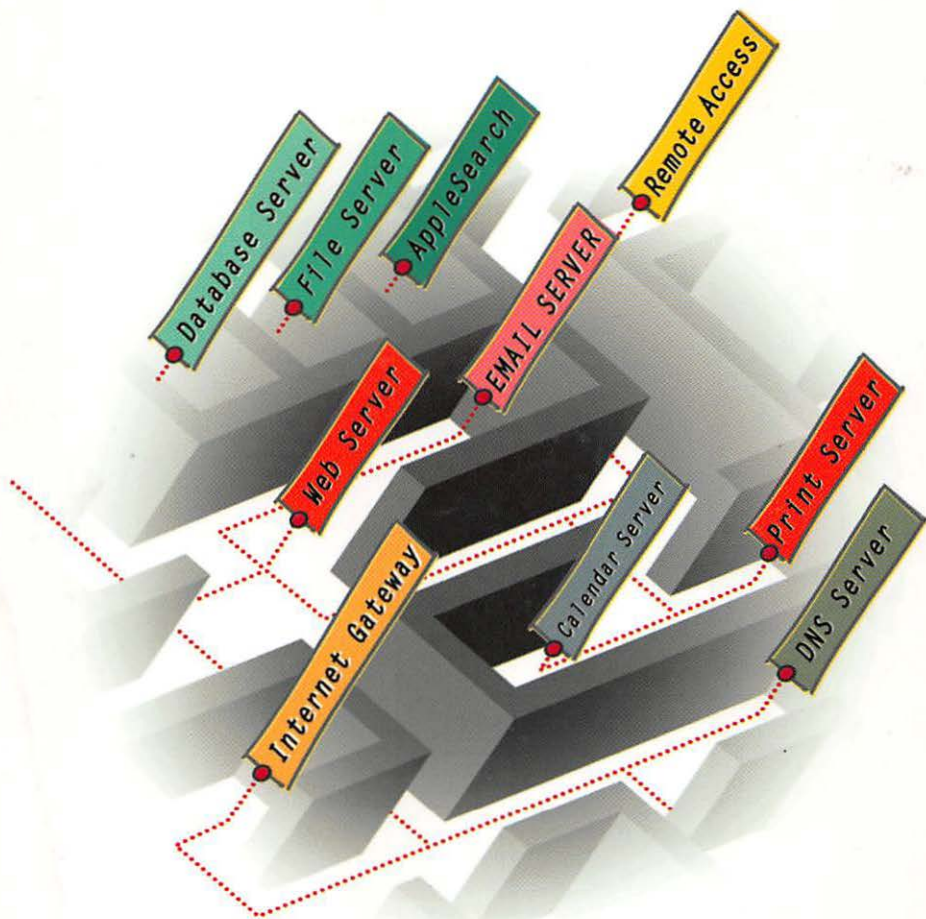




AppleTalk Network Services

by Dorian J. Cougias, Tom Dell, and E. L. Heiberger



- Services you can offer your users over the network — file sharing, printing, backup, e-mail, calendaring, forms, databases, digital workflow, Internal Web pages
- How network utilization is affected by the services you select
- Learn how to maximize network efficiency by distributing your services
- Macintosh-based Internet services: Web servers & CGIs, FTP, e-mail, list servers, DNS, and more



Includes free
network tools

AppleTalk Network Services

Network Frontiers Field Manual Series

LIMITED WARRANTY AND DISCLAIMER OF LIABILITY

ACADEMIC PRESS, INC. ("AP") AND ANYONE ELSE WHO HAS BEEN INVOLVED IN THE CREATION OR PRODUCTION OF THE ACCOMPANYING CODE ("THE PRODUCT") CANNOT AND DO NOT WARRANT THE PERFORMANCE OR RESULTS THAT MAY BE OBTAINED BY USING THE PRODUCT. THE PRODUCT IS SOLD "AS IS" WITHOUT WARRANTY OF ANY KIND (EXCEPT AS HEREAFTER DESCRIBED), EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTY OF PERFORMANCE OR ANY IMPLIED WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. AP WARRANTS ONLY THAT THE CD-ROM(S) ON WHICH THE CODE IS RECORDED IS FREE FROM DEFECTS IN MATERIAL AND FAULTY WORKMANSHIP UNDER THE NORMAL USE AND SERVICE FOR A PERIOD OF NINETY (90) DAYS FROM THE DATE THE PRODUCT IS DELIVERED. THE PURCHASER'S SOLE AND EXCLUSIVE REMEDY IN THE EVENT OF A DEFECT IS EXPRESSLY LIMITED TO EITHER REPLACEMENT OF THE CD DISC(S) OR REFUND OF THE PURCHASE PRICE, AT AP'S SOLE DISCRETION.

IN NO EVENT, WHETHER AS A RESULT OF BREACH OF CONTRACT, WARRANTY OR TORT (INCLUDING NEGLIGENCE), WILL AP OR ANYONE WHO HAS BEEN INVOLVED IN THE CREATION OR PRODUCTION OF THE PRODUCT BE LIABLE TO PURCHASER FOR ANY DAMAGES, INCLUDING ANY LOST PROFITS, LOST SAVINGS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PRODUCT OR ANY MODIFICATIONS THEREOF, OR DUE TO THE CONTENTS OF THE CODE, EVEN IF AP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, OR FOR ANY CLAIM BY ANY OTHER PARTY.

Any request for replacement of a defective disc must be postage prepaid and must be accompanied by the original defective disc, your mailing address and telephone number, and proof of date of purchase and purchase price. Send such requests, stating the nature of the problem, to Academic Press Customer Service, 6277 Sea Harbor Drive, Orlando, FL 32887, 1-800-321-5068. APP shall have no obligation to refund the purchase price or to replace a disc based on claims of defects in the nature or operation of the Product.

Some states do not allow limitation on how long an implied warranty lasts, nor exclusions or limitations of incidental or consequential damage, so the above limitations and exclusions may not apply to you. This Warranty gives you specific legal rights, and you may also have other rights which vary from jurisdiction to jurisdiction.

THE RE-EXPORT OF UNITED STATES ORIGIN SOFTWARE IS SUBJECT TO THE UNITED STATES LAWS UNDER THE EXPORT ADMINISTRATION ACT OF 1969 AS AMENDED. ANY FURTHER SALE OF THE PRODUCT SHALL BE IN COMPLIANCE WITH THE UNITED STATES DEPARTMENT OF COMMERCE ADMINISTRATION REGULATIONS. COMPLIANCE WITH SUCH REGULATIONS IS YOUR RESPONSIBILITY AND NOT THE RESPONSIBILITY OF AP.

AppleTalk Network Services

Network Frontiers Field Manual Series

Dorian J. Cougias • Tom Dell • E.L. Heiberger



AP PROFESSIONAL

AP PROFESSIONAL is a division of Academic Press, Inc.

Boston San Diego New York
London Sydney Tokyo Toronto



AP PROFESSIONAL

An Imprint of ACADEMIC PRESS, INC.
A Division of HARCOURT BRACE & COMPANY

ORDERS (USA and Canada): 1-800-3131-APP or APP@ACAD.COM
AP Professional Orders: 6277 Sea Harbor Dr., Orlando, FL 32821-9816

Europe/Middle East/Africa: 0-11-44 (0) 181-300-3322
Orders: AP Professional 24-28 Oval Rd., London NW1 7DX


Japan/Korea: 03-3234-3911-5
Orders: Harcourt Brace Japan, Inc., Ichibunan

Australia: 02-517-8999
Orders: Harcourt Brace & Co. Australia, Locked Bag 16, Marrickville, NSW 2204, Australia

Other International: (407) 345-3800
AP Professional Orders: 6277 Sea Harbor Dr., Orlando FL 32821-9816

Editorial: 1300 Boylston St., Chestnut Hill, MA 02167; (617) 232-0500

Web: <http://www.apnet.com/>

This book is printed on acid-free paper. 
Copyright © 1996 by Network Frontiers, Inc.
All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission in writing from the publisher.

All brand names and product names mentioned in this book are trademarks or registered trademarks of their respective companies.

United Kingdom Edition published by
ACADEMIC PRESS LIMITED
24-28 Oval Road, London NW1 7DX

ISBN 0-12-192570-6

Printed in the United States of America
96 97 98 99 CP 9 8 7 6 5 4 3 2 1

DEDICATION

Dorian Cougias's Dedication:

*To the gang at MCO
who really know the definition of the word "service"*

CONTENTS

Introduction	xi
Acknowledgments	xv

Chapter 1:

Introducing E-Mail Messaging with QuickMail.....	1
Mainframe Message Centers	3
Electronic Bulletin Board Systems.....	5
LAN Client/Server Electronic Mail Systems.....	7
Peer-to-Peer Electronic Mail Systems	10
Enterprise LAN/WAN Electronic Mail Systems.....	12
A Better Definition of E-Mail	16
Setting up the Basic QuickMail Server.....	18
Using QM Administrator	22
Setting up the QuickMail Client.....	31
How QuickMail Messaging Works	43
Setting up the Advanced QuickMail Server.....	48

Chapter 2:

Publishing E-Mail with ListSTAR	77
Installing ListSTAR for SMTP.....	79
Setting up Administrative Services	89
Setting up E-Mail-on-Demand Services	96
Go Forth, Brave Souls.....	108

Chapter 3:

Gathering Internet E-Mail with Mail*Link Gateways.....109
Where the Mail Goes 111
Setting up the Mail*Link Gateway 112
Establishing User Accounts..... 131
Queued Messages 136

Chapter 4:

Exchanging Documents on the Small LAN with PowerTalk139
PowerTalk on the Desktop 140
PowerTalk to the World..... 151

Chapter 5:

Using the Apple Internet Mail Server as Mail Host155
Installing the Apple Internet Mail Server 156
Setting up Your Users..... 165

Chapter 6:

Networking Art with a Cumulus Database167
Working with the Cumulus Server 169
Working with Cumulus*Client 176
Planning for the Cumulus System 180

Chapter 7:

Database Servers and CGIs.....185
Planning the Database..... 189
CGIs or ACGIs 201

Chapter 8:

Deploying an SQL Database with Butler213
Working with the Butler SQL Server..... 215
Creating a Butler SQL Database..... 225
Interfacing with a Butler SQL Server..... 229

Chapter 9:	
Making Your Forms Smart with Informed	239
Informed System Components.....	241
Informed in Action	243
Planning for Electronic Forms	256
Chapter 10:	
Finding Information with AppleSearch	259
Setting up the AppleSearch Server.....	261
Using AppleSearch via AppleTalk	271
Using AppleSearch via the Web	276
Other Good Things to Know about AppleSearch	280
Chapter 11:	
Working with the Domain Naming System and QuickDNS Pro.....	285
Navigating the DNS Database	287
DNS Resource Records.....	290
Putting up QuickDNS Pro Server.....	296
Host Name Dos and Don'ts	301
Chapter 12:	
Serving the World Wide Web with WebSTAR.....	303
HTML for Fast Typists.....	304
Putting up a WebSTAR: The Basics	308
Putting up a WebSTAR: Additional Considerations	318
Chapter 13:	
Erasing Barriers in Digital Publishing with Adobe Acrobat	331
Components of Adobe Acrobat Publishing Software.....	333
An Exercise in Digital Publishing.....	336
Converting Existing Documents into PDF Material	340
Using Acrobat on the WWW	351

Chapter 14:	
Selecting Printers.....	361
Simple Printer Checklist.....	370
Chapter 15:	
Adding File Sharing and File Servers.....	371
What Does It Take to Be a File Server?	372
Non-Dedicated File Sharing.....	373
Dedicated File Sharing—Hardware that Doesn't Move.....	375
Short Version of a Needs Analysis	381
Building a Server Plan from Scratch	384
Workgroup Server User Access Planning.....	386
Workgroup Server Users & Groups Planning.....	389
Building a Server Plan for an Inherited Network.....	393
Chapter 16:	
Planning for Backup Management.....	397
Planning for and Implementing Retrospect Backup Server.....	399
Chapter 17:	
Scheduling Users with Now Up-to-Date and Now Contact	407
Setting up the Up-to-Date and Contact Servers.....	409
Setting up the Up-to-Date and Contact Clients	426
Chapter 18:	
Service Performance.....	431
E-Mail	434
File Sharing.....	440
Printing.....	446
Web Servers.....	448
Databases	452
Backup Servers	455
All Together Now	456

INTRODUCTION

Network design consists of two parts. The first part concerns all you need to know about cabling, hubs, switches, and routers. It encompasses RFPs, contracts, and putting network pieces together. All of that was covered in the companion book to this one, *Designing AppleTalk Network Architectures*. It's all important stuff. But what you need to know in addition to that is what you want to run *on* your network. That's why networks are built in the first place, isn't it?

Networks are built so that users can print, so that users can share files with each other, and so that users can send each other mail. All the hubs and switches in the world aren't going to do you any good unless you have a service to offer. That's the second part of network design: knowing how to offer services with which the users want to work, services that make the users feel that having a network is worth something to them. So guess what this book is about? Yup, just like the title says, it's about services design and how to get those services running.

This book doesn't flow like some of our other books. This book is really broken down into five different "service types." They are as follows:

E-Mail It used to be that networks were built so that users could share incredibly expensive printers. Now most networks are built so that users can share e-mail. It's kind of funny, but a CEO recently called me up in the middle of the night and he seemed to be in a cold sweat. "Do you know that our e-mail system is the backbone and lifeblood of communications in our company?" he asked. Sure, I did. I was the one who designed the system years ago. He was just figuring that out. Well, it *is* important to a lot of people. To that end we cover three different types of e-mail systems that you might want to employ. We cover the easiest system out there, PowerTalk, which ships with each Macintosh. We then show two more types, the Internet-friendly (but not very robust) Apple Internet Mail Server and the corporate e-mail application QuickMail. We also cover some e-mail options that you definitely will want to use, such as StarNine's ListSTAR and Mail*Link.

Databases The second most popular network service these days is the database. FileMaker from Claris is the "everyman's" database and is probably the most widely used system in the universe. Well, we not only cover it, but we also cover how to tie it into the World Wide Web using CGIs. To that end, we also cover Butler and Cumulus as well. Butler is a very strong SGML queryable database and Cumulus is by far

the best art database that we've ever seen. We also cover one of my personal favorites, *Informed Designer* from Shana. *Informed* is a package that lets you build custom forms and then route those forms to other users through e-mail (see a tie-in here?). Many of our books ship with some of the forms that we've designed. Great package and a great way to capture information!

Internet Related The next suite of applications we discuss falls within the rubric of Internet-related applications, like DNS servers and WebSTAR. To us, these are the key Internet-related services your network should employ. If you are creating a network that has links to the Internet, these will definitely be the applications you'll use, and probably plenty of others as well. Check out AppleSearch as a part of this. If you have information that you need to get out to a lot of people, then tying in AppleSearch with the Web is a fantastic, and often overlooked, tool.

Printing There are only two chapters about printing. One of them is about how to choose a printer, and the other is about that great paperless printer, Adobe Acrobat. If you deal with documents in any way (that means all of you), this should be on your must reading list.

Da Rest of da Services File Server choices, more information about backup planning, and two can't-do-without-them applications—Now Up-to-Date and Contact—are lumped together here. Check out Now Up-to-Date. In the *Managing AppleTalk Networks* book there's a calendar I've created called the network manager's calendar. I swear by it. I swear by this product as well.

Network Traffic

At the end of this book is a chapter about network services and the network traffic that they cause. We put it at the end of the book because it was the last chapter we wrote. It's a chapter that you **need** to read because it tells you about how much traffic each of these types of applications causes when you run it on the network. A part of network services is knowing the types of traffic that the services cause and how the services will affect network performance. Don't believe what you read in the magazines about how much traffic something causes. Follow what we tell you here, and do some of your own investigation as well. *Before* you put any of the services that we mention in this book on your network, you need to know how it will affect the ability of the network to deliver already existing services to your users. If you know that, then you are in good stead. If you don't, then you are taking your chances.

What Didn't Make It

What didn't make it into the book is anything about video conferencing and Webcasting. This stuff is too new to write about. You'll see a new book come out from us that covers adding multimedia to your network. That development will be something pretty significant and will definitely change at least some of the ways you provide services and plan designs. And by the way, if you are one of those folks who are saying that you won't be adding video conferencing to your network, get ready to change your tune. Look at the Nokia monitors! They are shipping with video cameras in them. I'm sure that by the time this book is a year old there will be more and more people who are doing direct video conferencing.

As Always

If you read something in here that doesn't seem right, or you need to know more, want to take a class, or you don't agree with us, feel free to call, e-mail, or fax us, and we'll try to help or put you in touch with the folks who can.

From the team at Network Frontiers, who include

Cassandra Kovel, Editor in Residence

Tom Dell, Chief Writer for the Book

Lynn Heiberger, Writer

Tom Hessel, Writer

and yours truly, Dorian Cougias, Chief Prognosticator

thank you for purchasing and reading this book. Here's how to get a hold of us:

Phone (415) 896-2939

Fax (415) 896-1573

All of the following e-mail addresses are at @netfrontiers.com

dorian_cougias

tom_dell

lynn_heiberger

cass_kovel

tom_hessel

Introduction

Our Web page is at *www.netfrontiers.com*.

One last thing: Be sure to check out the digital audio tracks from the Love Zombies on the CD-ROM that comes with this book.

ACKNOWLEDGMENTS

Special thanks and appreciation go to Jeff Pepper of AP PROFESSIONAL, our infinitely patient publisher.

We also thank the staffs of the Apple Market Centers around the country who have always made us feel at home when we teach in them.

Thanks also go to Adobe, AG Group, Apple, Canto, Claris, Dantz, EveryWare, Farallon, Helios, Men & Mice, Neon, Now, Santorini, Shana, and StarNine.

Special thanks to Luis Cubero of Hewlett-Packard.

We would also like to thank Dick Skeie for coming back to CE Software.

Very special thanks to our support staff of Felix Barrientos and Maizie Gilbert. Thanks also to our artist, Monte Lewis, and our indexer, Steve Rath. Tom Hessel made his writing debut with this book. We thank him especially.

CHAPTER 1: INTRODUCING E-MAIL MESSAGING WITH QUICKMAIL

There are currently more than 19 million electronic mail—or *e-mail*—users, who send well over 15 billion messages per year in the United States alone. It is by far the cheapest form of communication, averaging around 16 cents per page to send cross-country. Compare this with facsimile messages costing \$1.86 per page and \$13 for overnight express mail carriers. In the computing world, electronic messaging is big business.

There are many definitions of e-mail. There are the technical definitions, there are the product-specific definitions, and there are the corporate communication definitions that have been appearing lately in the “reinvent your business” books. What we will try to do first is meld most of these definitions together into a single definition that, hopefully, expresses the correct concepts. After that, we will go into great detail about how it all works so that you will have a good base of knowledge for our discussions of QuickMail, Mail*Link, PowerTalk, ListSTAR, and the Apple Internet Mail Server.

E-mail is the term used for the exchange of text-based communication—perhaps including enclosures of data and voice messages—between a sender and designated recipients. It is delivered through computer networks and telecommunication links. In other words, it is a method of communication using computers and associated networking technology. It is asynchronous and usually fairly fast. It can be as simple as sending a text-based message—like “Hi. How are you doing?”—


from one computer to another within an office, or as complicated as sending data and voice enclosures to other users across the world on different systems.

E-mail can be sent between computers using a variety of different methodologies and technologies. Some companies employ only one of those methods or techniques, others employ more than one, and some employ all of them. For your benefit, we will go over each of them and then show you how Network Frontiers' e-mail system fits into the categories we have explained.

MAINFRAME MESSAGE CENTERS

The mainframe messaging center was the first design in which e-mail was employed. Systems like IBM's PROFs and the commercial bulletin board e-mail systems of today (such as America Online and CompuServe) act as hosts to e-mail user clients. These clients link up to the host via internal cabling schemes or through public access lines. Once connected to the system, the client can access waiting mail and send mail via a password-protected user account.

Mail *notification* is only performed when the user has connected to the system. The system has no notification method when the user is not connected. Because of this, most systems use the inbox/outbox metaphors for e-mail storage.



Mainframe Messaging Center

Characteristics of Mainframe Message Centers

Network Design

If you want to connect to a mainframe mail system, whether that mail system is internal to your organization or one such as America Online or CompuServe, you need to create a mail gateway from your LAN-based e-mail system to the mainframe mail system. All messages flowing from your organization to members of the mainframe system will pass through this single gateway. This means that the amount of mail in megabytes will affect your e-mail system's network performance and storage space requirements.

Costs

If you are connecting to a mainframe messaging system that is external to your organization, you must first purchase an account. These accounts can sometimes cost quite a bit of money over a year of regular usage. Some accounts have a fixed fee, but most of them are based upon the number of kilocharacters you transfer between your system and theirs. Up to \$0.05 per kilocharacter is common. It might not sound like much now, but add it up over the year and it could run into thousands of dollars!


Legal Issues

In this system, the host acts as a repository for mail awaiting transfer from the sender to the receiver. With many commercial systems, even after the mail has been downloaded to the local computer by the recipient, the message remains in the host's storage system for a certain amount of time before it is discarded. Users have options for removing the mail manually in most commercial systems.

Public electronic mail systems, such as America Online, CompuServe, AppleLink, and Source, all provide records of electronic mail transactions that could be used to disclose activity. These records and files are subject to United States laws and can be retrieved for litigation purposes, if necessary.

ELECTRONIC BULLETIN BOARD SYSTEMS

Electronic bulletin board systems are the most basic of all electronic messaging environments. In a bulletin board system, users *post* messages in the various areas of concern on the bulletin board. Once the message is posted, any other reader who enters that area of concern on the bulletin board can read and reply to the posted message.



Electronic Bulletin Board System

There is usually a person in charge of the bulletin board called a *System Operator*, or a *Sysop*. Larger commercial bulletin boards break down the responsibilities for different areas of concern and place individual Sysops in charge of each area. It is this person's duty to maintain all the bulletin board *chat* areas and direct posted messages to the correct areas of the bulletin board. Thus, this person has access to *all* the messages posted on the bulletin board and can eliminate any messages that he or she deems inappropriate. This gives the Sysop a lot of control over the management of messages. It also means that the Sysop is almost guaranteed to read most of the posted messages.

Bulletin boards are a great place to gather information. If a topic is hot, many replies are posted. In many cases, the replies themselves branch off into other discussions. These secondary discussions are called *threads*.

Characteristics of Electronic Bulletin Board Systems

Network Design

Much like the mainframe message systems, you can create a gateway to electronic bulletin boards, and they can act as e-mail repositories.

Some of the bulletin boards have “auto-download” features that can monitor conversations and threads in messages, and then forward those conversations and threads to a user’s mail account. From there, your LAN-based e-mail system can download the information through a gateway into an individual’s mail account.

Costs

Commercial bulletin boards cost money. Each connection costs something. You might want to ensure that the connections you establish will be worth what you will be paying for them.

Security


All messages posted to bulletin boards, whether they pertain to a topic in general or are strictly between two people, are open for public viewing and response.

Messages that are sent directly to users are generally private. However, any message sent to a user can be *forwarded* by that user to anyone else who is a member of the bulletin board system, or to whole groups of individuals.

Sysops are under certain legal obligations concerning data privacy and privileges.

LAN CLIENT/SERVER ELECTRONIC MAIL SYSTEMS

This is what you will have if you purchase the Macintosh platform's market leader in LAN client/server e-mail systems, CE Software's QuickMail. LAN *client/server* e-mail systems use the concept of an e-mail *server* that acts as the routing entity for e-mail being passed between sender and receiver. Users, or *clients*, receive accounts on the mail server. They can connect, either automatically or manually, to the mail server during the course of the work day. Once connected through the network, users can send e-mail to other users on the network via the mail server. In this case, mail is sent from the sender to the mail server only. The mail server then *notifies* the recipient that mail is waiting at the server to be picked up. The recipient may then request that the e-mail be forwarded to his or her computer by the server machine.



LAN Client/Server Electronic Mail Systems

LAN-based client/server systems require that there be a dedicated mail server running on the LAN. In the case of QuickMail, this service resides on a Macintosh. The e-mail server software can run concurrently with some other dedicated services, such as file services, but *not* other systems, like Dantz' Retrospect. User accounts and passwords are set up at the server or through remote server management applications.

All mail in this *store and forward* stasis is located at the server until it can be distributed. There are two main methods for distribution in a client/server model. One method is through *file-based* transactions. Here the server does not act as a transportation engine. It merely files mail in the appropriate "in" and "out" boxes and then uses the LAN's native file transport mechanisms to move the mail

between sender, server, and recipient(s). The other method is through a standard mail transport format, such as MHS or X.400, or through a proprietary engine similar to the one implemented by QuickMail.

Characteristics of LAN Client/Server E-Mail Systems

Network Design

QuickMail needs to run on a Macintosh. The machine on which the QuickMail server is running cannot be the network's backup machine because of network traffic interference during the backup process. It should not be a workstation. It can be both an e-mail and a file server.

Costs

There is a cost associated with the server machine. Fortunately, QuickMail runs well on low-end machines. An old SE/30 works just fine in most cases, but a base model PowerMac is preferable. Additional costs depend on how many users you will support and how many bridges and gateways you will have. The number of users will determine how large the hard drive needs to be, and the bridges and gateways will determine the RAM requirements.

You need a modem if you want to dial out to other services or networks, or if you wish users to be able to dial in to your network from remote locations. This is highly recommended. You may also need an Internet feed.

You need the initial QuickMail software, plus enough user accounts to support each of your users. You should also budget for training for yourself and your users.

Administration

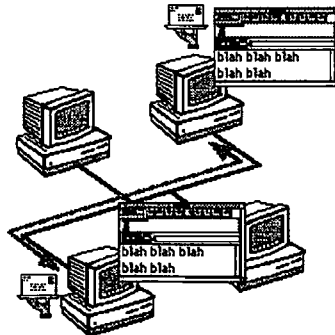
All messages are stored on the server until they are requested by the recipients. Most systems maintain a copy of files for a certain amount of time *after* they have been forwarded to the recipients.

Security

Most systems file the messages in their own proprietary format, but most of those formats can still be decoded using a standard word processor or text editor, unless an encryption option is available.

PEER-TO-PEER ELECTRONIC MAIL SYSTEMS

The Apple Open Collaborative Environment (AOCE) offers peer-to-peer electronic mail under the PowerTalk system. Instead of sending mail to a recipient *through* a mail server, the mail is sent directly from one computer to another. Outgoing mail is kept in the user's *Outbox* until it can be forwarded to the intended recipients. Mail in the *InBox* is actually located on the user's computer instead of on the mail server.



Peer-to-Peer Electronic Mail Systems

To the administrator, this means that mail that has been addressed but not sent is not located on a central server. Instead, it is distributed throughout the network on the senders' computers.

PowerTalk, which we discuss later in the chapter called "Exchanging Documents on the Small LAN with PowerTalk," can also use the client/server methodology through its personal gateways, but these do not make the computers true clients.

Characteristics of Peer-to-Peer E-Mail Systems

Network Design

Apple created the PowerShare server to act like the QuickMail server in implementing store-and-forward, but they have since stopped supporting it. When planning for PowerTalk, plan for a peer-to-peer system. That is its strength.

PowerTalk can be used with QuickMail through the QuickMail Personal Gateway from Apple.

Administration

Mail is decentralized throughout the corporate network, meaning that it is *much* harder to back up and protect than a centralized client/server system.

Tracking the trends of electronic mail usage is nearly impossible with today's tools.

Security


Sensitive data in store-and-forward mode on a mobile computer goes mobile when that computer leaves the office.

ENTERPRISE LAN/WAN ELECTRONIC MAIL SYSTEMS

In the real world (where most of us live), none of these systems stand alone. Instead, we usually have a mixture of all or part of these systems in the same corporate e-mail environment. Mail messages not only move from the sender to the receiver, they also traverse various systems that can be similar or as different as night and day. What we will do here is put together some of the more common arrangements, including a final arrangement that shows everything tied together.

LAN-to-Host Electronic Mail Systems

Many of the well-established companies with which we work have moved some of their e-mail users to LAN-based systems, but still have a great many electronic mailboxes on their legacy mainframe mail packages. For one reason or another, they will not move completely to a LAN environment. Because of that, a *gateway* from the LAN system to the mainframe system must be used to translate messages from one format to the other.




LAN-to-Host Electronic Mail Systems

LAN-to-LAN Electronic Mail Systems

It has long been said that “no Mac is an island.” Who said it? I don’t know, but it has been used to refer to the fact that Macintoshes work well on LANs. Sometimes it is used to refer to the fact that the Macintoshes need to communicate with the

PCs of the world. With respect to electronic mail it means that a great many of us end up sending mail across multiple LAN systems, whether they are local system to local system or local system to another local system over a dial-up connection.



LAN-to-LAN Electronic Mail Systems

Connect to Everything and Anything

The next illustration shows the way Network Frontiers communicates through e-mail. Starting from bottom right and working clockwise, we will explain with whom we connect and how.

Bottom Right

Our LAN-based electronic mail system consists of both a client/server system (QuickMail) *and* a peer-to-peer system (PowerTalk). They communicate with each other through the Apple Personal PowerTalk Gateway for QuickMail, formerly from StarNine Technologies.

Bottom Left

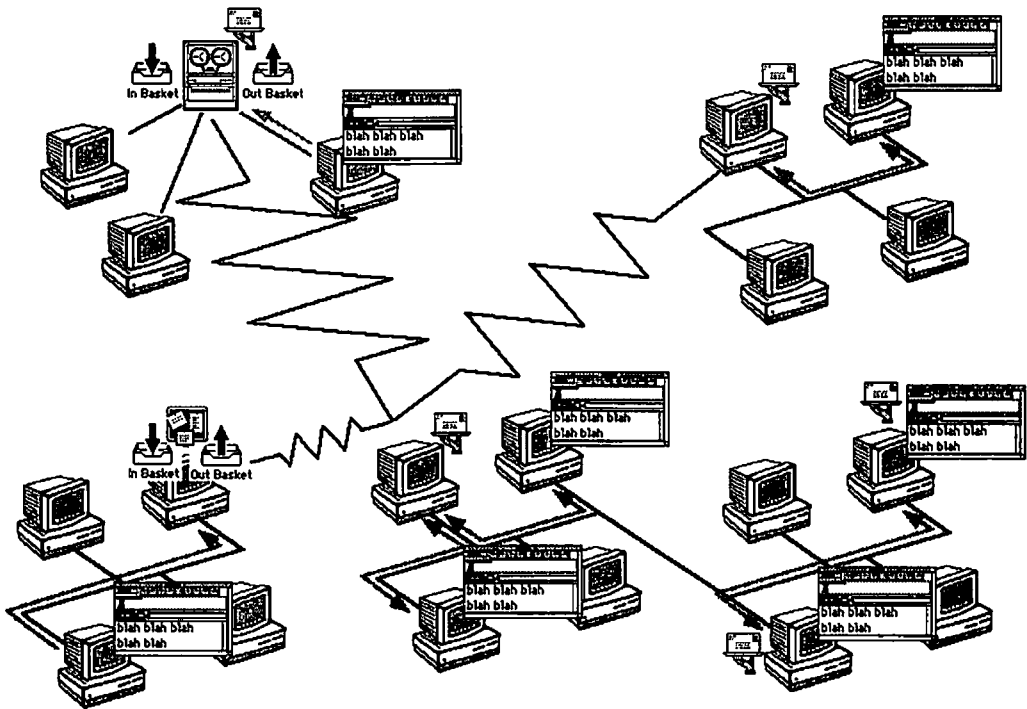
Through the use of QM-Script Gateway, we also connect to CompuServe's commercial bulletin board and electronic mail service. Not only do we download and upload electronic mail messages, we use QuickMail's scripting technology to gather information on topics of interest.

Top Left

Through the use of QM-Link Gateway provided with QuickMail, we connect to the AppleLink commercial bulletin board and electronic mail service. Since there is no scripting language associated with AppleLink, we use it only as a connection method to other AppleLink users.

Top Right

In addition to commercial services, we also connect to other mail systems through regular telephone lines. On an average, we send approximately 100 electronic messages per week to other organizations across the country.



Network Frontiers' E-Mail System

Not shown, because we cover it in so much detail elsewhere, is our Frame Relay LAN-to-Internet feed. This is serviced by two systems: QuickMail with the

Mail*Link SMTP gateway (*netfrontiers.com*) and the Apple Internet Mail Server (*pop.netfrontiers.com*).

So what? Well, we didn't show you this to prove how well we connect with the outside world. We showed you our electronic mail environment so that you can quickly understand what we came to realize after some time: Security in e-mail messages that flow through multiple formats defaults to the lowest common denominator. Even though AOCE allows the use of key signatures and QuickMail WAN is *very* secure, other paths are not.

Characteristics of LAN/WAN E-Mail Systems

Network Design

Consider throughput and capacity between users on a single mail server, across multiple mail servers on a single network, across multiple mail servers on separate dial-up networks, and between users on the LAN and users on the Internet.

Administration

Administration must be accomplished at several levels, including the LAN, any gateways to outside systems, and gateways to other, different internal systems.

Security

Plan for security at the level of the lowest common denominator.

Legal Issues

Mail using outside services must follow state, interstate, and international law.

A BETTER DEFINITION OF E-MAIL

Electronic mail is a relatively new form of communications that will change the way your organization operates. As the telephone changed the way we communicate, so too has e-mail. For some reason, a lot of people do not realize this. There is an interesting book on the market by Cheryl Currid called *Electronic Invasion: Brave New World of Business Communications* (Brady Publishing, 1993). Her book covers many aspects of business communications, one of which is e-mail. Her definition is about the same as ours, and she, too, sees e-mail as a culture-changing mode of communication.

Currid notes that as users become more comfortable with e-mail as a communication method, the communication culture changes within an organization. The way they use e-mail begins to change, as well as the way they communicate with each other. Currid speaks of three phases in this path of change: the adjective stage, the noun stage, and the verb stage. You can tell how comfortable people are with the technology by the way they talk about their communication system.

Adjective Stage

This is the beginning stage of using e-mail as a method of communication. Fred says to Joe, "I'll send you an e-mail message." At this point users are differentiating e-mail from other forms of communication. Currid points out that messages are somewhat formal in nature, with the following characteristics:

- They are long, with multiple topics per message.
- They are structured like formal letters.
- Users follow the rules of organizational behavior.
- Users do not copy other users without prior approval.
- Users print out their messages continually.
- Users do not trust the system very much.

Noun Stage

During this next stage, users are more familiar with the e-mail system. They no longer distinguish between this method of communication and other methods, such as voice mail and handwritten memos. Users make statements like, "I'll send you an e-mail on that." Here e-mail is becoming the preferred method of communication. Some of the behaviors you will notice are:

- Messages are not as formal.
- Messages are more direct and cover single topics.
- The paragraphs are shorter.
- The content is much more focused.
- Message sending is more frequent.

Verb Stage

Among the advanced users, e-mail is now a verb. You will hear your users saying, "I'll e-mail you about that." People make action-oriented statements referencing their e-mail system. Once e-mail is a verb, you will notice these e-mail behaviors:

- Message formality has gone out the window.
- Messages are short and to the point.
- The paragraphs have become one or two sentences in length.
- Action items are requested clearly, as are responses to questions.
- *Flaming*, or sending emotional statements, sometimes occurs.

SETTING UP THE BASIC QUICKMAIL SERVER

CE Software has made the process of installing the QuickMail system on the LAN a breeze. Run the installer on a dedicated Macintosh, restart the machine, and welcome to the wonderful world of e-mail (adjective stage)!

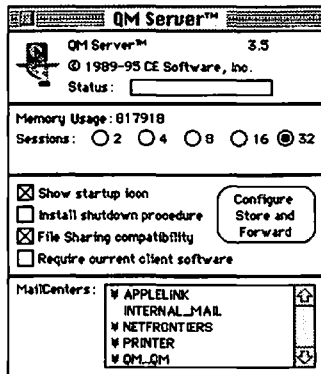
QuickMail Server Parts

What has been installed are QuickMail's essential, but as yet unconfigured, components, which are as follows:



QM Server

This control panel notifies network users when they have e-mail and distributes that e-mail to them. This is also where the number of allowable simultaneous sessions is set, each session being a single transaction between server and workstation.



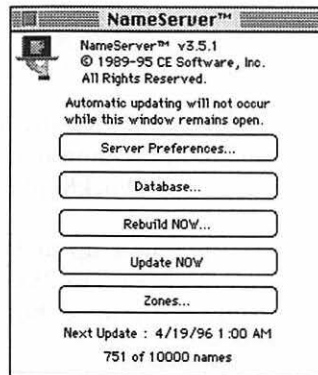
QM Server Control Panel

QM Server runs in the background using memory set aside in the system heap. It should be in the Control Panels folder of the System Folder. If it is not running, mail will not be distributed. Each QM-Server-equipped computer can handle up to 32 MailCenters.



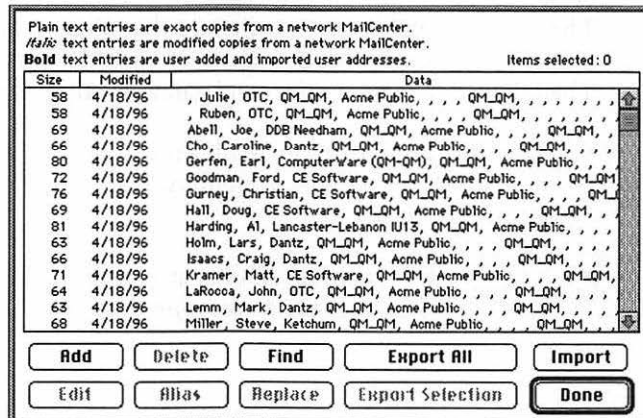
NameServer

The NameServer control panel usually resides on the mail server, where it runs in the background using memory from the system heap. This control panel gathers information from various MailCenters on the network, with which it creates a database of names and e-mail addresses. NameServer acts as a directory for users searching for others' addresses on the network and for routing messages that have come in through a bridge or gateway.



NameServer Control Panel

NameServer should be found in the Control Panels folder of the System Folder. Up to 16,399 names and addresses may be stored in each NameServer. Greater numbers are possible, but these addresses will not be visible in the Database window and are therefore difficult to administer.



NameServer Database



QM Administrator Application

QM Administrator is used to create and maintain MailCenters and the user accounts in these MailCenters. If bridge and gateway MailCenters are being used, or if users are dialing in from remote sites, this application *must* be running to handle communications.

When the QM Administrator application is opened, it will activate all the existing bridges and gateways, find the online MailCenters, and initialize any modems.

MailCenters

MailCenters are repositories for the mail associated with user accounts. *Online* MailCenters contain the accounts of users connected directly to the network—for example, Macintosh and QuickMail for DOS with AppleTalk users.

Each mail server can handle up to 32 online MailCenters. Each of these MailCenters can accommodate up to 254 user accounts, although no more than 100 accounts are recommended for all but the fastest server machines.

File-based MailCenters are repositories for mail associated with QuickMail for DOS with File Servers, QuickMail for OS/2, and QuickMail for Windows users' accounts. These are online MailCenters that reside physically on an AFP-compatible, file-based server volume. There can be only one file-based *server volume* per mail server.

The *bridge* MailCenter is software that allows communications between two like devices or protocols, such as one QuickMail server to another QuickMail server. The *gateway* MailCenter is software that allows communications between two unlike devices or protocols, such as a QuickMail server to an information service. The following are the bridges and gateways shipped with QuickMail.

