

MIT Project Athena

A MODEL FOR DISTRIBUTED CAMPUS COMPUTING

江苏工业学院图书馆

藏书章

GEORGE A. CHAMPINE

Copyright © 1991 by Digital Equipment Corporation.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the publisher.

Printed in the United States of America.

9 8 7 6 5 4 3 2

Order Number EY-H875E-DP

The views expressed in this book are those of the author, not of the publisher. Digital Equipment Corporation is not responsible for any errors that may appear in this book.

Trademarks and sources are listed on p. 282.

Production: Quadrata, Inc.

Editing: Jerrold Moore

Text design: Melinda Grosser

Index: Marilyn Roland

Jacket design: Dick Hannus

Composition: Compset, Inc.

Library of Congress Cataloging-in-Publication Data

Champine, George A.

MIT project Athena : a model for distributed campus computing /

George A. Champine.

p. cm.

Includes bibliographical references and index.

ISBN 1-55558-072-6

1. Massachusetts Institute of Technology—Data processing.

2. Electronic data processing—Distributed processing. I. Title.

T171.M44 1990

004'.36—dc20

91-10948

CIP

MIT Project Athena



A MODEL FOR DISTRIBUTED CAMPUS COMPUTING



FOREWORDS

I arrived at MIT in the fall of 1990, midway through the final year of Project Athena. With vision and determination, Professor Gerald Wilson, then Dean of Engineering, together with President Paul Gray and Provost John Deutch, had led this remarkable effort to create a new educational computing environment at MIT. The project challenged the entire MIT community, especially the faculty, students, and Project Athena staff. It succeeded in creating a system that has begun to transform the ways in which our students and faculty approach computing, learning, and information sharing.

I believe that the development of information technology will revolutionize the ways in which many of us work and interact. It is important that MIT prepare its graduates for such an environment and that we explore and implement ways to use information technology to enhance the educational process itself. With over 1000 networked workstations and servers located throughout the campus, Athena seems ubiquitous to this year's first-year students. By the start of classes in the fall, a typical freshman had likely become an Athena user, beginning with electronic mail and messaging and basic text editing. He or she might well have been enrolled in 3.091, Introduction to Solid State Chemistry, which uses the new On-line Teaching Assistant (OLTA) capability. In later years, a typical member of the class of 1994 will probably have created new software on Athena and will likely have taken several of the 200 subjects now using courseware on Athena. Some of the courseware developed

in Project Athena is already in use on other campuses around the country.

Project Athena was made possible by the partnership with the Digital Equipment Corporation and the IBM Corporation, and is a model of university–industry collaboration. Working side by side, people from the three organizations created new client–server software technology, including X Window, Kerberos, and many other system and application level programs, all of which were made available to the outside world.

The author of this book, Dr. George Champine, was the DEC Associate Director of Project Athena. His perspective provides valuable insight and a lasting account of the creation of the Athena Computing Environment.

Charles M. Vest

President

The Massachusetts Institute of Technology

From time to time, windows of opportunity open to make major advances in computing technology. One such window was the occasion to participate in Project Athena in 1983. Three unique factors came together at that time to make this an unparalleled opportunity. These factors were

- a major change in computing from centralized time-sharing systems to a distributed network of workstations;
- the desire by MIT to develop a campus-wide world-class computation system for education as well as research;
- the desire by Digital to participate in a very large scale implementation of networked workstations to test its strategy and products.

Nevertheless, the project was clearly very risky because of the many unknowns.

Now, eight years later, Athena has utilized this opportunity to make very significant contributions to distributed computing and has been successful beyond any of our expectations. Through the X Window System, Kerberos Authentication Service, and other similar developments, Athena has changed the world of computing. Athena is one of the first and one of the most successful implementations of

the client/server model of distributed computing. This model has now been endorsed by the Open Software Foundation for their Distributed Computing Environment. Athena has emerged as the model of choice for distributed campus computing, and I expect that it will serve as the preferred model of distributed computing for the 1990s.

