, October, 1991.

[14] OSF Distributed Management Environment (DME) Architecture, Open Software Foundation, May, 1992.