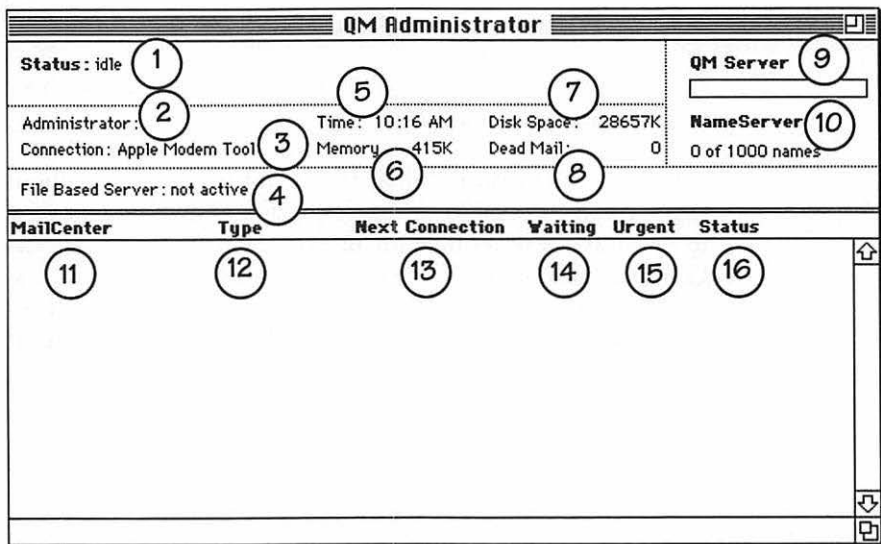
- The QM-Printer Bridge MailCenter is used by QM Administrator to collect mail for those without computers. It sends the mail to the printer in a batch at specified times.
- The QM-Direct Bridge MailCenter directs mail traffic to a MailCenter in another location, possibly across long-distance telephone lines.

- The QM-QM Bridge MailCenter is used to create a MailCenter that directs mail traffic to multiple MailCenters in other locations, usually across telephone lines.
- The QM-Script Gateway MailCenter directs mail traffic to information services such as MCI Mail, Sprint Mail, CompuServe, and other customized services. This software also has a number of associated prewritten scripts.
- The QM-Link Gateway MailCenter provides access to a specific commercial information service, AppleLink.
- The QM-AOL Gateway MailCenter provides access to a specific commercial information service, America Online.

The Store & Forward MailCenter is an integral but nearly invisible MailCenter that is created when QuickMail is installed. Any time a message needs to be routed to a mail server other than the one on which its sender has an account, this MailCenter is called upon. There is only one such MailCenter per server. There are few administrative functions that cannot be performed with QM Administrator. When you first establish your QuickMail service, you'll use it almost exclusively to set up your various MailCenters and user accounts. Later, you will use it in maintaining these accounts.

USING QM ADMINISTRATOR

After using QM Administrator to set up your online MailCenters, QM Server will distribute mail to and from your users with NameServer, which will keep track of their addresses. QM Administrator looks like the following screen shot.



The QuickMail Administrator Window

1. **Status.** This describes the current action QM Administrator is taking. Some actions may be completed so quickly that you cannot read the description, but you will see words flash on the screen.
2. **Administrator** or **MailCenter** or **User.** If no connection is being made, “Administrator” is listed. If a connection to a gateway or bridge MailCenter is occurring, the name of the connected MailCenter is listed. If a user is calling in using QM Remote or any other telecom package, the user’s name is shown.
3. **Connection.** This lists the tool used to handle incoming calls to this machine when a remote site or user connects to your site. Different tools may be used when a bridge or gateway calls from this site to another site.

4. **File Based Server.** If this machine is used as a server for file-based MailCenters, this line names the mounted AppleTalk File Protocol (AFP) volume that is used for server duties. If this machine is not a server for file-based MailCenters, it displays “not active.”
5. **Time.** This displays the current time as set on this computer’s system clock.
6. **Memory.** This is the amount of free RAM available to QM Administrator.
7. **Disk Space.** This is the amount of disk space available to QM Administrator. QM Administrator can use free space only on the volume containing the System Folder.
8. **Dead Mail.** This is the number of messages stored on the mail server that cannot be delivered to the addressee or returned to the sender.
9. **QM Server.** This status bar displays how much of QM Server’s processing power is being used.
10. **NameServer.** The number of names stored and the number of names that are allowed to be stored in the NameServer database are shown here.
11. **MailCenter.** This shows all MailCenters that QM Administrator can find in the zones searched. The icons vary and may look like any of the following.



These MailCenters are located on the local hard drive.



These file-based MailCenters are located on a file-based server volume.



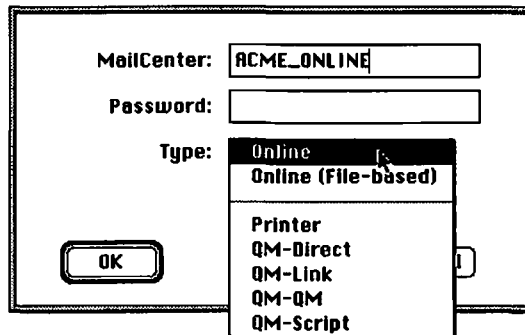
These MailCenters are located on a different mail server’s hard drive.

MailCenters listed in plain text can be administered from this computer. MailCenters listed in italics are “Not Handled.”

12. **Type.** This lists either “Online” or the bridge or gateway software used by this particular MailCenter.

13. **Next Connection.** This lists the scheduled time that a bridge or gateway MailCenter will activate and trade mail with another site. This is blank for on-line MailCenters or a bridge or gateway MailCenter with no calling schedule.
14. **Waiting.** This is the number of messages waiting to be sent from the bridge or gateway MailCenter.
15. **Urgent.** This is the number of top-priority messages that are waiting.
16. **Status.** This lists failure codes if this MailCenter is not connecting properly. It also may display “Not Handled” if this machine does not administer this MailCenter or “No bridge/gateway” if the necessary bridge or gateway software is not installed.

Before any local users can exchange mail, an online MailCenter must be created with QM Administrator to hold users’ account information. To do this select **New MailCenter** from the **File** menu. The MailCenter window appears.



Defining the MailCenter

In the **MailCenter** field, enter a name. All lowercase characters are automatically changed to uppercase and spaces are converted to underscore characters. Enter something in the **Password** field if you want. Although passwords are not required, they do prevent unauthorized access to a MailCenter. If the password is lost or forgotten, however, you are hosed! Select **Online** in the **Type** field. Clicking **OK** opens the first configuration dialog box for the new MailCenter.

Name: RCME_ONLINE
 Type: Online
 Password:
 Custodian: (None)
 Clear log every
 Send log every days
 Configure... OK Cancel

The Configure MailCenter Window

Enter a name and password for the user who will be custodian and click **Add**.

First Name: Tyrone
 Last Name: Johnson
 Password:
 Add Done

Entering the Custodian

So, what is a “custodian?” Although the terms *administrator* and *custodian* are often used interchangeably, they are not necessarily the same person. The custodian is usually someone responsible for a specific MailCenter. This person has probably been assigned these duties by the administrator.

If you are the QuickMail *administrator*, your duties probably include:

- **Supporting network users.** Answering questions about workstation software use, setting up online accounts, distributing QuickMail software, address books, groups, forms, and (perhaps) passwords. You might also be called upon to troubleshoot workstations.
- **Supporting remote users.** Answering questions about modem communications and distributing QuickMail software, address books, groups, and forms.

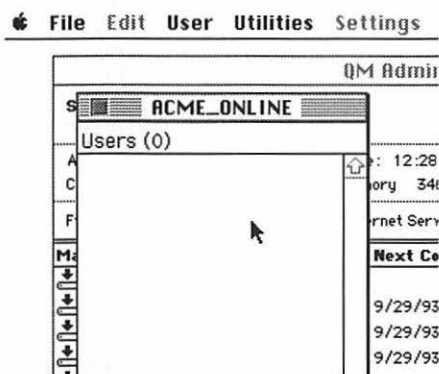
- **Maintaining the mail server.** Updating and upgrading server software when needed, regularly backing up the hard drive, clearing hard drive space of dead mail and inactive accounts, and troubleshooting hardware problems.
- **Coordinating efforts of other administrators.** If your network regularly communicates with another site through a QuickMail-to-QuickMail bridge, you will need to discuss policy and exchange user lists and phone books.
- **Coordinating efforts of custodians.** Support those people you have charged with maintaining MailCenters.
- **Helping to create policy.** This might involve simple decisions such as how many MailCenters will be created and who will use them, and more complex issues such as how your company will deal with e-mail privacy. Good planning at the beginning makes for easier going later. You will probably work with other network administrators and management on this.

If you are a QuickMail *custodian*, your duties probably include:

- **Maintaining MailCenters.** Helping end users with their accounts, reviewing logs for problems, and acting on notifications and alerts. You might also be charged with some or all of the duties of the administrator.

At smaller companies, you might assume both roles. Larger companies may have many people supporting QuickMail. Either way, administrators and custodians are vital system components.

Once you have created your first online MailCenter, you can add users' accounts by double-clicking on the MailCenter's name in the QM Administrator window.



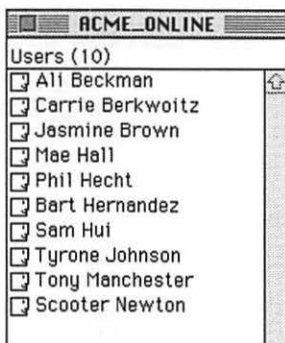
The MailCenter Open and Ready for New User Entry

Select **Create** from the **User** menu to open the User dialog box.

User Dialog Box

Enter the user's first name, last name, and password (if desired). When finished, click **Add**. The dialog box will remain on the screen, allowing you to enter multiple names without having to reselect to create a new user from the **User** menu. Click the **Done** button when finished.

The user names appear in the MailCenter's user list.

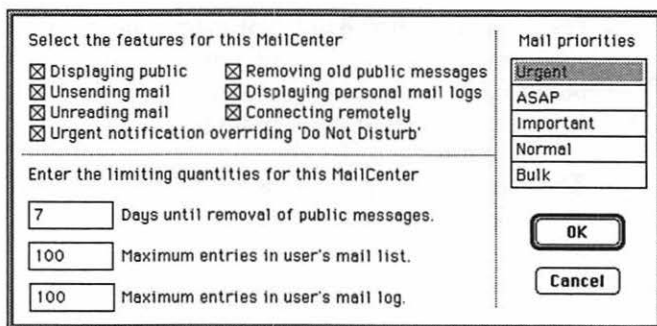


Online MailCenter User List

Close the MailCenter, and you are finished. If you are not using any bridges or gateways, you may close QM Administrator as well. Before doing so, however, there are a few more things you could do.

The way each MailCenter works can be tailored to the needs of your network environment. To change a MailCenter's configuration, double-click on the MailCenter in the QM Administrator window to open it, and then select **Configure** (*name of MailCenter*) from the **File** menu.

The Configure Online MailCenter window is displayed with the following available options.



Configure Online MailCenter Window

Leave the **Displaying public** checkbox selected if you wish to let the users of this MailCenter post messages visible to all users in the MailCenter. The **Unsending mail** option allows users to delete messages after they have sent them. The message

cannot be unsent after it is read by the recipient. If the **Unreading mail** checkbox is selected, users are allowed to change a message to Unread status after reading it. This can serve as a reminder for an important message. If the **Urgent notification override...** checkbox is selected, members of this MailCenter are notified when messages marked Urgent arrive for them, even if they have the Do Not Disturb feature of QuickMail turned on.


As your MailCenter activity increases, public mail messages will accumulate. If you select the **Removing old public messages** checkbox and specify the number of days next to **Days until removal of public messages**, messages older than specified are deleted automatically. This not only reduces the disk space taken up on the mail server, but also keeps only the most recent messages in the Public folder. A message with a subject beginning with an asterisk will not be deleted by this option. The **Displaying personal mail logs** checkbox allows a listing of the mail sent by individual users to be generated. This uses disk space on the mail server, but it allows users to track or unsend messages. You can limit the space consumed by each user's mail log by typing a number into the field next to **Maximum entries in user's mail log**. Users who exceed this number will see a dialog box that tells them they have exceeded the limit. They are asked to clear entries from their mail log, although they are not prevented from adding new entries. If the **Connecting remotely** checkbox is selected, users can call into this MailCenter using a modem and check, send, and read mail. (See the "Supporting Remote Connections" section of this chapter beginning on page 74 for more information.)

In the **Mail priorities** fields, you can change the words used to describe the five levels of QuickMail messages. The top level always has the highest priority in the QuickMail system regardless of the name. These settings only affect messages that are to be read by users of this MailCenter.


If the number entered in **Maximum entries in user's mail list** is exceeded, a dialog box appears asking or reminding the user to file or delete some messages currently in the user's mailbox.

Before we move on to the client end, we want to show you one more thing. You can communicate with a MailCenter's users by sending them a *newsflash*, which is a message that appears at the bottom of their QuickMail windows. This can be useful for informing users when a service interruption will occur or perhaps to direct them to documentation that you have posted in the Public folder, such as an e-mail privacy policy. To do this, open the desired MailCenter, and select

Newsflash from the **File** menu. A dialog box appears in which you may type your message, as shown in the next screen shot.



The Newsflash Window



The Newsflash in the User's Main QuickMail Window

SETTING UP THE QUICKMAIL CLIENT

While the QuickMail server requires a Macintosh, the QuickMail package comes with client software for computers running DOS, Windows, or the MacOS. Because of this, a cross-platform QuickMail system is possible as long as all the computers involved speak AppleTalk. Here are the essential components of the client software.

Windows Client

QMAILWIN.EXE

This is QuickMail's client software for Windows PCs connecting to the mail server through a file-based MailCenter.

QMNOTE.EXE

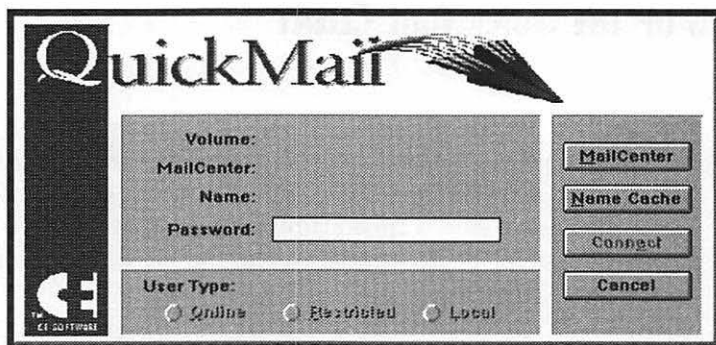
This is a notifier TSR that informs QuickMail for Windows users when mail has been received for them.

CE.INI

This is a file that supplies user information during workstation bootup.

CONFIG.EXE

This program allows workstation software preferences to be set up or changed.



QuickMail for Windows Connection Screen

DOS Client

QMAILATP.EXE

This is QuickMail's client software for PCs connecting to the network with an AppleTalk-compatible card.

QMATP

This is a notifier TSR that informs QuickMail for DOS with AppleTalk users when mail has been received.

QC.EXE

This is a notifier TSR that informs QuickMail for DOS with AppleTalk users when another online user wishes to hold a realtime QuickConference.

QCONF.EXE


This allows realtime conferencing between QuickMail for DOS with AppleTalk users and other online users.

CE.INI

This is a file that supplies user information during workstation bootup.

CONFIG.EXE

This program allows workstation software preferences to be set up or changed.



QuickMail for DOS Connection Screen

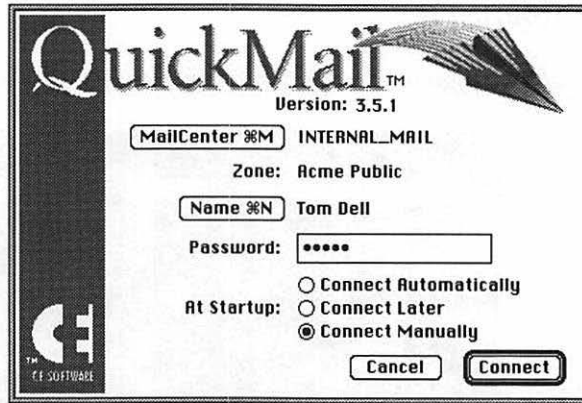
Macintosh Client

QuickMail Desk Accessory

This is QuickMail's client software in which messages are composed and sent and received and read. This software includes an INIT extension that informs users when mail has been received.

QuickConference

The QuickConference system extension allows realtime conferencing among online QuickMail users.

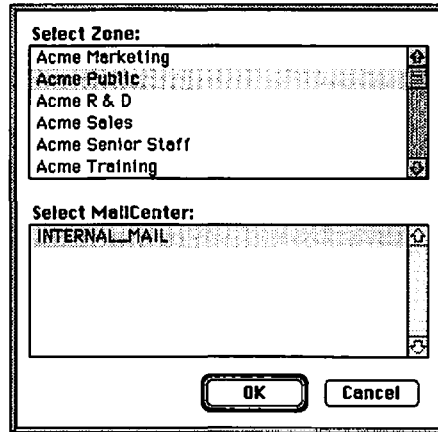


QuickMail for Macintosh Connection Screen

Since the Macintosh is the best computer in the universe, we will use the Macintosh to illustrate QuickMail client use.


After installing QuickMail, locate your new MailCenter and tell it who you are when you connect. You should see the QuickMail main connection screen when the computer restarts. If you do not, access it by selecting **QuickMail** from the **Apple** menu.

First, click the **MailCenter ⌘M** button to select from available MailCenters. A list of zones appears in the top half of the dialog box. Select the appropriate zone if applicable, and a list of its MailCenters will appear in the bottom section of the window. Highlight the MailCenter you created and click **OK**.



MailCenter Selection Dialog Box

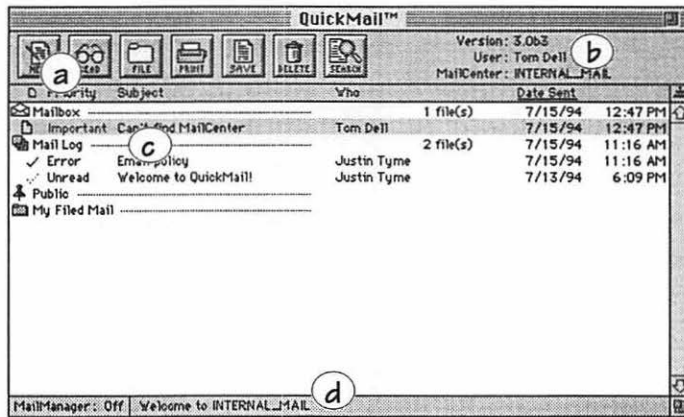
Next, click the Name \mathfrak{N} button. All the names belonging to the chosen Mail-Center are displayed. Highlight your name by clicking it or by typing the first letters of your last name. Click **OK** when finished.



Selecting a Name for QuickMail

Also type your password if you have one. Letters are replaced with bullets as they are typed to protect your password from straying eyes.

Once you have logged in, you are presented with the QuickMail Main window. This is where mail is created, read, filed, printed, deleted, saved, spell-checked, forwarded, and more.



The Main QuickMail Window

There are several components of this window that are critical to the way in which QuickMail performs:

- a. These are the function buttons. Clicking them performs the function identified by the icon.
- b. This is the version of QuickMail software you are using, your connection name, and the MailCenter to which you are connected.
- c. These are the listings of mail that have been sent to you (Mailbox), that you have sent to others (Mail Log), that are posted for public viewing (Public), or that have been filed on your local workstation in personal folders. Each message is tagged with the date and time it was sent.
 - The *mailbox* is your receptacle for up to 100 pieces of incoming mail. When you receive new mail, QuickMail flashes its icon above the **Apple** menu and plays a tone through the Macintosh speaker.
 - The *mail log* is where up to the last 100 messages that you have sent to others are stored for a time so that you may review them.
 - The *public* section of QuickMail is like a bulletin board. It's a place where anyone can post a message for everyone to read without addressing it to everyone. Public messages are controlled by the sender or the administra-

tor and are usually set to be automatically deleted after a specified period of time, the default setting being seven days.

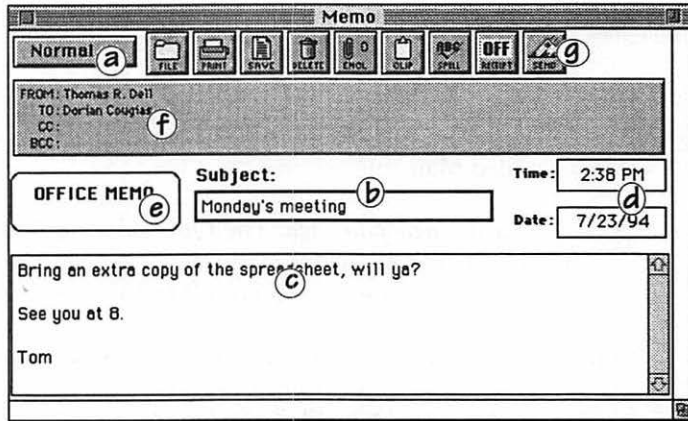
- You may create as many *personal folders* as you want to store messages locally for future reference. One has already been created for you called “My Filed Mail.”
- d. This is the *newsflash* area. The QuickMail administrator can put a one-line message here and change it any time. You can also determine if QuickMail’s MailManager function is on or off by looking here.

Here is how to send mail through the QuickMail system. To initiate the creation of a message, choose from one of the predesigned forms. To do this, use the **New** button in the main QuickMail window. A scrolling list of forms appears in a dialog box. Select one by double-clicking it.



New Form Dialog Box

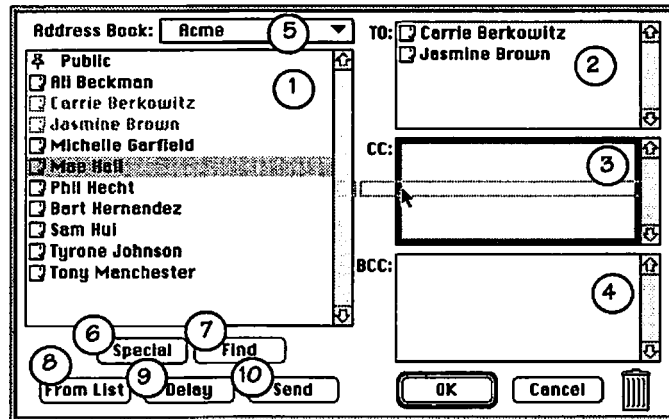
Forms serve a useful purpose when sending or receiving messages through Quick-Mail. They are a way to organize the information in an easy-to-read manner and can also portray a mood. Here are the parts of one of the most popular forms, Office Memo.



The Office Memo Form


- a. **Priority.** This pop-up menu allows you to choose from one of five priority levels for this message. This gives the recipient some indication of the urgency of your message.
- b. **Subject.** This is a text field where a brief meaningful title for this message should be typed.
- c. **Body.** This is where your message is entered. If you type more than will fit into this field's size, most forms will display scroll bars.
- d. **Time & Date Stamps.** Some forms automatically insert the date and current time into predefined fields when they are created. This text can be altered, but usually does not need to be.
- e. **Static Text.** Many forms include text that cannot be edited, called *static text*. The purpose of static text is to enhance the form or add to its functionality.
- f. **Address Label or Button.** After a QuickMail form is chosen and a message is entered, it is time to address the message. Click the **Address** label.
- g. **Send Button.** Once your form is complete, click this button to transmit the message to the QuickMail server, where it will be forwarded to all recipients.

After clicking the **Address** label or button (part f above), a dialog box appears, similar to the following screen shot, which allows you to choose the recipients.



Selecting Users in the Addressing Dialog Box

1. **Recipient List.** This is a scrolling list of all the users in your active *address book*. Users may be selected one at a time by dragging one name, or several at a time by holding down the **Shift** key and clicking multiple names.
2. **TO:** Drag users into this field to make them main recipients of this message.
3. **CC:** Drag users here to assign them as recipients of a “carbon copy” of this message. This means that the message was not targeted directly at them but contains information that they should know about.
4. **BCC:** Drag users into this field to assign them as recipients of a “blind carbon copy.” These users receive a copy of the message, but no other user will know that they have received it.
5. **Address Books.** If you have more than one address book, you can select a different one in this pop-up menu. This allows you to organize users into categories.



Changing Address Books

6. **Special.** This button is used to create a user “on the fly.” This could be someone you know exists but who is not in your address book. The following dialog box appears when you click this button.

The dialog box is titled "Create Special Address:". It contains the following fields and buttons:

- First:** A text field containing "Dorian".
- Last:** An empty text field.
- MailCenter:** A text field containing "INTERNAL_MAIL".
- Zone:** A text field containing "Acme Public".
- Address:** A large empty text area.
- Buttons:** "MailCenters", "OK", and "Cancel".

Creating a Special Address

If you simply enter the first and last name of the user to receive this message, QuickMail will search the NameServer for the proper MailCenter and zone. This user appears as an icon with dark glasses, since you are sending the message “blind,” hoping it will be received successfully.

 **Nathaniel DuPont**
A Special Address User

7. **Find.** Click this button if you want to locate a user in some other address book but are not certain of the spelling of his or her name. A dialog box similar to the following picture appears.

Enter Find criteria

First Name

Last Name

Department

Finding a User

Type the first or last name or just a few letters of either and click the **Find** button. QuickMail will attempt to locate this user in all available address books. Any matches will be displayed in a scrolling list. Select the name to which you wish to send this message and click the **Add** button. The user is added to the **TO:** field (but can be dragged anywhere) for this message.

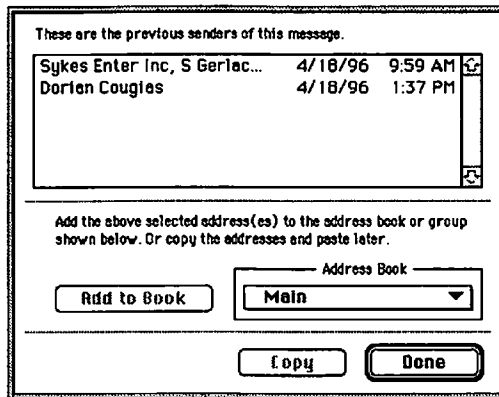
Names matching Dorien,*,*

Dorien Cougias INTERNAL_MAIL

Cougias,Dorien,COUGIAT,INTERNAL_MAIL,Acme
Public,,,INTERNAL_MAIL,,,,,

Results of a Find

8. **From List.** This is useful when replying to an existing message. You can have the previous senders added to one of your address books by clicking the **Add to Book** button. You can click **Copy** to store them in memory to be pasted into the **TO:** field of a new message.




Choosing a User in the List

9. **Delay.** Click this button to set the date and time you want this message sent. QuickMail can hold a message for any length of time before sending it to the recipients. You can also use this feature to send yourself mail as a reminder.
10. **Send.** Click this button when all other options have been completed to your satisfaction. This will transmit the message to the QuickMail server, where it will be forwarded to all recipients.

HOW QUICKMAIL MESSAGING WORKS

The process of delivering messages through the QuickMail system appears deceptively simple. When, as happens on occasion, mail is lost or misrouted and administrators attempt to determine why, they find out just how complicated it can be. Here is the path some users think that their mail takes through QuickMail.



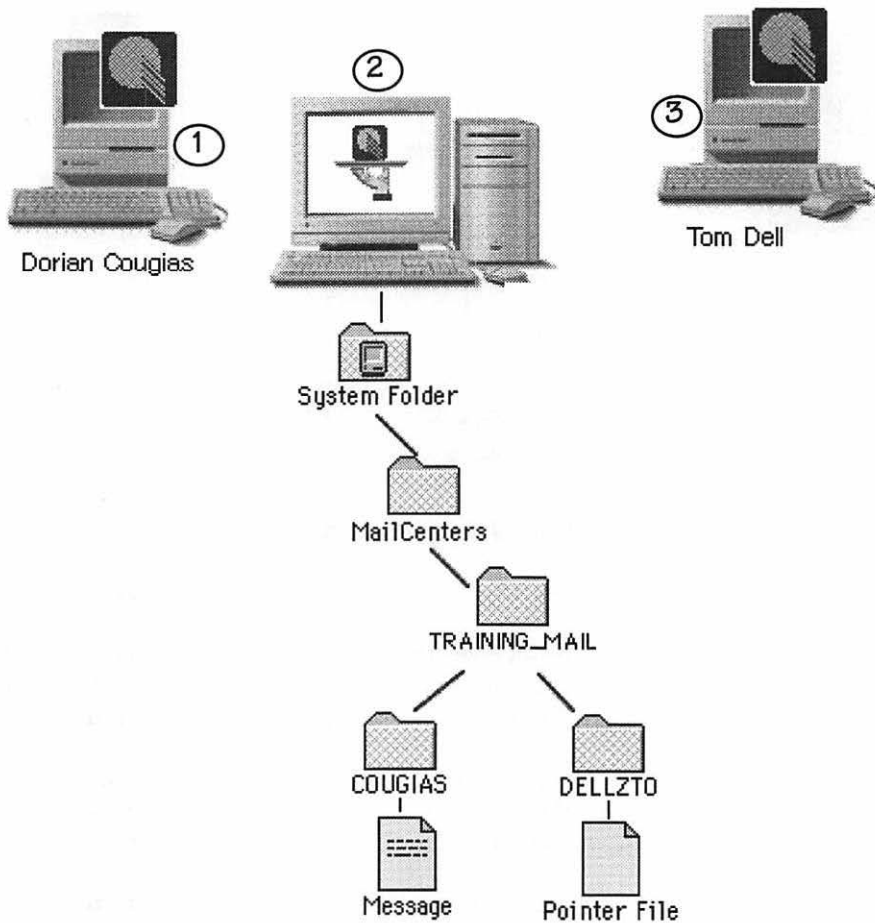
Is this the route a QuickMail message takes from sender to receiver?

1. The user creates the QuickMail message for another user.
2. Before sending the message, the user queries the mail server for address information. The user might not need the mail server. If the address information (recipient's name) is already known, the user just types it in.
3. With the address information provided by the user or the mail server, the sender launches the message onto the network, where the recipient's computer sees it and picks it up. The end user reads the message.

While this may seem logical, it is baloney. Mail is not sent directly from one workstation to another. In a simple transaction, mail is sent from a workstation to a

mail server, which in turn communicates with a recipient workstation. In more complex transactions, the sender's mail server may pass the information along to other mail servers that in turn forward the message to the recipient. If the recipient's mail server is offsite, bridges may be called into play. If the recipient's mail server is of a system other than QuickMail, gateways may be used.

Here is the path that a QuickMail message *really* takes between users, in this case between "Dorian Cougias" and "Tom Dell," both of whom are members of the "TRAINING_MAIL" MailCenter on a single server.



Path of a QuickMail Message from Sender to Recipient in the Same MailCenter

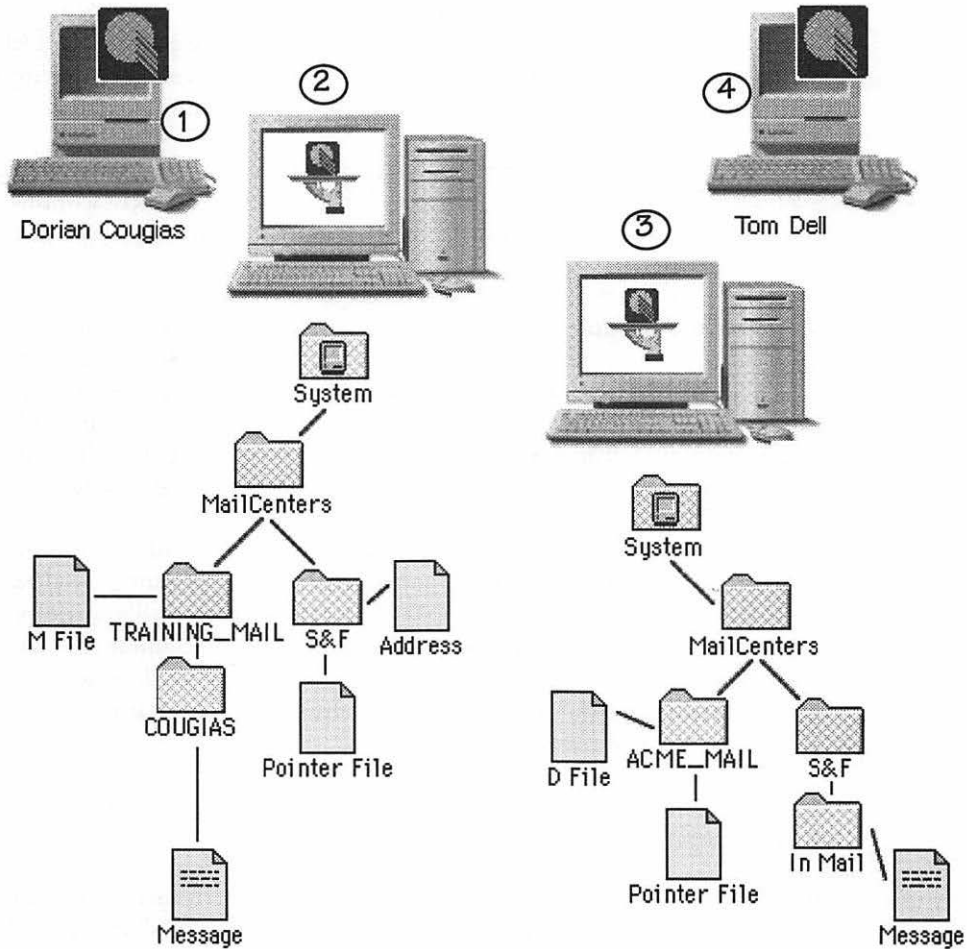
1. The first QuickMail user (Dorian Cougias) creates a QuickMail message, which he sends addressed to a user in the same MailCenter (Tom Dell). Although in this case the sender is using a Macintosh, he could be using an Intel-based PC equipped with QuickMail for DOS with AppleTalk.
2. QM Server receives the message and places it in the *sender's* online MailCenter folder. QM Server then places a *pointer* file in the recipient's online MailCenter folder. The small pointer file does nothing more than tell QM Server where to find the original.

Besides the unique online MailCenter folders associated with each user account, there are also files for each user's mail log entries (the *M file*) and mailbox entries (the *D file* or *Directory file*). QM Server makes an entry in the sender's M file so that he may track the message's status as "unread." QM Server also makes an entry in the recipient's D file so that he will see the message in his mailbox when he opens the QuickMail desk accessory.

3. QM Server communicates with the recipient's QuickMail notification INIT, which triggers a flashing icon over his **Apple** menu. Opening the desk accessory, the recipient (Tom Dell) sees the message in his mailbox and opens it. For this to happen, QM Server must again be called upon to look at the pointer file and locate the original message, which the desk accessory may then generate on the screen. QM Server also changes the status of the message in the sender's M file to "read."

This is the most simple mail transaction. It becomes more complicated when you add multiple QM Servers.

Here is the path that a QuickMail message takes between users of two different network mail servers, in this case between "Dorian Cougias" in "TRAINING_MAIL" and "Tom Dell" in "ACME_MAIL."



Path of a QuickMail Message from Sender to Recipient on a Different Mail Server

1. The first QuickMail user (Dorian Cougias) creates a QuickMail message, which he sends addressed to a user (Tom Dell) in another MailCenter that happens to reside on another mail server.
2. QM Server receives the message and places it in the sender's online MailCenter folder. QM Server then places a pointer file and a copy of the recipient's address information in the S&F MailCenter folder. This is the Store & Forward MailCenter that is called into play whenever the QM Server sees that

it needs to send mail beyond one mail server. QM Server makes an entry in the sender's M file so that the sender may track the message's status. QM Server also makes an entry in the S&F MailCenter's special D file. QM Server periodically calls out to the other mail servers on the network so that they may exchange mail. If the destination mail server is not available, the first mail server keeps trying for a given amount of time.

3. When contact is made, the *first* QM Server forwards the message to the second QM Server and updates the sender's M file and S&F MailCenter's D file. The *second* QM Server places the original message in its S&F MailCenter's special In Mail folder. It then places a D file entry and pointer file in the recipient's MailCenter folder.
4. QM Server on the recipient mail server communicates with the end user's QuickMail notification INIT, which triggers a flashing icon over his **Apple** menu. Opening the desk accessory, the recipient (Tom Dell) sees the message in his mailbox and opens it. QM Server is called upon to look at the pointer file and locate the original message, which the desk accessory may then generate on the screen. The next time the two mail servers communicate, the recipient QM Server notifies the sender QM Server of changes so that it may update the status of the message in the sender's (Dorian Cougias') M file.

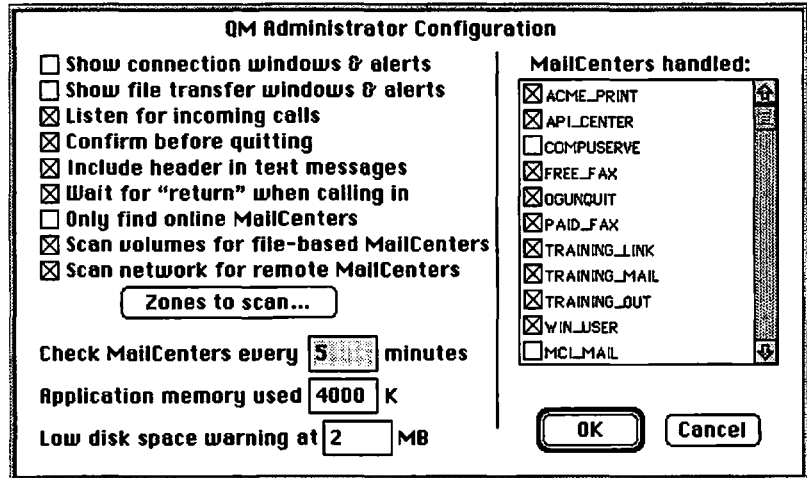
This is a more complicated mail transaction. Although the variables can change somewhat—for example, more mail servers or more MailCenters—these two processes provide the best overview of how mail is *really* transferred.

SETTING UP THE ADVANCED QUICKMAIL SERVER

Now that you have an understanding of how the QuickMail system looks and acts, we can give you a quick overview of how it can be expanded beyond the Macintosh and beyond your LAN through some of its many bridges and gateways.

Bridges and gateways are easy to set up in QuickMail and can extend your e-mail service to a global scale. Before attempting to install any bridges or gateways, however, there are some common rules that you *must* keep in mind.

- **Don't forget the passwords.** As with online MailCenters, bridge and gateway MailCenters can be password-protected. This prevents meddling by anyone at the mail server and, more important, by anyone with remote computers. If you forget the password, all you can do is delete the MailCenter and recreate it. *This will delete the entire user list and all the MailCenter's messages.*
- **Configure QM Administrator.** In the Configure QM Administrator window is the checkbox **Only find online MailCenters**. If this box is selected, any mail coming to you from another location will be delivered only to users in the online MailCenters of your mail server. If this box is not selected, your mail server will accept mail destined for others and forward it to users in other bridge or gateway MailCenters. You may or may not want to do this, as will be apparent later.



QM Administrator Configuration Window

- **Do not establish duplicate accounts.** When adding accounts to a bridge or gateway user list, make sure you are not adding users who already exist in another MailCenter. This will cause messages to be misrouted.