Project Athena also stands as an excellent example of what can be accomplished when the right kind of partnership is struck between industrial corporations and universities. What Athena accomplished could not have been done in the 1980s by MIT, Digital, or IBM alone. Working together, in the open environment of the university campus, we created the pioneering implementation of the vision of a coherent, powerful, distributed computing system.

On behalf of Digital Equipment Corporation, I would like to thank our two partners in Project Athena, MIT and IBM, for the chance to join with them in undertaking a project too large for any one of us to support alone, but whose benefits far outweigh the costs.

We congratulate the entire Athena team, knowing that the system they designed and built so well will serve as a platform for the major computing advances that will be accomplished in the 1990s.

Samuel H. Fuller

Vice President, Research
Digital Equipment Corporation



Creativity and innovation are the fuels that drive modern high technology companies. Just as the New England textile mills of the last century needed a steady flow of water to power their turbines, today's information-based industries require a steady flow of new ideas to compete in global markets.

Most people agree that innovation cannot be forced. Many believe it cannot be planned for or managed to a fixed schedule. In fact, there is evidence that most organizations are doing very well if they are able to encourage innovation rather than stifle it, and to recognize and reward creativity when it blossoms.

Developing collaborative research projects with universities is one successful strategy to encourage innovation that has been devel-

oped and used successfully for years by Digital Equipment Corporation. Project Athena at MIT was Digital's largest and longest running university project. During its eight-year history, Digital engineers spent more than 40 person-years working on site with MIT faculty, students and staff, and researchers from other companies. The trail of innovation leading from the people working together on this project is proof of the basic tenet of those of us who sponsor strong research connections with universities: When industry and academia pool their research talent, both sides prosper.

Project Athena was a success in several dimensions. Thousands of young engineers and scientists learned more about their core curriculum because of innovative learning software developed by faculty and staff. Perhaps even more important, they learned how to use networks of powerful workstations for collaborations in many endeavors. New technology that emerged from the project was turned into commercial products by Digital and other companies. Examples include the X Window System, Kerberos, and Hesiod. In addition, knowledge about how to design, install, manage, and modify an extremely complex network of distributed systems is influencing dozens of other organizations, both commercial and academic, that are on the forefront of computing.

Many of the lessons learned from Project Athena have to do with how bright people from different organizations, with different goals, and with very different cultures can work effectively together despite these differences. George Champine worked for nearly five years at bridging these cultures and contributing in many ways to the success of the project. This book is a great contribution to everyone interested in understanding the vision and the accomplishments of one of the most important university-industrial cooperative projects of the past decade. Equally important are his insights about how this remarkable partnership developed and functioned on a day-to-day basis.

Like most effective technical relationships, Project Athena was based on mutual respect, ongoing communication, a vigorous quest for excellence, and the challenge of growth and discovery. We hope that it can serve as a model for many other successful university-industry cooperative efforts.

John W. McCredie

Director, External Research Program
Digital Equipment Corporation

PREFACE

Tell me and I forget
Show me and I remember
Involve me and I understand

Ancient Chinese proverb



Project Athena at MIT has emerged as one of the most important models of next generation distributed computing in the academic environment. MIT pioneered this new systems approach to computation based on the client-server model to support a network of workstations. This new model is replacing time-sharing (which MIT also pioneered) as the preferred model of computing at the Institute. Athena is unique in that it is one of the first and one of the largest integrated implementations of this new model, perhaps paralleled only by the Andrew project at Carnegie-Mellon University. Athena's uniqueness has led to widespread interest in its design, implementation, and results. This book was written to supply information in a relatively concise form to meet this interest. For those who need more detail, a complete set (2800 pages) of Athena reports through September 1988 was published under the title "Project Athena: The First Five Years" and is available from Digital. This book is abstracted in large measure from those reports but with updated information where available and supplemental information to fill out the picture. This book also provides references to the more detailed papers describing Athena in the Bibliography.

The following characteristics, taken together, make Project Athena unique among the various higher education computing systems:

- magnitude, the largest educational project ever undertaken by MIT with a cost of about \$100 million;



Figure 1 Project Athena workstation cluster.

