# Project Athena: Supporting distributed computing at MIT

by J. M. Arfman
P. Roden

***Project Athena™ was an educational computing initiative at the Massachusetts Institute of Technology, undertaken in partnership with the IBM Corporation and Digital Equipment Corporation from 1983 to 1991. This paper gives an overview of the network-based distributed computing services, developed for a number of UNIX™-capable workstations. These services are extensions to the native operating systems of the workstations, and provide interoperability as well as systems administration facilities in a large heterogeneous workstation environment. Under Project Athena, a mature distributed computing environment was developed. Its organization and support structure may be used as a model when planning a new installation, whether on a university or commercial campus. A section of this paper deals with the support requirements for distributed computing environments, based on the Project Athena experience.***

Project Athena**, the Massachusetts Institute of Technology's widely known model of distributed computing, was launched in 1983 by MIT in partnership with the IBM Corporation and Digital Equipment Corporation (DEC) to address the Institute's need to provide widely available computing resources to the campus. The project partnership ended on June 30, 1991, eight years after it began. The Athena** computing environment has been officially adopted by MIT as its infrastructure for delivery of educational as well as selected research and administrative applications.

An MIT study in the early 1980s on the use of computer technology in undergraduate education observed that there was very little, if any, computer usage for educational purposes on the undergraduate level, compared to that in graduate education and research. The study originated in the School of Engineering, and DEC was approached as the single partner. However, the recommendation was soon made to create an extended project that would provide an undergraduate computing environment to benefit all of the MIT undergraduate curriculum. At the same time, additional sponsors were sought, and IBM and DEC officially became partners. Donations from many other companies and individuals were also received, particularly in the early years.

The designers of the Project Athena system chose to implement a UNIX**-based distributed workstation environment. This decision was made knowing full well that this technology was not available in the marketplace at the start of the project. The goals were to create a computing environment that would scale up to 10 000 workstations and accommodate heterogeneous hardware, but yet be "coherent." This concept means that a user could go to any workstation and access any files or applications without finding major differences in the user interface and service delivery. The Berkeley UNIX base was chosen originally to achieve this coherence on the IBM and DEC platforms, with IBM/4.3 on the RISC Technology (RT*) and 4.3 BSD** (Berkeley software distribution) for Digital's workstations.[1] The current and future environment will include Advanced Interactive Executive (AIX*) for IBM's (Personal System/2* (PS/2*) and RISC System/6000*, Ultrix** for DEC's workstations, A/UX** (Apple UNIX) for the Macintosh**, and others.

The concept of the public workstation was also important as it freed the users and their files from being associated with unique physical workstations. This system would be secure, "location-independent," centrally administered but physically disparate, and provide a consistent graphical user interface. It would be easy to install and update software. The concept had far-reaching implications for system design. It took almost five years to make it a reality, and was the reason for the extension of the original project from five to eight years. The last three years of Project Athena were dedicated to improving and solidifying the system infrastructure. As a result, computer usage at MIT has increased beyond all expectations, and the impact of this technology is growing and receiving worldwide recognition.

The system is totally transparent to the user. Project Athena's services make it possible to go to any workstation and access files and applications that can be resident anywhere in the system. The user provides a logon userid and password; the system services recognize and authorize the user to use the system.

MIT's students have embraced Project Athena enthusiastically. Over 95 percent of all undergraduates are regular users. Athena has had a significant impact on how they work and on how they interact with fellow students and faculty. A small but growing number of faculty feel that Project Athena also changed their teaching methods, making their face-to-face contact with the students more productive, and helping them develop intuitive learning in their students.

In retrospect, the challenges went beyond the scope that was originally envisioned. On the other hand, the significant accomplishments at the system design level were also far beyond anything expected at the time. For instance, X Windows**, which came from MIT, led to the establishment of the X Consortium, and is now an industry standard. More recently, the Kerberos** authentication mechanism, first implemented at Project Athena, has been adopted (with Hewlett-Packard extensions) by the Open Software Foundation (OSF) for inclusion in its Distributed Computing Environment (DCE**). Three other technologies were submitted by Project Athena to OSF in response to its call for Distributed Management Environment (DME**) technology: Moira** (service configuration management), Zephyr** (on-line notification), and Palladium** (print services), the latter together with IBM, DEC, and Hewlett Packard.

## The Athena computing environment

The following sections are devoted to describing the physical layout and architecture of Athena, as well as the services provided.

**The MIT campus network.** Project Athena does not manage MITnet (MIT network), the network on which it relies to distribute traffic. However, the impetus for MITnet did come at least in part from the reality that the success of Project Athena would depend on reliable transmission capabilities over the entire campus, as well as beyond it.

MITnet consists of a 100 Mb/sec fiber optic FDDI (fiber distributed data interface) spine, extending the entire length of the campus, or almost two miles. The spine does not have any systems directly attached to it, as shown abstractly in Figure 1. Rather, it has a number of routers to which various subnets are attached.

MITnet is connected to the rest of the world via the New England Academic and Research Network (NEARNET), which is one of the midlevels connected to the National Science Foundation Network (NSFNET). MIT has been awarded a contract to be the Nodal Switching System (NSS) for NEARNET, and is directly connected to the NSFNET backbone. A cluster of Athena workstations was installed 80 miles away from the MIT campus at the Woods Hole Oceanographic Institute, and is connected to MITnet by a series of microwave links. MIT's connectivity creates "communities of scholars" electronically. Students and faculty have embraced this method of sharing information on a global level. The approximately 4000 systems at MIT, ranging from personal computers to a Cray-2**, are connected via Ethernet subnets, attached to the spine via rout-