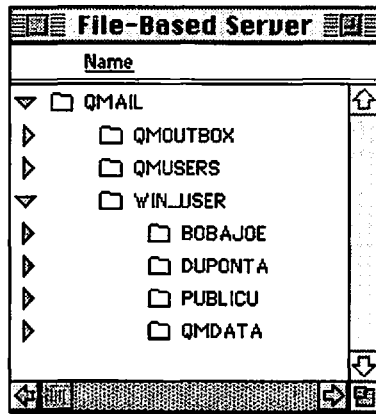
File-based MailCenters

Your network's users do not need a Macintosh to access QuickMail. Intel-based PCs need only have some type of network communication card installed—LocalTalk, Ethernet, or TokenTalk—along with the appropriate *AppleTalk* software drivers. Once networked, PCs have two options in communicating with the mail server. QuickMail for DOS with AppleTalk allows these computers to communicate directly with the Macintosh mail server and other Macintosh computers if the PC has been loaded with the AppleTalk protocol drivers. The second option is to use QuickMail for DOS with File Servers or QuickMail for OS/2 software. This software allows these machines to communicate with file-based servers that are in turn communicating with the mail server.

Here is the path a message takes between a Macintosh workstation and a PC running QuickMail for Windows.

1. The QuickMail Macintosh user creates the QuickMail message, which is sent to a PC user in a file-based MailCenter (WIN_USER).

2. QM Server inspects the message and places the original file in the S&F MailCenter, as it is destined for a different server. As regularly scheduled, the S&F MailCenter activates and passes the message on to QM Administrator. In the meantime, QM Administrator places the file in the FB Stuff folder on the mail server.
3. Eventually, QM Administrator transfers the message to the recipient's account directory on the file-based server volume (inside WIN_USER). QM Administrator "sees" what is depicted in the next screen shot.



File-Based MailCenter as Seen by QM Administrator

QM Administrator also updates the sender's M file (on the mail server) and the recipient's D file (in the QMDATA directory). On the Windows PC, the program QMNET periodically examines the file server in search of new mail. This program "sees" something like the following screen shot.

```

MS-DOS Prompt
E:\QMAIL\WIN_USER>dir

Volume in drive E is FILE-BASED
Directory of E:\QMAIL\WIN_USER

.                <DIR>
..               <DIR>
BOBAJOE          <DIR>      07-23-93   8:58a
DUPONTA          <DIR>      08-16-93   7:38p
PUBLICU          <DIR>      07-23-93   8:58a
QMDATA           <DIR>      07-23-93   8:58a
6 file(s)        0 bytes
622804992 bytes free

E:\QMAIL\WIN_USER>

```

File-Based MailCenter as Seen by QMNET

Having seen new mail, QMNET informs the QuickMail Notifier. The Notifier alerts the user.

4. The QuickMail for Windows user opens and reads the message.

Your network *does not* require a special Intel-based file server to accommodate cross-platform QuickMail. If your network has a Macintosh server running AppleShare, one of its volumes can be dedicated as a file-based server volume that both the mail server and PCs can access.

Each mail server can handle up to 32 file-based MailCenters. Each of these MailCenters can accommodate up to 254 user accounts, although no more than 100 accounts are recommended. Again, there can be only one file-based *server volume* per mail server.

To set up a file-based MailCenter, create, if it doesn't already exist, the file server volume in which the file-based MailCenter will reside. Do this with whatever file server software you are running that supports AppleTalk AFP transactions.

If you have not done so already, use your file server software to create an AppleShare client on the AFP file server. Another way of looking at it is to create a user on the AFP file server that is actually the Macintosh mail server. The mail server must have a permanent account on the file-based server with maximum privileges.

(For more information on this topic, refer to Appendix E of the *QuickMail Administrator Manual*).

Mount the file-based server volume from the Macintosh mail server using the new Macintosh server account. Use the Chooser desk accessory to do this.

Tell the QM Administrator application where to look for the file-based server volume. To do this, select **Configure File-based** from the **Utilities** menu. In the window that is generated, type in the requested information.

AFP File Server

Zone name: Acme Admin

Server name: Acme Ethernet Server

Volume name: File-Based Server

User name: Fax Mail Center

User password: *****

AFP File server access delay: 30 seconds

Low disk space warning at: 5 MB


Cancel OK

Configure File-based Server Volume Window

Zone name pertains to the network zone in which the file-based server is publishing itself. **Server name** corresponds to the network name under which the file-based server is publishing itself. **Volume name** pertains to the name given to the file-based server volume. **User name** is the client name the file-based server has given the mail server that it will use to log in. The mail server must have a permanent account with maximum privileges. **User password** relates to the client password the file-based server has assigned to the mail server's client account. **AFP File server access delay** applies to how long the mail server will wait for a response from the file-based server when logging in before assuming it is not up on the network. **Low disk space warning** pertains to how much disk space will be used on the server volume before QM Administrator warns file-based MailCenter custodians that room is becoming scarce.

Click **OK** when finished. QM Administrator will set up the necessary mail server files on the server volume. You may now configure a file-based MailCenter by selecting **New MailCenter** from QM Administrator's **File** menu.

When the first window appears, enter the desired name in the **MailCenter** field and any desired password in the **Password** field. Then, select **Online (File-based)** from the MailCenter **Type** pop-up menu. Selecting this MailCenter type causes a new choice to be generated. Select the **Drive**, or mounted AFP volume, that is to contain file-based MailCenters.



Selecting the AFP Server Volume

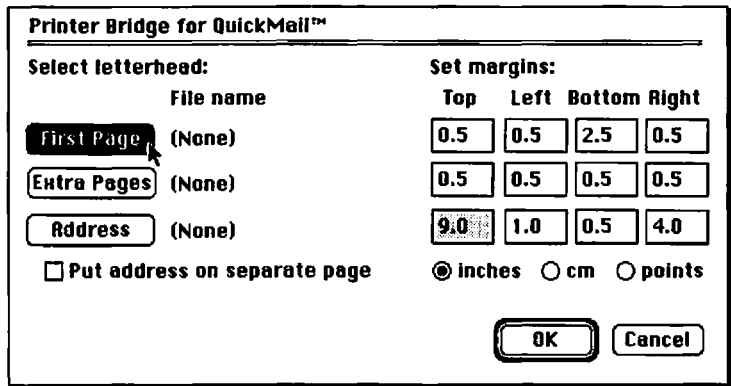
Clicking the **OK** button creates a new MailCenter. In the windows that follow, proceed as you would in creating any other online MailCenter.

Printer Bridge

The Printer Bridge allows QuickMail messages to be sent directly to a single network printer. This allows the administrator to set up a collection point for employees who do not have computers but who still want to receive electronic mail. It can also be used to send out postal mailings to groups of offsite employees or even clients. Someone just needs to pick up the printed messages and stuff the documents into envelopes. In this way, you could consider the Printer Bridge to be an e-mail-to-“real”-mail gateway!

You may have up to 32 Printer Bridge MailCenters on a server, although the server can print to a single printer only. The Printer Bridge MailCenter is one of the easiest bridges to create and maintain. To set up this service, select **New MailCenter** from QM Administrator's **File** menu.

When the first window appears, enter the desired name in the **MailCenter** field and a password in the **Password** field. Next, select **Printer** from the MailCenter **Type** pop-up menu. Click **OK** and, in the next window click **Configure**. The following window appears, containing the default margin and letterhead settings.



Configuring the Printer Bridge

The **Select letterhead** side of the window is for setting the path to letterhead files for the different pages on which the various messages will be printed. The letterhead can be any PICT or EPSF file. They must be stored in the QA folder within the mail server's System Folder. **First Page** is the actual first page of the memo, not a cover sheet. If a graphics file is chosen with this button, that picture will be placed on the first page. **Extra Pages** are for memos that require more than one page for all their contents. Whichever graphics file is selected for extra pages will be printed on all pages after the first page of the mail message. **Address** is for a cover page to be printed *before* the first page. If a cover page is going to be used to hold the user's address information, the **Put address on separate page** box should be selected. If not, the user's mailing address information is placed automatically on the first page.

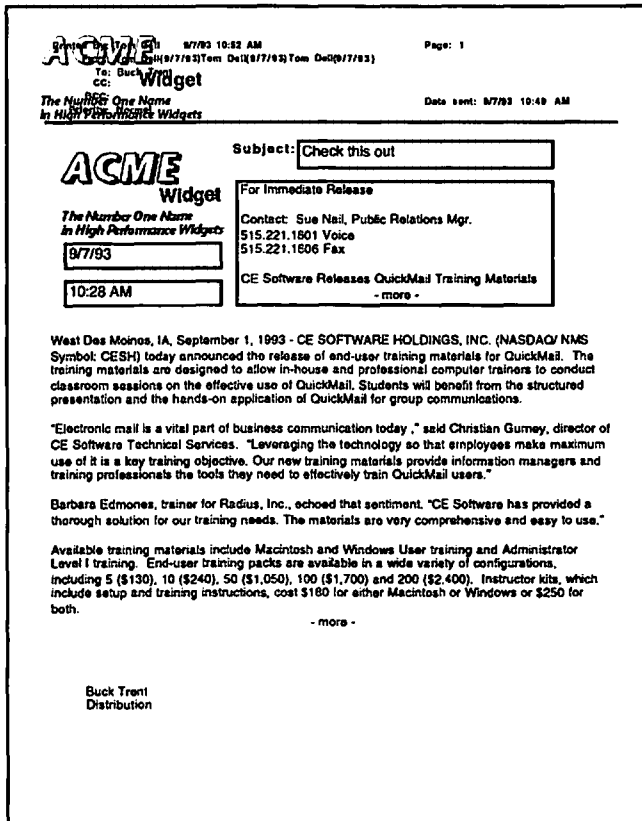
The **Set margins** side of the window can be used to establish dimensions for the *contents* of the first page and any extra pages. The e-mail's address header, subject, text, and any graphics visible in the QuickMail form chosen by the sender will be placed within the parameters set here. Any graphics added on the **Select letterhead** side of the window will ignore these margins.

You can also determine where the user's address will be placed. If a cover page is going to be used, the address information will appear on a separate page. *No mat-*

ter on which page it is placed, the address information will follow the address margins exactly.

There are a few things worth pointing out about graphic letterhead. The first is that the graphic will print in the same location as it was drawn in the original file. Wherever the image is created in the EPSF file, that is where the image will be placed as letterhead in the printed memo. In other words, if the picture is created in the top left of the page in the EPSF file, it will be printed in the same location in the printed memo. If the picture is created dead center in the EPSF file, that is exactly where the file will be printed in the mail message. This does *not* apply to PICT images. They print at the top left of the page, no matter what.

Also, if the margins are set incorrectly, the addresses or the text of the memo will print *over* the graphic letterhead. The following picture shows a form created with a logo and the Printer Bridge's default settings. The logo was saved as a PICT. Notice where the logo is printing on the page—at the very top left corner.



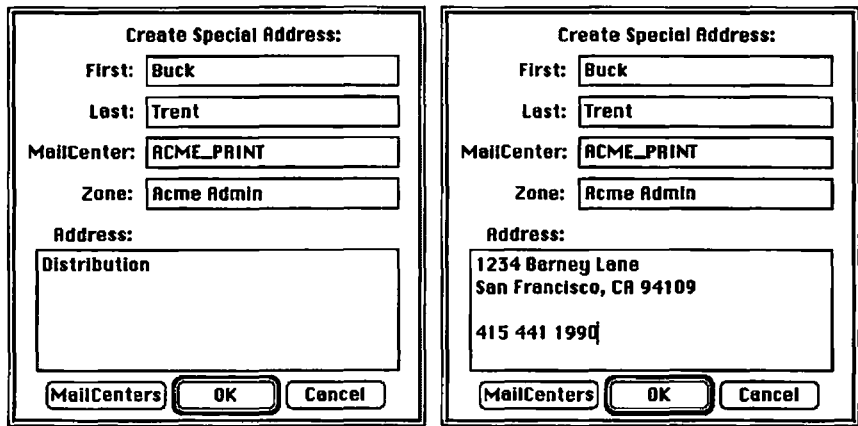
Incorrectly Placed Letterhead

Default margins for the message addressing and body are a half an inch on top, left, and right. The bottom is set to two and one-half inches. On the extra pages, the e-mail header is repeated with the company logo again overwriting it. The mailing address is tacked onto the lower left corner of the first page. Finally, we have a redundancy since the memo form used also bears the company logo. To make the printed messages more appealing, here is how we reset the defaults.

We decided not to use graphics on the extra pages. We again clicked on the **Extra Pages** button in the Printer Bridge for QuickMail configuration window, but this time we chose **Cancel** in the standard file selection box that came up. We reset the margins for the first page to two inches on top, zero inches for the bottom, and one-half inch for left and right. On the extra pages, we went for a half-inch all

Center is being used to send mail out of house, rather than if it is just being used for in-house employee information memos.

Printer Bridge MailCenter user accounts are created in the same manner as are those of other MailCenters: directly, by using the **Create** item from **QM Administrator's User** menu, or by importing the information from a database. Depending on how you use the Printer Bridge MailCenter, the address information in the user accounts might be different. In the following set of screen shots, for instance, on the **Address** field at the left only the department for an in-house employee was used. In the case on the right, the user's postal address was used.



Account with Department Address vs. Account with Postal Address

Either way, this is the information that relates to the **Address** button in the Printer MailCenter Configuration window.

QM-Direct Bridge

The QM-Direct Bridge allows your users to send and receive mail through a direct connection with one other QuickMail site. This other site might be another company division or perhaps another organization with which your company has constant contact, such as a supplier. The other site might also serve as a hub, permitting your users to reach more than just the other site by connecting your site to a Wide Area Network star configuration.

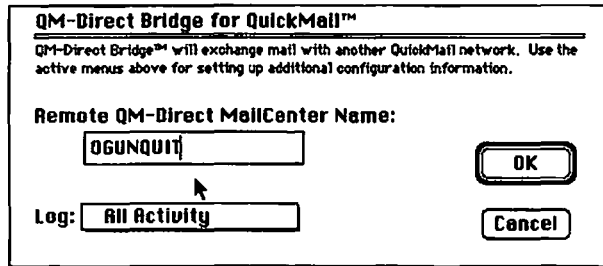
CE Software describes the QM-Direct Bridge as a phone on a “hot line.” When used, it dials to one other place only. The QM-QM Bridge would be more like a standard phone that can dial many numbers. There can be up to 32 QM-Direct MailCenters on a server. Each QM-Direct MailCenter contains users of the other site’s MailCenters.

Setting up this type of connection requires that the QuickMail Administrators at both ends of the bridge coordinate their efforts. Before anything can be accomplished, both sites must agree on the following:

- Which site will initiate calls, if not both, and what the phone numbers are.
- What the QM-Direct MailCenters at both ends will be called.
- When the QM Administrator application will be running at both sites, if not running at all times.
- The QuickMail account designations for users who will be exchanging mail through the bridge.
- The Comm Toolbox connection and file transfer tools both sites will be using. *They must be the same.*
- The times and conditions under which the connection will be made.

To set up this service, select **New MailCenter** from QM Administrator’s **File** menu. When the first window appears, enter the desired name in the **MailCenter** field and any password in the **Password** field. Then, select QM-Direct from the **MailCenter Type** pop-up menu. Use a name that relates to the given site or organization. Click **OK** to create a new MailCenter serving as the QM-Direct Bridge.

Click **Configure** in the next window. The following window, which contains instructions for how QuickMail should interact with the other site, appears.



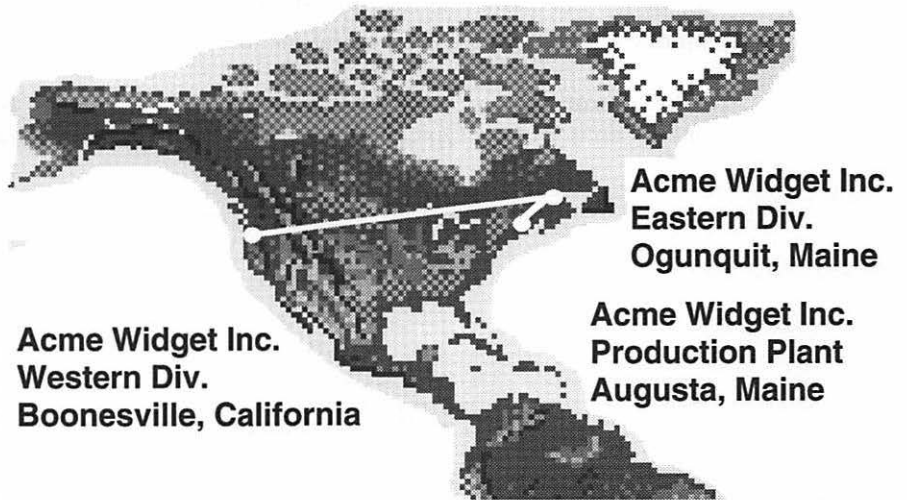
Configuring a QM-Direct Bridge

In this window, type in the name of the remote MailCenter. To finish setting up the QM-Direct Bridge MailCenter, enter the other site's phone number in the Connection Settings window.

QM-Direct Bridge MailCenter user accounts are created in the same manner as are other MailCenters: directly by using **Create** from QM Administrator's **User** menu, or by importing the information from a database. As with other online MailCenters, leave the **Address** field blank when typing in these names.

If there are more than a few users with which your site will be communicating, ask the administrator of the other site to send an exported list of users from its mail server.

There are situations in which your company might want to use a QM-Direct Bridge star configuration instead of QM-QM Bridges. In one scenario, several outlying offices (spokes) dial into one mail server (hub) using QM-Direct. The hub mail server forwards any mail it receives from one site to the other site or sites.



QM-Direct Link between West Coast and Production Divisions Routed through East Coast Hub

In terms of phone charges, this might make more sense than setting up numerous long-distance dialing QM-QM Bridges at outlying sites.

This bridge is fairly secure. Since the QM-Direct Bridge connects only two sites and connects only at agreed upon times, you will not receive unwanted mail from unknown sites.

QM-QM Bridge

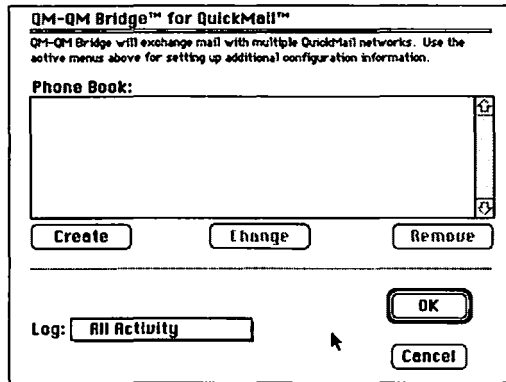
The QM-QM Bridge allows your users to send and receive mail through a connection with one or more other QuickMail sites. These other sites might be other company divisions or perhaps other organizations with which your company has constant contact, such as suppliers.

If your office needs to connect to one other site only, you might want to set up a QM-Direct Bridge instead. If more than one site is to be contacted, the QM-Direct Bridge method would require a new MailCenter for each site. One QM-QM Bridge MailCenter can communicate with many sites.

To set up this service, select **New MailCenter** from QM Administrator's **File** menu. When the first window appears, enter the desired name in the **MailCenter**

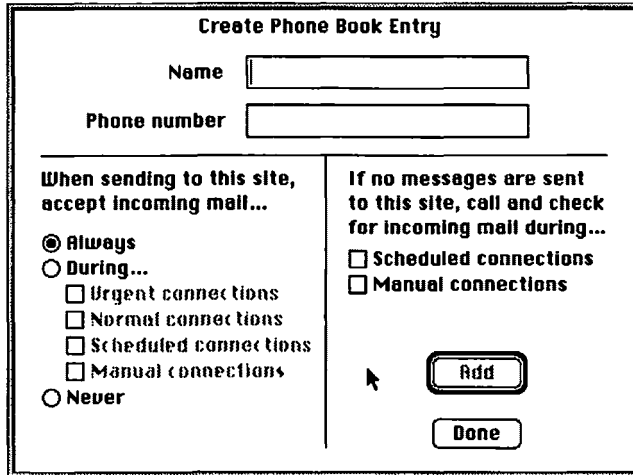
field and any desired password in the **Password** field, and then select **QM-QM** from the MailCenter **Type** pop-up menu.

Click **Configure** in the next window. Another window appears containing instructions for how QuickMail should interact with other mail servers.



Configuring the QM-QM Bridge

Begin setting up another mail server's QM-QM Bridge phone book entry by clicking the **Create** button.



Setting up Another Mail Server's QM-QM Bridge Phone Book Entry

Enter the other site's **Name** and its associated **Phone number**. In the bottom of this window, some account preferences can be set that work in conjunction with the QM-QM Bridge's Connect Times window settings.

When sending to this site... is used to determine under what conditions mail is downloaded from the remote site. If the first radio button is selected, all the mail stored for this account is accepted whenever a QuickMail-to-QuickMail connection has been made. If the second radio button is selected, all the mail stored for this account is accepted only if a QuickMail-to-QuickMail connection has been triggered by the chosen criteria. If the last radio button is selected, your mail server will not accept mail from the remote site at all.

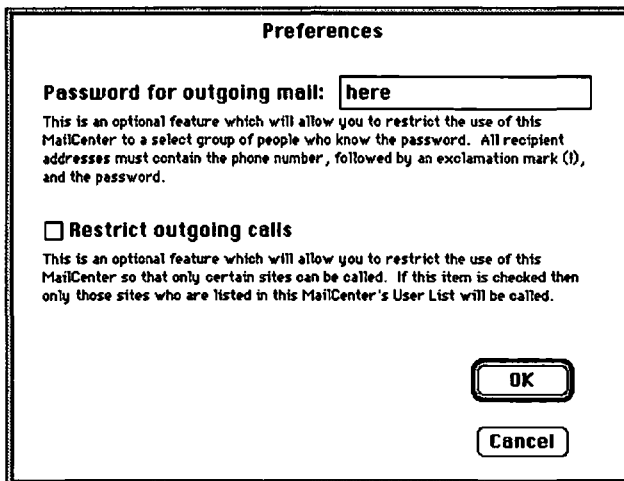
The **If no messages are sent...** options are used to determine under what conditions your mail server will call this site with the sole intent of downloading mail from the other mail server. If the **Scheduled connections** box is selected, QuickMail will call this mail server whenever it calls out to any other site as regularly scheduled. If the **Manual connections** box is selected, when you force QuickMail to call out, by selecting the MailCenter icon in QM Administrator's window and choosing **Connect** from the **File** menu, it will call out to this mail server, too, even if it does not have mail waiting for it.

Click **Done** when finished. The new account will be added to the phone book entries list in the Configure QM-QM Bridge window.

The screenshot shows a window titled "QM-QM Bridge™ for QuickMail™". Below the title is a brief instruction: "QM-QM Bridge will exchange mail with multiple QuickMail networks. Use the active menus above for setting up additional configuration information." The main area is labeled "Phone Book:" and contains a list with one entry: "CE Software" with the phone number "1,515-555-1234". Below the list are three buttons: "Create", "Change", and "Remove". At the bottom left, there is a "Log:" label followed by a text box containing "All Activity". To the right of the log box are two buttons: "OK" and "Cancel".

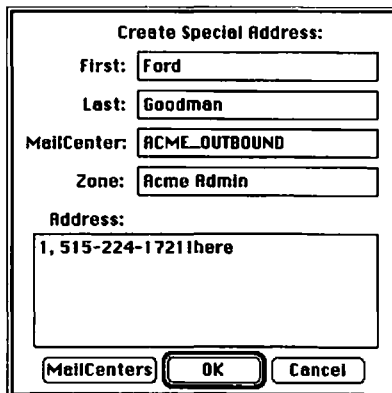
Newly Created QM-QM Bridge Phone Book Entry

Before the QM-QM Bridge MailCenter is ready to go, there are some preferences that should be set. To do this, first choose Preferences from the QM-QM menu. Here you may control who can use this bridge connection.



The QM-QM Bridge Preferences Window

Type in a password if you want only certain users to be able to send mail through this MailCenter. An unauthorized user trying to send mail to someone at this site by using his or her workstation's special address feature would still be out of luck, even if he or she knew the right phone number. He or she would also need to type the password into the Address field.



Using a Password for Outgoing Mail

Select the **Restrict outgoing calls** checkbox if you do not want this MailCenter to be used to dial out to any other sites but those listed in the MailCenter's phone book entry list.

To finish setting the QM-QM Bridge MailCenter to handle mail, enter the correct modem information in the Connection Settings window.

To finish setting up the QM-QM Bridge MailCenter for connecting with other MailCenters, choose **Setup** from the **QM-QM** menu. Here you tell your mail server and other sites about *your* office's telecommunications environment.

Setup

This information should describe YOUR site and this MailCenter's phone system. When two sites connect, some of this information is exchanged between sites.

Company Name:

Domestic long distance dialing prefix:

International long distance dialing prefix:

Country code:

Area code:

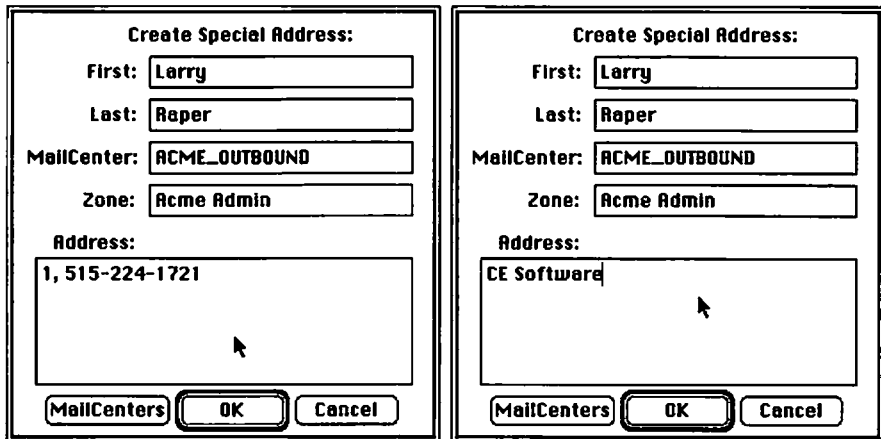
Local phone number:

Prefix for outgoing calls (optional):

Suffix for outgoing calls (optional):

The QM-QM Bridge Setup Window

QM-QM Bridge MailCenter user accounts are created in the same manner as other MailCenters: directly by using **Create** from QM Administrator's **User** menu or by importing the information from a database. The QM-QM Bridge MailCenter user list differs from those of other MailCenters, however, in that it contains recipients' QM-QM Bridge account names or phone numbers, not typical QuickMail addresses.



Address Boxes for Recipient of QM-QM Bridge Mail

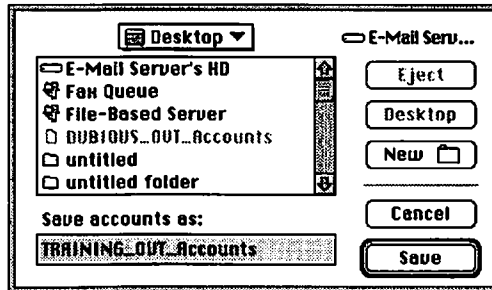
It further differs in that this particular bridge also uses a phone book that acts similar to but is different from the MailCenter user list. The QM-QM Bridge is the only MailCenter shipped with QuickMail that makes use of both a user list and a phone book. These are not the same.

The QM-QM Bridge user list, like other MailCenters, contains specific user accounts. Since there can be hundreds of users at each mail server that the QM-QM Bridge contacts, however, it is impractical to create user accounts for all of them. Your user list could become huge! Instead, create user accounts for only those with whom you are in regular contact.

The QM-QM Bridge phone book automatically keeps track of all other mail servers with which your mail server comes in contact. When a connection is made, the QM-QM Bridge exchanges information about phone numbers and MailCenter names, and both mail servers maintain this list.

As with user lists, the easiest way to exchange large amounts of such information with another site is to export it into a database file and send it to that site's Administrator. It may then be imported into the phone book database.

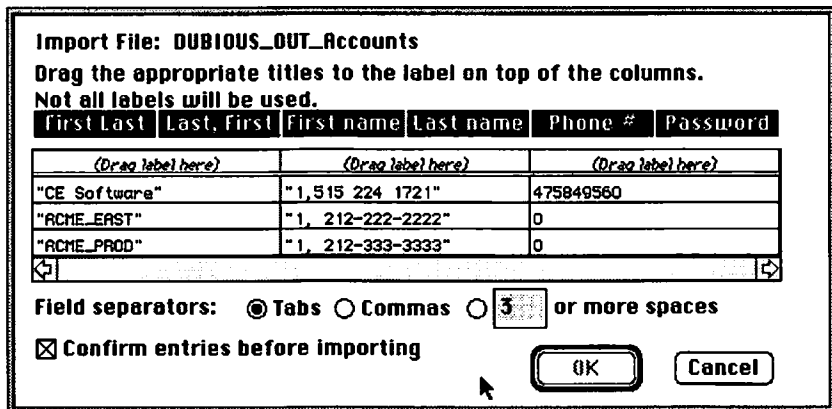
To export the phone book from a QM-QM Bridge MailCenter, open it in QM Administrator, and select **Export Phone Book** from the **QM-QM** menu. The standard Open File dialog box appears, prompting you to choose a location for the exported records. Note that the predetermined name for *all* exported phone book lists is the list's MailCenter name followed by "Accounts."



Exporting Phone Books from a QM-QM Bridge MailCenter

When the file is exported, it is exported with tabs separating the fields and with quotes surrounding the fields.

To import phone book lists into a QM-QM Bridge MailCenter, open the desired MailCenter in QM Administrator and select **Import Phone Book** from the **QM-QM** menu. The standard Open File dialog box appears, prompting you to select the desired file for importing. After the file has been selected, the Import File window appears with the names appearing in rows and columns below.



The Import File Window

Study the way the names are formatted, and then drag the labels **First Last** and **Phone #** over the title position for that column.

Last, First	First name	Last name	Phone #	Password
(Drag label here)	(Drag label here)	(Drag label here)		
"CE Software"	"1, 515 224 1721"		475849560	
"ACME_EAST"	"1, 212-222-2222"		0	
"ACME_PROD"	"1, 212-333-3333"		0	

Selecting a Title

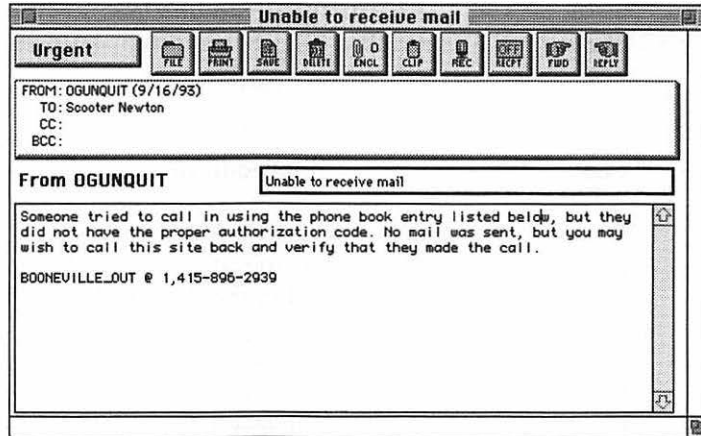
The title for the column has been selected (above) and dragged into place (below).

First Last	(Drag label here)	(Drag label here)
"CE Software"	"1, 515 224 1721"	475849560
"ACME_EAST"	"1, 212-222-2222"	0
"ACME_PROD"	"1, 212-333-3333"	0

Dragging the Title into Place

This sets the order of the names as they are imported into the MailCenter phone book. If you checked the **Confirm entries...** box, you are prompted for verification on each record that is imported into the phone book.

So, what will stop some QuickMail site with which you do not want to have contact from learning your mail server's phone number and dialing in? The QM-QM Bridge associates a coded password with each site with which it has contact. When a connection is made, both sites must agree on this code. They may then exchange mail. Should an unknown site dial in, it will not know the code. Your mail server will allow it to upload messages to your site, but will not download any mail to it. The QM-QM Bridge MailCenter custodian will then receive a notification such as the one in the following screen shot.



Security Alert from QM-QM Bridge MailCenter

The other site will also receive a notification informing its custodian that your MailCenter is no longer authorized to receive mail from it.

Sometimes this feature can cause a break in the exchanges of an authorized connection. This can happen, for instance, if one of the site's phone numbers has changed. Should you receive the previous notification from a site with which you *do* want to receive mail, locate that site's phone book entry in the QM-QM Bridge Configuration window, make note of its phone number, and delete it using the **Remove...** button. Next, create a message asking the administrator at the other site to remove *your* mail server's phone book entry from *his* mail server's phone book and reply to your message. To send your message, use the **Special Address** button in the desk accessory along with the phone number you copied down previously. The other administrator's reply will renew the authorization codes and permit mail exchanges to resume.

QM-Script Gateway

The QM-Script Gateway allows your users to send and receive mail through such information service providers as CompuServe, MCI Mail, Sprint, and GENie using QuickMail user software. A company can enable its QuickMail users to send mail to users of any of these information services by opening an account with the service and setting up a QM-Script Gateway MailCenter to connect. All that your users need to know is the account ID the service has assigned the user with

which they wish to communicate. Information service users need only know the account ID the service has assigned your company and the name of the employee with which they wish to communicate.

There can be up to 32 QM-Script MailCenters on a server. Each QM-Script MailCenter can contain hundreds of addresses.

To set up this service, select **New MailCenter** from QM Administrator's **File** menu. When the first window appears, enter the desired name in the **MailCenter** field and any desired password in the **Password** field, and then select **QM-Script** from the MailCenter **Type** pop-up menu.

Click **Configure** in the next window. A window appears containing instructions for how QuickMail should interact with the information service account.

The screenshot shows a window titled "QM-Script Gateway for QuickMail™". Below the title is a small text box: "QM-Script Gateway uses scripts to exchange mail with other mail systems. Use the active menus above for setting up additional configuration information." Below this are two input fields: "Account #:" and "Password:". Underneath is a section for "Multiple Enclosure Uploading:" with a dropdown menu set to "Separate Files". Below that is a "Log:" dropdown menu set to "All Activity". The "Script:" field is set to "(None)". At the bottom left are three checkboxes: "Display connection" (unchecked), "Include QuickMail headers" (unchecked), and "Strip 8th (high) bit" (checked). At the bottom right are "OK" and "Cancel" buttons.

Configuration Window for QM-Script Gateway

In the bottom of this window, several QM-Script Gateway preferences can be set.

Multiple Enclosure Uploading is used to determine how QuickMail messages containing more than one file enclosure are handled. If you choose **Separate Files**, each enclosed file will be uploaded to the service individually. If you decide to choose **PackIt I**, all the files will be combined in a Macintosh transfer format to be uploaded together. Many compression utilities can unpack this format.

Script shows the file from which the gateway is obtaining its instructions on how to dial into the information service. It might be one of those included with the QuickMail software or one that you have written yourself.

Display connection causes a window to be generated whenever QM Administrator logs into the information service. The flow of communications and the contents of messages are displayed in the window.

```
User ID: 76 170, 431/GO:EMAIL
76 170,431/GO:EMAIL
Password:

One moment please ...

% Checking Membership Information...
```

Connection Displayed (CompuServe)

If you are concerned about privacy or security, do not select this option.

Include QuickMail headers causes the QuickMail message header information—as in Date, From, and To—to be added to the information service’s native message header information.

```
2
Date: 30-Aug-93 22:56 PDT
From: Scooter Newton [724113,274]
Reply to: Artist
Date      8/30/93
Subject   RE:Artist
From      Scooter Newton
To        Dorian Cougias on CompuServe
```

QuickMail Header Included in CompuServe Message

Strip 8th (high) bit pertains to the number of bits that compose a character. Using the Apple Modem Tool, the mail server can send and receive whatever the computer system at the other end of the modem requires. The default is 8-bit,

which is not required by some systems. Using this checkbox removes this extra bit. If you are planning to transfer files using the XMODEM protocol, do not check this box.

After you have decided on these options, type in the **Account #**, or ID, and the **Password** at the top of the window.

QM-Script Gateway for QuickMail™

QM-Script Gateway uses scripts to exchange mail with other mail systems. Use the active menus above for setting up additional configuration information.

Account #: 76170,431

Password: ●●●●●●●●

Newly Created QM-Script (CompuServe) Account

To finish setting up the QM-Script Gateway MailCenter, enter the information service's phone number in the Connection Settings window.

QM-Script Gateway MailCenter user accounts are created in the same manner as are other MailCenters: directly by using **Create** from QM Administrator's **User** menu or by importing the information from a database.

The QM-Script Gateway MailCenter user list differs from those of other MailCenters, however, in that it contains information service mail *recipients*, not originators. The address fields contain the information service account IDs for people to whom you are sending e-mail.

Create Special Address:

First:

Last:

MailCenter:

Zone:

Address:

Address Box for Recipient of QM-Script (CompuServe) Gateway Mail

Do not type exactly the same account names of users from one of your network MailCenters into a QM-Script Gateway MailCenter. This name duplication could cause mail to be misrouted. Notice that in the earlier illustration, the account designation is for **First Name** "Dorian Cougias" and **Last Name** "on CompuServe." That is because this user is already a user in an online MailCenter as **First Name** "Dorian" and **Last Name** "Cougias."

Normal FILE PRINT

FROM: Scooter Newton
 TO: Dorian Cougias on CompuServe
 CC:
 BCC:

2
 Date: 30-Aug-93 22:56 PDT
 From: Scooter Newton (724113,274)
 Reply to: Artist

Date 8/30/93
 Subject RE:Artist
 From Scooter Newton
 To Dorian Cougias on CompuServe

Address in QuickMail Form vs. Address in CompuServe Message

When a message is received from an information service, QM Administrator will examine the file for the name of your onsite recipient. It will attempt to match it with a user in the NameServer. If no match is made there, QM Administrator will search all the user lists in the MailCenters it is set to scan. If QM Administrator can't match a user with a message, the mail will be routed to the QM-Script MailCenter's custodian.

No one can dial into your network through the QM-Script Gateway, since it is designed to dial out only. This greatly reduces any security risk.

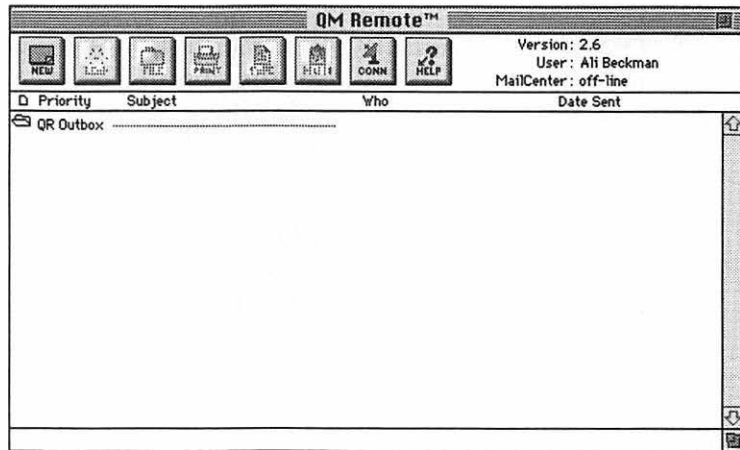
Supporting Remote Connections

Users at home or on the road need not miss out on important e-mail conversations if their remote computers are equipped with modems. Many times, however, you as QuickMail administrator will need to “walk them through” the dial-in process.