- workstation-based almost exclusively;
- campuswide, with one integrated system that serves the entire campus;
- coherent, in that all applications run on all workstations, providing the same user interface independent of hardware architecture and easy sharing of programs and data; and
- centrally managed and supported.

In the late 1970s the information processing needs of higher education were rapidly moving beyond the capabilities of conventional pedagogical techniques. These techniques—pencil, paper, books, blackboards, and lectures—have all been used for hundreds of years. At the same time, major computer manufacturers were developing high-performance graphics workstations and networks that they expected students could afford within a few years. The MIT faculty believed that a network of workstations with good graphics and communications could overcome the well-known obstacles to learn-

ing by providing capabilities for

- easier visualization of abstract concepts;
- powerful simulation tools for classroom demonstration and homework problems;
- productivity aids to allow the student to accomplish more in less time and with better quality; and
- improved education in the area of design, by solving realistic problems using sophisticated computer-aided techniques.

Figure 1 shows part of Athena's largest workstation cluster.

Ultimately, the system was intended to help students obtain new conceptual and intuitive understanding of disciplines. A one-time investment in training and development would become applicable to all uses of the system.

In May 1983, MIT announced the establishment of a five-year program to explore new, innovative uses of computing in the MIT curriculum. This program was Project Athena.

Athena was undertaken to create a computing environment that would make a significant and lasting improvement in the overall quality of education at MIT. Athena's mandate was to explore diverse uses of computing in support of education and to build the base of knowledge needed for a longer term strategic decision about how computers fit into the MIT curriculum.

Athena has a major role in the educational mission of MIT, with the objective of contributing to the quality of the educational process. This role includes instruction, research, and activities outside the laboratory and classroom.

The primary technical objective of Project Athena, as stated early in 1983, was: "By 1988, create a new educational computing environment at MIT built around high performance graphics workstations, high speed networking, and servers of various types" [Lerman 87]. This objective has remained constant over the life of the project. Additional technical objectives were to

- support a heterogeneous hardware configuration;
- provide a user interface independent of hardware;
- provide a software development environment independent of hardware;
- maximize the exportability and importability of software; and
- accomplish these objectives at an ongoing operations cost, including hardware, not to exceed 10 percent of tuition.

These objectives have largely been met.

The support of heterogeneous hardware is necessary because Athena's two major sponsors, Digital Equipment Corporation and IBM, have products with substantially different architectures. The attainment of the cost objective was aided to a considerable extent by having a single user interface and single software development environment independent of hardware. The computational environment includes hardware, data communications, physical facilities, systems software, documentation, user support, applications software, and information resources, such as the libraries and outside databases.

Project Athena has now been in existence for seven years. Today, it is a production system, having fulfilled its objective of providing campuswide access to workstation-based computing at MIT.

Athena is a distributed workstation system in use 24 hours per day with 10,500 active accounts and 1300 workstations. The Athena model of computing is that of a unified distributed system where a single log in provides universal access to a variety of authenticated network services. The distributed services supported by Athena convert the time-sharing model of computing supported by Unix into a distributed system operating environment.

Network services supported include authentication, name service, real time notification service, a network mail service, an electronic conferencing service, a network file service, and an on-line consulting service. The system has extensive failsoft capabilities that permit continuous operation despite equipment failures. The system is designed to support 10,000 workstations by minimizing the use of network bandwidth, mass storage, and labor—all scarce resources. The centralized management approach also helps to minimize support and operations cost.

The workstation types presently supported are the Digital VAXstation and DECstation, the NeXT computer, the IBM RT/PC and RS6000, the Apple Macintosh II, and IBM PS/2 personal computers. Communications is provided by a campuswide local area network. The backbone is a 10 million bit/second fiber optic token ring. Attached to the backbone are 23 IP routers, which support 41 Ethernet subnets using TCP/IP.