**Figure 1 The Project Athena computing environment**

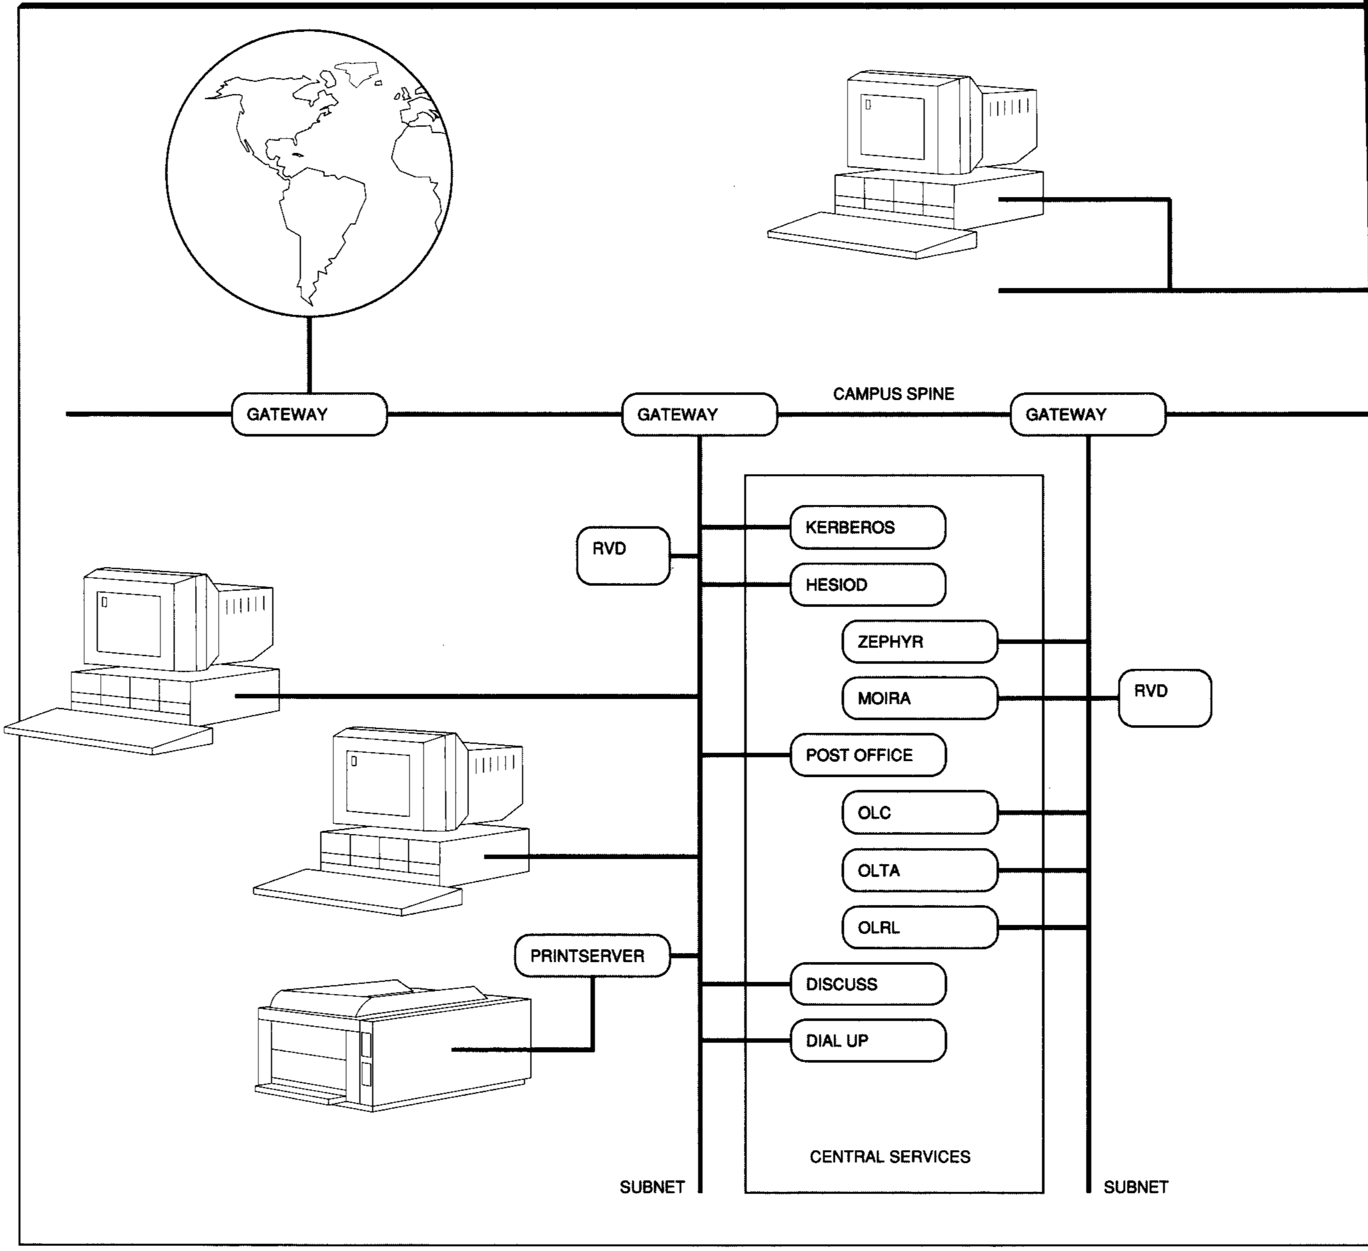


ers. While the number of systems managed by Project Athena is about 1500 out of the total 4000, the traffic generated by Project Athena accounts for over 50 percent of spine usage.

MITnet is partitioned into 40 Ethernet 10 Mb/sec subnets, which support Athena's more than 1300 advanced function workstations, 80 file servers, 40 print servers and 125 associated printers, as well as other MIT systems. As Figure 1 shows, every one of the subnets has at least one system software delivery server. Since only the absolute minimum of boot and communication software is kept on the workstation, the bulk of system soft-