Macintosh users connecting to the office network using AppleTalk Remote Access (ARA) software or a similar third-party product, such as Shiva’s NetModem with Dial-in, are the easiest remote users to support. They can use the same QuickMail desk accessory software and network services as any user on the premises. Depending upon the speed of the modems over which they are communicating, users will scarcely know the difference.


To use QuickMail over ARA, the remote user need only install the usual QuickMail workstation software on the remote computer. Tell users that, whenever possible, they should compose outgoing mail in the QuickMail desk accessory prior to logging in with ARA. This can save on phone charges and more quickly free up the mail server line for other uses. Tell users to remember that mail they file is *saved to the hard drive of the computer on which they are working*. Should they wish to call up the same mail from an office workstation, they should not file that mail on their remote computers.

QM Remote is a handy utility users can employ from any Macintosh equipped with a modem to dial into the QuickMail server and retrieve or send messages. The interface closely resembles that of the actual QuickMail desk accessory.



Main QM Remote Window

Users might not realize that they can read their mail after having disconnected from the mail server, thus saving phone charges and freeing up the mail server line for other uses. To do this, users select all the messages they wish to read by holding down the **Shift** key and clicking the **Read** button. Any messages that have not already been transferred to the remote computer will be downloaded to them.



Receiving Mail Remotely with QM Remote

After all the messages have stacked up on the screen, the user may log off. QM Remote can also be set to log off automatically once messages are retrieved. Tell users to compose outgoing mail prior to logging in whenever possible.

Finally, warn users that while they are connected to the QuickMail server with QM Remote, they will not receive any new mail that is sent to them. The only time QM Remote checks the status of their mail is when they first connect. If users want QM Remote to update its list of mail waiting for them, they need to choose **Refresh** from the **Mail** menu. Otherwise, they might miss important messages. Unlike with the desk accessory, there will be no automatic notification.

It is also possible to dial in to a QuickMail server from either a Macintosh or a PC using plain-text-based communication software. Although there are many such programs, the interface will vary little.

To dial in, users should type the following command string, replacing the *Xs* with the actual phone number:

ATDT XXX-XXXX

Depending on the program, this might be accomplished automatically. If the connection proceeds normally, they will see this message:

CONNECT 9600

QuickMail(tm) Copyright 1988-93 CE Software, Inc.

The users must type in a name and password as they are prompted for each.

Name:Ali Beckman

Password:*****

Next they will be welcomed to the MailCenter and notified of the total number of messages waiting for them, as well as the number of unread messages.

Welcome to QuickMail, Ali Beckman

Welcome to ACME_TRAINING

You have 30 messages waiting (3 unread)

L(ist) R(ead) W(rite) D(ownload) U(pload) B(ye) > L

CHAPTER 2: PUBLISHING E-MAIL WITH LISTSTAR

The World Wide Web and products such as StarNine's WebSTAR have captured the public imagination and are being credited with spurring the Internet publishing revolution. Credit would be better given to the academic community and their e-mail-based *list servers*, since they really propelled the Internet speech movement. Realizing this, StarNine makes a Macintosh-based list server product as well called ListSTAR. We take a look at its various flavors in this chapter.

ListSTAR performs two major functions. It can act as a traditional list server, processing the e-mail from many Internet users for the list server's subscribers. It can also be used as an automated e-mail response system (dubbed *e-mail-on-demand*), sending out information to users who specifically request it. There are several ways it can do these things. Each method and function is referred to as a *service*.

ListSTAR can be used in conjunction with QuickMail, StarNine Mail (formerly Microsoft Mail), Post Office Protocol (POP), and Simple Mail Transfer Protocol (SMTP) servers. Although we have chosen to speak of QuickMail in our discussion of LAN-based e-mail systems, we will concentrate on the SMTP version of ListSTAR as it ties in better with the other sections in this book. The same basic rules apply between the versions.

So, what specifically is a list server? It is simply a machine that processes e-mail-based discussions by allowing an organization to send identical mail messages to

all members of a given *mailing list*. The users can “sign up” for, or *subscribe* to, the service automatically by submitting an e-mail message containing a certain code. When used in conjunction with the Internet, these users can be located anywhere in the world.

While they function similarly, this is not really the same as an e-mail-on-demand system. E-mail-on-demand systems produce an automatic e-mail response to a specific user request, which can include both text and binary information. This is more like a telephone-based fax-back system.

ListSTAR also permits the use of a *direct mailing list*, a synthesis of the two. Here you can distribute a given e-mail to a list of people you explicitly create. Ever get junk e-mail? This is how it was sent to you.

Since I do not want to kill any more trees than are necessary to make the point, I will go into just one implementation of ListSTAR here: e-mail-on-demand. This will touch on enough of the important concepts that you should feel comfortable in understanding ListSTAR’s other functions as well.

INSTALLING LISTSTAR FOR SMTP

We recommend that you install ListSTAR for SMTP on a Macintosh equipped with at least a 68030 processor and a 200-MB hard drive. You also need System 7.x, MacTCP, AppleTalk (or Open Transport), AppleScript 1.1 or later with the scriptable Finder, and 1,200K of free RAM. I like to put a Dilbert doll alongside the CPU, but that is optional.

Also, make sure that you are connected to an IP network and that TCP/IP is configured correctly. See our sections elsewhere in this book if you need help.

The above is true for ListSTAR for SMTP and ListSTAR for POP. You can dispense with the IP connection in favor of just AppleTalk if you use ListSTAR for QuickMail. Ditto for ListSTAR for StarNine Mail. You also need an e-mail account on these respective systems for each ListSTAR service you implement to receive e-mail, so plan ahead.

You do not have to run ListSTAR on the same Macintosh that runs the QuickMail or StarNine Mail server software. In the QuickMail version, ListSTAR only needs QuickMessenger to log in to a MailCenter. In the StarNine Mail version, it uses the client software to log in.


You could run different versions of this software—say, both ListSTAR for SMTP and ListSTAR for POP—on the same Macintosh. You could also stick your tongue in a wall socket, but I would not recommend that either. StarNine has heard of cases in which the user interface becomes corrupted and it is difficult for AppleEvents to differentiate between the two applications. As for the wall socket . . . oh right, like you never tried that!

By the way, the difference between ListSTAR for SMTP and ListSTAR for POP is important. The former operates as its own e-mail server, although you do not need to go in and set up user accounts. The latter can dial up and log into a POP3 server that holds e-mail accounts for each of ListSTAR's services.

Do not run ListSTAR for SMTP on the same Macintosh on which you are running a Mail*Link SMTP gateway or the Apple Internet Mail Server. They both listen for SMTP messages on port 25. It is a crap shoot which one will receive a given message. It is fine to run ListSTAR for POP on the same Macintosh on

which you are running the Apple Internet Mail Server. ListSTAR can log into this POP server as long as you make sure **Ping before opening connection** is not selected in ListSTAR's TCP/IP preferences. This prevents ListSTAR from pinging its own host.

When you have chosen your destination Macintosh, you can install ListSTAR using the old familiar installer script. Be aware of this: If you attempt to install ListSTAR on a hard drive other than the startup disk, a System Folder will be created on that disk into which the installer copies the MacTCP 2.0.6 control panel and the Thread Manager extension. Avoid this by deselecting MacTCP and Thread Manager during a Custom install. Do this as well if your Macintosh is equipped with Open Transport.



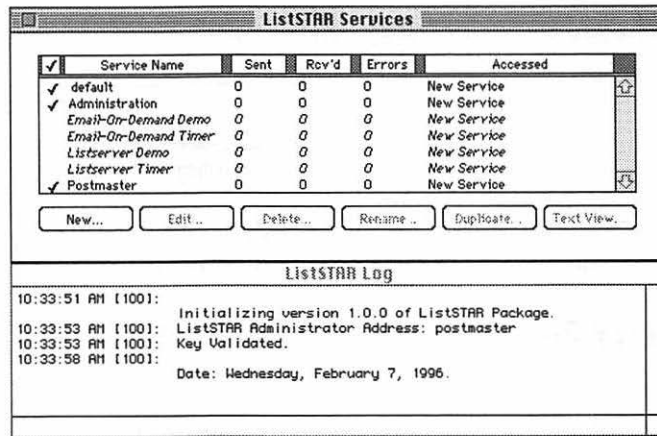
ListSTAR for SMTP Folder

Once the installer has finished, you will have a folder containing the ListSTAR Server application and five subfolders. The ADMIN folder contains preferences, e-mail account information, logs, and other administrative software. The Message Spool/SMTP folder is where ListSTAR stores incoming and outgoing messages. The Services folder contains a subfolder and preferences for every ListSTAR function you set up. The other folders contain examples and technical information. The loss of any of these components will result in a major headache for you once you have taken the time to configure ListSTAR. Back up the work immediately after you are finished with your configurations, and at least daily thereafter.

Before you actually launch ListSTAR, you should verify your network connection. Because we are working with ListSTAR for SMTP here, which can act as its own mail host, we need a little more than just a correctly configured MacTCP control panel. First, give ListSTAR a host name that is unique to your domain, such as *listserv.netfrontiers.com*. Second, you will want the name of another SMTP mail host, such as *smtp.netfrontiers.com*. It is to this machine that ListSTAR can

forward mail that it has been unable to deliver itself or, if you prefer, all of its mail. Since this will involve a change to either your network's domain name server or to the tables of your Internet service provider's (ISP's) name server, take care of this before you launch ListSTAR.

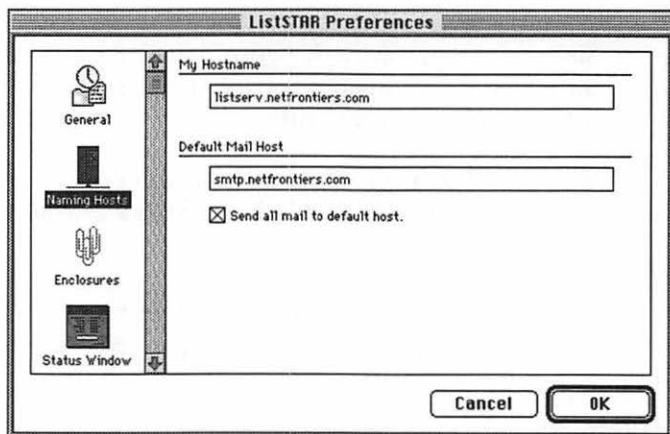
When you have these things ready, launch ListSTAR and type in its serial number. You are greeted with one window containing demonstration services and another containing the log.



ListSTAR Windows

Naming Host Preferences

Select **General Preferences** from the **Windows** menu. This generates the Preferences window. Here select the **Naming Hosts** icon. In the window that appears, add the list server's domain name to the **My Hostname** field and the mail host to which you will be forwarding in the **Default Mail Host** field. If you want to off-load the work of all mail delivery onto this other mail host, select the **Send all mail to default host** checkbox.

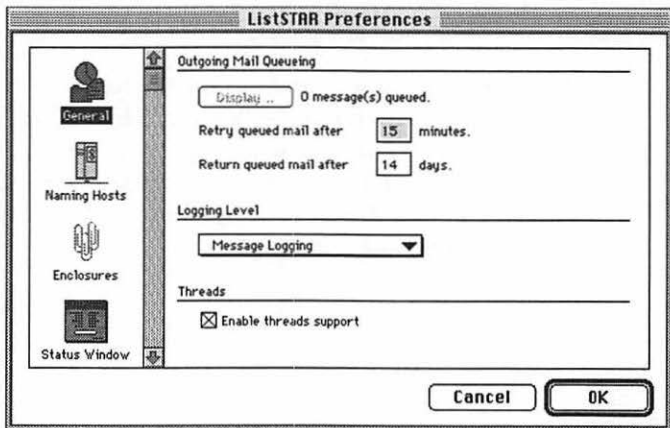


Setting up Host Names

Now you are more or less ready to put this server into action. While we are at it, however, a review of ListSTAR's other preferences is in order.

General Preferences

This category contains instructions for how ListSTAR queues mail, how much detail it goes into when logging sessions, and whether or not it will run in a threaded environment.



General Preferences Window

The choices in the **Outgoing Mail Queuing** pane pertain to how long ListSTAR attempts to forward undelivered mail; this pane is fairly clear. The mail might not go through the first time because the host ListSTAR to which it is sending is down. If you know the host server will be down regularly for periods longer than those set in the defaults, change these preferences.

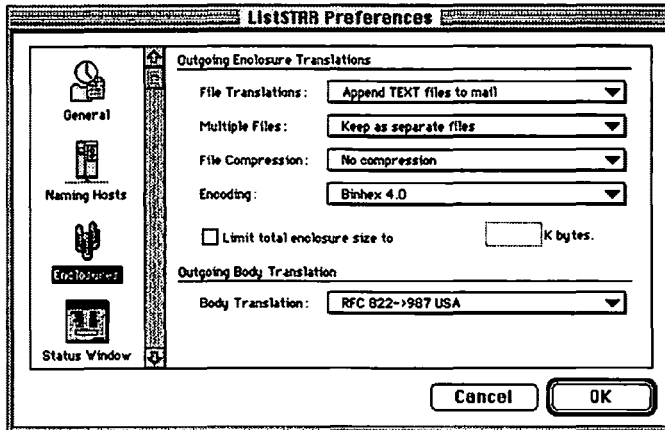
Below the **Outgoing Mail Queuing** pane is the **Logging Level** pop-up menu, in which you can decide how much information you want ListSTAR to gather about each session. The chances are good that you will never need more than the default **Message Logging**. If you do have connection problems, however, selecting higher levels will provide you or StarNine with invaluable information.

Below the **Logging Level** pop-up menu is the **Threads** section. You will almost certainly want to keep this checkbox selected. When running under System 7.5 and later, or earlier System 7 with the Thread Manager Extension installed, ListSTAR can open a number of simultaneous connections. Without it, ListSTAR can open more than one connection simultaneously, but it will only be able to move e-mail over them one at a time. We will discuss this further on.

Enclosures Preferences

ListSTAR can send file enclosures in addition to plain text e-mail, although those enclosures need to be translated into a *binary* format—in other words, computer-speak. Because the formats of files created by various applications can be hard to read by remote mail systems, it is conventional to *encode* file enclosures as more universal formats. By default, ListSTAR automatically translates outgoing enclosures of any size using the *BinHex* format.

You can change this and other settings here. We recommend that you stick with the defaults for Internet use. If you are using ListSTAR internally, however, some modifications might be helpful. Read the *ListSTAR Administrator's Guide* for complete information. Be aware that these settings affect *all* mail leaving ListSTAR.



Enclosures Preferences Window

The **File Translations** pop-up menu determines how ListSTAR will deal with straight text file enclosures in outgoing mail. The default of **Append TEXT files to mail** directs ListSTAR to add the contents to the bottom of the e-mail message.

The default **Keep as separate files** in the **Multiple Files** pop-up menu directs ListSTAR not to pack different file enclosures together. Suppose you had one text file and one program (binary) file you wanted to send to a mailing list. ListSTAR would handle the text file based on the criteria set elsewhere in this window under **File Translations** and would handle the binary file using criteria established under **Encoding**. Two enclosures would leave the list server.

The default is to have nothing in the **File Compression** pop-up menu. If you are distributing files to *Macintosh users only*, however, you could save them transfer time by choosing to enable Aladdin's StuffIt compression. Of course they will need at least the free StuffIt Expander application to decompress the files. (You should have StuffIt. It compresses and decompresses almost everything.)

While there are many options that you can choose in the **Encoding** pop-up menu, there are only three that most of us need to consider. If you choose **AppleSingle—UUencode**, ListSTAR translates the Macintosh file into a file format that permits both its data and resource forks to be recreated on the receiving system, then encodes this resulting file using the UNIX standard *UUencode* format.

The data fork of a Macintosh file contains kernel information, such as text. The resource fork contains the format and icon information used by the MacOS to

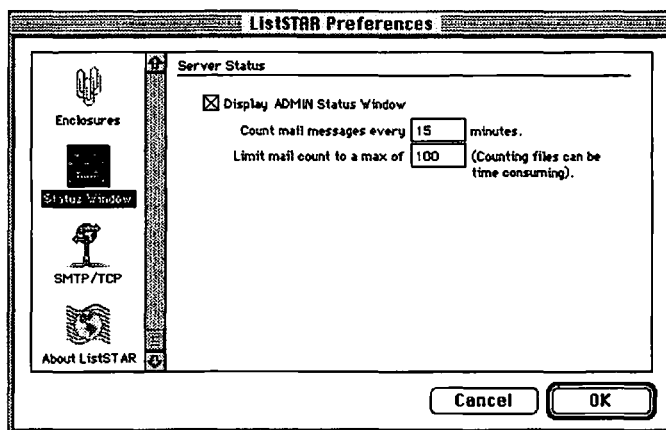
make it so friendly. PCs and other machines only use a data fork. This method is the one to use when you have Macintosh files that are going to God-only-knows-what kind of system.

The default choice is **Bihex 4.0**, which is used to turn binary data into ASCII text and back again. This method is the one to use when you have Macintosh files that you know are going to other Macintosh users.

Choosing Data fork Only—UUencode directs ListSTAR to apply the UNIX standard UUencode algorithm on the data fork of a Macintosh file, but also directs it to chuck the resource fork. If the ultimate destination of these enclosures is other Macintosh users, this will deprive them of resource information needed by the MacOS. If, however, the ultimate destination is users of older PCs or UNIX machines, this is a good choice.

You can limit the size of outgoing messages by setting the **Limit total enclosure size to** option, but we do not see much reason to do this with a list server. You can determine how the Macintosh's 8-bit characters will be translated for systems using 7-bit character sets in the **Body Translation** pop-up menu. We suggest you stick with the default of **RFC 822->987 USA**, an accepted standard that is observed even on really old e-mail systems.


Status Window Preferences



Status Window Preferences

Selecting **Display ADMIN Status Window** here generates a handy monitor of the application's activity. There's nothing like kicking back in front of a warm list server and counting the bits as they go by, right? The numbers you set here determine how much time ListSTAR will spend counting messages to show you that it has been working hard. Every 15 minutes is plenty often, unless you are doing troubleshooting. The gateway will stop counting messages after it reaches the number you specify in the **Limit mail count to a max of** field. Make this number smaller than 100 to conserve processing resources. Really, do you need to enable this at all? I doubt it.

SMTP/TCP Preferences



SMTP/TCP Preferences

Here you select how ListSTAR will work with TCP/IP and the threaded environment we mentioned earlier.

If you go with the default of 5 in the **# of Incoming Connections** pop-up menu, you will have a total of eleven concurrent processes, or *threads*: five listen for mail, four send mail, one is the *main* thread, and one more can also act as a sender. We suggest that you leave this setting alone.

The **Tcp Command Timeout** pop-up menu controls how long ListSTAR waits for a response from a mail host before assuming it doesn't exist. The default is 127 seconds, but we have had better luck with a setting of 255.

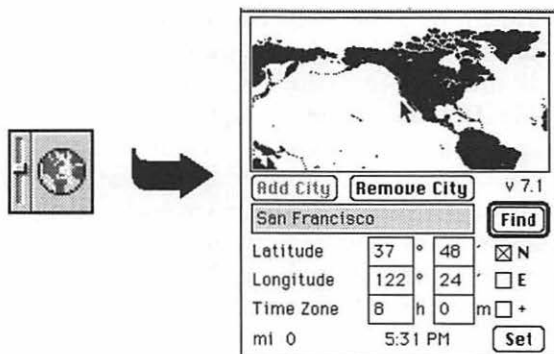
Finally, we come to the **NameResolver** pop-up menu. When ListSTAR wants to find the IP address that goes with a given host name, it can ask MacTCP (or TCP/IP under Open Transport) or its own *StarNine Resolver* for the information. If **MacTCP only** is the only option you can set here, then the mail server's MacTCP control panel does not know about any DNS. It must fall back on a Hosts file you create in the System Folder. If you have other choices, then MacTCP does know about DNS and can ask the servers for resolution. Unfortunately, MacTCP only knows a few DNS records, so it cannot gather as much routing information as it might. That is where the StarNine Resolver comes in. It knows more—specifically, it knows about MX records.

In my opinion, the best way to set this preference is to choose the **StarNine then MacTCP** option. It provides the best chance for making a connection. Of course, if you set ListSTAR to forward its mail to another mail host in the **Naming Hosts** pane, MacTCP's resolver is all that is necessary.

If you select **Ping before opening connection**, ListSTAR attempts to confirm that a remote host is available before attempting to communicate with it. It does this by bouncing a packet off the mail host like a destroyer bounces sound off the hull of a lurking submarine . . . sorry, I've been playing too much Harpoon again. The trouble with using this setting is that even if a mail host is there, your e-mail might not go through if that mail host is so busy that it cannot respond to a ping query. We have found that America Online's servers have this characteristic. If you start receiving administrative e-mail messages warning you that a given host cannot be found on the Internet when you know darn well it *is* there, deselect this option.

And Another Thing . . .

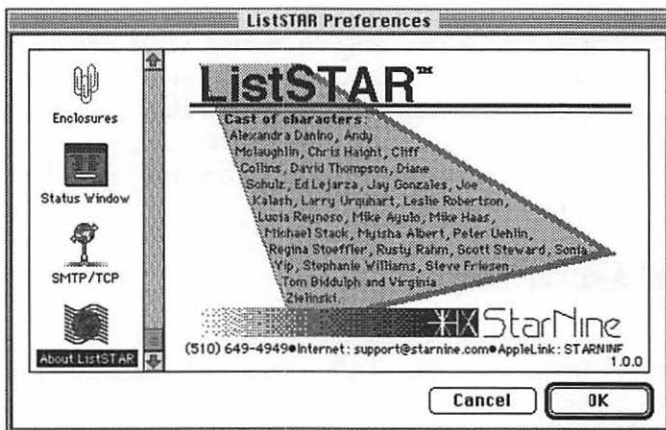
To make sure ListSTAR generates the correct time in its RFC-compliant headers, you must tell it what time zone it is in relative to Greenwich Mean Time. This is easy on a Macintosh. Double-click the Map control panel, found in the Control Panels folder of the System Folder.



Map Control Panel

Click on the map to show the Macintosh where you live. (No, there is no map for Mars.) If there is a big city nearby, it will appear. Click the **Set** button to save this setting, and then close the control panel.

I am going to show one more little preference pane—the **About ListSTAR** pane.



Technical Support Numbers

Here you can find StarNine's technical support number and e-mail addresses. StarNine's technical support is great, so do not be afraid to avail yourself of it. While you are at it tell Tom and the other folks how great you think ListSTAR, WebSTAR, Mail*Link, and their other products are. Dorian and I have known these folks since before we formed Network Frontiers, and they have always been great to work with.

SETTING UP ADMINISTRATIVE SERVICES

At the simplest level, ListSTAR's services are composed of individual e-mail accounts that have sets of *rules* associated with them. For a good example of this, take a look at the list server for a technical support discussion on StarNine's own products. ListSTAR sends and receives mail via an e-mail account dedicated to that service. A series of rules is applied to that mail that does such things as automatically adding and deleting subscribers to the mailing list, forwarding mail to the subscribers, and compiling a digest of the list discussion. To subscribe to this list, send an e-mail to *liststar-request@starnine.com* with the word "subscribe" in the Subject field.

ListSTAR services can be one of two types: *Mailer* or *Timer*. Mailer services, such as list services and fax-on-demand, react to e-mail, as in "when an e-mail comes in that says *this*, do *that*." Timer services, such as list server digest distribution and log forwarding, react to the clock, as in "when it is *that* time, do *this*."

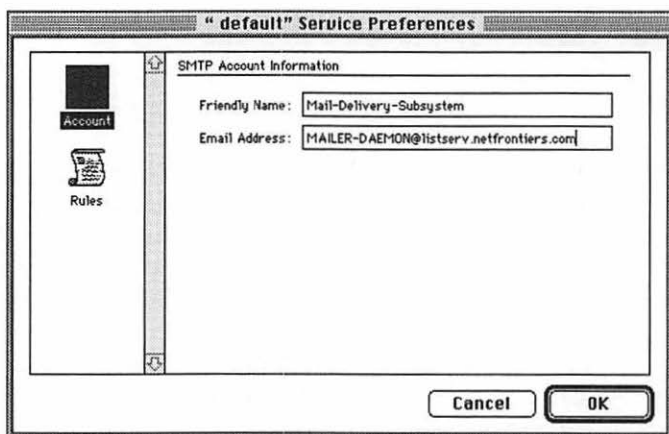
Since a Mailer service must both send and receive e-mail, it requires its own e-mail account. Since a Timer service needs to send e-mail only, it does not require its own e-mail account.

When you first launch ListSTAR, several administrative services are already active. These are in plain text. In addition, several demonstration services are inactive. These are in italics. The active services are Timer and Mailer services set up to automate ListSTAR administrative functions, so work with them first.

Default Service

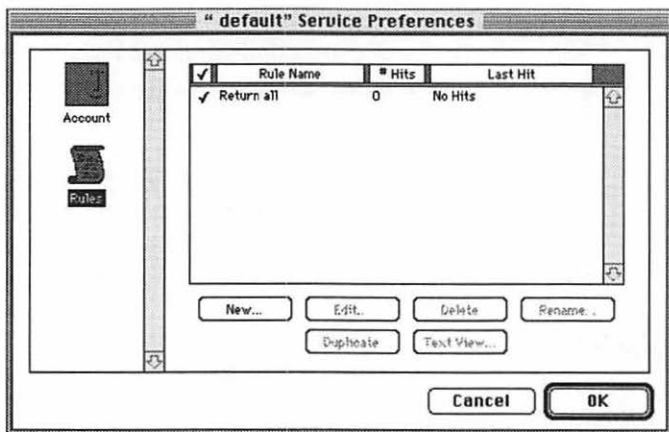
It is the job of the "default service" to handle all mail not addressed to any other service. Remember the Elvis song, "Return to sender; address unknown"? That is what it sings.

To configure the default service, double-click on its line item. In the Service Preferences window you will find panes for both the service's e-mail account information and associated rules. In the Account pane, edit the Email Address field to reflect your site, such as *MAILER-DAEMON@listserv.netfrontiers.com*.



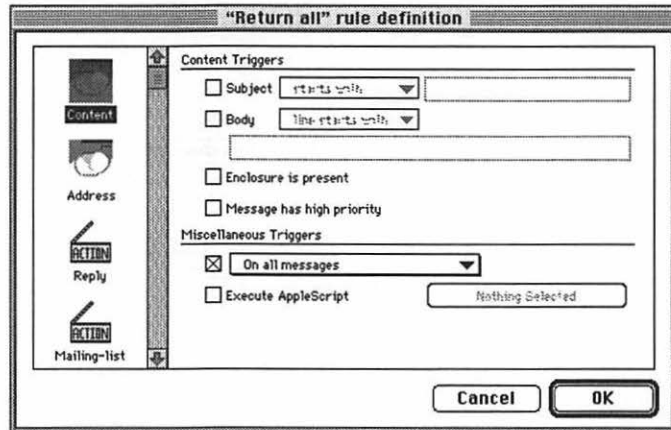
Default Service Preferences Account Pane

Now take a look at the series of rules that pertain to how misaddressed e-mail will be handled in the **Rules** pane.



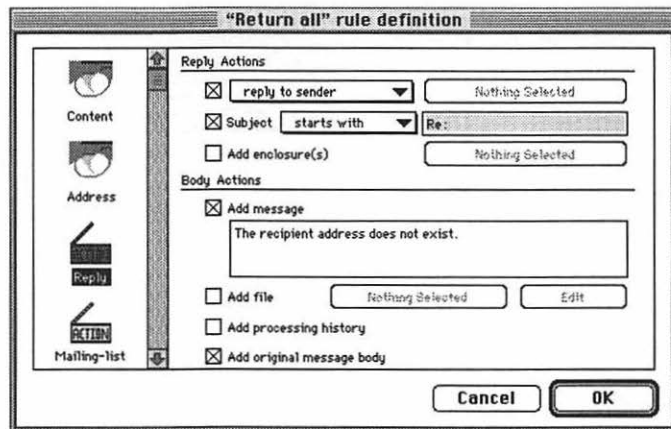
Default Service Preferences Rules Pane

In this case, there is one rule: "Return all." Double-click its line item to examine it.



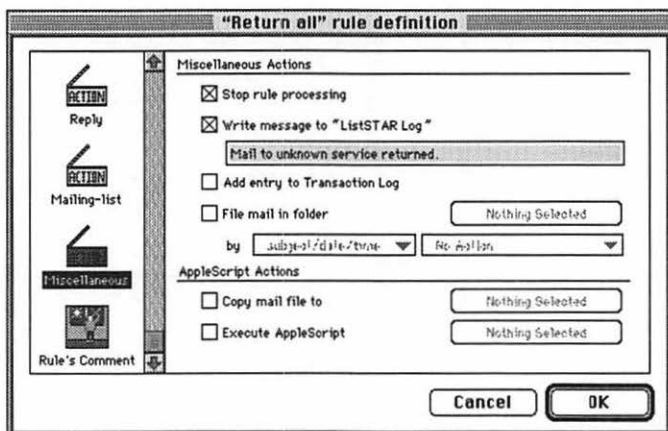
Default Service Content and Miscellaneous Triggers

It is all very logical. Rules are nothing more than *if-then* instructions: *If* the light is red, *then* stop (a logic that escapes most San Franciscans). In this case, it is *if* you receive a message (or rather, **On all messages**) . . .



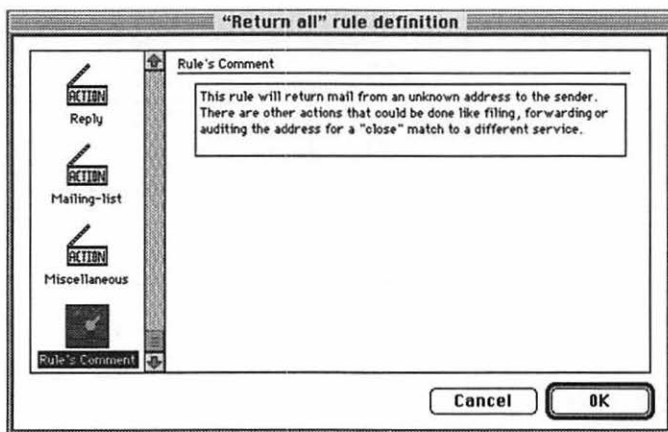
Default Service Reply and Body Actions

. . . *then* reply to the sender by appending the subject with Re:, adding the message "The recipient address does not exist," and adding the original message body to the reply.



Default Service Miscellaneous Actions

The instructions go on to say: Once this is done, stop processing the e-mail and record your actions in the “ListSTAR Log” by saying “Mail to unknown service returned.”



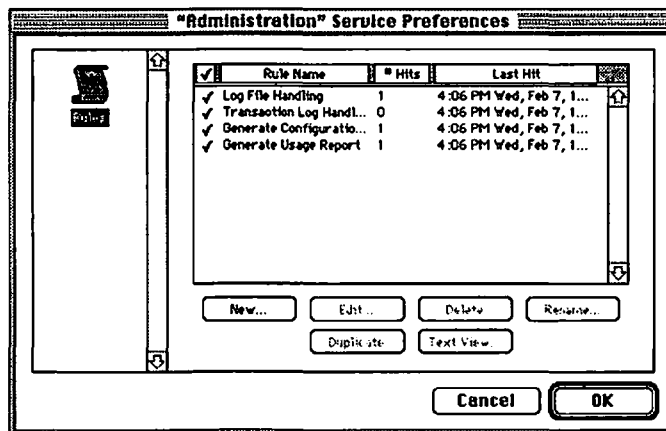
Default Service Rule's Comment

That is all that is really going on here, except that you must use the pop-up boxes and input fields to communicate with the computer and can't just talk to it (yet). To communicate with other humans who follow in your wake, explain your directions in the Rule's Comment pane.

By the way, if you change the name of this service, include a blank space before the new name. It needs to be at the top of the list to execute at the right time.

Administration Service

The Administration service is an example of a service that does not handle incoming e-mail. Thus, it does not require its own e-mail account and so has no Account pane. What it will do is file ListSTAR'S various log files within the subfolders of the Administration folder within the Services folder.

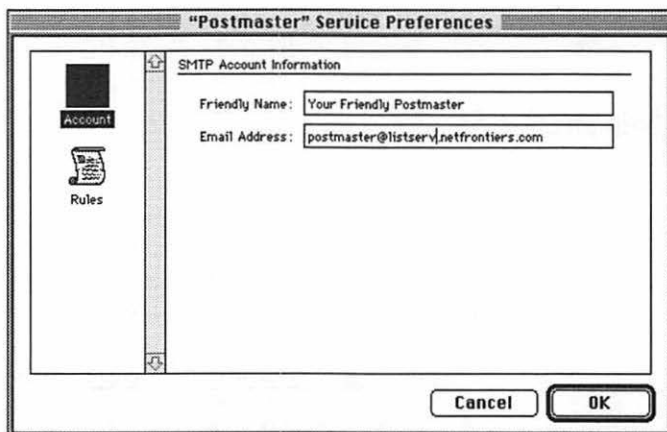


Administration Service Window

You can go in and change the folder to which the logs are saved or have the service forward some or all the logs to your e-mail account on another system.

Postmaster Service

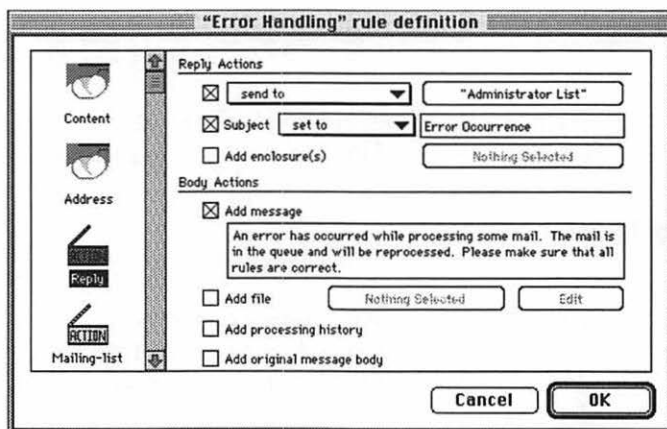
The Postmaster service supports the Postmaster account, which is designed to receive error messages and user reports of problems. This account is *required* for all mail hosts by RFC 822, the Internet *Standard for the Format of ARPA Internet Text Messages*.



Postmaster Services Window

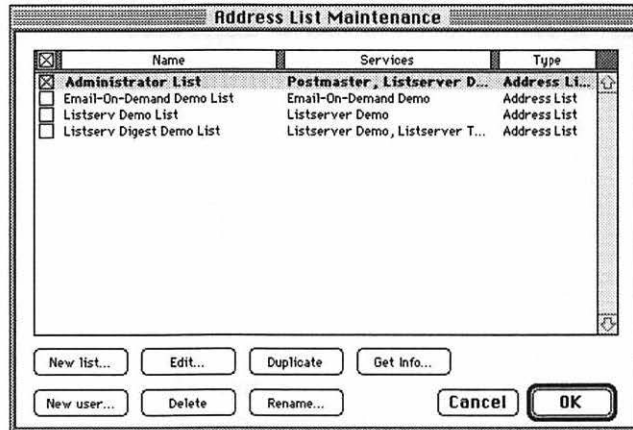
In the **Account** pane, edit the **Email Address** field to reflect your site, as in *postmaster@listserv.netfrontiers.com*. Also, go into the **Rules** pane and make sure the service is doing what you want. Take a look at the **Error Handling** rule, for example. This one goes to work when an error occurs during e-mail processing. It sends the mail message to the Administrator List and writes a message in the log file.

Click the button labeled **“Administrator List”** in the **Reply Actions** section.



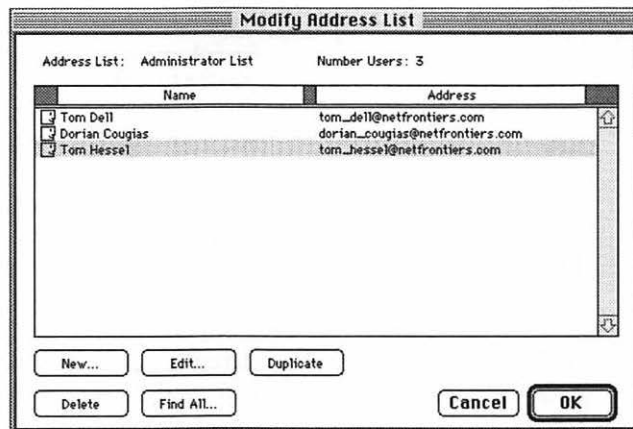
Error Handling Reply Action Pane

This generates the Address List Maintenance window, in which you place e-mail account information.



Address List Maintenance Window

Double-click the Administrator List line item to see who will be receiving these error messages. You don't see very much, do you? You should be in here! Click the New... button to give ListSTAR your e-mail address, as well as those of any other network administrators you wish to torment.




Administrator Mailing List

One nice thing here: Once created, these lists are *global*, meaning that they can be used in more than one rule. Creating the list is tedious, however. There is no “import” function, so unless you use a subscribe method, these types of lists must be created one user at a time. That bites? You might very well think so. I couldn't possibly comment.

SETTING UP E-MAIL-ON-DEMAND SERVICES

Now that you have had a look at how to use ListSTAR's rules, let's move on to setting up a new service. This can be accomplished quickly thanks to StarNine. They have thoughtfully created a number of demonstration services that need only be modified to get your site up and running. We begin with the one called "Email-On-Demand Demo."

Select this service in the ListSTAR Services window, and then click the **Duplicate** button to make a copy. Always keep the original demonstration services untouched for future reference.



Duplicating Email-On-Demand Demo

Call the copy whatever you like, and then click **Create**.

Here is what I intend to show you. I tell you ahead of time because you should always have these things thought out before you jump into ListSTAR. Network Frontiers teaches numerous classes on such subjects as network design, network management, servers, backup systems, connecting to the Internet, Webmastering, and speaking with a very loud voice. We often receive phone calls from would-be students enquiring about the latest list of cities in which we are teaching and the dates of those classes. We also wrote the Apple Certified Server Engineer (ACSE) program based largely on what is in this Network Frontiers Field Manual Series of books. We field many phone calls from students about that, too. To better serve this audience, I want to create a service here that will:

- Go by the e-mail account name of "Schedule."

- Accept e-mail addressed to that account containing the words “class info” in the Subject field and respond with an e-mail form containing an index of available files.
- Based on selections made in the form above, as returned by the user, enclose the requested files and reply.
- Automatically update the index daily.
- Collect the e-mail addresses of those asking for information.
- Log these transactions.

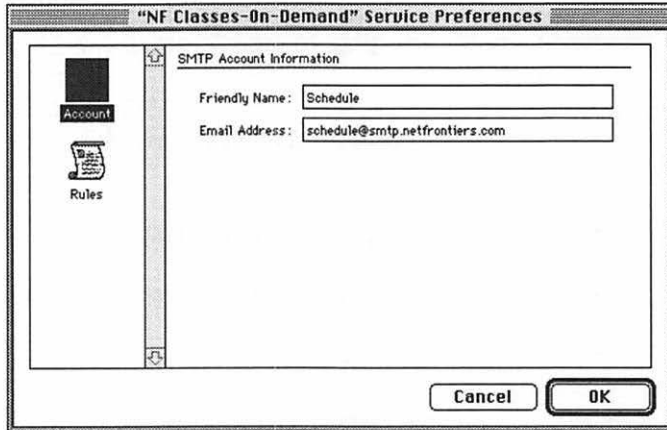
ListSTAR Services					
<input checked="" type="checkbox"/>	Service Name	Sent	Rev'd	Errors	Accessed
<input checked="" type="checkbox"/>	default	0	0	0	New Service
<input checked="" type="checkbox"/>	Administration	0	0	0	1:03 PM Thu, Feb 8, 1996
	Email-On-Demand Demo	0	0	0	New Service
	Email-On-Demand Timer	0	0	0	New Service
	Listserver Demo	0	0	0	New Service
	Listserver Timer	0	0	0	New Service
	NF Classes-On-Demand	0	0	0	New Service

New... Edit... Delete... Rename... Duplicate... Text View...