The system provides a coherent model of computing in which all applications can run on all supported workstations independent of architecture. Because of the strong level of coherence, the human interface to the system is independent of the type of workstations

being used. Differences in equipment are invisible to the user. Only one training course set and one documentation set are needed, no matter what workstation hardware is used. Students need to learn only one system during their stay at MIT, thus improving their productivity. The ultimate objective of coherence is to make the transfer of information limited only by its usefulness, not its availability.

Currently, 16 public clusters of workstations are available for use by students. These public clusters include 10–120 workstations each, and are open 24 hours per day. One of the public clusters is an electronic classroom. In addition to the public clusters, there are 25 department clusters and four projection-equipped facilities. Workstation clusters have been installed in five student residences, with some having workstations in the bedrooms and others having them in common areas. Figure 2 shows the locations of Athena workstation clusters.

At present, 38 Network File System (NFS) file servers, 8 Andrew File System (AFS) file servers, 36 Remote Virtual Disk (RVD) file servers, 135 Postscript printers, three each of name servers and post office servers, and two authentication servers are utilized. In addition to the “generic” Athena monochrome workstations, 15 multimedia workstations support full motion color video and sound. The system has 100 gigabytes of rotating storage in the workstations and an additional 50 gigabytes of rotating storage in network file servers.

The current 10,500 active user accounts (from a total of about 15,000 total accounts) generate about 5000 log ins and about 13,000 mail messages per day. The average student uses the system eight hours per week. In aggregate, the users generated 24,000 questions for the on-line consulting system and printed 5 million pages during 1990. About 96 percent of the undergraduate students and about 70 percent of the graduate students use the system. Athena usage is increasing substantially each semester.

The design of the Athena system was led by Professor Jerome Saltzer of MIT, in association with the Athena staff, which includes among others Ed Balkovich, Tony Della Fera, Steve Dyer, Dan Geer, Jim Gettys, John Kohl, Steve Lerman, Steve Miller, Cliff Neuman, Ken Raeburn, Mark Rosenstein, Jeff Schiller, Bill Sommerfeld, Ralph Swick, and Win Treese.

Initially, about 1500 Institute-owned workstations will be deployed. Later, all students and staff will be allowed to purchase private workstations, and plans call for ultimate system extension to