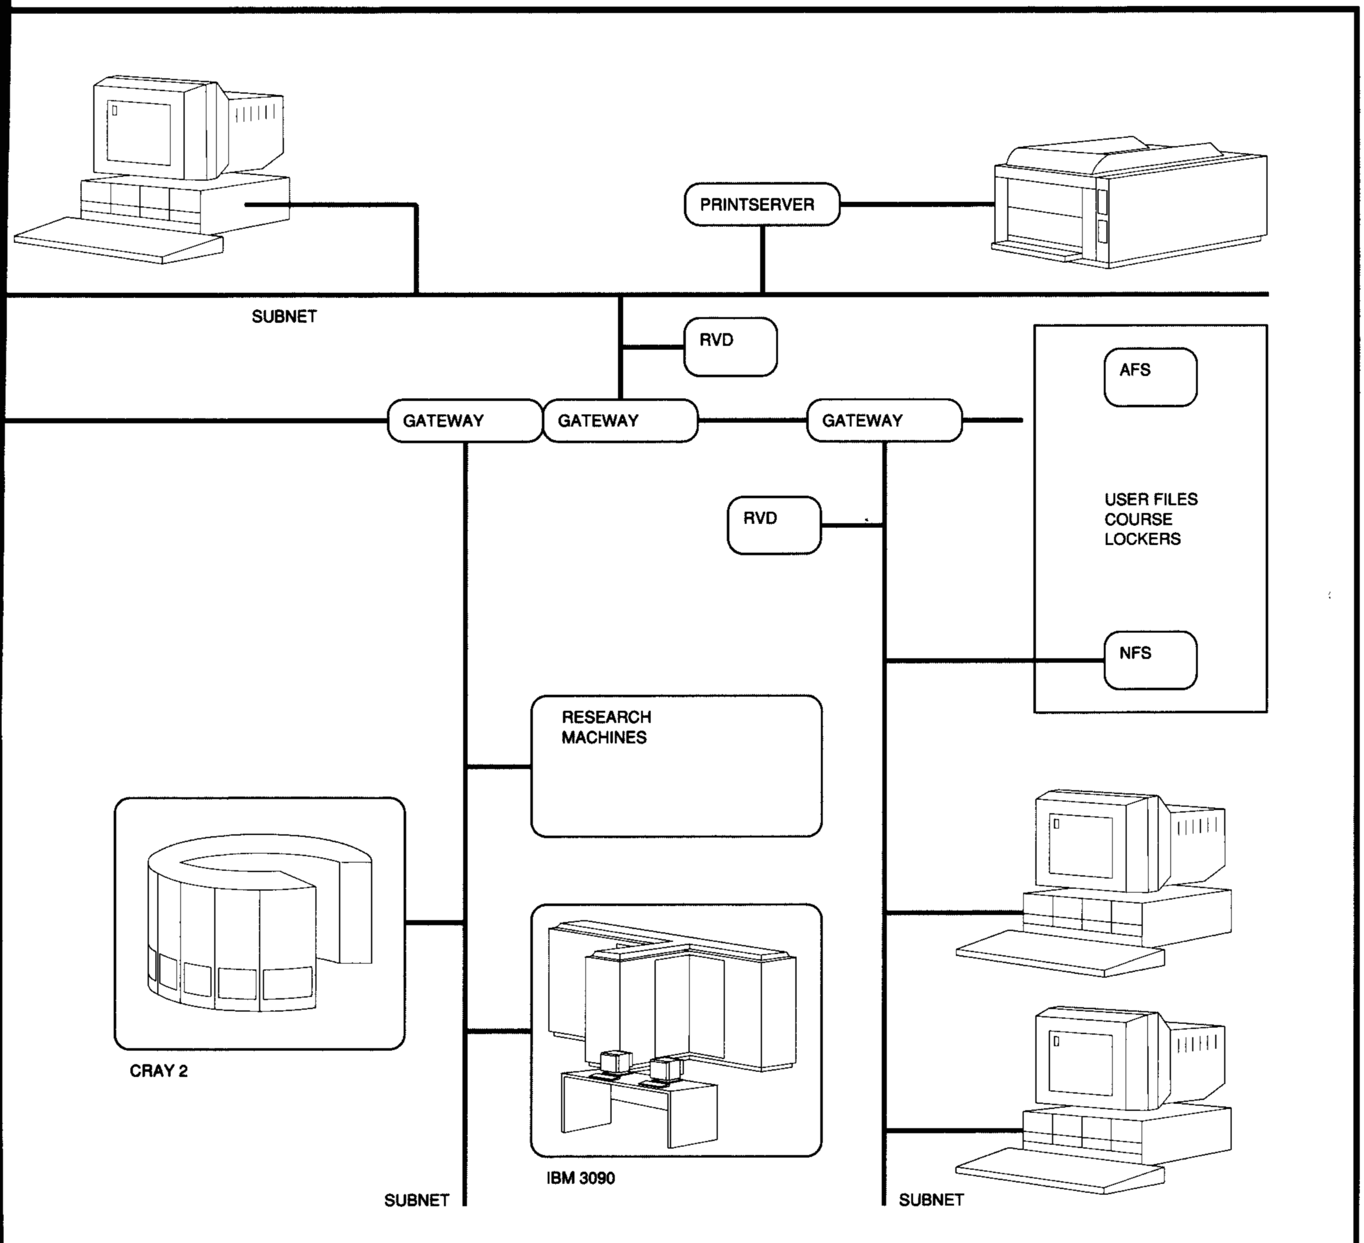


Redrawn with permission of the Massachusetts Institute of Technology

ware has to be provided rapidly on demand. Each subnet has one of these servers, with the binaries for every type of workstation present on that particular subnet.

While the initial focus of Project Athena was on undergraduate education, the implementation now serves the whole MIT campus as a delivery system. In 1989, a decision was made to grant graduate students Athena accounts. As of November 1990, about 70 percent of all graduate students had accounts, in addition to 98 percent of all undergraduates. The IBM and Cray systems, which are used for administrative and research purposes, are also accessible from Athena workstations. Someone who has an account on one of these systems can access mainframe applications within an Athena window. MIT's strategic goal for distributed computing is to make the Athena computing environment the service delivery mechanism for all instructional, research, and administrative applications, such as the library catalog, the Registrar's Student Information System, and others.

**The "dataless" workstation model allows a user to log on anywhere on the network.**

**File servers.** The user and courseware file servers, as well as those for central services, can be located anywhere on the network. Project Athena initially adopted, with security modifications, the Network File System (NFS**) from Sun Microsystems, Inc., an industry standard. Athena also uses the Andrew File System (AFS**) from Transarc Corporation, and a migration to AFS is under way. This file system, originally developed by Project Andrew at Carnegie Mellon University, has superior distributed file management, performance, and backup characteristics. The third type of server, the Remote Virtual Disk (RVD), was developed at MIT in the Laboratory for Computer Science to provide a fast, read-only, binary server for each subnet. Typically, the subnet RVD contains the code for the operating systems of the workstations in that particular subnet.

Since Project Athena has adopted the "dataless" workstation model, and a user can log on anywhere on the network, the location of the user's files is irrelevant to the individual and is managed through a mapping from logical to physical addresses by the Hesiod** name server. The *attach* command is an abstraction that allows the user to be unconcerned about the actual physical location and structure of the files.