Selecting the Copy of a Demo Service

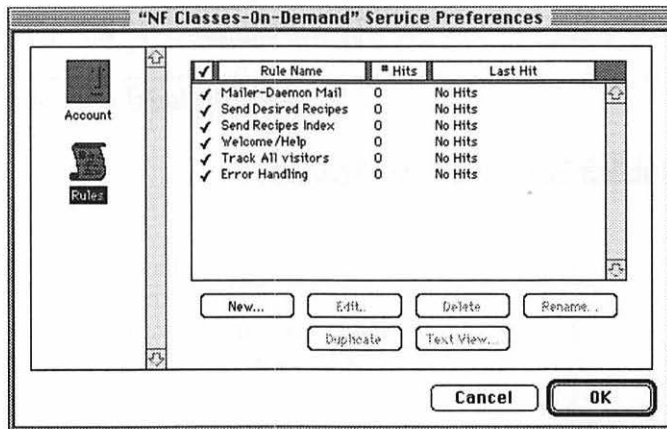
Establish Service E-Mail Account

The first step in doing this is to double-click the new service, in this case “NF Classes-on-Demand.” I then give it a name in the fields of the Account pane. This will not just copy the service’s line item, but will also copy all the files associated with that service in the Services folder of the ListSTAR folder.



Establishing the Account Name

Next, I select the Rules pane so that I can modify StarNine's settings. I have never been one to make more work for myself than necessary.



Modifying the Rules

Handle Reserved Addresses

These rules are executed in the order in which they are listed. The first one is "Mailer-Daemon Mail," which handles any e-mail from *reserved addresses* like "mailer-daemon." Such e-mail comes in when ListSTAR sends mail to an address

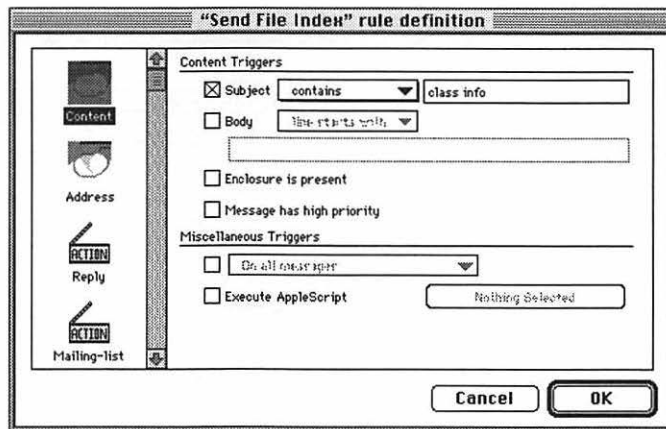
that is invalid or unavailable, which is bound to happen from time to time. Star-Nine's logic, as written in the **Rule's Comment** field, is this:

If the from address is another list server or control address, then we send the message to the Administrator List. We also want to discard it and stop processing the e-mail message.

I see no reason to change anything here. In fact, this function is vital and will be part of every service I create.

Send File Index

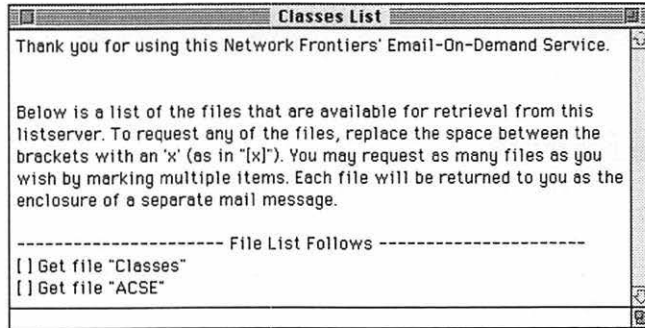
The second rule with which I want to work (the third in the list) is "Send Recipes Index." I use the **Rename** button to make this "Send File Index." This one is more complicated, since it involves an advanced ListSTAR feature called the *AutoResponder*. This feature uses AppleScript to produce a selection form in which users can see what files are available and choose only the ones they want to have e-mailed to them. It is like fax-back: "For sales, press one. For tech support, press two. To hear a Greek myth, press three."



Accepting E-Mail

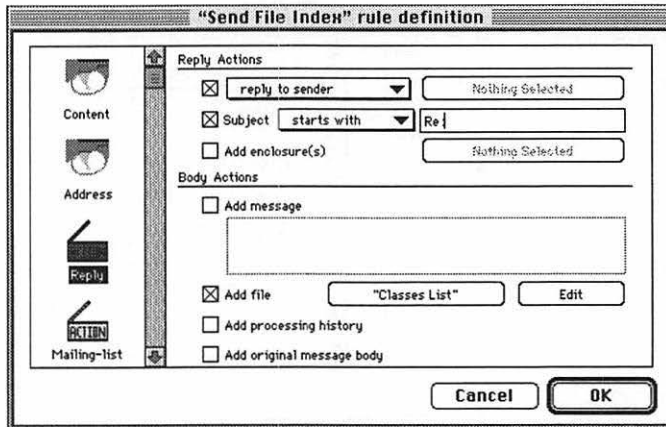
In the **Content** pane, I modify the rule to accept e-mail that contains **class info** in the **Subject** field. I then move on to the **Reply** action.

In this pane, I direct ListSTAR to reply to the incoming message with the words **Re: class info** in the **Subject** field. In StarNine's demo, ListSTAR is also directed to enclose a file entitled "Recipe List." That would not be appropriate for my purposes. Thus, I go into the automatically created NF Classes-On-Demand folder within the Services folder, rename the file, and modify it to serve my needs.



Modified File Index

Who said cheaters never win? Next, I click the button in the **Add file** field to select the modified file.



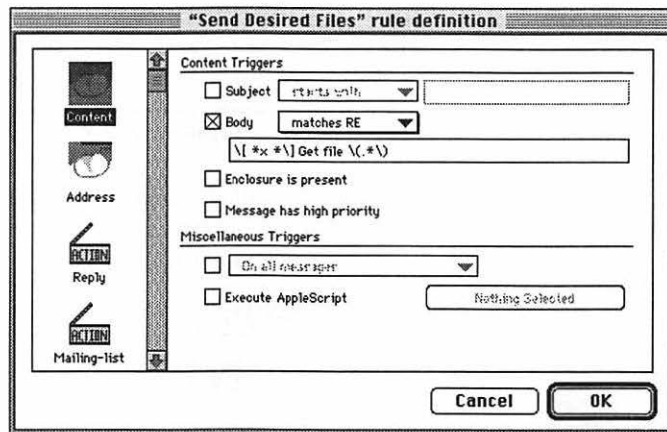
Selecting File Index Enclosure

I also find a Filed Recipes folder within the NF Classes-On-Demand folder within the Services folder, which I ruthlessly rename Filed Schedules. Within this folder are a bunch of text files StarNine puts there to illustrate the demo rules. I throw

out their files with such titles as “Beer Cheese Soup” and “Calzone Di Broccoli” and coldly replace them with two of my own, “Classes” and “ACSE.”

Send Desired Files

The third rule with which I want to work (the second in the list) is “Send Desired Recipes.” I use the **Rename** button to make this “Send Desired Files.” In the **Content** pane of this rule, you see something scary, as shown in the following picture.



Uh-oh, it looks like some kind of code!

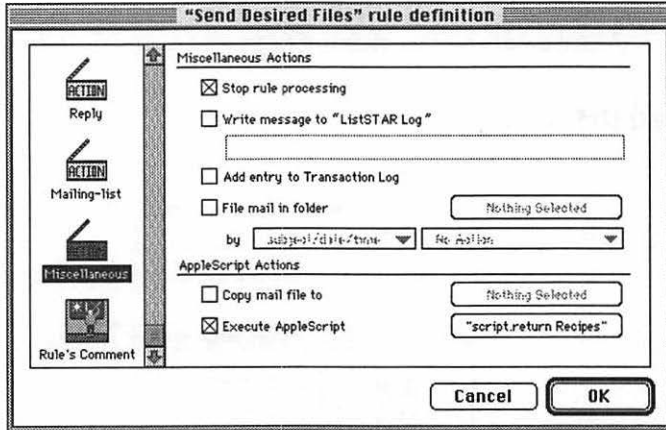
Don't panic. It's okay. That gibberish is well documented in the *ListSTAR Administrator's Guide*. What this means to ListSTAR . . .

```
\\[*x *\\] Get file \\(.*)
```

. . . means something like this to humans:

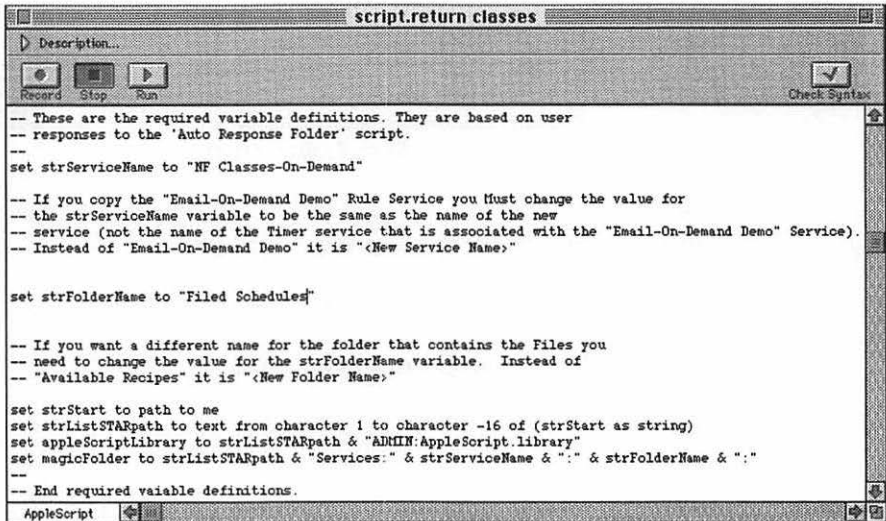
```
[x] Get file "Classes"
[x] Get file "ACSE"
```

This action tells ListSTAR that when a match occurs between data in the incoming form and the Classes List I just created, it should hang onto the data. If an AppleScript is launched later in the rule, ListSTAR will pass the data on to that script. I click on the **Miscellaneous Triggers** field to find that, indeed, an AppleScript is needed.



Selecting an AppleScript

Hmmm . . . any AppleScript called “script.return Recipes” is probably not going to work for me. I rename it “script.return classes,” and then open it in Script Editor to modify any other inaccurate values.

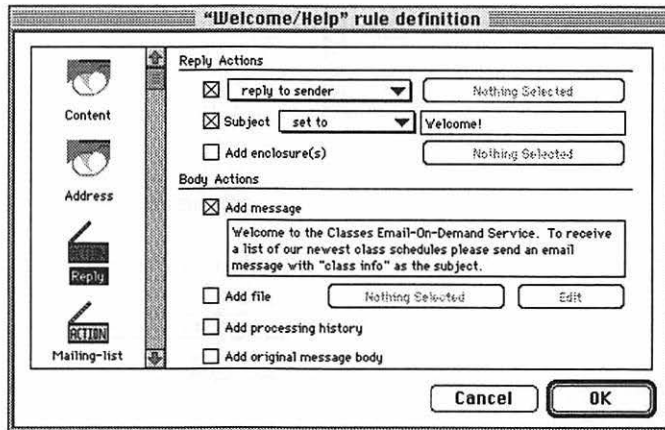


Changing AppleScript Variables

I do not compile this new modified script, as it must remain text. I do click the button in the **Execute AppleScript** field to choose the modified file.

Welcome/Help

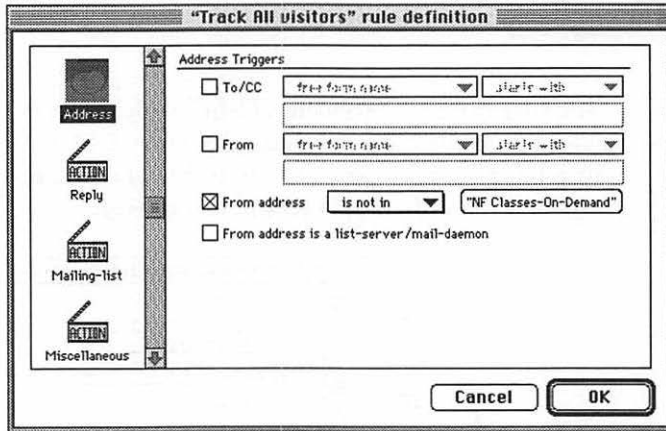
The fourth rule is “Welcome/Help,” which I do not need to rename. It will go into action only if the previous rules were unable to process an incoming message. Its job is to give a user instructions by sending back an informative “Welcome” message. I modify the content of this message in the **Reply Actions** pane.



Adding a Welcome Reply

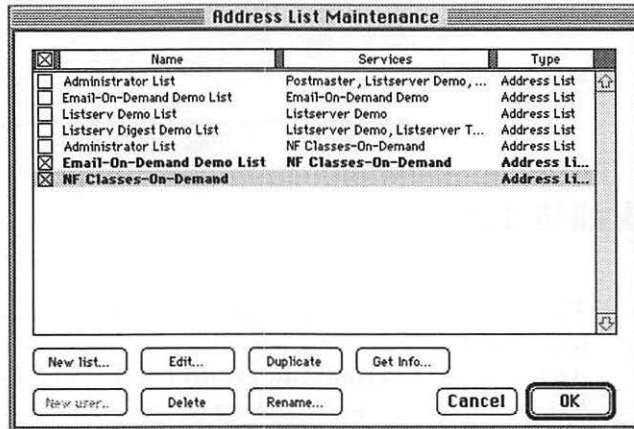
Track All Visitors

The fifth rule is “Track All Visitors,” which I also do not need to rename. Its job is to record the e-mail addresses of every user who requests service. First, I click the button in the **From Address** field of the **Address** pane to pull up the Address List Maintenance window, where I create a new mailing list for this service.



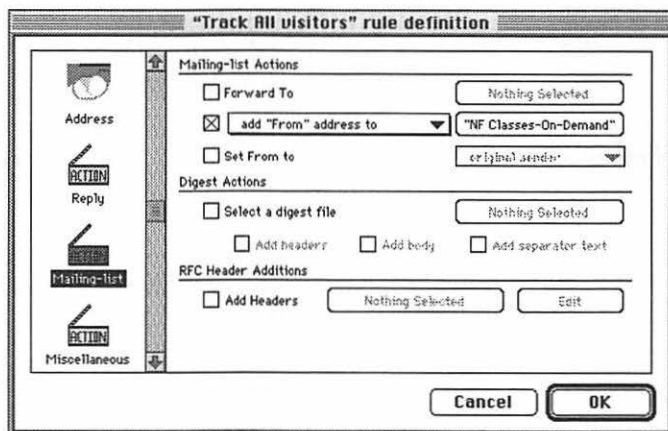
Selecting the "From" List

I also deselect the "Email-On-Demand Demo List."



Creating Service Mailing List

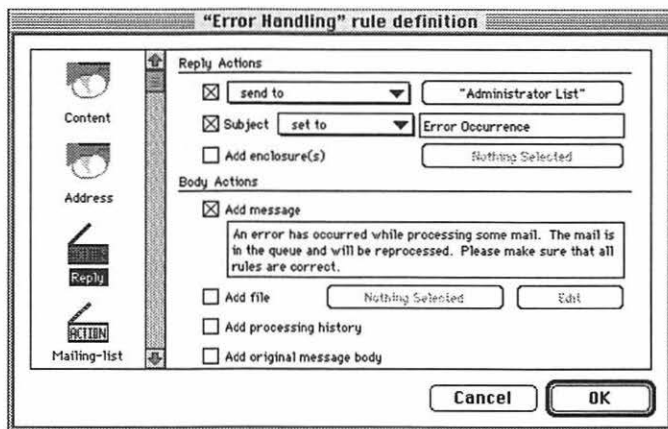
I then select the newly created list in the button by the **From address** field of the **Address** pane. I also select the list in the button by the add "From" address to pop-up menu in the **Mailing List Action** pane. I again deselect the "Email-On-Demand Demo List."



Selecting the "From" List Again

Error Handling

The sixth and final rule is "Error Handling." I do not rename this either. In fact, I do not need to change a thing. This rule allows me and others in the Administrator List to know when there has been an error in message processing.




Telling Me about Errors

At this point, I click the **OK** button to return to the ListSTAR Services window. There I click next to the "NF Classes-On-Demand" line item to add a checkmark

and start up the service. This completes the steps necessary to use the Mailer service for this e-mail-on-demand system. There is an additional service to set up though: the Timer service that updates the "Classes List" index file.

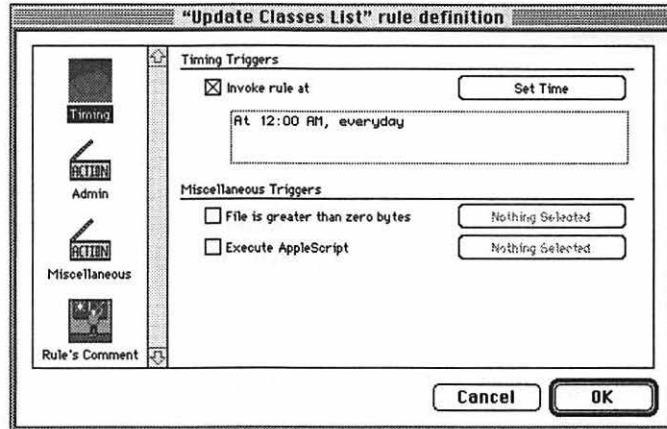
E-Mail-on-Demand Timer

As I did previously with the Mailer service, here I copy and rename the StarNine demo service. Upon opening it, I find but one rule, "Update Recipes List," which I rename "Update Classes List."



Renaming Update Rule

In the Timing pane, I find that this function will execute every day at 12:00 A.M. I do not bother to change this.



Setting the Timing

In the **Miscellaneous** pane, I find that the service is to execute an AppleScript called “script.create Recipes List.” As with the earlier AppleScript, I modify and rename this one. Its purpose is to update “Classes List” whenever additional files have been added to the Filed Schedules folder.

Once I have finished any other modifications I want to make, I activate the service. With minimal effort, I now have a robust, self-contained list server system.

GO FORTH, BRAVE SOULS

This short chapter can in no way replace StarNine's *ListSTAR Administrator's Guide* for detail, nor does it really go into AppleScript. It should have made you familiar enough with what is going on here, however, to make further reading easier and to help you come up with new ideas.

The most important parting comment I can offer is this: *Test the thing!* Peg the list server from several different external e-mail accounts before even thinking of putting it up for real. There are a lot of steps here, and any one of them that is out of place can make a real mess. You need to catch any mistakes *before* your ListSTAR receives heavy usage.

CHAPTER 3: GATHERING INTERNET E-MAIL WITH MAIL*LINK GATEWAYS

Once you are familiar with how LAN-based e-mail works, you are ready to move on to the grand-daddy of all e-mail transports, the Internet. Gathering e-mail from the Internet and placing it in the LAN-based e-mail accounts of your users is easy, requiring only that you install a *gateway* to act as a translator between local and remote systems. If your users are not clamoring for an Internet gateway yet, they soon will be. It has been estimated that one trillion e-mail messages were sent in 1995, while the U.S. Postal Service made “only” 180 billion deliveries!

Since we have chosen CE Software’s QuickMail in our example of a LAN-based e-mail system, we will use QuickMail-to-Internet gateways in our discussion here. The two best known are both made by StarNine Technologies. Mail*Link Remote allows your QuickMail users to exchange mail with users of QuickMail, Quarterdeck Mail (formerly Microsoft Mail), or Internet mail via UUCP (UNIX to UNIX Copy Protocol). Mail*Link SMTP allows your QuickMail users to exchange mail with Internet users via SMTP (Simple Mail Transfer Protocol). Although each does its job a little differently, they are both configured in much the same way.

Why a choice? In the 1980s, the Internet community chose SMTP as its standard mail transfer protocol. Before this, however, AT&T had already developed and adopted UUCP and its protocols as the preferred method of mail transfer between UNIX systems. Not surprisingly, AT&T wanted a protocol that worked well over

analog telephone lines. Despite not having been accepted through the consensus process used to choose SMTP, UUCP protocols are nevertheless a de facto standard, especially for dial-up connections.

Before you try to set up the Mail*Link gateway to work with a UNIX-based mail system, you need to provide information *to* and gather information *from* that system's administrator. This might be a person at your own company or someone at a commercial Internet service provider (ISP). These people will have an e-mail *host*, which is a machine that sits on the Internet and gathers mail for your users by proxy. The most important question to ask these people initially is what protocol you should use to connect to the host and with which to exchange e-mail. If the host makes use of UUCP, you want Mail*Link Remote. If the host makes use of SMTP, you want Mail*Link SMTP. For the rest, the basic considerations are these:

- What will you name the gateway?
- What is the name of the remote host?
- When should QM Administrator call the remote host?
- What IP address will the gateway use (Mail*Link SMTP)?
- What domain name will the gateway serve? A domain name is the textual representation of a numerical Internet machine address. For instance, we use *net-frontiers.com*. Domain names cannot contain spaces, tabs, parentheses, commas, brackets, semicolons, colons, backslashes, or braces, and are always in lowercase.
- What are the login procedures and passwords for the remote system?

You have one important option. If you have a dedicated mail server Macintosh connected directly to the Internet with its own IP address that is mapped to your domain address, you can dispense with the remote host altogether and just let Mail*Link SMTP act as a host. If your ISP does not maintain a mail server, this is your *only* option, unless you use the Apple Internet Mail Server that we speak about in the chapter entitled "Using the Apple Internet Mail Server as Mail Host," beginning on page 155.

As it can be a lot to think about, we suggest that you gather this information before going further. The time you lose in preparation will be offset by the time you save when troubleshooting later.

WHERE THE MAIL GOES

The following steps show an example of one path a message can take through the Mail*Link gateway:

- Step 1:** The QuickMail user creates a QuickMail message, which is addressed to a user in the Mail*Link gateway MailCenter. The address field associated with the recipient does not contain a QuickMail-created designation but rather a remote system account address, such as an Internet domain name.
- Step 2:** The Mail*Link gateway is located on the mail server and is active when QM Administrator is running. Messages are queued here.
- Step 3:** When a previously set time arrives, QM Administrator passes control of the system to the gateway, which will dial out through a modem to the remote host in the case of Mail*Link Remote, or will establish a connection with the remote host via MacTCP in the case of Mail*Link SMTP. Once logged into the remote system, the gateway uploads outgoing mail to the remote system. It also asks for any incoming mail. Incoming messages have either a UNIX-compatible version of your users' QuickMail account names or an alias for the same.
- Step 4:** The recipient's mail system forwards the message received from your gateway.

There are variables to this simple model. For instance, as the mail goes out over the Internet it could actually be routed through many mail systems before reaching the recipient. Also, the mail forwarding might not be initiated by the local system. Mail*Link can be configured to pass on mail only when a remote system calls it. Finally, mail going through the Mail*Link SMTP gateway does not go straight out to the Internet. It is passed over to the application `smtp.daemon` first. For more information about these nuances, refer to StarNine's documentation. It is cumbersome but comprehensive.

SETTING UP THE MAIL*LINK GATEWAY


The way you set up the Mail*Link SMTP and Mail*Link Remote gateways is similar, so we will talk about both here.

Configuring the Mail Server Macintosh

Before you do anything with the Mail*Link software you have installed, you need to tell the world where your mail server lives.

Set the Time Zone

To make sure Mail*Link generates the correct time in its RFC-compliant headers, you must tell it what time zone it is in relative to Greenwich Mean Time. This is easy on a Macintosh. Double-click the Map control panel, found in the Control Panels folder of the System Folder.

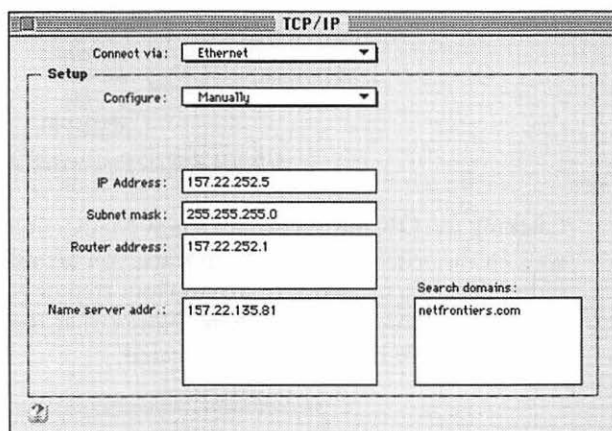


Map Control Panel

Click on the map where your city should be. Its name, or the name of a city nearby, will appear in the control panel. Click the **Set** button to save this setting, and close the control panel.

Configure TCP/IP (Mail*Link SMTP)

If you are using Mail*Link SMTP with a direct Internet feed, configure the mail server Macintosh's TCP/IP control panel (formerly the MacTCP control panel under Classic Networking) with the machine address you wrote down earlier.

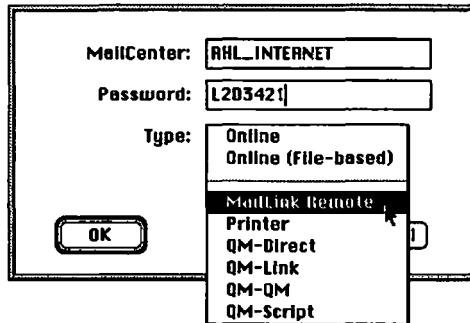


TCP/IP Control Panel

Restart the mail server when you are finished. That is probably all *you* need to worry about as far as IP addressing is concerned. Your ISP has one additional task, however. The unique IP address of your mail host must be mapped in the tables of the ISP's domain name server to some logical domain name—as in *smtp.netfrontiers.com*—so that other mail hosts on the Internet can learn of it.

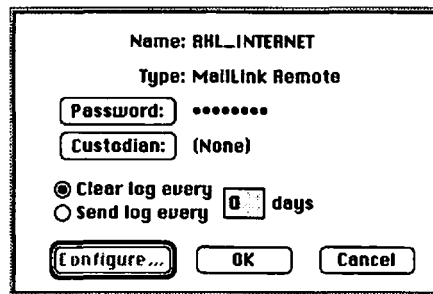
Configuring the Gateway MailCenter

To set up the Mail*Link Remote or Mail*Link SMTP gateway, select **New MailCenter** from QM Administrator's **File** menu. When the first window appears, enter the desired name in the **MailCenter** field and any password in the **Password** field (optional). Select **MailLink Remote** in the MailCenter **Type** pop-up menu.



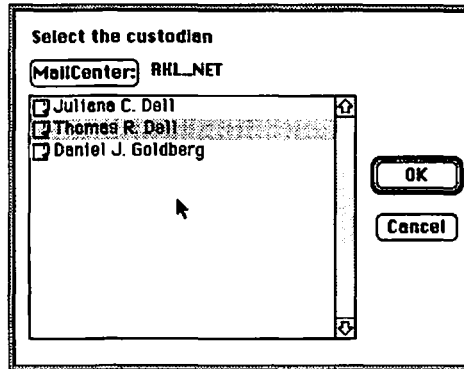
Selecting MailLink Remote MailCenter Type

Clicking the **OK** button creates a new MailCenter that serves the Mail*Link gateway. In the window that next appears, you are asked to select a custodian. If you did not select a password in the previous step or you wish to change the password, you have a chance to do so here. Notice that the **Configure...** button is grayed out when there is no custodian selected.



Selecting Custodian (Configure... Button Grayed Out)

To continue, click the **Custodian** button. A scrolling window appears in which you can choose the name and online MailCenter of the user who will receive maintenance logs and notifications for the new MailCenter.

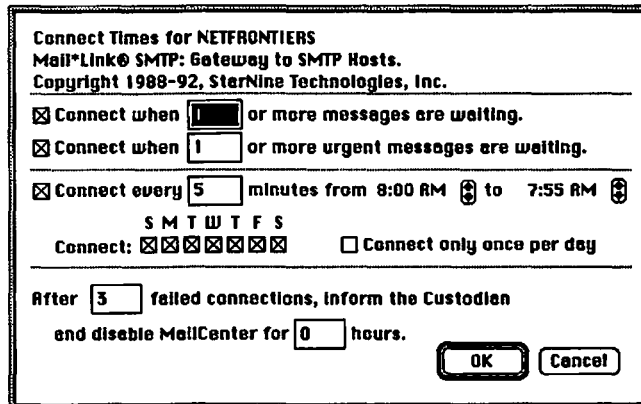


Choosing the Account Holder to Be Custodian

Click **OK** when finished, and you are returned to the previous window. Now, however, **Configure...** is an option. Click this button.

*Mail*Link SMTP Preferences*

If you are setting up the Mail*Link SMTP gateway, clicking this button takes you to the Connect Times window. This window is identical to those for other Quick-Mail gateways.



Connect Times Window

As with other QuickMail gateways, you should reset these defaults to take into account the volume of mail generated by your users, the urgency of that mail, and the most advantageous (non-peak) times for connecting to remote systems. Other

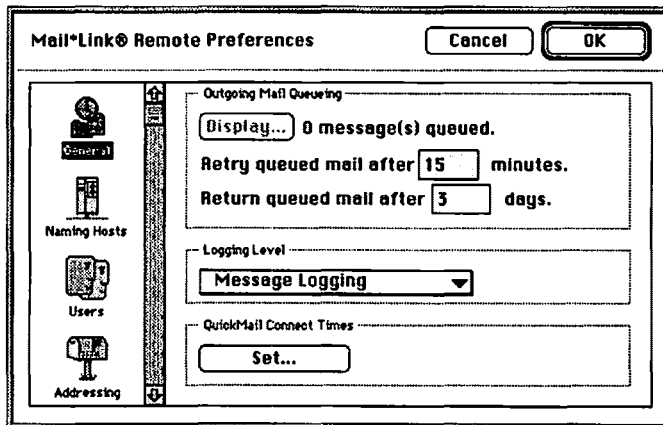
Mail*Link SMTP preferences are set using the smtp.daemon application, which we discuss on page 126.

Mail*Link Remote Preferences

If you are setting up the Mail*Link Remote gateway, clicking this button will put you into the Mail*Link Preferences window. There are many features to be configured, each represented by an icon in the scroll box on the left. The important ones are described next.

General Preferences

This category contains instructions for how Mail*Link queues mail, how much detail it goes into when logging sessions, and when QM Administrator will be permitted to initiate a session.



General Preferences Window

The choices in the **Outgoing Mail Queuing** pane pertaining to how long Mail*Link will attempt to forward undelivered mail are pretty straightforward. The mail might not go through the first time because the host server is down. If you know the host server will be down regularly for periods longer than those set in the defaults, you will want to change these.

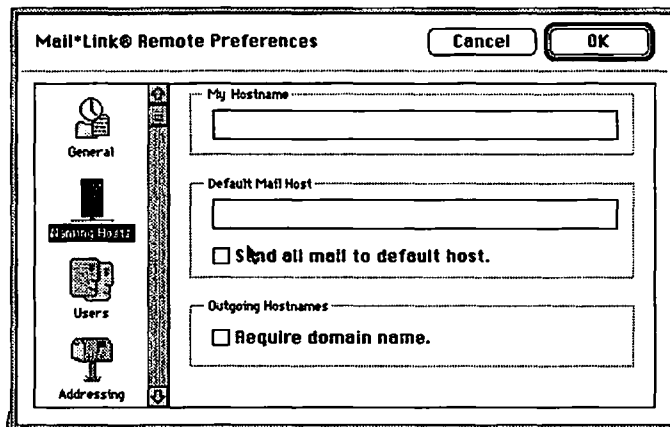
Below the **Outgoing Mail Queuing** pane is the **Logging Level** pop-up menu in which you can decide how much information you want Mail*Link to gather about each session. The chances are good that you will never need more than the default **Message Logging**. If you do have connection problems, however, select-

ing higher levels will provide you, or those you call upon for technical support, with invaluable information about the intricacies of what is going on.

Below the **Logging Level** pop-up menu is the **QuickMail Connect Times** section. The **Set...** button generates a window identical to the Mail*Link SMTP window we just talked about. Keep in mind that what you set here does not determine when the Mail*Link Remote gateway makes a connection, but rather, at what times it is *allowed* to make a connection. *When* a connection is actually attempted is set in the System-to-Call Configuration window that we discuss on page 123.

Naming Hosts Preferences

This category contains instructions for how Mail*Link presents itself to remote systems and to which remote host it will look to forward outgoing mail with unknown destinations.



Naming Hosts Preferences Window

In the **My Hostname** field, enter the name that will identify your local gateway to remote systems. As these remote systems are UNIX-based, you must use a name that does not use characters they cannot read—such as non-alphabet characters, spaces, and punctuation—and which is preferably less than eight characters long. The name you enter here will be used in logging into remote systems. The administrator of these remote systems must make sure that the UNIX password file and UUCP system file have entries for the gateway.

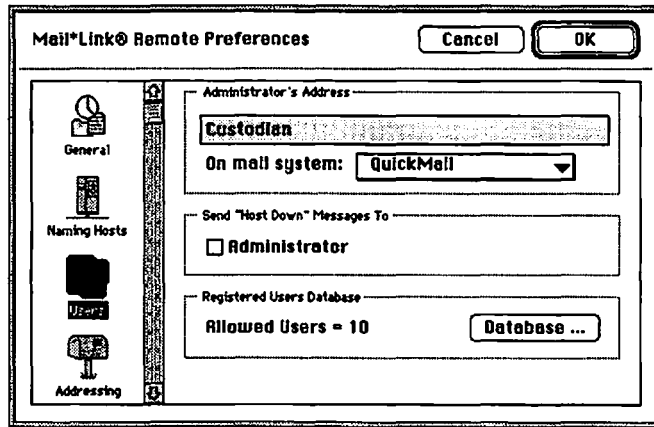
The local host name is also used in the return addresses the gateway puts on mail going to remote systems, and in the header information of mail coming from remote systems.

In the **Default Mail Host** field, enter the name of a mail host to which Mail*Link will forward mail that is meant for a destination unknown to the gateway. The default mail host will most likely be a UNIX mail host that can forward mail through many other systems until it reaches its destination, such as that of your ISP. If you wish to forward *all* mail through such a remote system, select the box in the **Send all mail to default host** line.

Some mail hosts require that a domain name, such as *netfrontiers.com*, be present in all mail addresses. If you are connecting to such a system, selecting the box in the **Outgoing Hostnames** pane causes any mail that does not contain at least one period to be returned to the sender. Selecting this *also* requires you to have at least one period in the names in the **My Hostname** and **Default Mail Host** fields.

Users Preferences

Settings having to do with the administrator and users of the Mail*Link gateway are accessed here.



Users Preferences Window

The default for the **Administrator's Address** field is **Custodian**, the name you selected when first creating the Mail*Link MailCenter. The proper format for this entry is the custodian's QuickMail account address appended with "@" and the name of the MailCenter to which the custodian belongs. For instance, *tom_dell@netfrontiers.com*. Note that underscores are used in place of spaces. UNIX again!


The administrator can also reside on an Internet-based mail system. In this case, choose **UUCP** or **SMTP** instead of **QuickMail** from the **On mail system** pop-up

menu and enter the user's correct Internet address above it. Any mail addressed to "postmaster" is forwarded to the user's address in this field. If you select the **Administrator** checkbox in the Send "Host Down" Messages To pane, the gateway sends an urgent message to the person you designated as administrator after it has failed in three attempts to contact a host system.

Like QuickMail, Mail*Link is sold for a specified number of users. The number of users that your system is licensed to serve appears in the **Registered Users Database** pane. Clicking the **Database...** button generates the Registered Users window. We will talk more about that window starting on page 131.

Addressing Preferences

UNIX mail systems require that account names adhere to certain characteristics. QuickMail account names may or may not meet these requirements. Mail*Link can create its own UNIX-compliant addresses for QuickMail users, called *aliases*, automatically. How the gateway will create these addresses and how it will display them on incoming and outgoing mail is set up here.



Addressing Preferences Window

Selecting the **Auto-register users on outgoing mail** checkbox in the **Alias Creation** pane directs Mail*Link to add users to its database whenever they send mail through the gateway. It will add users until it reaches the maximum number for which you are licensed. If you do not select this, you need to add users explicitly through the Users Preferences window. Unregistered users will have their mail returned to them.

It is sometimes better to create your own aliases. The way Mail*Link translates a QuickMail account name to be UNIX-friendly can leave you with something pretty ugly. For example, QuickMail user *Thomas R. Dell* becomes *Thomas_R#d#_Dell*. You can specify a simpler alias such as *trdell* or just *tom* in the Users Preferences window, or you can let your users do it by selecting the **Allow-user to change their alias** checkbox. I have a friend whose name is so frequently misspelled that she uses aliases to add over nine ways to spell her name, just to make sure she receives any mail destined for her.

By selecting the **Use alias as reply address** checkbox in the **Outgoing Mail** pane, you direct Mail*Link to use a presumably simpler alias for a user's QuickMail account name, as we have already seen, in the reply addresses. You do not *have to* use aliases. If you deselect this box, Mail*Link uses the actual QuickMail account name. However, to make it consistent with RFC 822, the standards for this type of UNIX mail addressing, Mail*Link will perform a translation on the name that could change some characters and make what the recipient sees in the **From:** field fairly nasty. Should you prefer to go with the real QuickMail account names, as modified, you can also have the sender's MailCenter information tacked on by selecting **Use MailCenter in reply address**. Of course, it is translated, too.


When mail travels through the Internet, it can go through many different mail systems before reaching yours. In so doing, it picks up RFC 822-compliant address information from these systems that is placed in what is called the *header*. These headers can become quite long. By making a choice in the **Place RFC header at pop-up menu**, you can designate where this header appears in the QuickMail messages that are ultimately seen by your users, if they are seen at all.

Whether or not you select **Strict incoming address validation** depends upon the size of your organization. When mail comes in, Mail*Link calls on QuickMail's NameServer to associate a QuickMail user and MailCenter address with the address on the message. If this checkbox is not selected, it sends the mail to the first account name that matches. If it is selected, it waits until both the account name and MailCenter address match. If you have a big site where there could be people of the same name in different MailCenters, you will need this.

Enclosures Preferences

Besides the information that users type into QuickMail forms, they can also send file enclosures if the site with which you are exchanging mail permits this. Because the formats of files created by various applications can be hard to read by remote mail systems, it is conventional to *encode* file enclosures in more universal formats, usually as straight ASCII text. Mail*Link automatically scans for and translates

incoming enclosures using the UUencode/AppleSingle and BinHex formats. You can determine which will be used for outgoing mail in this window.