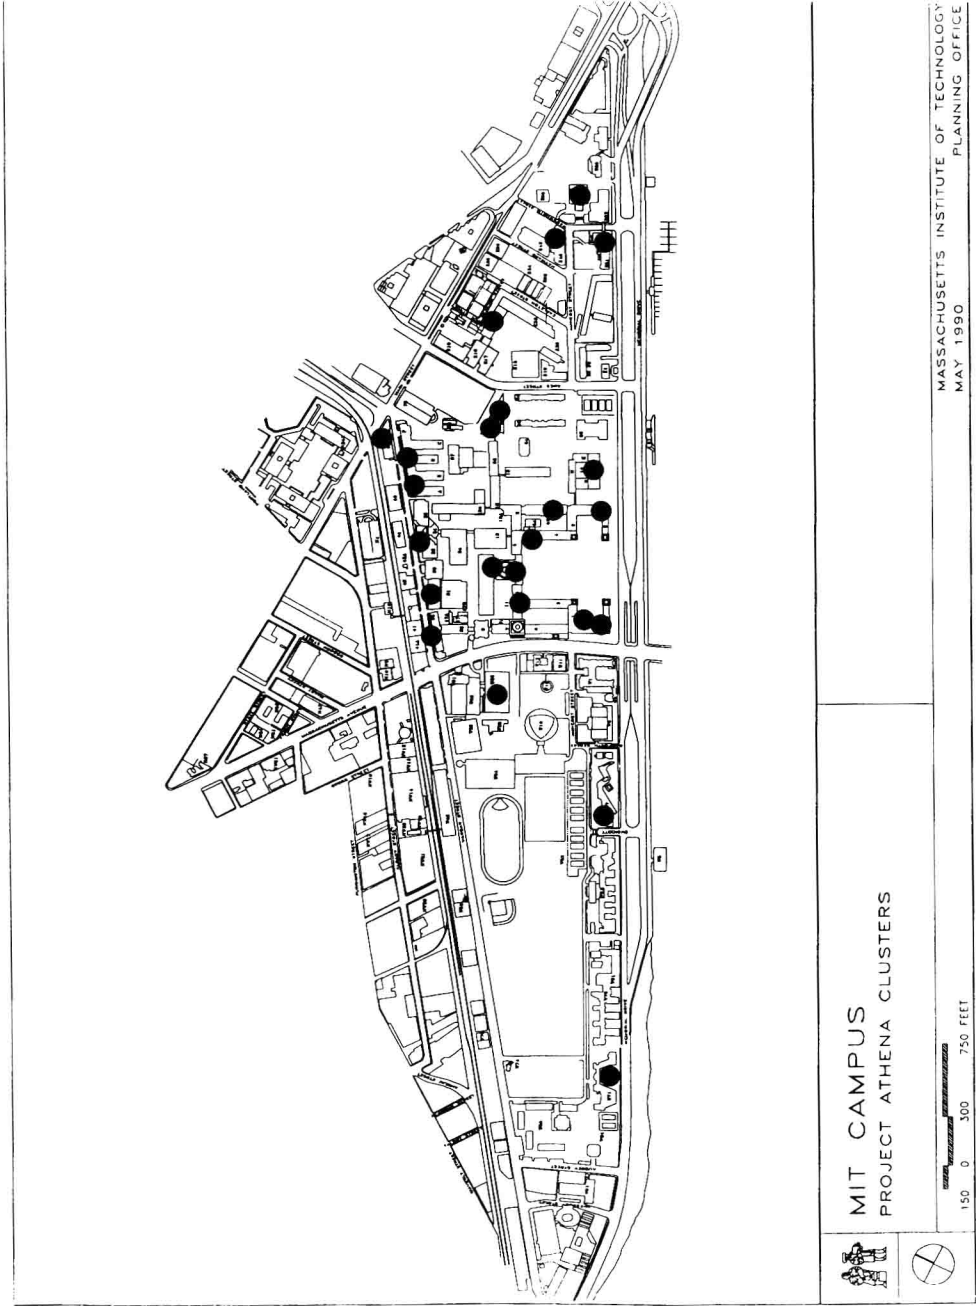


Figure 2 Locations of Athena clusters.

support approximately 10,000 workstations. System software is based on Unix and includes additional MIT-developed software that implements a client-server distributed system.

A significant part of the project is development of instructional software; to date 125 projects have been funded for that purpose. Students use Athena for most courses in some manner. Often Athena is used with personal productivity software, such as word processing, spreadsheets, graphics, and mathematical application packages. In addition, about 200 courses are presently being taught that use MIT-developed instructional software. Some courses require the use of Athena in mid-semester and final exams. Considerable anecdotal evidence shows that Athena has improved the quality of education at MIT, fulfilling its primary objective, even though the instructional software developed to date falls short of what ultimately can be developed.

Another part of the project is development of a multimedia workstation that uses interactive video disk technology to deliver full motion video in an X window simultaneously with conventional computer-generated text and graphics. The multimedia workstation project includes development of an advanced authoring environment, called MUSE, to reduce the cost and skill level needed to develop instructional software.

Athena continued as a free-standing project until June, 1991. At that time it was integrated into the MIT administrative system.

In many ways, Athena mirrors larger campus issues. Some of the more fundamental issues are the proper role of computers in higher education, the role of student housing as integrated into coursework versus a defense against the demands of coursework, and the inherent conflict between the priority given to teaching and the priority given to research (and tenure). To the extent possible, these larger issues are described in this book, with Athena being only one point in the spectrum of tradeoffs that can be made. All campuses are different, and strategies that work well on one campus could end in disaster at another. No claim is made that the approach selected by MIT is appropriate anywhere else. However, the lessons learned may be helpful to others designing similar systems.

Not all sections of this book are likely to be of equal interest to all readers. To help you find the areas of interest to you, we divided the book into four parts: Development, Pedagogy, Technology, and Administration.