User files, courseware, and licensed software are stored on the file servers in "lockers," hierarchical directory structures, e.g. NFS mountpoint or AFS volume. Eventually, both the RVD and NFS file service for users' home directories and system software delivery will be replaced by AFS. The user file and course locker servers are shown at the upper right in Figure 1. The total amount of backed-up storage on these servers exceeds 40 gigabytes. At present, students are given 1.2 megabytes of backed-up file space, with more available if the user can get approval or is willing to do a private backup. This is not so much a limitation in the system design as a financial decision based on the number of users and the number of file servers available, taking into account the available backup facilities, technology, and personnel.

**Central services.** A client-server application is considered a "service" offered on the network. Some services are required for the system to function, such as authentication, and others are value-added services, such as the on-line help system. These services are centrally administered, but are physically located so as to minimize impact on users, in the event that portions of the network are unavailable.

The Central Services, shown in the left half of Figure 1, consist of the suite of Athena network services, developed over the lifetime of Project Athena. Together, they constitute a complete distributed systems environment, which is scalable up to at least 10 000 workstations. The architecture of the Athena system is shown in Figure 2, layered on top of the operating system and communication protocol. While the systems services are shown together in Figure 1 for the sake of schematic clarity, this does not imply that there is only one server of each type, or that they are all located on the same subnet. On the contrary, it is necessary for continuous availability to have at least the crucial services replicated across the network. Other services require multiple servers because the number of servers grows with the number of users.

**Kerberos.** A consequence of the public workstation model is that authentication, or identification of the user, is no longer relevant at the workstation level. All public workstations are considered "untrusted," to the extent that the root password is published in all documentation. To authenticate

**Figure 2 Project Athena system architecture**

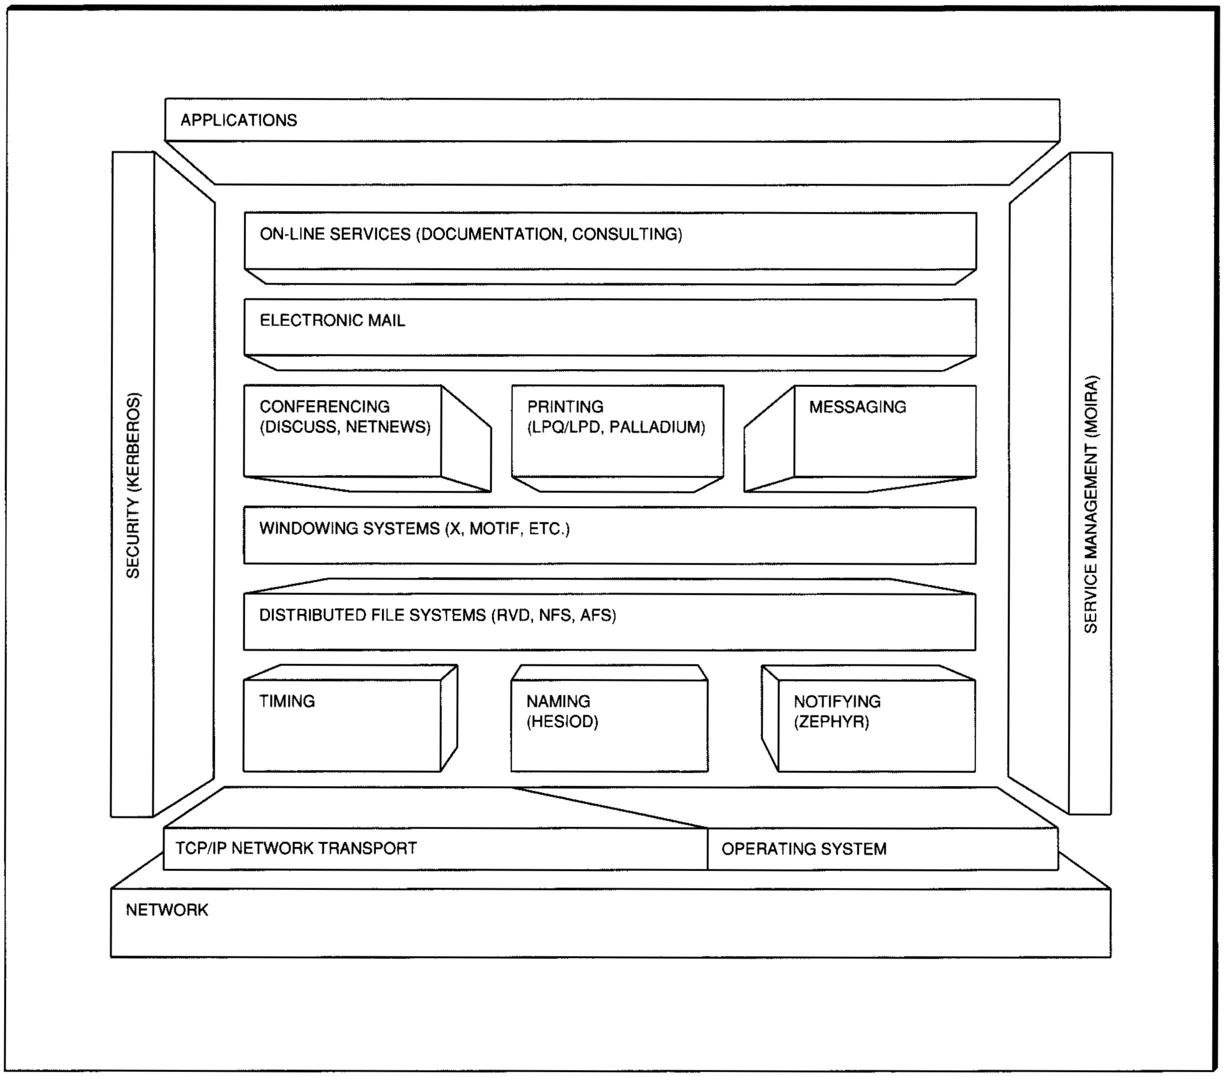


Redrawn with permission of the Massachusetts Institute of Technology

the user, rather than the system, the Kerberos authentication system, based on the trusted third-party model,[2,3] has been designed and implemented at Project Athena. The user's password is encrypted, so that no password will enter the network in the clear. The secure Kerberos server then grants encrypted tickets, valid for a few hours, which identify the user as legitimate to Athena network services such as file servers. All of this is transparent to the user beginning with the logon process.

**Hesiod.** In order to associate the user with his or her own tailored environment and home directory, a mapping function is required that translates logical into physical addresses. Since the user can be at any workstation, only the authenticated userid for that session is of importance. The Hesiod[4] server updates its information from the Moira database (described later), and performs the required mapping when the user signs on. As a consequence, the user is always assured of a completely personalized workstation environment, with home directory and other relevant files attached as if they are stored locally.