Enclosures Preferences Window

By selecting **Treat each enclosure as separate file**, you direct the gateway not to pack different file enclosures together. Say you had one text file and one program, or binary, file enclosed. The gateway handles the text file based on the criteria set elsewhere in this window under **TEXT enclosures**, and handles the binary file using criteria also established elsewhere in this window. Two enclosures would leave the gateway. By selecting **Treat multiple enclosure as 1 PACKED file**, you direct the gateway to cram the different file enclosures together into one file, then encode it. Take the same text file and binary file. The gateway will handle both using the criteria established under **Binary enclosures** elsewhere in this window. One enclosure would leave the gateway.

The **TEXT enclosures** pop-up menu determines how the gateway will deal with straight text file enclosures in outgoing mail. Choose **Use Binary Method** to apply the criteria set under **Binary enclosures** to text as well. Choose **Append to Mail** to have the content of the message tacked onto the bottom of the e-mail message itself.

The **Binary enclosures** pop-up menu determines how the gateway will encode most, if not all, file enclosures in outgoing mail. If you choose **AppleSingle-UUencode**, Mail*Link translates the Macintosh file into a file format that permits both its data and resource forks to be recreated on the receiving system, and then encodes this resulting file using the UNIX standard UUencode format. The data

fork of a Macintosh file contains “the guts,” such as text. The resource fork contains format and icon information used by the Macintosh operating system. (PCs only use data.)

Another common format is BinHex, which can be used to turn binary data into ASCII text and back again. To use it, select **Binhex 4.0**.

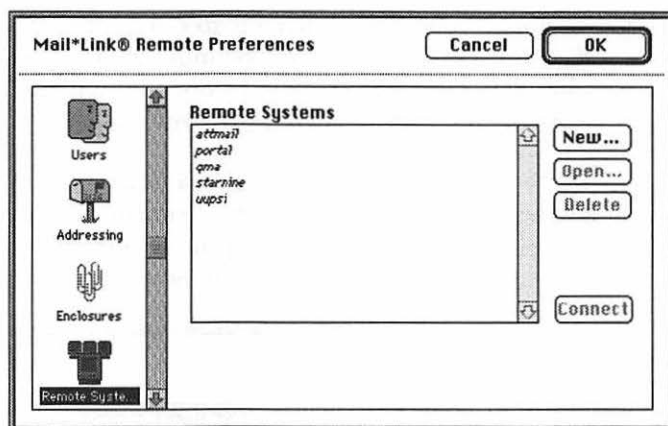
Choosing **Datafork only-UUencode** directs gateways to apply the UNIX standard UUencode algorithm on the data fork of a Macintosh file and just discard the resource fork. If the ultimate destination of these enclosures is other Macintosh users, this will leave them without resource information needed by the Macintosh operating system. If, however, the ultimate destination is PC users, they only need the data fork.

Which of these options should you choose? That is one of the things you need to discuss with administrators of the remote sites. One drawback to consider is that the choices you make here apply to *all* remote systems with which you wish to exchange mail. Because of this, you probably want to end up with methods that are as widely compatible as possible.

Depending on how you feel about telephone bills and server hard drive usage, select the **Limit total enclosures size to...** checkbox and enter a value. How large this number is depends on the types of file enclosures your users need to send out, how often they do it, and the relative cost of tying up the modem line long enough to send them. Naturally, if you have a fast modem, for example 28,800 bps, you might permit larger files than if your modem is slow, such as 9600 bps. If a user tries to send a file that is larger than the limit you place here, it will be returned. If you receive many complaints, you can always adjust the size. Be aware that the size you enter here also pertains to files as they come into the gateway. After they are encoded, expect their size to increase by about a third.

Remote Systems Preferences

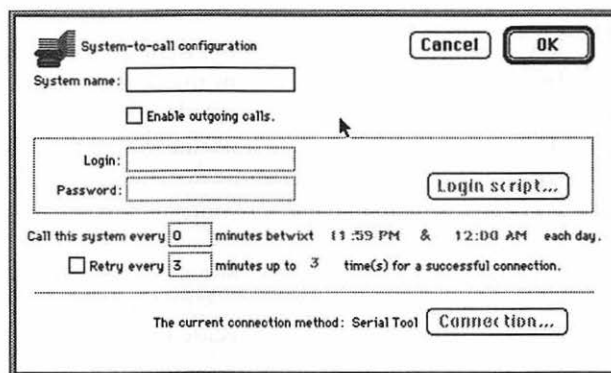
Mail*Link Remote knows about remote systems based on what is entered in this window. Each remote system with which the gateway will connect is listed here and is associated with instructions on how these connections take place.



Remote Systems Preferences Window

StarNine provides a few login scripts to commonly used systems with Mail*Link Remote. You will see these displayed in italics. To use any of these login scripts, follow the steps outlined in the following paragraphs using the **Open...** button instead of the **New...** button.

Click the **New...** button to set up your own remote system. The System-to-Call Configuration window appears.

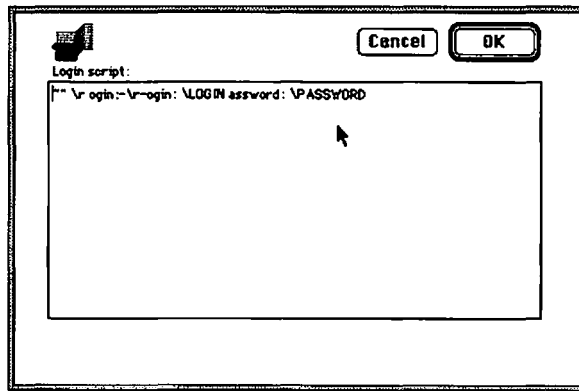


System-to-Call Configuration Window

Enter the name of the remote system's host in the **System name** field. This name will now appear in the Remote Systems list. If the name is the same as what you typed in earlier for the default mail host name, it will appear in bold. Even if you

do not want Mail*Link Remote to call out to other systems, but rather want to have them call your system, you must enter the remote system names here. You can still restrict outgoing calls by not selecting the **Enable outgoing calls** checkbox. The name in the Remote Systems list will appear in italics.

In the **Login** field, enter the login name and associated password assigned to you by the administrator of the remote system. Clicking the **Login script...** button brings up a window in which you can add other communications parameters established between you and the remote system's administrator.



Login Script Window

The login script is nothing more than the set of instructions you would use to log in manually, scripted in a form the computer can use to do it automatically. It is similar to a modem script, a DOS batch file, or an AppleScript. In many cases, the default script provided by StarNine will work. Mail*Link Remote knows to substitute the proper information for `\LOGIN` and `\PASSWORD`. In other cases, you need to work with the remote system's administrator to fine-tune it.


Like the **Connect Times** window in **General** preferences that determines under what conditions a communications session is permitted to take place, when a session actually takes place is decided through the **Call this system every...** field. Unlike the **Connect Times** settings, however, what you set here causes Mail*Link Remote to initiate calls at those times whether there is outgoing mail in the queue or not. In this instance, Mail*Link Remote just asks the remote system for any incoming mail.

If Mail*Link Remote is to call out to another Mail*Link Remote gateway, additional transfer methods are made available by clicking the **Transfer** button. As we are only talking about UNIX mail here, however, stick with the **Send mail with built-in UUCP** option.

The **Connection** button generates the same window of Apple Communications ToolBox options used by other QuickMail gateways and QM Administrator itself. For details on these settings, see the *QuickMail Administrator Manual*.

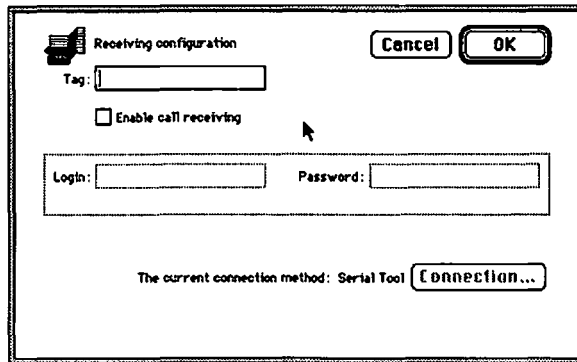
Remote Receiving Preferences

The Mail*Link Remote gateway can wait for calls from remote systems as well as initiate them. How these connections are handled is configured in this window. Use this window *only* if you are dedicating a modem or serial connection to the sole use of Mail*Link Remote. Otherwise, let QM Administrator handle the incoming calls.



Remote Receiving Preferences Window

Click the **New...** button to prepare for a remote system to dial into the Mail*Link Remote gateway. The Receiving Configuration window appears.



Receiving Configuration Window

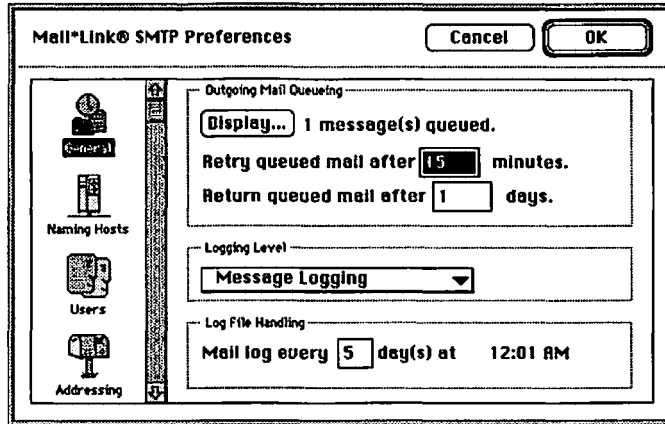
What you enter in the **Tag** field is for your own reference and is particularly useful if you are using multiple connections for mail transfers. You might tag one modem connection as “modem on port 1” and another connection as “ADSP to server,” for example. Enable or disable the various connection methods by selecting or deselecting the **Enable call receiving** checkbox.

What you type in for login name and password is what Mail*Link Remote will require of any remote system dialing in. In addition, you must have entered the calling system in the Remote Systems Preferences window. If not, the call will be terminated even if the login and password are correct.

*Mail*Link SMTP's smtp.daemon Preferences*

To configure Mail*Link SMTP completely, you must launch the smtp.daemon application that comes with it and select **Preferences** from the **Edit** menu. These preferences are like those of Mail*Link Remote, but with these differences:

General Preferences

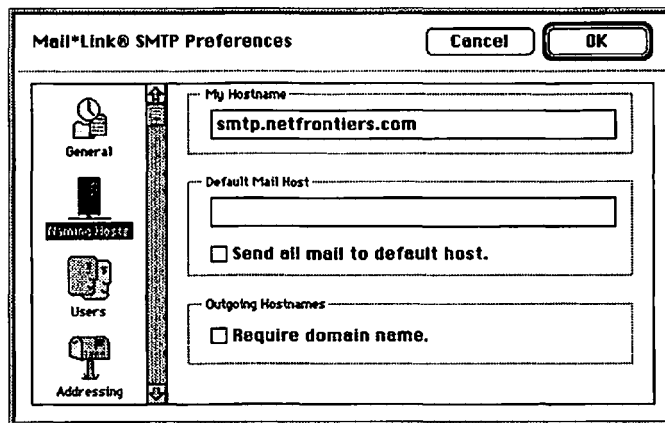


General Preferences Window

In the General pane of smtp.daemon's Preferences window, the **Connect Times** pane is replaced with a **Log File Handling** pane in which you may designate when smtp.daemon's log is sent to the Administrator. This is not the same as the Mail-Center log, which is sent according to criteria established in QM Administrator.

Naming Hosts Preferences

The Naming Hosts pane of smtp.daemon's Preferences window is the same as that of Mail*Link Remote. It is configured the same way, unless you plan to do without a remote host and just put your mail server directly on the Internet.



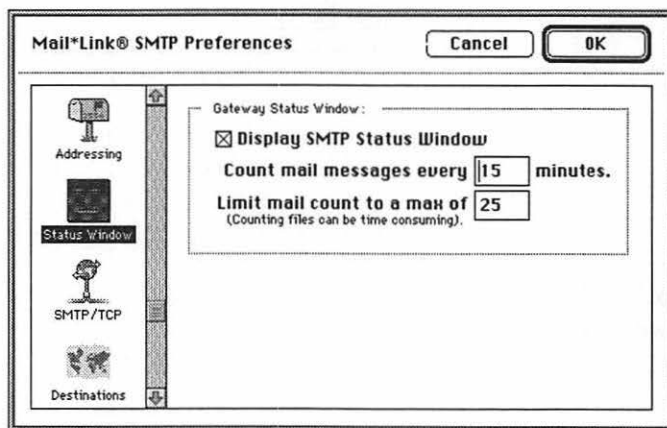
Naming Hosts Preferences (Gateway Also Host Setup)

If you wish to have the Mail*Link gateway send and receive its own Internet mail directly, make sure your network's DNS, or that of your ISP, is configured to map the IP address of your mail server Macintosh to an appropriate domain, such as *smtp.netfrontiers.com*. Type this domain into the **My Hostname** field. You can leave the other fields blank, or add the domain of another SMTP mail server to which your machine can forward mail if it cannot deliver it directly.

Do not select the **Send all mail to default host** checkbox if your gateway is acting as its own host. That would, of course, defeat the purpose.

Status Window Preferences

Unlike Mail*Link Remote, *smtp.daemon* has a Status Window preferences pane.



Status Window Preferences

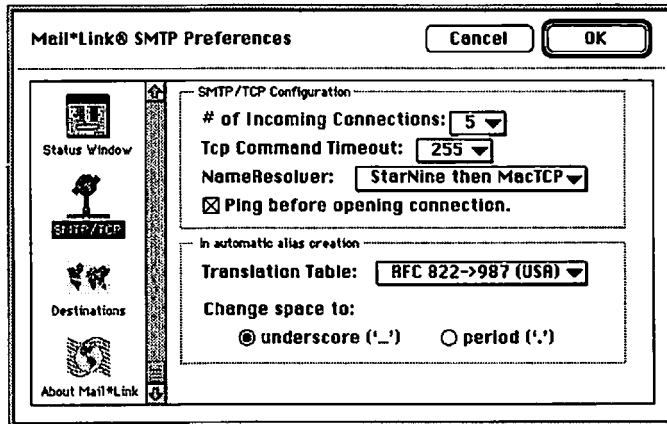
Selecting **Display SMTP Status Window** generates a handy monitor of the application's activity.



SMTP Status Window

The numbers you set here determine how much time smtp.daemon spends counting messages to show you their status. Every 15 minutes is often enough, unless you are doing troubleshooting. The gateway will stop counting messages after it reaches the number you specify in the **Limit mail count to a max of** field. Make this number smaller than 100 to conserve processing resources.

SMTP/TCP Preferences



SMTP/TCP Preferences

Since smtp.daemon uses a direct Internet connection to exchange mail and not a dial-up, its preferences do not include panes for modem connections. Instead, you simply select a few options for how Mail*Link works with TCP/IP.

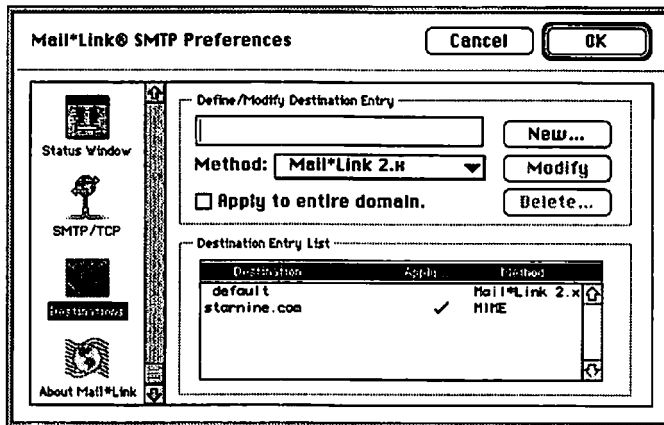
When running under System 7.5 and later, or earlier System 7 with the Thread Manager Extension installed, smtp.daemon can open a number of simultaneous connections. If you go with the default of 3 in the **# of Incoming Connections** pop-up menu, you will receive a total of *seven* concurrent processes, or *threads*: three will listen for mail, two will send mail, one is the *main* thread, and one more can also act as a sender. I do not really understand this. I am not that flavor of geek. I do know that each thread takes up 24K of RAM, so if you crank this number to 5, you should also move smtp.daemon's memory allocation up to 1,500K in its **Get Info** window, as you will be supporting 11 threads.

The **Tcp Command Timeout** pop-up menu controls how long smtp.daemon will wait for a response from a mail host before giving it up for dead. There is a lot of traffic out on the Internet right now, so I recommend going with the maximum setting of 255 seconds here. In networking, that is an eternity.

The **NameResolver** pop-up menu is important. When `smtp.daemon` wants to find the IP address that goes with a given host name, it can ask the TCP/IP (formerly MacTCP) control panel or its own StarNine Resolver for the information. If **MacTCP Only** is the only option you can set here, the mail server's TCP/IP (MacTCP) control panel does not know about any DNS servers and is falling back on a Hosts file to forward its mail to machines listed there. If you have other choices too, TCP/IP (MacTCP) does know about DNS servers and can ask them for resolution. Unfortunately, TCP/IP (MacTCP) only knows about a few DNS records, so it cannot gather as much routing information as it might. That is where the StarNine Resolver comes in. It knows more—specifically, it knows about *Mail Exchanger (MX)* records. The best way to set this is to choose the **StarNine then MacTCP** option, covering all your bases. You *must* do this if you are setting up Mail*Link as an independent host.

The rest of the settings you can leave alone. If you select **Ping before opening connection**, Mail*Link will try to make sure a remote host is “out there” before attempting to communicate. We usually do not select this since Mail*Link will not forward to mail hosts that do not respond to a ping but are there nevertheless.

Destination Preferences



Destination Preferences


Chances are good that you will never need to mess with `smtp.daemon`'s **Destinations** pane. If you go with the defaults, you will be all set to exchange mail with methods that are consistent with MIME, PC, and UNIX systems. If you need to set particular transmission characteristics for some sites, you can do that here. If so, read the StarNine manual, since this is beyond the scope of this book.

ESTABLISHING USER ACCOUNTS

Unlike many other QuickMail MailCenters, Mail*Link users are not established directly through QM Administrator. Instead, they are added to the Mail*Link user database automatically whenever they use the gateway, as we discussed earlier, or are added explicitly by the administrator. As with QuickMail itself, the number of users who can be part of your Mail*Link Remote database is limited by the number of licenses you purchased for them.

Setting up Accounts Manually

To set up user accounts manually, access the Users Preferences window. Clicking the **Database** button generates the Registered Users window.

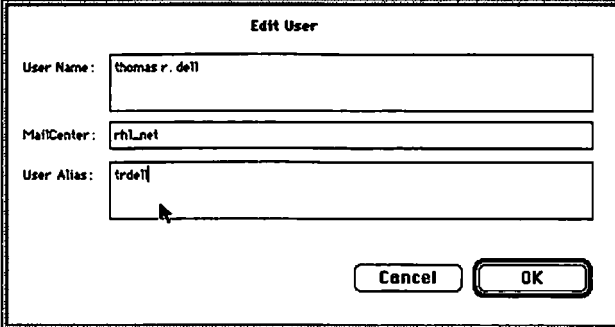


Registered Users Window

You have the following options in working with the accounts of Mail*Link users:

- **New:** Use this button to explicitly create a user account.
- **Purge:** Use this button to force Mail*Link to recreate its database from previously saved information, if necessary. This could be perhaps after—God forbid—a crash occurs when you are working within the database.
- **Import:** Use this button to accept a list of users from an ASCII text file.

- **Edit:** Use this button to edit the account of a user. It generates the following window. Why bother? Mail*Link generates its own user *alias* that is its way of presenting a QuickMail account name in a form UNIX understands. In doing so, it might translate special characters in the user's QuickMail name in somewhat unpleasant ways. For instance, *Thomas R. Dell* becomes *Thomas_R#d#_Dell*. Here you can specify a simpler alias, such as *trdell* or *tom*.



Edit User Screen

- **Delete:** Use this to remove a user, perhaps one who has left the company.
- **Export:** Use this button to create an ASCII text file that can be imported into another Mail*Link Remote gateway elsewhere.

The **Sort by** pop-up menu lets you change the criteria under which the accounts are ordered in the window. The default **Unsorted** displays the accounts in the order they were created.

Setting up Accounts Automatically

If you have purchased enough licenses and would rather let Mail*Link do the work, simply make sure the **Auto-register users on outgoing mail** checkbox is selected in the **Addressing** pane of the Preferences window. Whenever a user sends mail through the gateway, that user is added to the database.

Of course, when that user is added, Mail*Link creates its own funky alias in the database as well. You can change this alias to something more human yourself, or let the user do it. If you make sure the **Allow-user to change their alias** checkbox is selected in the **Addressing** pane of the Preferences window, users can establish

their own aliases by sending a message to "MAILER-DAEMON." More specifically, the following are the steps the user needs to take:

Step 1: Open a message and use QuickMail's **Special** address function to create an address with "MAILER-DAEMON" in the **First** name and **Address** fields.

Create Special Address:

First: MAILER-DAEMON

Last:

MailCenter: RHL_INTERNET

Zone: No Zone

Address:
MAILER-DAEMON

MailCenters OK Cancel

Special Address Window

Step 2: Type "Change alias" in the **Subject** field.

Step 3: Type the new alias in the format "alias:name." For example, "alias:tom." Do not use special characters, like punctuation.

Memo

Normal FILE PRINT SAVE DELETE ENCL CLIP REC REPT SEND

FROM: Thomas R. Dell
TO: MAILER-DAEMON
CC:
BCC:

OFFICE MEMO Subject: Change alias Time: []
Date: []

alias:tom


QuickMail Memo Form

Step 4: **Send** the message.

Users can find out what aliases Mail*Link might have previously assigned them by going through the same procedure but typing “Request alias” instead of “Change alias” in the **Subject** field.

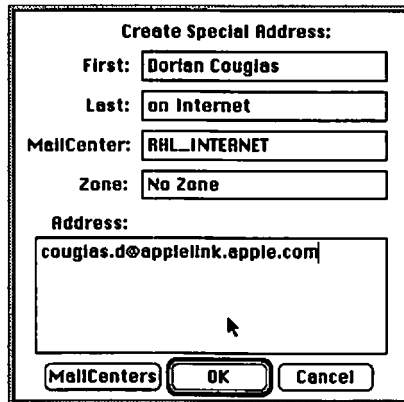
Mail*Link MailCenter Users List

The Mail*Link gateway MailCenter accounts are created in the same manner as are other MailCenters: directly by using **Create** from QM Administrator’s **User** menu or by importing the information from a database.



Adding a User to the MailCenter Users List

The Mail*Link MailCenter Users list differs from other MailCenters, however, in that it contains UNIX mail *recipients*, not local network originators. The address fields contain Internet addresses for people to whom you are sending mail.



*Address Box for Recipient of Mail*Link Remote Gateway Mail*

Importing Internet Lists

To import Internet address lists into a QuickMail MailCenter, open the desired MailCenter in QM Administrator and select **Import** from the **User** menu. The standard Open File dialog box appears, prompting you to select the desired TEXT file for importing. The TEXT file can be in a comma-delimited, space-delimited, or tab-delimited format. This format is set by the database or spreadsheet in which the file was originally created or maintained. After the file is selected, the Import File window appears, with the names appearing in rows and columns below.

Import File: Exported Names

Select the proper field separator so the information (below) is in separate columns.

Field separators: Tabs Commas 3 or more spaces

Select an inverse label which describes a column's contents and drag it to the title position of that column.

Not all labels will be used, some columns may not have titles.

First Last	Last, First	First name	Last name	Password	Address
(Drag label here)	(Drag label here)				
"on Internet, Dorian Cou..	"couglas. d@applelink.ap...				
"on Internet, Lynn Heibe..	"heiberger. l@applelink...				
"Rocha, Max"	"rocha@bearriver.com"				

OK Cancel

Import File Window

Study the way the names are formatted, and then drag the appropriate label over the title position for that column. This sets the importing order for the names as they are brought into the MailCenter Users List.

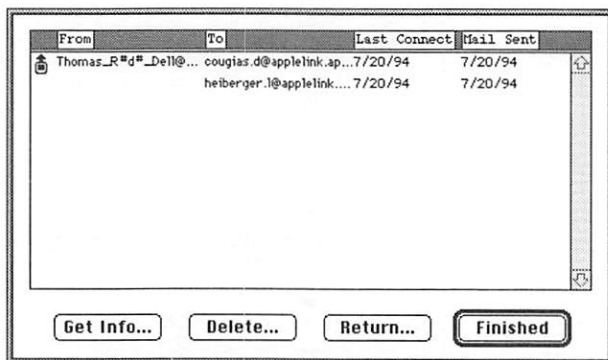
First Last	First name	Last name	Password	Address
(Drag label here)	(Drag label here)			
"on Internet, Dorian Cou..	"couglas. d@applelink.ap...			
"on Internet, Lynn Heibe..	"heiberger. l@applelink...			
"Rocha, Max"	"rocha@bearriver.com"			

Last, First	Address
"on Internet, Dorian Cou..	"couglas. d@applelink.ap...
"on Internet, Lynn Heibe..	"heiberger. l@applelink...
"Rocha, Max"	"rocha@bearriver.com"

The titles for the columns have been selected (above) and then dragged into place (below).

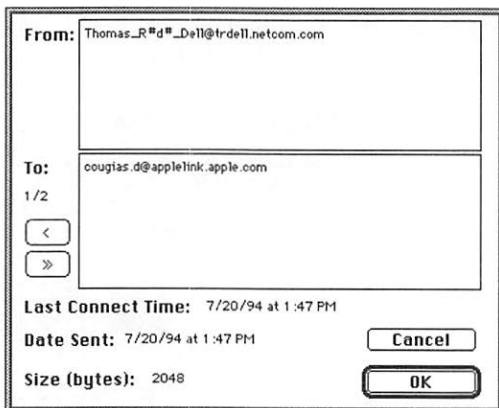
QUEUED MESSAGES

Mail*Link gives you greater access to queued messages than most other Quick-Mail gateways. To see queued messages, access the **General** pane in the Preferences window and click the **Display** button. When queued messages are present, this generates a window and the following choices.



Queued Messages Window

- **Get Info:** In addition to the sender and recipient information seen in the main window, this button provides individual files' sizes in bytes.



Get Info Window

- **Delete:** This button removes the message from the queue.
- **Return:** This button sends the message back to its originator.
- **Finished:** This button returns you to the Preferences window.




Each message is also accompanied by the icon with an upward arrow. This icon denotes mail that has been forwarded to Mail*Link Remote from QuickMail and is waiting to go further. The icon with a downward arrow denotes mail that has come in from a remote system but has not been passed on to QuickMail yet. The icon with an upward arrow and an “x” through it denotes mail that has been subject to a problem in delivery.

CHAPTER 4: EXCHANGING DOCUMENTS ON THE SMALL LAN WITH POWERTALK

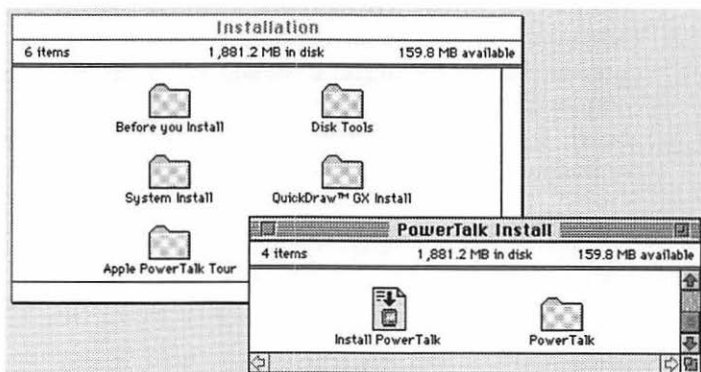
Beginning with System 7 Pro, Apple has provided a useful way for networked Macintoshes to exchange electronic documents with each other as part of the features set in PowerTalk. Although it is robust enough to be used in a client/server fashion—an implementation that involves specially designated PowerShare server Macintoshes—its immediate value to the small office LAN is its ability to be used peer to peer. In this section, we will talk about how your small LAN's users can employ PowerTalk to exchange e-mail and set up contact catalogs. We'll also look briefly at how those users can extend PowerTalk to communicate with commercial online services, the Internet, fax machines, and pagers.

POWERTALK ON THE DESKTOP



A Scene from Apple's PowerTalk Tutorial


Apple began giving away PowerTalk with System 7.5. It is installed separately from System 7.5x, so you will have to go back and find the system installer diskettes to put it on your users' Macintoshes.



PowerTalk—Included in, but Separate from, System 7.5x

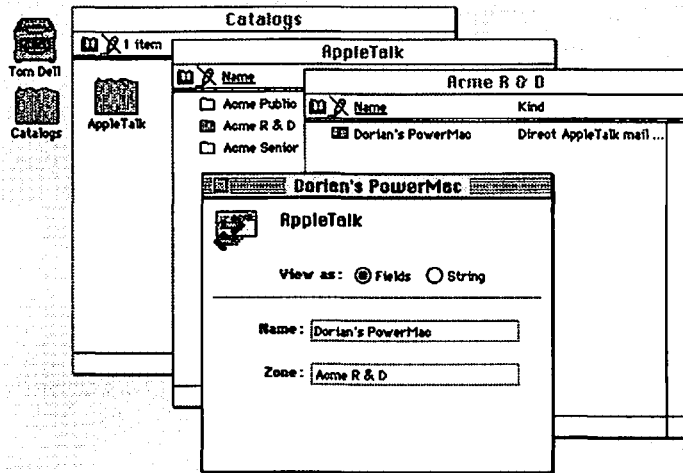
It is even easier to install than System 7.5. You don't have to worry about **Easy** or **Custom** in the Installation window. Your choices are to do it (**Install**) or not to do it (**Quit**). Alas, if only the rest of life were that straightforward.

The changes this installation makes to your users' desktops are immediately obvious. An in/out tray icon, called the *mailbox*, appears. This is where users can access their incoming mail and check on the status of their outgoing mail. Also on the desktop is an icon of an open book called *Catalogs*. This is where users can store information about those on the LAN with whom they exchange e-mail. Under the **Apple** menu, the **Mail and Catalogs** item appears. Among the items aliased here is *AppleMail*, a stand-alone application that can be used to create and read e-mail. Let's take a look at each in turn.




Rifling through the Catalogs

There are a few different types of PowerTalk catalogs. The one that appears on the user's desktop we think of as "live." Opening it will first show you the types of PowerTalk services available on the LAN. Opening one of these, such as *AppleTalk*, will show you the Macintoshes partaking of that service.




A Peek in the Desktop Catalog

Opening the desktop catalog causes the item **Catalogs** to appear on the Macintosh's menu bar. Use this menu to create the next type of catalog.



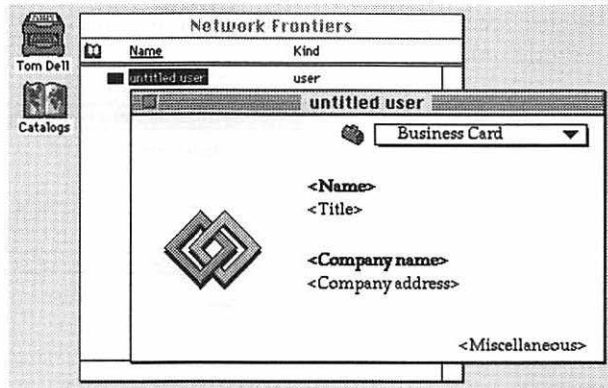
Catalogs Menu Item

The second type is more static. It is called a *personal catalog*. Use the **New Personal Catalog** command to generate one of these, and then give it a name indicative of the contacts it will contain, for example "Network Frontiers."



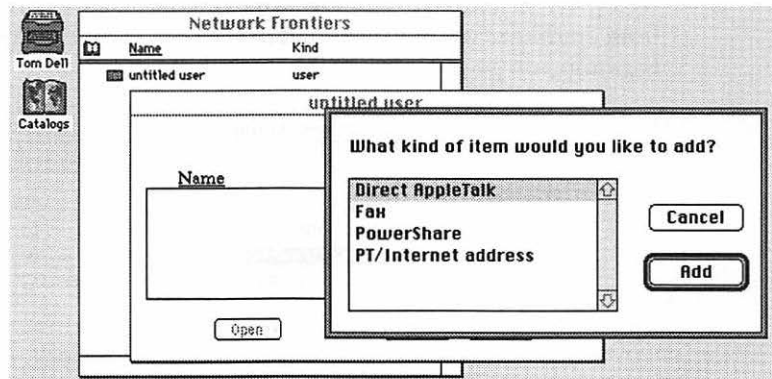
New Personal Catalog

Within this personal catalog, you can store contact information for people both *on* and *beyond* your network. To begin, open the new catalog and choose **New User** from the **Catalogs** menu.



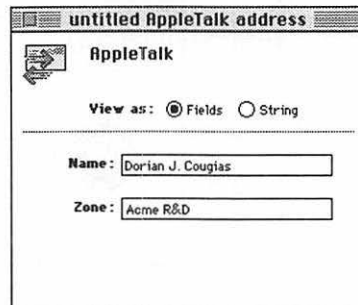
New Information Card

This generates a generic information card with four parts accessed through a pop-up menu. The **Business Card** part is just that, complete with a place for company logos or the user's picture. Click on the pop-up menu in the upper right to switch to **Personal Info**, where you can add the user's home address and any peculiarities you have noticed. Switch to **Phone Numbers** to add, well, phone numbers. Switch to **Electronic Addresses** and click the **Add** button to record the address or addresses at which the user can receive e-mail.



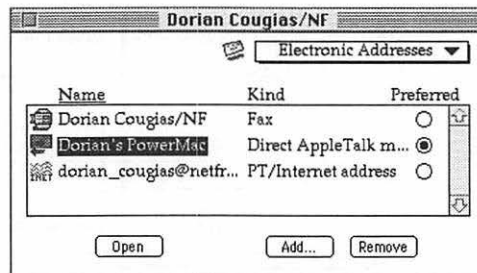
Adding Electronic Addresses

If your PowerTalk is the “stock” version, you will see **Direct AppleTalk** and **PowerShare**. PowerShare was a store-and-forward server mechanism Apple no longer sells or supports. For our small network peer-to-peer implementation, Direct AppleTalk is all we care about, and the e-mail addresses here will be the same as those seen in the desktop catalog. These are other PowerTalk users who are currently online.



Basic User LAN Address

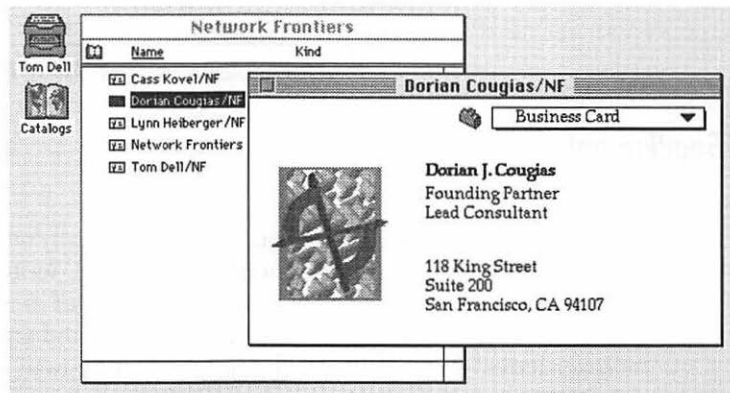
Note that there are a few other choices in the next illustration, which was taken from my PowerBook. These are provided by third-party gateways. With the PowerFax gateway from STF Technologies and my fax modem I can send electronic data to fax machines. I have a CompuServe gateway that lets me exchange mail with that commercial online service, and a StarNine Technologies Mail*Link gateway that interfaces with our LAN’s CE Software QuickMail server (not shown). My most frequently used add-on is the Apple PowerTalk Personal Internet Gateway, previously called the Mail*Link Internet Gateway for PowerTalk, available from StarNine. We show you that in the section called “PowerTalk to the World” beginning on page 151.



Multiple Electronic Addresses

Which of the electronic addresses PowerTalk will use to contact an individual is determined by selecting the **Preferred** radio button and can be changed any time. Look at the previous illustration, for example. When I'm in the San Francisco office, I can communicate with Dorian through "Direct AppleTalk." When I'm on the road, I can still reach him via "PT/Internet"—whether he likes it or not.

Create information cards for all the PowerTalk users on your LAN. You can get pretty creative with these. I have put the cards illustrated in this chapter on the CD-ROM that comes with this book. That way, you have a few samples and you can send us your comments on the book, if you want. Now, be nice!




Dorian's Information Card

Distribute the information cards you create to the rest of your LAN's users, and you're in the e-mail business!

Here's a good way to do this. Create a shared folder called "Install PowerTalk" on one of the networked Macintoshes. You can put everything the user will need to get up and running in this folder. To start, if PowerTalk is not already installed on your users' computers, create another folder inside the first called "Install This!" Copy the contents of the PowerTalk installer floppies into this folder. Be sure not to rename the resulting folders so that the installer can be used to add PowerTalk over the network.

Next, create another folder within the Install PowerTalk folder called something like "Our Users." Put the information cards you have created in this folder. Finally, create a folder called something like "Learn This!" Here you can put Apple's PowerTalk tutorial, a copy of HyperCard Player (in case the users trashed theirs), and an implementation guide like the one we made for you on the CD-

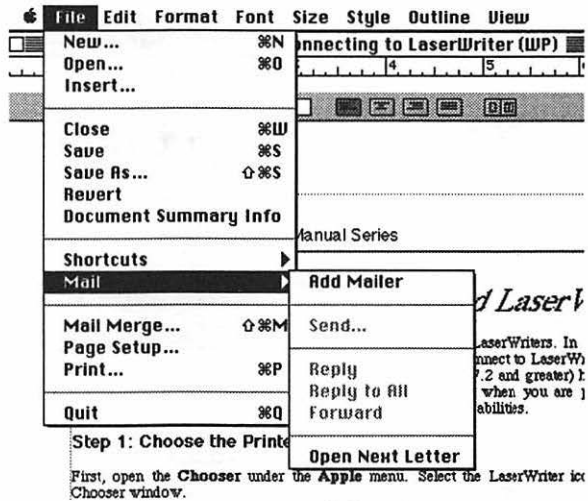
ROM. Now it should all be automatic, except that you'll need to tell people that this shared network volume exists!



PowerTalk Installation Folder with Address Cards

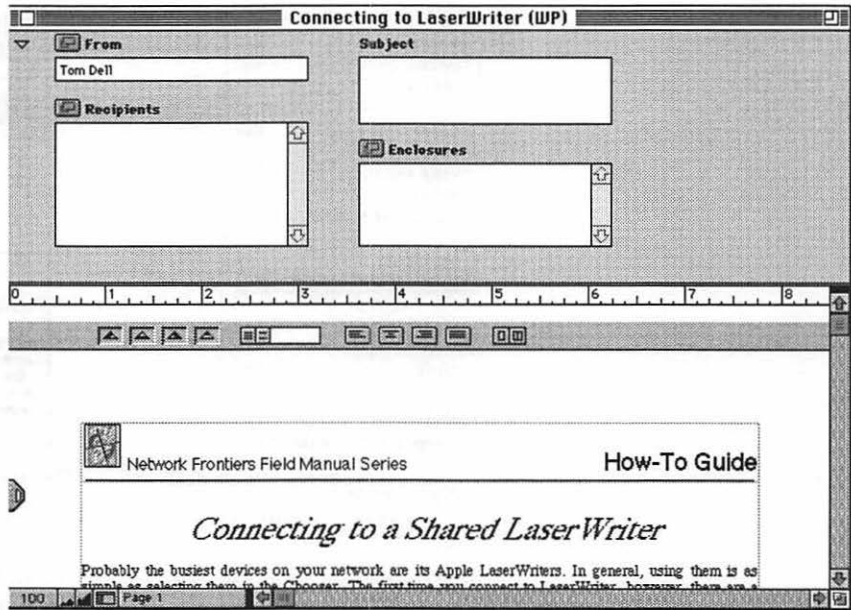
Sending out the Mail

When we teach novice users about e-mail, we usually tell them about its dual character. At times it can be thought of as a postcard. Write down a message, address it, and it's off. At other times it can be thought of as an envelope. Enclose a file created in some application or other, address it, and send it off with or without a "cover letter." PowerTalk can be used to send e-mail either with the aid of the AppleMail application or directly from PowerTalk-savvy applications.