The Development part includes Chapters 1 and 2. Chapter 1 describes the process of creating the vision that became Athena, the

negotiations necessary to achieve some level of consensus about that vision, and the trigger necessary to start implementation. Chapter 2 then describes the project undertaken to implement the vision, starting with the activities to develop the project's infrastructure.

The Pedagogy part includes Chapters 3 and 4. Chapter 3 describes the instructional software developed for Athena, including the process by which it was funded, and some of the more widely used packages. Chapter 4 describes project-faculty relations. An important aspect of Athena's success, these relations have not always been smooth.

The Technical part includes Chapters 5, 6, and 7. Chapter 5 summarizes the technical aspects of the system, and Chapter 6 presents them in much greater detail. Chapter 6 deals with the system at a technical level and assumes a knowledge of computer systems software and Unix. It also describes the human interface work at Athena and the campus communications system. Chapter 7 describes multimedia workstation development and use.

The Administration part includes Chapters 8–11. As part of the Athena experiment, workstations were put in five student housing facilities; Chapter 8 describes these installations and assesses the results. Chapter 9 describes the financial and organizational aspects of the project. Chapter 10 provides an assessment of the project. Chapter 11 looks at possible futures for the project and the direction chosen by MIT.

Three documents of interest to implementors of campus computing systems similar to Athena are included in the Appendixes: Guidelines for Installation (Appendix II) and Principles of Responsible Use of Project Athena and Athena Rules of Use (both in Appendix IV).

Acknowledgments

Project Athena would not have been possible without the vision of many individuals from MIT, Digital, and IBM and their willingness to take risks. At MIT, the primary driving forces behind getting Athena started were Gerald Wilson, Dean of Engineering, Joel Moses, Head of Electrical Engineering and Computer Science, and Michael Dertouzos, Director of the Laboratory for Computer Science. Moses and Dertouzos gave much of the early direction to technical considerations, including the concept and definition of coher-

ence. Steve Lerman, Professor of Civil Engineering, was Project Director for the first five years and provided the managerial direction necessary to start the project and keep it on track. Jerome Saltzer, Professor of Electrical Engineering and Computer Science, provided technical leadership for the first five years; without him the project might have failed. James Bruce, Vice President of Information Systems, provided continuing support of Athena in the implementation of the campus communication system and in significant contributions in serving on the Athena directors committee. Earll Murman, Professor of Aeronautics and Astronautics, accepted the position of Director of Athena for the last three years and managed the difficult transition from development to delivery of highly reliable and quality computing services.

The responsible people at Digital include Kenneth H. Olsen, Founder, President—and graduate of MIT—who was willing to bet on a vision. They also include Winston R. Hindle, Senior Vice President, who provided continuing support for the pedagogical aspects of the project; Samuel H. Fuller, Vice President of Corporate Research, who argued the case within Digital to fund the project; Dieter Huttenberger, the first Director of External Research, John W. McCredie, Director of External Research, who staffed the Digital part of the project and integrated it into the larger Digital External Research Program; Edward Balkovich, the first Digital Associate Director at Athena who contributed substantially and crucially to its management and design; and Maurice Wilkes (inventor of the first programmable electronic digital computer), who was my immediate predecessor. I joined the project in August 1986 as Digital's Associate Director and participated in its management through its mid-life to the end.

Key people at IBM who made significant personal contributions to Athena include Ralph Gomory, Senior Vice President for Science and Technology; Richard Parmelee, first IBM Associate Director at Athena; and Les Comeau, Manager of Technical Computing Projects.

People who contributed significantly with ideas or in the review of this book include Jim Bruce, Ed Balkovich, Jim Gettys, Dan Geer, Anne LaVin, Jack McCredie, Jerry Saltzer, Naomi Schmidt, and Win Treese. Joel Moses and Mike Dertouzos provided material from their personal files concerning creation of the Athena vision and the early days of the project. Earll Murman provided significant help with the chapter on instructional uses of Athena. Ben Davis provided sub-