**Post Office.** In addition to Kerberos and Hesiod, a number of other central services are available to the users via the network. The Post Office[5] servers provide the store-and-forward capability required for mail, since the user no longer has a permanent physical workstation address. At any time during a session, the mail that has been kept stored in the Post Office server may be requested and transferred into the user's home directory.

**Zephyr.** In addition to mail, there is an instantaneous notification service called Zephyr,[6] which allows users or services to send a message to one or more individuals who are logged in at the time. A small Zephyr window, or Zephyrgram, pops up, containing the message. It is also possible to subscribe to "instances" or classes of Zephyr messages, so that a user who has subscribed to all messages on a given subject can receive these automatically. Those not subscribed will not receive the messages.

**On-Line Consulting.** On-Line Consulting[7] (OLC), which is built on top of the Kerberos, Hesiod, Zephyr, and Post Office services, allows the user to send an instant message to a group of consultants, to ask for assistance or information about anything regarding usage. More than 3000 of these questions are answered every month, of which 95 percent are answered in under five minutes. These questions are also logged and edited into the "stock answers" database. If the user is no longer logged in when the answer becomes available, the system reverts to regular electronic mail for delivery. Consultants can also be contacted via telephone or in person, but experience has shown that electronic assistance is far more popular.

**On-Line Teaching Assistant.** A variant of OLC is On-Line Teaching Assistant (OLTA). This service was inaugurated in September 1990, to provide students with direct access to the teaching assistants for the courses in which they are enrolled. It has been an immediate success, with over half the freshmen using it. About ten of the large freshman courses utilized OLTA during the 1991 academic year.

**Moira.** Moira's service management system[8] provides the central management facility for all of Project Athena. Once a distributed environment reaches more than 200 installed workstations, it is essential to provide an economical way to install and update client and server software. Moira makes use of a central relational database, where all system control and configuration information is kept. It contains the authoritative data about files, users, services, hosts, mailing lists, and group lists. Changes in the database are propagated as service-specific configuration files to appropriate servers on a regular basis. Project Athena's model is influencing the direction and standards for the industry.

**Dialup.** A number of workstations are configured as Athena time sharing systems. Any terminal emulation program, such as PC communication or network remote software for logging in, may be used to access this service. Analog connections are established by dialing into the modem pool of the MIT digital telephone switch. These are then routed through the telephone system to an X.25 PAD (Packet Assembly/Disassembly) and onto one of MITnet's Ethernets. Any available dialup server is assigned, unless a specific type is requested. Authentication is provided at the host workstation level. Only those applications and services that do not require X can be accessed. Dialup is popular for remote access to mail and other common services.

**Other common services.** A discussion of other common services follows.

*Discuss.* Electronic conferencing, or Discuss**,[9] allows for both public and private meetings. It provides for authentication, moderators, mail archiving, and transaction flags. There is no equivalent supported product available today on heterogeneous hardware platforms.

*Printing.* Distributed and authenticated print spooling, and location-independent printer configurations are provided. In addition, user quotas for printing have been instituted. The current quota is 1200 free pages per year. Students must pay for any excess printing through the regular student bill. This prevents the use of printers as free copying machines. The Palladium print system, jointly developed with IBM and DEC, was accepted by OSF for DME, and may eventually replace Athena's current print system.

*On-line documentation.* Information about all aspects of the systems is carefully maintained and available on line. Searches can be done by either menus or keywords. The data are stored in a hierarchy of small, logical modules. There is a single documentation source pool for both on-line and hardcopy viewing.

*Map-based applications.* Campus resources, such as available workstations in public clusters, or offices, buildings, and libraries, can be located on electronic maps. The application Xcluster displays a map of the campus and the location of each public cluster. When a particular location is clicked on with the mouse, the number and type of workstations currently free, as well as the number in use, is displayed. The status of each printer in that cluster is also shown. This is very useful because as students move from class to class, they can plan where they could next use a free workstation (of a specific hardware type if desired) to continue their work.

*Global messaging.* This provides a mechanism for disseminating messages to various classes of users. For instance, a message could go to all users in the freshman class to alert them to new course offerings.

## Multimedia projects

The Visual Computing Group (VCG) at Project Athena occupied a very special place within the project for the last half of its existence. Members of the group were looking ahead to the time when capabilities of the hardware and advances in telecommunications would allow distribution to the individual workstation of multimedia objects and applications. They envisioned the need for a multimedia authoring environment that would be easy to use and flexible enough to accommodate a wide variety of applications and authoring styles. The prototype language or (maybe more accurately) environment they developed is called Athena MUSE**.[10] This prototype was tested against a great variety of real applications. It encompassed education and other applications. As a result, Athena MUSE has been enriched to the point where it is necessary to re-architect and add functionality, particularly in the area of multimedia editors. There has been worldwide interest in this work, particularly because it has been demonstrated to thousands of visitors, and several courses at MIT are now using applications developed with Athena MUSE.[11-13]

The strength of Athena MUSE is that the language is extensible, i.e., the higher levels are written in terms of the lower levels of language. Also, it is possible, depending on the author's interests and level of expertise, to enter Athena MUSE at the lower levels. To the greatest extent possible, Athena MUSE has been made independent of the multimedia platforms on which it resides, making it a relatively simple matter to port it to new UNIX operating systems and hardware.

Athena MUSE is not dependent on the Athena distributed environment, although it has been implemented to support Athena at MIT. The hardware included IBM RTs and DEC workstations, with the Parallax** video digitization board, and IBM Megapel displays. Video is supplied by the campus video cable or video disks. In 1991, work was completed by Bob McKie[13] at the Cambridge Scientific Center to port the Athena MUSE capability to the IBM PS/2 with AIX and M-Motion.

The work done by the VCG is the basis for the Athena MUSE Consortium established at MIT as part of the new Center for Educational Computing Initiatives, under the direction of Professor Steve Lerman. The MIT members of the VCG form the core of the Athena MUSE Consortium staff.

## Usage

As of June 1991, there were about 11 000 users registered, with over 3500 individuals accessing the system at least once a day. Overall usage has risen dramatically over the last two years, with each semester's statistics surpassing those of the previous one. An average of 100 regularly scheduled courses each semester now require the use of Project Athena. About two-thirds of the 1300 or so workstations are available in public and departmental clusters, the rest are deployed as private workstations. As an early experiment, some of them were placed in student living quarters; others are used in electronic lecture halls and classrooms.