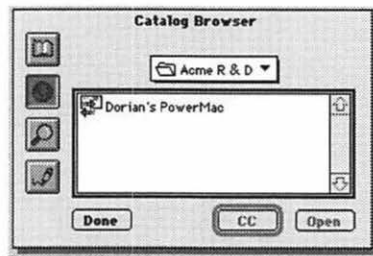
Mailing a Document Directly from ClarisWorks

In the following illustration, we show one of the user How-To Guides we made for the CD-ROMs that come with our books. Choosing **Mail** from the **File** menu in PowerTalk-savvy ClarisWorks lets you generate a *Mailer*. Click the **Recipients** button to open a catalog.



Document with Mailer Added


The Catalog window can be used to search through various catalogs. You can search for specific names as stored in information cards, or you can type in an address for recipients for whom you don't have an information card.



Finding Recipients

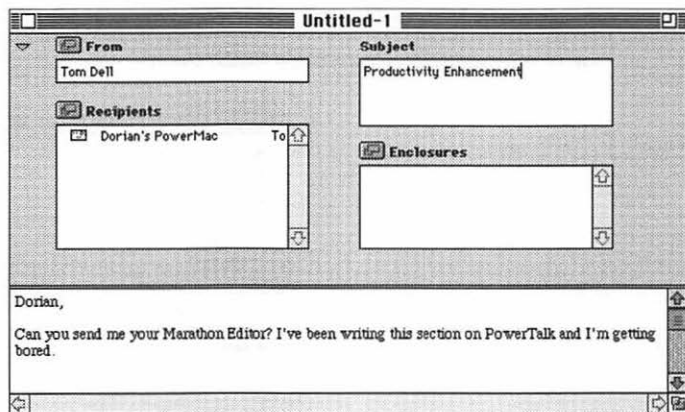
Once you have selected a recipient and typed in a subject, choose **Send** from the **Mail** menu. This generates a window in which you can choose the format your e-mail will adhere to in the **Send as** pop-up menu. If the recipient has ClarisWorks, you can send the document in its native format. If the user doesn't have ClarisWorks, the document can go out in AppleMail format or as a snapshot. Users can

edit text, graphics, sounds, and movies in the AppleMail format, while snapshot gives them a non-editable picture. But that's not all—Claris' file filters can be accessed here so the e-mail could go out as, say, Microsoft Word or WordPerfect!



Choosing the E-Mail's Format

This is a good example of e-mail-as-envelope. The ClarisWorks document goes out as an *enclosure*. To see an example of e-mail-as-postcard, choose **Mail and Catalogs** from the **Apple** menu and then select **AppleMail**.

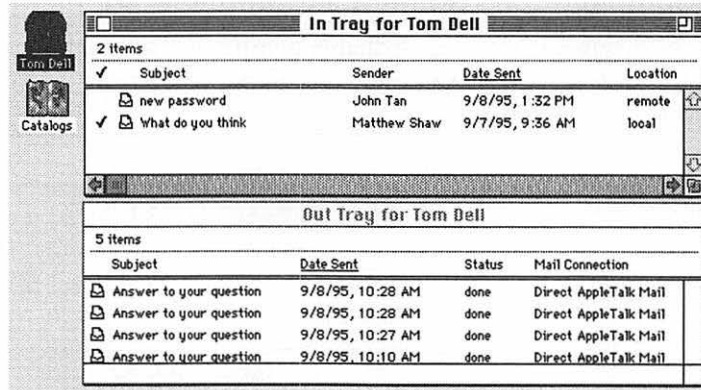


Quick AppleMail Message

Click the **Enclosures** button to select the files on your hard drive that you wish to send to send someone else. Your e-mail message becomes a sort of cover letter.

Checking the Mailbox

Users can view mail sent to them by double-clicking on the mailbox icon and then double-clicking on a piece of e-mail therein to launch AppleMail.

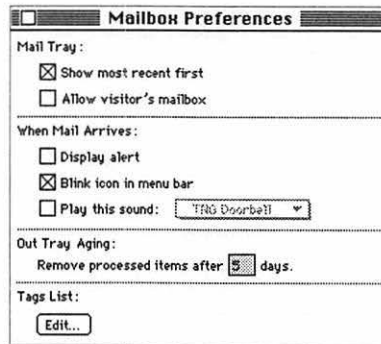


Open Mailbox

When the mailbox is open, a **Mailbox** menu bar item appears.



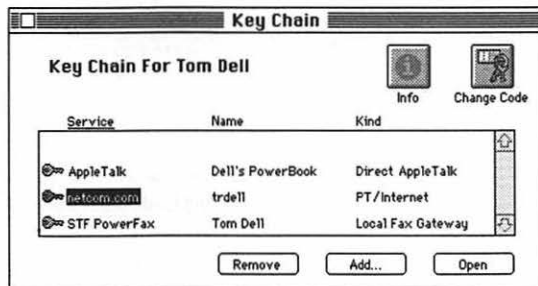
By selecting **Preferences** from the Mailbox menu, users can establish how they'll be notified when new mail arrives: an alert dialog box, a flashing menu bar icon, or maybe a really annoying sound! Neat!



Setting Notification Options

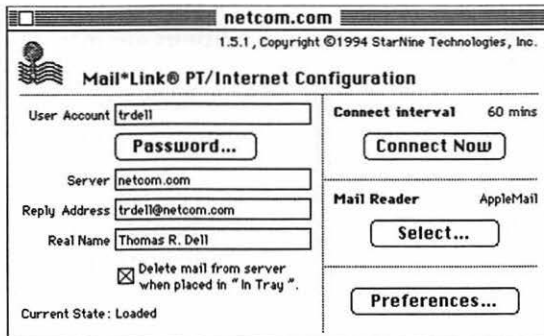
POWERTALK TO THE WORLD

On larger LANs equipped with an e-mail server or direct Internet access, the easiest way to give users Internet mail capabilities is through a special mail system-to-Internet gateway. For users with a desktop modem and an Internet dial-up account, however, this is not a necessity. These users can receive their Internet mail from all over the world delivered to the same place where they find their LAN-based e-mail, the PowerTalk mailbox. This is made possible by the Apple Personal Internet Gateway, previously called the Mail*Link Internet Gateway for PowerTalk, from StarNine.



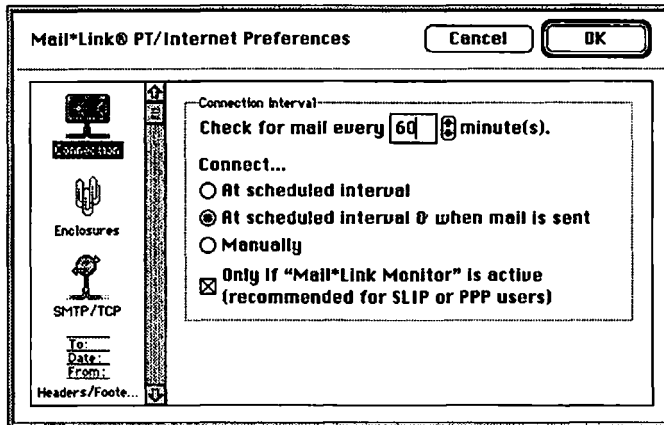
PowerTalk Key Chain with Internet Gateway Installed

This gateway is easy to configure once installed. It can be accessed by selecting the PowerTalk **Key Chain** item in **Mail and Catalogs** from the **Apple** menu.



Entering Mail Account Information


In the main window, enter the user's e-mail account name, the password, and the name of a Post Office Protocol (POP)-compatible mail host such as the Apple Internet Mail Server. Also enter the user's e-mail address and real name. Next, select the **Delete mail from server...** checkbox. Most access providers charge you for mail that piles up in their machines.



Enabling Connection Interval in Preferences

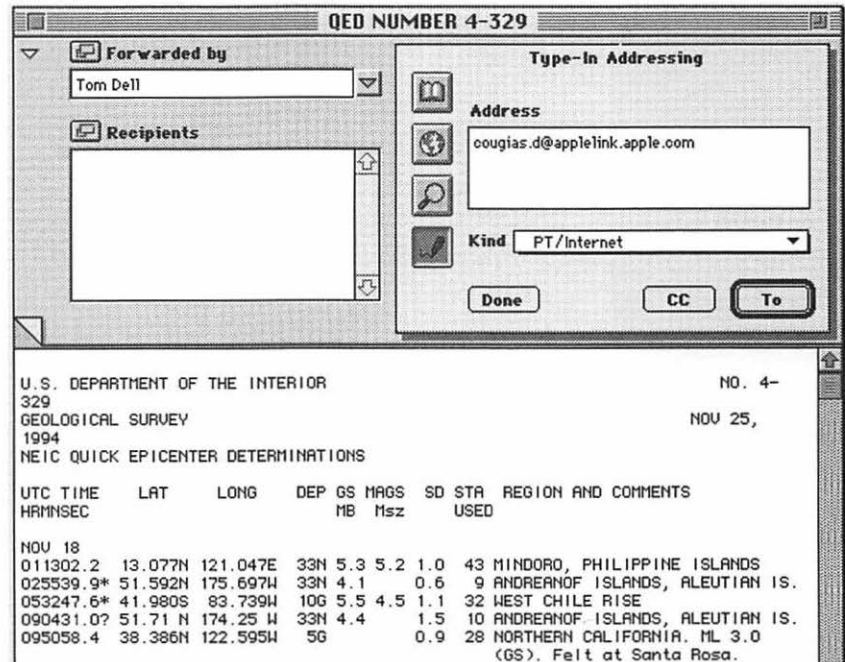
Finally, click the **Preferences...** button and select the **Connection** icon. There is a choice to permit connections only if Mail*Link Monitor is active. This is a useful utility that shows you the status of your mail host connections. I suggest you use it. Also make a choice about when the user's gateway should connect. This will depend upon the urgency of the user's mail and how often the user receives mail. If you are connecting via a modem, the conditions you establish here will cause MacPPP or MacSLIP to be launched automatically and a dial-up connection made whenever these conditions are met. On a LAN, these criteria determine when PowerTalk will use the MacTCP or TCP/IP control panel to make a connection. It probably won't be necessary, but if you want to mess with any of the other settings, consult StarNine's extensive documentation.

With the gateway in place, the user is insulated from most of the nastiness of Internet mail. Mail labeled as **remote** in the Location column of the Mailbox's in tray has not yet been read or downloaded, while mail labeled **local** has.



Internet Mail in the In Tray


Sending mail is also simple, permitting the user to access addresses previously recorded in a catalog or to enter them “on the fly.”



Forwarding Internet Mail to an Address Entered “On the Fly”

Unless the user is too busy to pay attention to such things, we recommend setting the connection interval to **Manually** in the Preferences window and requiring Mail*Link Monitor to be active. Make an alias of Mail*Link Monitor and put it in the Mail and Catalogs folder that appears under System 7.5’s **Apple** menu.

Tell users to launch this application when they want to send Internet mail. Once Mail*Link Monitor is running, they can initiate a connection with the **Connect now** choice from the **Administration** menu.



*Watching the Connection with Mail*Link Monitor*

This enables users to see how the connection proceeds. If there is a problem, they can tell you at what point the transfer broke down. By having users only connect manually, you make sure that Internet connections are not being established and charges are not being incurred when a user is not around.

CHAPTER 5: USING THE APPLE INTERNET MAIL SERVER AS MAIL HOST

The Apple Internet Mail Server, formerly known by its shareware name of MailShare, is a Macintosh-based Post Office Protocol (POP) and Simple Mail Transfer Protocol (SMTP) host created by Glenn Anderson. This application is simple to set up and use. It permits you to pluck mail from the Internet that is destined for your users, as well as letting them send mail. This application gives you the abilities of a mail host without forcing you to have a UNIX machine of your own or an Internet service provider (ISP) with POP3 services.

On the client side, Apple Internet Mail Server works nicely with Eudora (from Qualcomm, Inc.), which comes in both Macintosh and Windows versions, and with the shareware application NotifyMail (by Scott Gruby). It can also be used with Netscape Navigator and with the Apple PowerTalk Personal Internet Gateway we talk about in another section of this book called "PowerTalk to the World," beginning on page 151. But that's not all! Apple Internet Mail Server can be used with an SMTP gateway, such as StarNine's Mail*Link, to deliver mail from your LAN's existing e-mail system, and it has some ability to act as a list server as well.

The Apple Internet Mail Server requires a Mac Plus or better with System 7.x and MacTCP 1.1.1 or later. We suggest you run it on at least a 68030 machine with System 7.5.1 and MacTCP 2.0.6 or later.

INSTALLING THE APPLE INTERNET MAIL SERVER

This application installs in mere minutes—provided you have done some planning ahead of time.

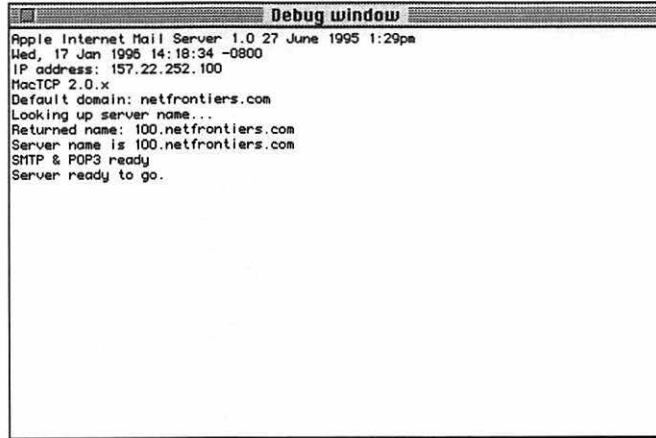
DNS Is Your Friend

First of all, to use the Apple Internet Mail Server as your network's mail host, you need to fall back on one of the Internet's most important support systems: the *Domain Naming System (DNS)*. When you send e-mail over the Internet, you send it to both a user and a mail host. This address is expressed as a domain name such as *tom_dell@netfrontiers.com*. The user is "tom_dell." The mail host is "netfrontiers.com."

That is a straightforward address format for people, but it is not so great for computers. A computer sending an e-mail message to such an address must first *resolve*, or look up, a unique machine number associated with the would-be recipient's mail host. In this case, "netfrontiers.com" is associated with, or mapped to, the IP address 157.22.252.75. To make this determination a mail host must contact a *name server*.

When setting up a mail host, make sure that information is added to the domain naming system that properly reflects your installation. For most of us, this means nothing more than telling your ISP that you want to dedicate a given machine's IP address to your domain's mail service. The ISP can update the DNS.

When you first launch Apple Internet Mail Server, it generates a Debug window. During the launch, it does a reverse domain name lookup for the address of the machine on which it is running. You will see this in the Debug window. It should match the information you have exchanged with your ISP.

*Debug Window*

If you see something in the Debug window similar to what is pictured here, you are ready to proceed.

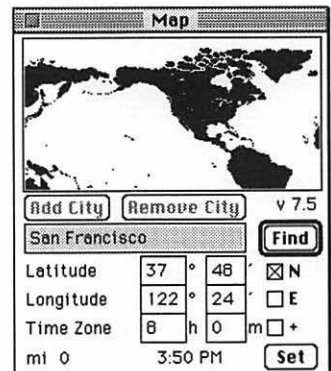
Configuring the Mail Host

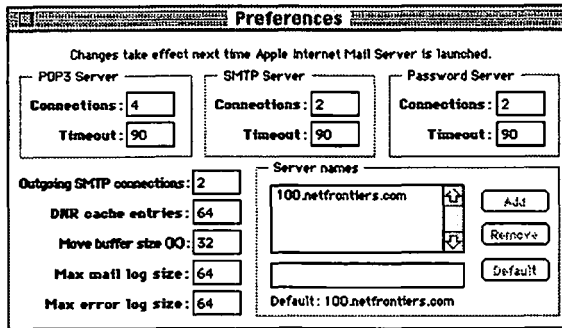
Date and Time

The Apple Internet Mail Server retrieves its date and time information from the Date & Time control panel. It finds its time zone information from the Map control panel. Make sure these are correctly set on the mail host Macintosh before you start sending mail.

Preferences

Choose **Preferences** from the **Server** menu.

*Verifying Time Zone Information*



Preferences Window

The address that the Apple Internet Mail Server found when it did its reverse domain lookup will be seen in the **Server names** field. This is the domain name that will be used for your site’s Internet mail. In the illustration above, it is “100.netfrontiers.com.”

That’s fine, but let’s say you also want to receive mail for “user@train.netfrontiers.com.” In that case, you need to type “train.netfrontiers.com” in the lower pane and click the **Add** button to put it in the list. Doing this tells the mail server to accept messages for users in this domain as well. However, your ISP must add the additional domain name mapping to your DNS for such mail to be properly routed to your mail server. As long as this is the case, your single Apple Internet Mail Server can be used to service multiple e-mail addresses. It will identify itself to other mail hosts by whatever domain appears in the **Default** field.



Be aware of this: The user name space spans domains. That means that any account you create for one domain will also be valid for another domain served by the same copy of the application. You might have a “netfrontiers.com” and a “train.netfrontiers.com” as separate domains, but “tom_dell@netfrontiers.com” is equal to “tom_dell@train.netfrontiers.com.” Mail addressed either way will end up in my account if the same application is handling both accounts.

Now take a look at the **POP3 Server**, **SMTP Server**, and **Password Server** panes. For each of these services you can define an upper limit on the number of simultaneous connections allowed, within MacTCP’s limit of 64 (Open Transport does not have this limitation). Upon startup, Apple Internet Mail Server grabs one connection for domain name resolution and another for sending finger (NotifyMail) requests. That makes 62 connections that can be distributed among the POP,

incoming and outgoing SMTP, and password services, along with any other TCP/IP applications you might have running on the mail host Macintosh.

The maximum number of allowable outgoing SMTP connections is set beneath these panes in the **Outgoing SMTP connections** field. If you notice that the outgoing mail queue is often occupied by more than a few messages, increase the number in this field. If you have a lot of e-mail traffic, you might need to increase the other allocations as well.

Personally, I like to set **Password Server Connections** to zero. All these connections do is allow users with clients that support password services to change their passwords. I like to assign passwords because I know I can make up *good* ones! Setting this to zero denies users the ability to change the passwords I assign.

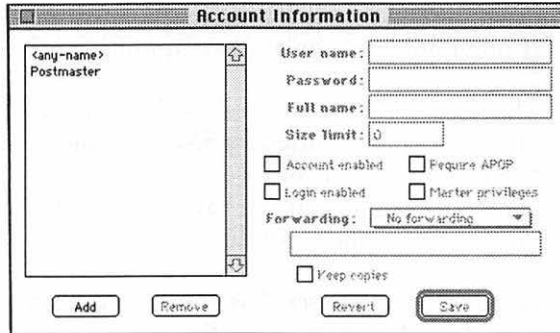
If the Apple Internet Mail Server crashes after you have made these changes, give it more memory by selecting **Get Info** from the **File** menu and jacking up the number in the **Preferred size** field.

If you have users dialing into their accounts over modems, or if you have an overburdened network, you will probably want to increase the number in the **Timeout** fields from 90 seconds to something like 180 or 240 seconds.

There is only one other field worth noting in this window: **DNR cache entries**. The Apple Internet Mail Server has a built-in domain name resolver that replaces the one in MacTCP. It keeps a cache of its successful lookups, which increases performance and reduces the load on the name server. By default, it remembers the last 64 lookups. Increase this number if you like, but note that each entry requires about 280 bytes of additional application memory.

User Accounts

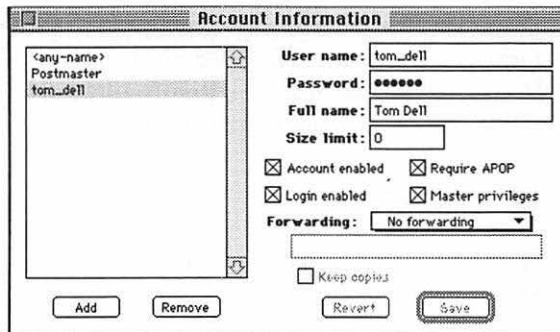
Now you can add your users. To do this, select **Account Information** from the **Server** menu.



Initial User Accounts

Notice that before you have configured anything, you already have two default accounts: `<any-name>` and `Postmaster`. The `Postmaster` account receives error messages and user reports of problems. This account is *required* for all mail hosts by RFC 822, the Internet *Standard for the Format of ARPA Internet Text Messages*. It is most likely your job to regularly read the mail sent to this account. The `<any-name>` account is a catchall. It accepts mail for any user name not defined in your domain's user list. It ensures that you do not lose mail because of bad addressing.

The next step should be to add *your* account, as you are presumably the person who will be managing the system. To do this, click the **Add** button.

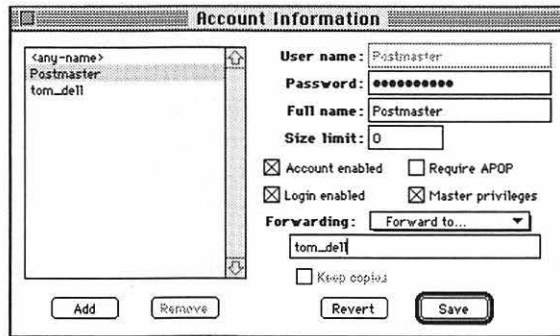


Setting up Your Account

Next, type in an account name in the **User name** field and a good five-digit alphanumeric nonsense password in the **Password** field. Type your real name in the **Full name** field. We strongly recommend that you also select **Require APOP**, or *Authenticated Post Office Protocol*. This ensures that your password travels across

the network in an encrypted format, not as clear text. You can choose **Master privileges** to permit remote administration, but this option does not work yet (version 1.1)! Bug Apple about it. Click **Save** when you are finished.

Now, here is a nice trick: Go back and open the Postmaster account. Click the **Forwarding** pop-up menu and select **Forward to...** Then type your account name in the field beneath it.



Forwarding Mail

Do the same with the <any-name> account. Now all your administration-related mail will come to the same mailbox. This is also a nice trick if you are traveling. For instance, when I am on the road I forward my “tom_dell@netfrontiers.com” mail to “trdell@aol.com.” That way, I can receive my mail through the America Online commercial service, a local call from most places in the metropolitan U.S.

Go ahead and add the accounts for the rest of your users. If you are worried about them hogging too much space on your mail host’s hard drive, you can type a number in the **Size limit** field. Any mail coming in that would push a user’s account over the limit will be rejected. I do not recommend trying to enforce account tidiness in this way—or at least not for your boss’s account.

Before we move on, let’s go back and look at the forwarding options.



Forwarding Options

- **No forwarding** Incoming messages are spooled in a mail drop and retrieved by the user with a POP mail client.
- **Forward to...** Incoming messages are sent on to another e-mail account.
- **Save as archive...** Incoming mail goes into a UNIX mail format text file. This is the same format Eudora uses to store its mail. If you place an archive in the Eudora folder in the System Folder of a machine that has Eudora installed on it, the e-mail client should recognize it as an additional mailbox the next time it is launched.
- **NotifyMail to...** NotifyMail is an extension that listens for finger requests. When it hears one, it sends an AppleEvent to Eudora that checks the user's mail. If this option is chosen, Apple Internet Mail Server sends a finger request to the specified domain name.
- **NotifyMail to last IP** Like the option just mentioned, this one tells Apple Internet Mail Server to send the finger request to the last IP address from which a user checked his or her mail.
- **Mailing list...** This is a nifty option if you do not have good list server software like StarNine's ListSTAR. You can create a simple mailing list in which a message sent to a single address is redistributed to multiple recipients. The mailing list is in the form of a plain text file with one address per line. Each line is in the format "user@host (optional comment)." An example of such a mailing list might be:

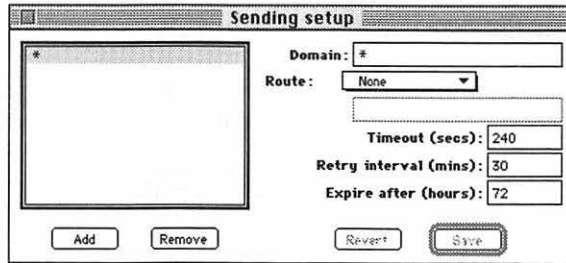
```
tom_dell (local user)
trdell@netcom.com (Internet user)
dorian (local user)
tom_hessel@eworld.com (Internet user)
```

If you supply only a file name, Apple Internet Mail Server looks in the Mail folder of the System Folder for this file. If you want it to look somewhere else, provide a full path name.

- **Save as files...** Incoming messages are spooled to a folder that you designate as separate text files with the line feeds deleted. If you do not specify a folder, the Apple Internet Mail Server saves the files in its own folder. This is a handy feature if you want to use the text files with, say, a Web server script.


Sending via Other Hosts

The last setup parameter, which you may not need to worry about, is for routing mail through other mail hosts. Select **Sending setup** from the **Server** menu.



Sending Setup Window

You may want to use this window if you need to handle outgoing BITNET addresses that end in *.BITNET*. There are still a few LISTSERV addresses out there that use this format. To accommodate these, you can direct the Apple Internet Mail Server to reroute addresses that end in *.BITNET* by using the wildcard character "*" matching any domain plus BITNET in the **Domain** field, as illustrated below. Next, choose **Via host...** in the **Route** pop-up menu and type in the name of a mail server that is connected to both the Internet and BITNET.



Setting Alternate Sending Option

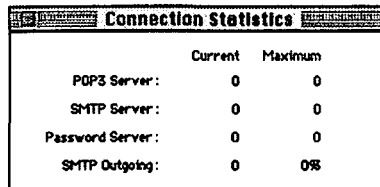
Click **Save** when you are finished.

This would also work if, for some reason, you wanted the Apple Internet Mail Server to act solely as a mail relay, forwarding all mail to another machine, such as your ISP's mail host, for actual delivery. In this case use the wildcard character "*" to match any domain, and use the **Route** fields to establish the address of the sending host.

You could choose **Save as files** from the **Route** pop-up menu to save the mail that matches a particular criterion. Then you could use a batch script to deliver all the queued mail for a specific destination at a given time of day. This is handy if you have dial-on-demand Internet service rather than a full-time connection.

Connection Statistics

Once you have all this stuff set up, keep an eye on how the server is acting by choosing **Connection statistics** from the **Server** menu.



	Current	Maximum
POP3 Server:	0	0
SMTP Server:	0	0
Password Server:	0	0
SMTP Outgoing:	0	0%


Connection Statistics

The current number of connections and the maximum number of simultaneous connections allowed for each service are displayed in the Connection Statistics window. If you find the maximum number for any service is equal to the maximum number of allowed connections you have set, try increasing the values. Ideally, you want to be able to support more connections than are necessarily called upon so that no user ever gets a “server is busy” message.

A Scary Thing Back up the hard drive of your mail host Macintosh immediately after setting this up. If your Apple Internet Mail Server Preference file becomes corrupted, you will lose your work—unless you are real handy with ResEdit.

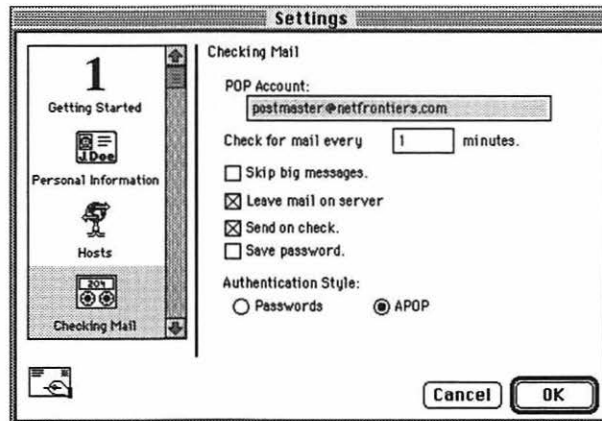
SETTING UP YOUR USERS

Setting up your users to access the Apple Internet Mail Server is not difficult, but there is something you should know about. If you have selected the **Require APOP** checkbox when setting up your users' accounts, they will need an APOP-compliant mail client application to receive their mail. Eudora, written by Steve Dorner of Qualcomm and available all over the Internet, is one such e-mail client.



Eudora Password Dialog Box

If you have done this, be sure to tell your users that they must enable APOP support in their mail client. With Eudora, this is accomplished through the **Checking Mail** pane of the Settings window.



Enabling APOP


CHAPTER 6: NETWORKING ART WITH A CUMULUS DATABASE

Canto Software's Cumulus is a client/server database for digital art images. Its sole function is to store and categorize graphics files and QuickTime movies on a server for distribution to Macintosh network clients. It does this with impressive speed, providing an extensive number of details about each image and a great search engine.

You might think this type of specialized database would be useful only to an organization dealing in graphic arts with *a lot* of images. We are not so sure. It seems that many of our clients who have no direct connection to the graphic arts industry still have a large amount of company artwork that would be well suited to this service. An example that comes immediately to mind is one of our clients involved in real estate brokerage. The CIO is learning Photoshop. Why? Because, she explained, with the desktop publishing revolution, it is not enough to just create good work, it must also *look* like good work. That means semi-professional graphic design and, in many cases, art. The typewriter doesn't cut it any more.

Cumulus does not actually store art in its image databases, it stores *references*; that is, it stores information about a given graphic file along with a pointer to wherever it is physically located. The logic in this is that even though graphics files are among the largest files, you don't end up with the mother of all database files—each Cumulus record takes up only about 5K of hard drive space. The beauty of this, in a networking sense, is that the art can reside anywhere, not just on one

server. The art can even be on offline media, such as those cheap CD-ROMs you can buy that come crammed full of stock art. Users need only select the records with which they wish to work, and Cumulus gathers the art from wherever it is located! If it's from offline media, users will be asked to insert the disk in a drive on the *client* Macintosh.



Cumulus's Client/Server Art Database Catalog System

Overall, Cumulus works in much the same manner as other client/server databases, but with some important distinctions. The *Cumulus* application has the functionality of both server and client. The *Cumulus•Client* application is only that, a client. The full version of the Cumulus application can be a server to *Cumulus•Client*-equipped computers when it opens a database for their access. It can then change hats and be a client to a database shared from another machine. Besides differing in purpose, Cumulus and *Cumulus•Client* differ in function. Cumulus is used to *administer* databases, allowing you to enter art, organize images, assign privileges, and perform maintenance. *Cumulus•Client* *edits* databases, allowing you to categorize and retrieve images.

In the following pages, we show you one way to set up the Cumulus system. Our goal is simplicity, so we will be adhering to more centralization than is required. We will be talking about a solution for the average office environment, not a graphic arts firm. Nevertheless, the principles we apply are fairly universal.

WORKING WITH THE CUMULUS SERVER

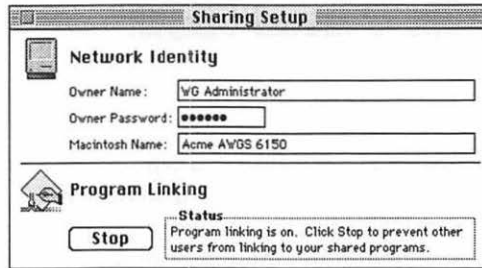
The first thing to consider when setting up the Cumulus server has little to do with Cumulus itself. You need to decide which machine will host your graphics database. If it is to be a Macintosh, it must be a color-capable computer—640 by 400 pixels—running System 7 or later. If you anticipate database usage will be light, you could get by on a non-dedicated machine: that is, not specifically a server. In the example we present here, we add Cumulus to the services available through an AppleShare server. This will be the Apple Workgroup Server 6150 used by our fictitious company, Acme Rubber Chicken, Ltd.

Most planning issues are the same as those we have already addressed in this book for other services. For instance, taking traffic into consideration, where should the database be placed on the network? Remember, Cumulus is primarily a catalog of art *references*, so even though art files can be huge, the traffic generated by clients browsing the server is no greater than it would be for another type of database.

Once you have decided on the would-be Cumulus machine, gather together the art that you plan to categorize. That's right: company logos, reusable advertising art, symbols, maps—everything except the JPEG smut you've been pulling down from that certain Web site after hours. For ease of administration, put the graphics in a folder on the designated Macintosh and name it something logical, such as "Acme Art Gallery."

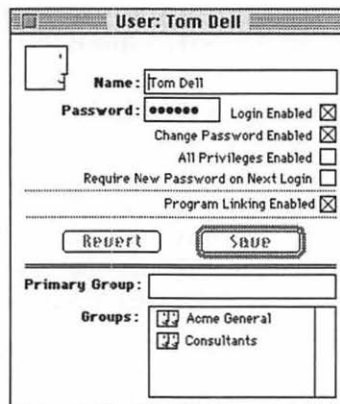
Next, make sure you have user accounts set up for all the users on your network who are to have access to the database. If the database is on a Macintosh using Personal File Sharing, do this through the Users & Groups control panel. If you are using an AppleShare server, use AppleShare Admin or, better yet, Server Manager.

The ability of the Cumulus database to share its contents depends on the Macintosh's Program Linking capability. Make sure this is enabled in the server's Sharing Setup control panel.



Enabling Program Linking

Also, ensure that each user can use the database by selecting the checkboxes that allow Program Linking in their account windows. (For more information on this sort of thing, see our book *Managing AppleShare & Workgroup Servers.*)



Program Linking Enabled

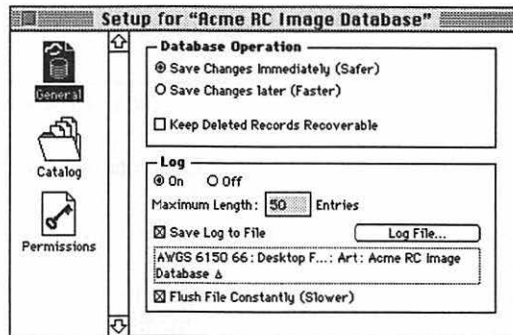
Now you are ready for the Cumulus server itself. Copy its folder from the CD-ROM to your server's hard drive. In the Extensions subfolder you will find system software that may or may not be installed already on the server machine. Drag those elements that are not already installed into the System Folder and restart.

When the server comes back up, launch Cumulus and register it when prompted. Print and send in the registration to be eligible for Canto's tech support.

Click **Cancel** when prompted to supply the location of a database file. At this point, choose **New** from the **File** menu to create a database file. Name it some-

thing that makes sense, such as “Acme RC Image Database,” and an empty database window appears.

Before adding any art to this database, set a few preferences. First make sure that **Allow Sharing** is selected under the **Database** menu bar item. If it is not, the database will be accessible only locally. Next, choose **Database Setup** from the **Database** menu.

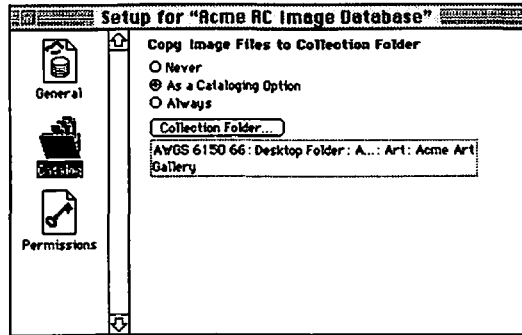


General Preferences

In **General** preferences, direct Cumulus to save changes to this database immediately (while the database is open) or later (when the database is closed). We recommend the former as it is safer, mostly because the shared database may be closed infrequently if it’s on a dedicated server. If you do choose the latter, you won’t be totally out of luck if the server crashes, however. Under this setting, data such as the indices isn’t written immediately to the hard disk, but records are.

We recommend that you do not select the **Keep Deleted Records Recoverable** checkbox, which is selected by default. When this is not selected, Cumulus writes new data over the areas once occupied by deleted files. Otherwise, it sort of hides records rather than deleting them so that they can be “recovered,” or made visible again. We figure it is more worthwhile to recatalog a graphic that has had its record deleted accidentally than it is to deal with the hard disk fragmentation that this type of failsafe incurs.

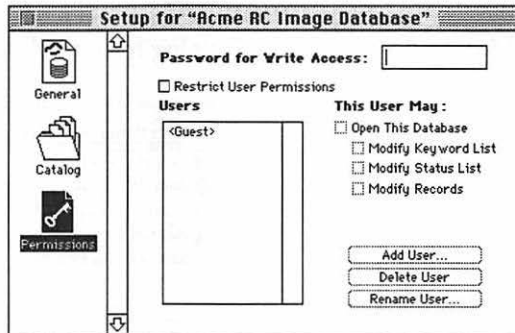
As with other server-based systems, we strongly recommend that you keep a log of Cumulus activity. Select **On** in the **Log** pane, and select **Save Log to File** to create and select a file to which tab-delimited entries will be written whenever the database is closed. If you are going to have the database open most of the time, select **Flush File Constantly (Slower)** so that the log is kept up to date.



Catalog Preferences


In **Catalog** preferences, click **Collection Folder...** to select the folder we created earlier, into which you have already dumped all your company's art files. In our example, this is the Acme Art Gallery folder.

Above this button is one of Cumulus's neatest features. If you choose the **Always** button, Cumulus copies a graphic file that has been referenced in its database from any source on the network to the folder you designated. This is a great way to centralize your art, assuming your server has sufficient hard drive space. Users then need not worry about keeping cataloged images on their hard drives. If you choose **Never**, the graphic file is left where it was originally found, and where it is expected to be found when a user tries to access it from the catalog. This is neat in that it lets you keep art all over the network and not just on the server, but it's a pain in that if the source of a cataloged entry is unavailable (for example, Lou took home the PowerBook), so is the art. If you make the differences and their implications plainly known to your users, you can probably leave it up to them on a record-by-record basis by selecting **As a Cataloging Option**. We have made you a user How-To Guide to help with these choices and placed it on the CD-ROM that ships with this book.



Permissions Preferences

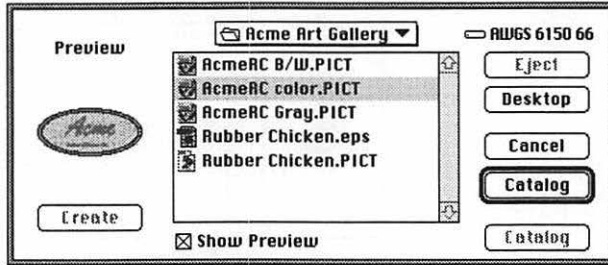
In **Permissions**, you can set read/write access privileges similar to those used by AppleShare and other database programs. If this is the only Cumulus database you will be running, you might find this a bit redundant, as you have set up AppleShare or Personal File Sharing passwords already. Cumulus allows you to have many databases on the same machine, however, and you may want to give some users access to some of them but not others. In this example, we are dealing only with one, the Acme RC Image Database.