While studies have been made of student usage patterns, there have been none in the past two years, during which the utilization time has increased rapidly. Past questionnaires have shown that word processing and electronic mail tended to predominate initially, accompanied by simulation and games. The actual use of course materials was not as heavy as had been hoped. This can be accounted for partly by realizing the magnitude and difficulty of writing and maintaining courseware for a noncommercial operating system, since third-party tools are not available. Second, more funding for courseware development was available during the early years of Project Athena, when the operating system was still under development. Much of the courseware was written by students, and therefore the original developers were often no longer available at MIT to make the necessary changes when a new release of the system software required them. Third, importation and porting of suitable software was not always possible or affordable.

There has been recent, qualitative evidence that the usage is starting to shift somewhat from writing papers, sending mail, and playing games toward scientific modeling, visualization, and courseware use by students, indicating a higher level of maturity in the integration of Project Athena into the undergraduate curriculum. Another observation is that many students like working in the workstation clusters, where they can ask questions, work together, and learn from each other.

Athena resources are made available to the faculty to help keep courseware current. The migration to vendor-supplied operating systems, the expanding suite of third-party software, and the stability of the authoring environment, will additionally have a positive influence on courseware import, creation, and usage.

## The future

MIT made the crucial decision in the fall of 1990 to make the Project Athena infrastructure part of MIT's basic service delivery organization, thereby insuring that the legacy of Project Athena lives on at MIT. This decision was made as a result of a very comprehensive and detailed year-long study by the Committee on Academic Computing in the '90s, composed of a broad selection of MIT faculty, which assessed the Institute's academic computing needs for the next decade. In addition to educational computing, access to libraries and administrative services is now provided. Under Steve Lerman, there is a focus for new academic computing initiatives at MIT, which includes industrial partners in many of its projects.

Much of the distributed computing technology developed at Project Athena has been adopted or submitted to OSF and will become available in vendor products over the next few years. Universities and commercial customers are beginning to look for this technology to solve their growing needs for system management and interoperability in a heterogeneous distributed computing environment.

## Support requirements for DCE

The objective of this section is to provide DCE planners with realistic support estimates, based on the customer's own general requirements and environment.

**Organization.** Using the current Project Athena organization as a model, reasonable assumptions can be made concerning which parts would be necessary in order to support the basic installation. These assumptions were arrived at with the assistance of the MIT staff. Project Athena went through extensive self-examination in order to determine what parts were to be merged with its Information Systems organization after the formal end of Project Athena in June of 1991. Four functional areas are required to support the infrastructure: requirements and initiatives, user services, systems and operations, and administration. The titles of the functional areas are generic, and do not necessarily represent actual departments.

*Requirements and initiatives.* This area consists of faculty or management liaison, third-party software and application support, authoring tools, and productivity tools.

- Faculty or management liaison—Users and their management are provided with information via newsletters, meetings, etcetera. The people in this group are the catalysts for change—whether acting as ombudsmen for user representatives, or presenting new ideas and facilities to the users. They teach and assist the faculty, and give advance warning of significant changes and improvements. When necessary, software development support and re-

sources are provided to ensure that user applications are portable after system software upgrades or changes.
- Third-party software and application support—Commonly used applications are selected, tested, and ported if necessary. These applications may need to be replaced or upgraded from time to time. This group makes sure that the right applications are available to the users, using known user requirements as criteria for selection. Examples are CAD/CAM applications, technical training packages, foreign language instruction, GIS (Geographical Information System) software, or database management systems.
- Authoring tools—In most environments, users want to or need to write their own software. It is the responsibility of the requirements and initiatives group to support these efforts by selecting, testing, and installing authoring tools that will be easy to use and that will result in quality software packages.
- Productivity tools—These are a special set of applications, whose goal is to increase the user's personal and professional productivity. Examples are found in office systems, word processing, and spreadsheet applications.

*User services.* This area consists of user accounts, consulting, training, and documentation.

- User accounts—The management of the user accounts may differ depending on the customer's environment and policy. For instance, in the university, about one quarter of the users leave and about the same number enter during the academic year. Adding a significant number of users can be done quite mechanically. For instance, MIT manages this by extracting information from tapes provided by the registrar and the payroll office. All the user has to do is to access the system by entering an ID, establish a password, and wait about a day for the network resources and configuration to be set up. Deleting users turns out to be much more of a challenge, since many may come back as graduate students. In any case, private files would be destroyed prematurely if users were removed as mechanically as they can be added.

  In a commercial environment the user group may be more stable than at universities, and adds, deletes, or changes usually do not happen at a special time of year or all at once. On the other hand, control over the accounts may be much more important and require more timely action because of security considerations. It may well be that the resources required in both environments turn out to be about the same when calculated proportionally.

- Consulting—The degree to which consulting is available to the users depends on management philosophy. In principle, all the tools are available, from personal contact, to assistance by telephone, to on-line consulting and on-line help. The latter does not require a person to be at the other end, and is generally a menu-driven search program for information and answers to the most frequently asked questions. The human consultants are trained in a wide variety of skills, with some necessary specialization. They can range from full-time paid professionals to part-time volunteers. These may also be distributed directly into the software support organizations of vendor suppliers.

- Training—User training can be implemented in many different ways. In universities, it is most efficient to schedule the initial training when students arrive, and it pays to have a live instructor give out information to large groups. For faculty, training is often done on an individual basis. While new users in the commercial environment need training on demand, a case could still be made for periodic group instruction, as well as self-study or on-line education.

- Documentation—Documentation, both on-line and hardcopy, is one of the top user requirements. This function requires a staff of people who understand the technical environment as well as the user requirements for lucid, well-organized documentation. Although much of it can be made available on line, some basic manuals have to be available in hardcopy, to initially train the user, or to provide information in case of system malfunction.

*Systems and operations.* This area consists of customer service, system support, release engineering, and software engineering.

- Customer service—A staff has to be available to plan, install, and repair the systems. They are not only reactive, but proactive, and constantly monitor the overall and individual status of the installed equipment. This is done in several ways. The staff is capable of monitoring from a central point. In addition, "cluster patrols" visit the workstations about once a day, to solve obvious problems where possible, and report those that require further attention. They are the interface to vendor service organizations, if repairs or replacements cannot be done in-house.

- System support—The system support staff takes care of systems administration and ensures availability and performance of central network servers. File server capacity, systems and user file backup are monitored.
- Release engineering—This function is required even if only one operating system is installed. The Athena services have to be reintegrated for each new release or version of the operating system. When distributing software updates to a large group of users, the potential for disaster is always present. Not just one user, but all users will be affected by any error or code defect introduced by a new software release. Software updates must be done in a controlled and tested way. When more than one operating system is involved, release engineering becomes one of the most essential systems functions. The staff does the systems integration and testing, and ensures that everything remains interoperable.
- Software engineering—New versions and releases of software are a fact of life. As more and more of the operating systems are provided by vendors, the task of integrating the Athena software and building new releases is a constant one. To some extent, each Athena-like installation will have to face this, since local modifications will exist, and a wide variety of workstations and associated operating systems may be installed. In addition, it is important that new releases of the Athena suite of software be incorporated as they become available, in order to take advantage of fixes and functional enhancements.

*Administration.* Other critical but often overlooked aspects of DCE support requirements are the areas of budgets, inventories, and license and contract negotiations, as well as other administrative issues. Management and personnel for this functional area have to be planned for, and resources allocated. The responsibilities include finance, property management, contracts, software licensing, and administrative support.

**Assumptions.** Certain assumptions had to be made in order to make the rules we developed generally usable.

Telecommunications and network support has not been included. There are many university and commercial customers who already have a telecommunications network and associated support structure in place. It is therefore assumed that the network infrastructure lies outside the scope of these estimates.

The proposed DCE installation is on one campus. The assumption is that all workstations are within a reasonably small geographical area and are all served from the same high-speed network. For instance, it should be possible to walk or bicycle to a workstation location within a few minutes, not hours. An example of this is a single university campus, or a group of commercial buildings which are within the same city and which have the same access to the company's high-speed network.

There is no application development. Commercially available applications satisfy the need of the customer. Of course this is not a realistic assumption for most customers, but it would be impossible to estimate the resources to support every customer's application requirements.

The space for workstations and servers has been allocated and all facilities arrangements are in place. It may be difficult to find and build suitable space for the DCE clients and servers. The estimates in this paper do not take into account the very considerable initial work required to find and fund the space required, both for the workstations and servers, and the people who are going to support them. The assumption here is that, once having made the decision to install the DCE, these resources are the same whether or not the Project Athena model is chosen. This assumption may have to be carefully examined by the customer for accuracy.

**How to estimate support requirements.** At first, it might seem simplistic to design a "formula" to estimate support requirements. In talking with the experts on the subject, managers, and personnel at Project Athena, it became evident that although everyone had an intuitive understanding about the human resources required, it was difficult to translate this into real numbers when a new installation was discussed. Resources are defined in terms of full-time employees. The following variables are relevant.

*Workstations (clients only).* The number of workstations affects the resources required for tracking and inventory, hardware maintenance, installation/re-booting/updating, and system-wide services.

Resource estimate: W/500, where W is the number of workstation clients.

*Users.* The number of users affects accounts administration, user training, "how to use" consulting, documentation, and server configuration and sizing.

Resource estimate: U/1000, where U is the number of users.

*Clusters.* A cluster is a physically co-located workgroup, with systems in open areas and/or private spaces. Clusters are characterized by distance between clusters, size (i.e., number of clients per cluster), and printer and system server requirements. These factors affect resources required for planning, deployment, paper and other consumables, space management and security, and server administration.

Resource estimate: C/15, where C is the number of clusters.

*Supported applications.* An application is defined as a licensed product or its equivalent if it was developed in-house. Examples are: word processing, a CAD application, or a courseware package. The number of applications provided and supported centrally, exclusive of development, affects the resources required for compilation (if porting is required), building the software distribution, integration and testing, installation, updates, software maintenance, application-specific user support, and higher-level documentation.

Resource estimate: A/50, where A is the number of supported applications.

*Licenses required.* A license is defined as the right to use the application software for potentially multiple users on multiple platforms. In all cases, at least one contract has to be negotiated, renewed, etc. The number of licenses may reach several hundred. The number of contracts to be negotiated, and contract and license renewals can be significant.

Resource estimate: L/25, where L is the number of licenses.

*Vendor operating systems and applications (VOSAP).* VOSAP means the vendor operating system supplied for a specific platform. New versions and releases are a fact of life. One school of thought is that there is a nonlinear relationship between the number of operating systems and the resource required to support them, with a "wall" at about four or five. The activities associated with VOSAP support are release engineering, client configuration, testing, documentation, and defect tracking and reporting.

Resource estimate: V, where V is the number of distinct vendor operating systems and their platforms.

*Combined requirements.* When combined, the total human resources (HR) required are expressed in Equation 1:

$$HR = W/500 + U/1000 + C/15 + A/50 + L/25 + V \qquad (1)$$

One additional piece of information is needed in order to use this formula for estimating resource requirements. There are five sets of human resource skills required to support the total environment. These are roughly related to each term in the formula, with A and L combined. The minimum value of HR is therefore five. Depending on the environment and how resources are counted, customers are reporting anywhere from 10 to 25 workstations supported per support resource in a non-Athena installation. This brings us to the conclusion that a Project Athena-like environment only starts to pay off in terms of resources required to support it when the number of client workstations is somewhere around 150.

An important factor to keep in mind is that these estimates are based on support requirements for a mature installation. It is quite possible, when starting a new installation, to get by with far fewer support people. In this case the favorable factors are that all hardware was provided by a single

**Table 1 Test of the support formula parameters, showing actual and calculated resources**

| Parameter | Divisor | MIT | NCSU | UCSC | UMD | ISU | UMICH |
|---|---|---|---|---|---|---|---|
| | | Profiles of Six Athena Installations | | | | | |
| Workstations | 500 | 1300 | 200 | 20 | 350 | 265 | 850 |
| Users | 1000 | 11000 | 3050 | 300 | 2000 | 1550 | 8750 |
| Clusters | 15 | 40 | 2 | 1 | 22 | 26 | 12 |
| Applications | 50 | 200 | 100 | 75 | 200 | 200 | 200 |
| Licenses | 25 | 40 | 4 | 20 | 40 | 25 | 20 |
| VOSAP | 1 | 4 | 1 | 2 | 5 | 1 | 4 |
| **Support Resource Requirements** | | | | | | | |
| Actual | | 25 | 3 | 3 | 15 | 10 | 23 |
| Calculated | | 26 | 7 | 5 | 15 | 10 | 20 |

vendor, they started with all new equipment (i.e., no maintenance), they installed one current version operating system, they had low user-support expectations initially, and they got their initial installation support, planning, and training from the vendor. This is not a viable situation in the long run.

**Verification of concept.** We had the opportunity to test our assumptions and modify our earlier estimates after the First International Athena Technical Conference in April 1991. We obtained profiles of five Athena-like installations at universities: North Carolina State University (NCSU), University of Maryland (UMD), Iowa State University (ISU), and the University of California at Santa Cruz (UCSC). In addition, we obtained data from the University of Michigan (UMICH) in an informal discussion at the Engineering Workstation Conference held in Boston in August 1991. Using these profiles, and making reasonable assumptions in those cases where information was incomplete, we constructed Table 1. Note that the formula works well for the larger installations with many clusters. North Carolina State University has been able to manage with a very small staff up to now because they have a single vendor installation, very few clusters, and new equipment. As they diversify equipment, distribute it to more locations, and expand their user services, more staff will likely be needed. Similarly, the University of California, Santa Cruz, installation is in the initial stages and presents an anomaly.

After spending some time working on this concept and discussing it with a number of managers and staff at Project Athena, we feel quite comfortable with this approach. Until more data are available from other installations that use the Project Athena software, and until more installations have matured, this is our best estimate.

An interesting extension of this work would be to see if one can correctly predict the server configurations required to support these environments. Similar variables will be required, augmented by others that are related to the type of specialized central servers, such as Kerberos, Zephyr, Moira, etc., as well as customer requirements for user file space servers, backup requirements, and others.

## Conclusion

The Project Athena partnership between MIT, IBM, and DEC has proved to be very fruitful. It has given rise to new standards, such as the X Window System, and continues to make significant contributions to open systems architectures, in particular to OSF's DCE and DME. Until these architectures have been implemented by the workstation manufacturers, many universities, companies, research laboratories, and government agencies are looking to elements of the Athena suite of network services to manage their distributed systems installations.

Our participation in Project Athena has demonstrated the value of joint projects and partnerships with universities in an open systems environment. Workstation installations from a single vendor are the exception in the real world. It is therefore exceedingly important that we take advantage of the opportunities we have in the establishment of open standards. At the same time, it allows us to incorporate strategic architectures into our products at an early stage.


## Acknowledgments

We want to thank the management and staff of Project Athena for their thoughtful comments. We particularly want to express our gratitude to Earll Murman, Ron Orcutt, Catherine Avril, Naomi Schmidt, and Don Davis. Thanks also to George Hebert of IBM for his feedback.


*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of the Massachusetts Institute of Technology, UNIX Systems Laboratories, Inc., the Regents of the University of California at Berkeley, Digital Equipment Corp., Apple Computer, Inc., the Open Software Foundation, Cray Research, Inc., Sun Microsystems, Inc., Transarc Corp., or Parallax Graphics, Inc.

## Cited references


1. G. W. Treese, "Berkeley UNIX on 1000 Workstations: Athena Changes to 4.3 BSD," *USENIX Association* (February 1988).
2. R. M. Needham and M. D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers," *Communications of the ACM* **21**, No. 12, 993–999 (December 1978).
3. J. G. Steiner, B. C. Neuman, and J. I. Schiller, "Kerberos, an Authentication Service for Open Network Systems," *USENIX Association* (February 1988).
4. S. P. Dyer, "The Hesiod Name Server," *USENIX Association* (February 1988).
5. J. G. Steiner and D. E. Geer, Jr., "Network Services in the Athena Environment," *USENIX Association* (July 1988).
6. C. A. DellaFera, M. W. Eichin, R. S. French, D. C. Jedlinsky, J. T. Kohl, and W. E. Sommerfeld, "The Zephyr Notification Service," *USENIX Association* (Winter 1988).
7. B. Anderson, T. Coppetto, D. E. Geer, Jr., and G. W. Treese, "OLC: An On-Line Consulting System for UNIX," *USENIX Association* (February 1989).
8. M. A. Rosenstein, D. E. Geer, Jr., and P. J. Levine, "The Athena Service Management Systems (Moira)," *USENIX Association* (February 1988).
9. K. Raeburn, J. Rochlis, W. Sommerfeld, and S. Zanarotti, "DISCUSS: An Electronic Conferencing System for a Distributed Computing Environment," *USENIX Association* (February 1989).
10. M. E. Hodges, R. M. Sasnett, and M. S. Ackerman, "A Construction Set for Multimedia Applications," *IEEE*, 37–43 (January 1989).
11. "Report Card on Technology: MIT Computer Network Offers Look at Future Education," *IBM Multimedia Solutions* **4**, No. 10, L. J. Baratto, Editor (October 1990).
12. "MIT Breaks New Ground with Interactive Multimedia," *IBM Multimedia Solutions* **4**, No. 11, L. J. Baratto, Editor (November 1990).
13. R. A. McKie, "Converting from RT to PS/2: The Fun Is Just Beginning," *IBM Multimedia Solutions* **4**, No. 12 (December 1990).


## General references


C. R. Avril and R. L. Orcutt, "Athena: MIT's Once and Future Distributed Computing Project," *Information Technology Quarterly* (Fall 1990).

G. A. Champine, *MIT Project Athena, a Model for Distributed Campus Computing*, Digital Press (1991).



*Accepted for publication January 20, 1992.*



**Josina M. Arfman** *IBM Corporation, Route 100, P.O. Box 100, Somers, New York 10589.* Dr. Arfman was IBM program manager and MIT Associate Director of Project Athena at the Massachusetts Institute of Technology from November 1988 until June 1991. Her current IBM position is program manager in the Advanced Workstation Division. She joined the Research Division in 1976, and has had a number of management positions, including those in Corporate Ph.D. Recruiting and Worldwide Internal Telecommunications Strategy. She has a Ph.D. in X-ray crystallography.

**Peter Roden** *Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139.* Mr. Roden is the Manager of Distributed Systems in Information Systems at MIT. Before the merger with Information Systems, he was the Manager of Systems and Operations at Project Athena. Prior to his coming to MIT, he was the Manager of VMS Systems for Harvard University's Faculty of Arts and Sciences. He has a degree in computer science and is a member of ACM and IEEE.



Reprint Order No. G321-5486.