Creating Keywords

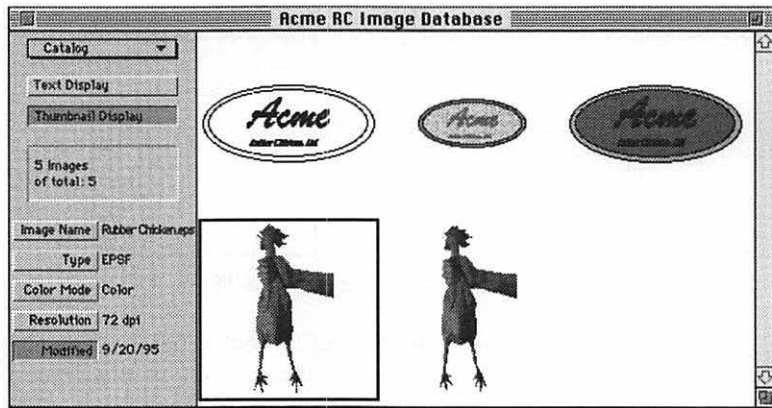
One of Cumulus' real "powers" is that it has an extensive ability to *catalog* your network art. This lets you find the *right* art for the job and find it *fast* through the use of "keywords" that you designate. You could set up all or some of them now by choosing **Show Keywords** from the **Keywords** menu to generate an index window. Next, select **New Keyword** from the same **Keywords** menu and type the word into the highlighted box that appears. In our initial list, we added words for different file formats, such as EPS, PICT, and TIFF, and for advertising media, such as magazines, newspapers, and flyers. We cover more specific contents of keywords further on.

Now you are ready to catalog the art you have gathered together in the collection folder. To do this, choose **Catalog** from the **Record** menu.



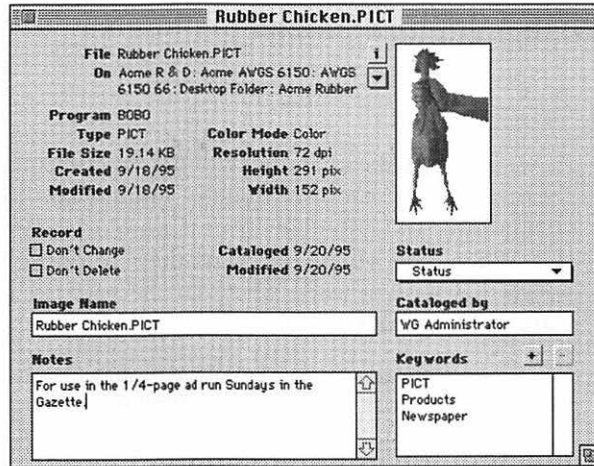
Choosing an Image to Catalog

Note when you open the designated folder that the Open File dialog box has *two* Catalog buttons in it. The bottom Catalog button allows you to bring in the contents of the entire folder. The top Catalog button is for importing individual files only. The images you choose will appear in the database's main window, which, although part of the server, acts like the Cumulus client software. For our example, we will import one picture for now.



Cataloged Art in Database Main Window

Double-click on a *thumbnail* picture to generate a window featuring information for that record only.



Database Record Window

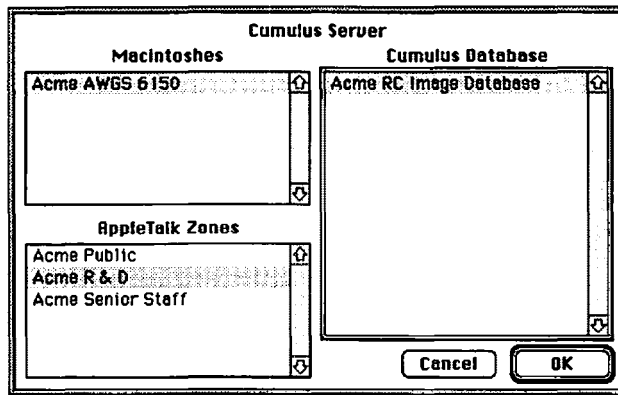
Hey, wait a minute! There were two rubber chickens in that illustration! What's the problem? No problem, my friend. One picture is a low-resolution PICT file for, say, black and white newspaper ads. The other picture is a high-resolution EPS file for, say, full-color magazine ads. The fields in a Cumulus record help you to avoid confusing seemingly like files left "loose" on the network. Records can also be differentiated by their **Notes** and by **Keywords**.

Now that you have this thing up and working, let's take a run out to the workstation end of the system before we come back to working with Cumulus as a whole.

Isn't that a fine rubber chicken we have pictured? You know, when you need a really good rubber chicken—and who among us can honestly say they haven't at one time or another—the best place to buy them is from Oriental Trading Company of Omaha, Nebraska. Give them a call at (800) 327-9678 and they will send you a catalog full of fun stuff. The pictures in those catalogs were organized for production in a Cumulus database, a high-end implementation of the product.

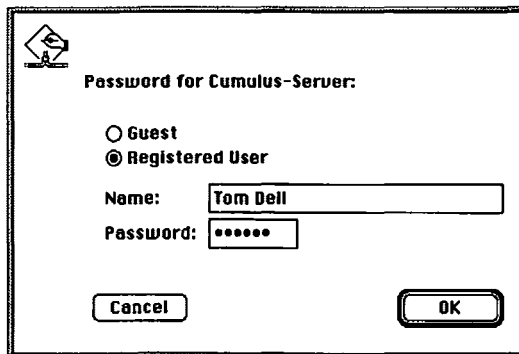
WORKING WITH CUMULUS•CLIENT

As we mentioned earlier, many of the functions that can be performed locally on the image database server through Cumulus can also be performed from the workstation through Cumulus•Client. Cumulus makes use of Program Linking to do this. When users launch the client, they are asked the location of the database with which they wish to work. They can also generate this window by selecting **Share Catalog** from the **File** menu.



Selecting the Shared Database

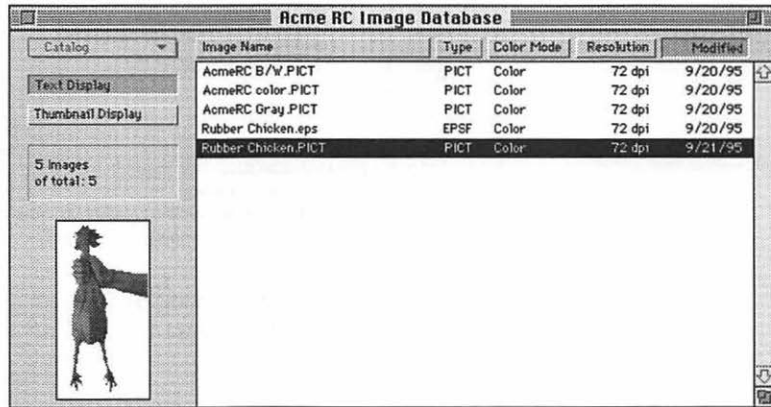
After they highlight the relevant Macintosh, AppleTalk zone, and Cumulus database, users can click **OK**. They will be prompted for their AppleShare or Personal File Sharing user names and passwords, after which they can click **OK** again.



Enter User Name and Password


If you enabled password protection when setting up the Cumulus server, users next need to enter a password to be able to make any changes to the database.

When the database is opened, they have two choices on how to view it. By clicking the **Text Display** button on the left, they see a list of the available art.



Text Display View

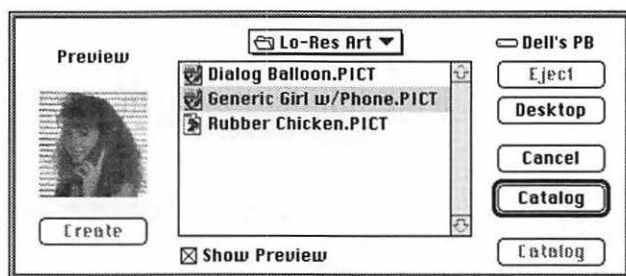
By clicking the **Thumbnail Display** button, they see the art illustrated.



Thumbnail Display View

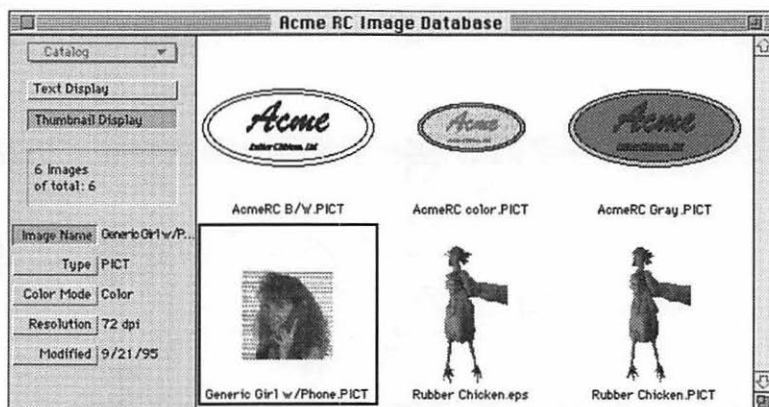
The art can be sorted by the **Image Name**, **Type**, **Color Mode**, **Resolution**, and **Modified** buttons.

Art is added to the database remotely from the Cumulus•Client in the same way it is added locally through the Cumulus server: Users must choose **Catalog** from the **Record** menu.



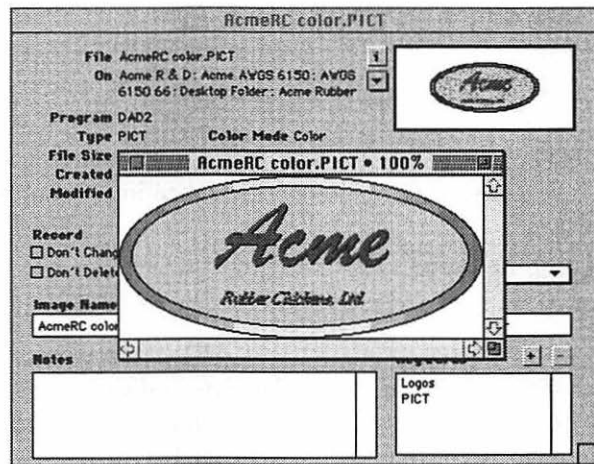
Adding Art to Database from Client

Next, they choose a file they wish to catalog. If they are not sure which file of several it might be, they can click the **Create** button and select the **Show Preview** box to see what it looks like first.



New Art in Database Window

The new record is now visible to all network users. When users need a given file, they can double-click its entry to see its related information, keywords, and notes.



Record with Preview Open

Double-clicking the preview in the upper-right corner of the record pulls up the actual file. Using the **Zoom In/Zoom Out** choices from the **Edit** menu will provide different views. Using **Copy** from the **Edit** menu allows users to pull pictures into any working documents via the Macintosh Clipboard.



Art Copied and Pasted from Cumulus

PLANNING FOR THE CUMULUS SYSTEM

The previous section was meant to familiarize you with Cumulus. Now let's talk about some things to consider when implementing its service on your network.

One Lump or Two?

One of the first questions you should ask yourself is whether you should lump all your art into one image database or distribute the records through multiple databases. The easiest course of action is probably the one we have already shown you: a single database cataloging a single primary server folder full of art.

There are several reasons why multiple databases, perhaps distributed across multiple servers, might be the better way to go. Do you have tons of images? If you do, multiple databases make more sense. The number of records your image database can have is ultimately limited only by memory, but Canto recommends that you not exceed 50,000 records. We feel that the limit is realistically much lower than that, under 10,000. First, can you imagine the number of keywords needed to permit users to quickly find what they want in a database of that size? It would be better to segregate like art into like databases. For instance, one image database might be for "technical" art and another for "marketing" art. This is especially helpful if art can be segregated by department. Why should a marketing person have to slog through the blueprints cataloged by the technical people? Also, you notice performance slowdown when a single database creeps past 10,000 records.

Whether or not you take the next step and spread image databases across multiple servers will depend on the numerous factors already discussed in relation to network design. How powerful are the servers? How big are the databases? How many people are using the databases? How often are they using the databases? These types of questions will point you toward the most efficient deployment of Cumulus. You might need to experiment a bit to find the perfect configuration.

If you are going to need multiple databases, it is best to set them up at the outset. If you do lump everything together now and decide later that the multiple database approach is needed, you won't be out of luck, however. It's easy to divide a Cumulus image database.

First, select the records you wish to move to another database while holding down the **Command** key. Next, use **Export Records** from the **File** menu. Move over your keywords as well by selecting all the listings in the Keywords window and choosing **Export Keywords** from the **File** menu. Reverse the procedure after creating a new image database. As we said earlier, Cumulus can have more than one database open at the same time.

Opening Her Up

Another thing to consider is how important speed of access is and for whom. In most cases, it is hard to imagine that anyone would complain about the speed of Cumulus, but heavy users might. For these users, ask yourself this:

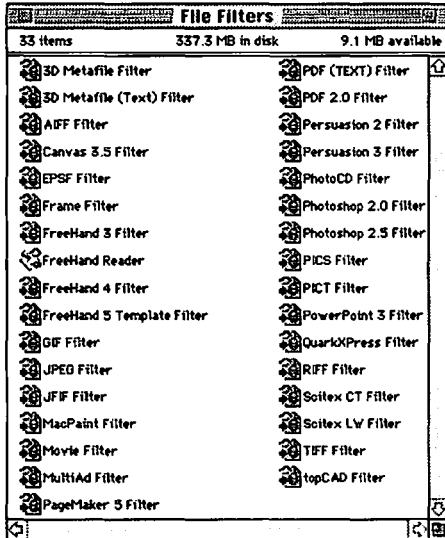
- Do they have the best possible network connection? If your heavy users are dispersed throughout the network, it might be better to place them and the Cumulus server together and segment their traffic on your fastest network.
- Should I dedicate the server to Cumulus? If you are running many services alongside Cumulus, see how much performance is gained by relocating them one by one to another machine. Maybe you should let Cumulus have a Macintosh all to itself.
- You might want to upgrade the server machine. The speed gains in moving from one Power Macintosh to another probably won't impress you a whole lot, but those made from moving from a 68K machine to a PowerMac will often blow your socks off! Don't forget about accelerator card upgrades. Don't forget the workstation, either. If you have Cumulus running on a 9500 but your graphic artist is using an LC, the slowdown is going to have nothing to do with the server!
- Should I jack up the RAM? You can gain some performance by allocating more memory to Cumulus, which might also entail buying some more. Don't try to cheat with Connectix' RAM Doubler. While it is a way cool product, it is best applied to occasional memory needs and not for perpetual use on a dedicated server. The best gains come from making sure the memory available to Cumulus is greater than the memory it requires when it has all its databases open. Take a look at what Cumulus requires by choosing **Administration** from the pop-up menu at the upper left of the Database window. Again, don't forget the workstation. Try allocating more RAM to Cumulus Client.

Acme RC Image Database				
Administration		Records Deleted: 0	Memory Available: 5344 KB	
Database Activity		Database Memory Free: 1667 Byte	Required: 80.00 KB	
		Disk Space Available: 9136 KB	Recommended: 512.0 KB	
User	Date	Time	User	Action
All	11/16/95	12:27 PM	WG Administrat	Added: AcmeRC Gray.PICT
WG Administrator	11/16/95	12:27 PM	WG Administrat	Added: ACSE Logo.tif
(Guest)	11/16/95	12:27 PM	WG Administrat	Added: Backup T
	11/16/95	12:27 PM	WG Administrat	Added: Backup T.GIF
	11/16/95	12:27 PM	WG Administrat	Added: Backup.rgb
	11/16/95	12:27 PM	WG Administrat	Added: BUM1.gif
	11/16/95	12:27 PM	WG Administrat	Added: Calendar2.rgb
	11/16/95	12:27 PM	WG Administrat	Added: CD Artwork.rgb
	11/16/95	12:27 PM	WG Administrat	Added: DAN.gif
	11/16/95	12:27 PM	WG Administrat	Added: Designing2.rgb
	11/16/95	12:27 PM	WG Administrat	Added: Dialog Balloon.PICT
	11/16/95	12:27 PM	WG Administrat	Added: FIN.gif
	11/16/95	12:27 PM	WG Administrat	Added: Generic Gir1 w/Phone.PICT

Cumulus Memory Usage

- Should I obtain a faster hard drive? Cumulus reads from and writes to the hard drive a lot. Naturally, the faster the access time of the hard drive, the faster Cumulus runs. This is less the case when the entire database fits neatly into available RAM, but it is still an issue.

Standardizing Formats



Different File Formats

One of the strengths of Cumulus is that it can open many graphics file formats. The filters that come with version 2.5 are shown in the illustration, and as one guy at Canto told me, they add new ones all the time. While this is a welcome feature, I would argue that it is one the average office should not overuse. I say this because while Cumulus can handle pretty much whatever you throw into it, many of the applications your users have can't. I suggest that you support at most three basic formats for new files. These are detailed below.

High Resolution

The high-resolution format is for sophisticated publications, such as full-color brochures and professionally printed material. Anywhere high-quality reproduction is required, use these formats. Most often this is EPS or TIFF. These will usually be strictly for the use of your artists and marketing people. They are often quite large.

Low Resolution

The low-resolution format is for documents that need a little dressing up but do not require as many levels of detail for printing. For instance, it might be the company logo that someone wants to slap on a report they just wrote. The report is probably going to be printed out at 300–600 dpi, so a high-quality image would be wasted. Low-resolution files should be available to everyone, so follow two rules: Make them compatible with as many user applications as possible, and make them small, under 200K. The format that fits the bill best here is PICT.

Screen Resolution

You may or may not need this category at all. This is for art that will be seen only on the computer screen. For instance, it might be a picture of the company building that someone wants in a presentation. The presentation will be viewed at the resolution of the monitor (72 dpi). Most of the detail that a high-resolution image

for printing contains will never be discernible, and the more detail in the image, the slower it is drawn on the screen. Audiences have been known to doze off while waiting too long! As the goal here is small size, it is best met with the PICT format.

A subcategory might be online resolution. This is for art that is to be “published” through online services or in HTML documents. In this instance, small size is vital, as it translates into more productive online time. Platform independence is also necessary, as you don’t know what kind of machine is viewing and downloading the document. The format to use is GIF or JPEG, and the latter is arguably the best choice.

Enough, already! Cumulus comes with a clearly written manual, so you can learn everything you need to know there. The main thing to realize is that if your organization makes use of a lot of art in its day-to-day operations, Cumulus will greatly improve your organization’s productivity by making it easier for users to find what they need—in one place, in the right version, in the right format, and filed under an appropriate keyword.

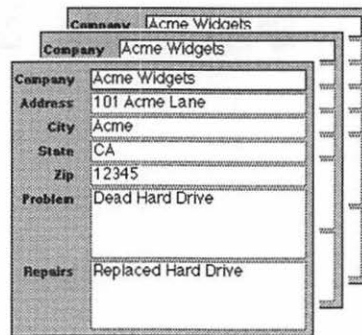
CHAPTER 7: DATABASE SERVERS AND CGIs

This chapter was written by Tom Hessel

All electronic information, whether it is a little or a lot, is somehow stored, organized, updated, compared, searched, sorted, and output. There are many different ways that these operations can be performed. Some people do them manually by filing, cutting, pasting, and printing. More clever people tend to use applications, such as databases, that aid in these special operations. I guess you can say that *databases* are collections of data that, once entered, can be easily organized, updated, compared, searched, sorted, and output. Databases can be used as both organized storage facilities for information and as utilities to perform comparative analysis tasks between different types of data. This data can be rearranged and reformatted so that people can have a better understanding of its significance through charting and customized reports. The beauty of today's database applications is that they allow for easy development and customization by the end user's own set of rules. I can remember when databases were mainframe-based systems only, and it took a multitude of tiny little programmers to create and update these databases. The users were constantly at the programmer's mercy. Well, for the most part, those days are behind us.

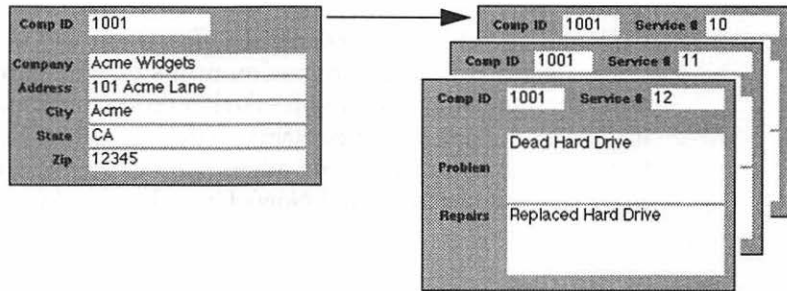
A database can consist of one file that contains all its information. This is more commonly known as a *flat file* database structure. This may be a database that holds information regarding technical services performed for clients, for instance. In a flat file structure, there are sometimes a multitude of fields within each record. The problem with a flat file database structure is that each individual record

within the database has no relationship with any of the other records. Because of this, your database starts to fill up with redundant information. If there need to be multiple entries in the database for the same company, the company information must be re-entered time and again, even though their problem and solution changes. Flat file design decreases data access speeds when performing such functions as searches and sorts. However, database development time is minimal when working with a flat file structure, because the databases are often less complex.



Flat File Database Containing Records and Fields for Each Single Entry or Activity

The second type of database structure allows you to share or associate information between two data files. This allows the creation of multiple data files. For example, a company could have one data file containing all client information and one containing all the work performed for the clients. This is more commonly known as a *relational database* structure. A data file that contains information needed by another file is called the *related file*. The connection between these two is known as a *many-to-one* relationship. This means, in our example, that there are many service entries or activities for a single company. With a structure like this, you need to enter the company information only once, which decreases the amount of redundant information in your database. Relational databases increase data access speeds when performing such functions as searches and sorts. However, due to their complexity, database development time increases when working with relational file structures.



In this relational database, the company file on the left has many related service entries, shown on the right.

The beauty of a relational database structure is that whichever file I use, I am given immediate access to all the information relating to that file. If I access the record for Service #12, the database also loads the related information that tells me to which company that service record belongs. If I access the record for Company ID 1001, the database also loads all the related service records. In the example shown in the preceding picture, they are service records 10–12.

Databases have become powerful tools that can be used as simple stand-alone packages on your personal computer, thus allowing great speed in accessing your information. You no longer have to depend on the reliability of an additional computer holding the main database. However, with few changes, you can use most single workstation packages to develop a fast, reliable, multi-user database on a local file server for a small workgroup or business.

Current database developments enable us to access data files over multiplatform and multiprotocol networks. FileMaker Pro v3.0, for example, allows you to share data files simultaneously over AppleTalk, IPX/SPX, and TCP/IP. This allows FileMaker Pro clients to select their own flavor of computer—PC or Mac—along with the protocol of their choice. Other databases also enable users to *host* databases from different platforms, but all platforms must use the same protocol. I will explain this in more detail shortly.

FileMaker Pro, in my opinion, is a wonderful, low cost, user-programmable relational database development package. Wow, am I a salesman or what? FileMaker Pro will form the foundation for our discussion of databases in this chapter. While we won't include much information about setting up FileMaker Pro or FileMaker Pro Servers, we will cover design issues and other related tips for placing database servers and *CGIs*, or *Common Gateway Interfaces*, on your company's networks.

We will discuss the relationship between FileMaker Pro databases and WWW CGIs and how client/server databases affect traffic and workstation performance. Other database development environments that are currently available include ACI's 4th Dimension, Microsoft's FoxPro Professional, and one of the older original database development packages on the Mac, Helix Express. Elsewhere in this book we will talk about some others: Apple's AppleSearch, Canto's Cumulus, EveryWare's Butler SQL, and Now's Up-to-Date and Contact applications.

PLANNING THE DATABASE

Whether you are going to develop a database package that is a single-user or a multi-user, client/server application, there are many steps to consider before creating the database. Although FileMaker Pro allows you to easily make modifications to the design at a later date, try to minimize time and effort by planning carefully before creating the database. Remember that the more you plan ahead, the fewer problems you will have in the future. Here are some things to consider before you create your database:

- Decide whether this database will be accessed throughout the entire corporation or if it will be a single-user application. If your decision is client/server, then you need to decide which type of computer best suits your server needs.
- Determine the types of information with which you will be working. To decide whether you need to design a relational or flat file database, ask yourself a very simple question here: Will our database have information that must relate to other information within the same database? If so, then your design will be a relational one. When using a relational database, you will want to design a separate data file for each type of information with which you will be working. If you decide that you need only a flat file, it will obviously be a single file structure.

Using the example shown in the picture on page 187, we would need one data file for the company information and a second data file for the service information.

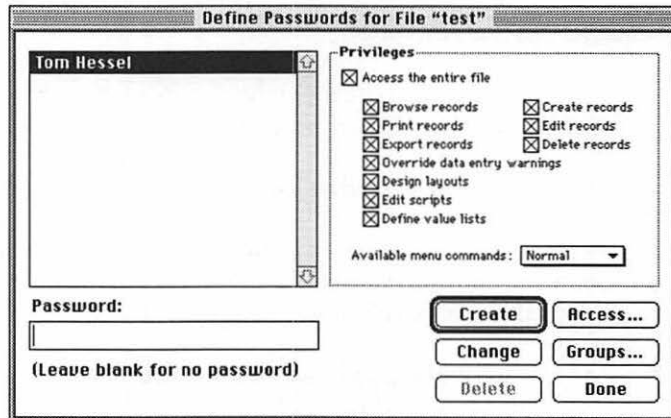
- Determine what data each of the data files will be holding.

Using the example shown earlier, we would need fields for the company ID, company name, address, city, state, and zip code for the company data file and company ID, service ID, problem, and repairs fields for the service data file.

- If this is going to be a relational database, determine which field in one data file relates to which field in the second data file. FileMaker lets you build a flat file database, which will let you easily add relationships to already existing or newly created data files in the future. Remember that using the relational feature enables you to use information that resides in a separate data file by de-

fining a relationship between the two data files. Earlier I showed the relation between company information and the services performed for that company.

- Look at how your company's current information is being handled. This will help you determine how to format your data for output. How will you reproduce your company's current reporting system, invoices, service reports, and mailing labels with the new database application?
- Based on the data file and field decisions you have made, design your layouts for the database, how the data will be entered, and how the data will be output. If you will be printing invoices, mailing labels, or service reports, you should design them on paper first.
- If this will be a multi-user database, and the information needs to be secure, set up some Users & Groups information. Determine whether users will have "super user" privileges; read only; or read, add, and delete privileges. Just about every database development package allows for user and group access privileges. As an additional reminder, there must be at least one user within your setup who has the ability to *access the entire file*. This user should be the server machine that will be hosting the data file. I will explain this further in just a moment. The following picture shows FileMaker's access privileges and password security features. Other packages may be different.




User Creation with Associated Privileges

Picking the Proper CPU

When talking about selecting the appropriate computer to serve the needs of your database for your company, we will discuss solely client/server multi-user operations. When setting up a database for single-user, single-workstation operations, the solution is obvious—bigger, faster, stronger, and with more RAM! The only person degrading the performance of the database and the computer on which it resides would be that single user. Simple enough?

Trying to choose a home for a multi-user database is not quite as easy. The best solution would be to dedicate a machine to the database and apply the rules of bigger, faster, stronger. Unfortunately, we live in a world with budget constraints that can make this impractical. For those of us without a bottomless pocketbook to pick, let's take this discussion further (congress members do not apply here).

There are a few different ways in which FileMaker Pro can be configured as a multi-user database application. Without using additional software such as Claris's FileMaker Pro Server software, the only way that FileMaker-designed databases can be accessed by multiple users across the network is through a feature called *peer-to-peer hosting*. This is kind of like using System 7's Personal File Sharing, but it uses AppleShare for the sharing of files across a network. Hosting dictates that whoever opens the database file first will then *host* it to everyone else needing access to it. I guess you could define *hosting* as the ability of the personal computer on the network, whether it is a PC or a Macintosh, to allow access to its data file. This can occur across the Internet or within an internetwork. Hmmm . . . let's take a look at this in the following picture.




Felix is hosting the database for Lynn and Maizie.

There is just one problem here: If, for some reason, Felix decides to restart or turn off his computer, or maybe go home early, he can no longer host the database file to the other users, Lynn and Maizie. I call this type of hosting *nondedicated* hosting. No one is dedicating a computer to the sole purpose of serving our database. Whoever opens the database first will then host it to everyone else. Today it could be Felix, tomorrow it could be Maizie, and so on. Are you starting to see the pattern? The catch is that in order for Felix, Maizie, and Lynn to all be able to host the database based on whoever opens the database first, they all must have a copy of the database file on each of their machines. That copy must be current so that everyone has access to the same and most current information each day. This is the point of multi-user database files: the sharing of information that is common to more than one person. How would you like to be responsible for making sure that every night whoever was the last person hosting the database ensures that everyone else has a copy of this current file in case they host the data file the next day? Not me!

Both the beauty and the beastliness of FileMaker Pro is that right out of the box FileMaker's software is both a client and server, or hosting, application. When hosting data files using the AppleTalk protocol, FileMaker does not use the *Apple Filing Protocol* (AFP) for sending data from client to host, unlike AppleShare. When transmitting data across your networks for Macintosh computers using Personal File Sharing or AppleShare file servers, you are using AFP. FileMaker, however, uses a protocol within the AppleTalk protocol stack called the *AppleTalk Data Streaming Protocol* (ADSP). This protocol provides an advanced two-way messaging system between two nodes on your network. ADSP is a symmetrical protocol that supports full-duplex streams of data allowing both clients at either end to maintain equal control over data exchange sessions. It is not necessary to share the database data file across your networks using Apple's Personal File Sharing or AppleShare. When clients want to sign on to a hosted database, FileMaker will do a network scan for available databases using ADSP, not AFP.


So, what if you put the database file on your company's AppleShare file server? That would be a good idea in that it will keep the database file resident on the one machine whose responsibility it is to never go down. This is because file servers never go down, right? Please answer, "Yes." In the following picture, however, notice that even though the data file is now placed on the server and the company maintains just one copy of the data file, because of FileMaker's hosting technology, the same problems we just discussed still apply—the first person opening the data file is still hosting that data file to everyone else.



The data file resides on the server and Felix opened it first.

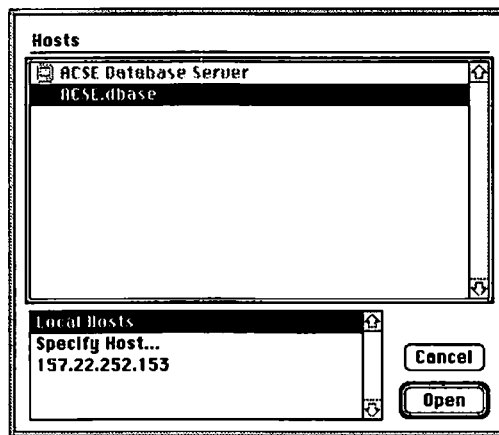
In the picture above, Felix was the first person to double-click the data file that resides on the server. In the chain of transaction commands, the data file is sent from the AppleShare file server to Felix's machine using AFP. Then, all hosts will sign on to Felix's machine using ADSP. This is what I call *ghosting hosting*—whoever double-clicks the data file first will become the host to everyone else, even though the data file resides on an AppleShare server.

So, how the heck *should* you plan for the location of the database so that no one other than a dedicated machine is hosting the database? Since FileMaker is both your client and your hosting software, when planning for the use of a company-wide, multi-user database, you need to make sure that the data file resides on a central server and that the data file has first been opened on that file server from the FileMaker Pro application on the server. Users wanting to access the database should open the file through the **Hosts** button in the Open dialog box from within FileMaker. The big mistake here happens when people select **Open** from the **File** menu and then select the **Open** button instead of the **Hosts** button. They proceed to navigate through files until they find the data file as it sits on a file server. Again, that is opening the data file using the AFP protocol (bad), instead of allowing FileMaker to open the data file from a host computer using ADSP (good). If everyone logs on to the AppleShare file server and double-clicks the data file instead of using the open hosts command, you will again experience the *ghosting hosting* effect.



FileMaker Pro's Open Dialog Box


The database must already be running somewhere before you click the **Hosts** button, or no host will be displayed. If you did not dedicate a machine for this database, you will still click the **Hosts** button, since this will display whoever is hosting the data file. Since we want a file server to host the data file, it should already be running on your server.



Dialog Box Displaying Available Hosts on the Network

The picture above shows the dialog box that is displayed once you have clicked the **Hosts** button. All of your available hosts or data files will be displayed. Highlight the appropriate one and finalize the procedure by clicking the **Open** button. You have now successfully logged in to a hosted data file using the ADSP portion of the AppleTalk protocol stack. Congratulations.

With the file server acting as the host to the rest of your database clients, your database will always be prepared to host its services. Also, since there is no one sitting in front of the file server all day long, you won't have to worry about degrading the performance of the workstation that is hosting the file.



FileMaker on a Server Hosting the Data File to Network Users

Management applications report back that most AppleShare file servers are used for a total of about 45 minutes per day per person. Yeah, that's right, 45 minutes, got something to say about it? People log on, people retrieve data, and people remain idle for the remainder of the day. With activity like this, it makes sense to use your AppleShare file servers as hosts for your FileMaker Pro databases. However, please keep in mind some of my recommendations when doing so:

- Use some common sense. Will hosting a 500-MB data file with approximately 25 clients degrade the performance of additional services that computer may be offering? What do you think?
- To ensure that the data file is always open on the server, try putting an alias of the data file—*not* the FileMaker Pro application—in the Startup Items folder on your server. The Startup Items folder is located in your System Folder and whatever item resides in that folder will open automatically when the computer starts up or is restarted.
- For every user who will be a client to FileMaker's data file, increase the RAM allocation for FileMaker Pro on the host computer by 500–1,000K (all users accessing it should have their memory requirements set to the suggested size). Base this number on the activity of your users and the complexity of the database they are accessing. Set it lower for less complex databases with lower amounts of usage, and higher for more complex databases with very active us-


ers. Still, you should not have to raise this number any higher than 10 MB. This means you would be hosting this data file to 20 users, giving 500K of application RAM to each user. This is done so FileMaker can add records to a data file by more than one user simultaneously, while other users are searching for records that already exist. FileMaker stores added records in RAM until there is no activity in the database for 90 seconds or on a user-definable interval starting at ten minutes. Then the records are written to the data file. Write aheads can go to RAM while reads are pulled from the hard disk, which increases performance. Without this functionality, people would read and write to the hard drive at the same time, decreasing database performance. Does this mean that if one user adds a record to a remote database, all other users have to wait 90 seconds before they can retrieve the file? No, FileMaker is intelligent enough to know where a record is, whether it's still in RAM or on the hard drive. If another client to the database requests a record that is still in RAM, FileMaker will go to RAM for that record.

- AppleShare or Personal File Sharing is not needed when *hosting* database files to multiple users. While FileMaker Pro uses ADSP (part of the AppleTalk Protocol), TCP/IP, or IPX to communicate between client and server, AppleShare uses AFP. If you are putting your FileMaker Pro databases on the same machine as your AppleShare file servers and you do not want users double-clicking the database file itself, do not make that database file a shared point. This ensures that users receive the best performance from the FileMaker hosting technology by using ADSP and not AFP.

When to Use FileMaker Pro Server

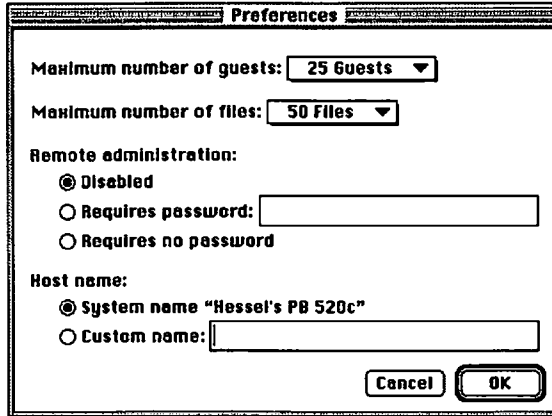
Remember that FileMaker Pro is targeted to support the database needs of small to medium sized groups of users. If you need to host a data file to more than 25 clients using Claris' peer-to-peer hosting technology, purchase FileMaker Pro Server 3.0, which can host up to 100 database files and 100 network users. Each client of the server can open up to 50 files at a time, as long as the combined number of open files does not exceed 100. FileMaker Pro Server 3.0 also supports asynchronous I/O and multithreading operations, allowing for multiple database operations to take place simultaneously. FileMaker Pro Server 3.0 does not allow database formatting from within the server; the software is a dedicated data file hosting application only. Like FileMaker Pro 3.0, the server software allows for a maximum database file size of 2 GB.

FileMaker can transmit data using one of three protocol suites: AppleTalk, TCP/IP, and IPX/SPX using peer-to-peer hosting. The only problem is that all *hosts* must use the same protocol. FileMaker Pro Server, on the other hand, allows different clients using different protocols to access its shared data. This allows you to share data from one central location to multiple types of platforms—PCs and Macs. Ask yourself one very important question: Do we have clients on different platforms using different protocols who need to access the data on a FileMaker Server? If your answer is yes, then you do need to use the FileMaker Pro Server software to share this common data.



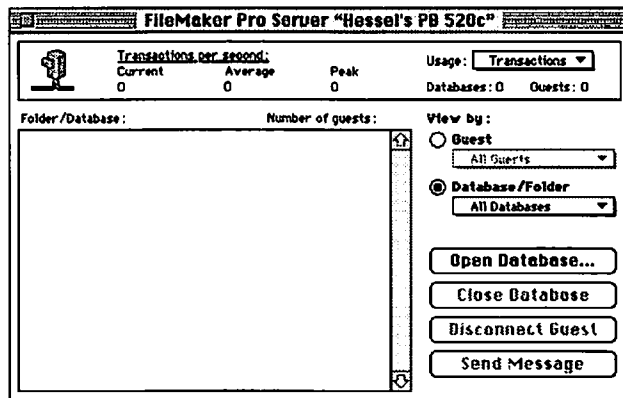
FileMaker Server with Different Protocol Clients

Upgrading to FileMaker Pro Server is an extremely painless operation. It allows you to use the same data files that you were using when hosting with the FileMaker peer-to-peer technology. In this way, you can upgrade very inexpensively and easily as your company's needs grow. Before you can start using the server software to host your data, make sure that from within the application you have configured your FileMaker database file as a multi-user file and that your Users & Groups privileges have been set. FileMaker Pro Server's default preferences are set to be the same as that of FileMaker Pro, with the maximum number of guests set to 25 and the maximum number of files opened as 50. Change these to suit your needs. The next picture shows the default preferences within the server.



FileMaker Pro Server's Default Preferences

Once your preferences have been set, it is time to open your data files from within the server. As you can see, there are not a heck of a lot of settings within the server software. This makes setting up the application very, very simple. The **Administer** option in the **File** menu brings up the window you see in the next screen shot. Administrators can see information about their data files and the performance of the server in a few different ways—transactions per second, network kilobytes per second, hard drive kilobytes per second, cache hits, swaps per second, available databases, and number of connected guests. This provides different levels of analytical information for planning upgrades and growth. This information can provide baseline server statistics showing how your server is performing on typical days. Dramatic increases or decreases in these numbers are out of the normal scheme of things and will help you with troubleshooting and upgrades.



Administration Window as Displayed in FileMaker Pro Server v3.0

By clicking the **Open Databases...** button, you are prompted to select the data file that is to be hosted. It's as simple as that. The backup scheduling features and additional options within FileMaker Pro Server are self-explanatory. However, I would like for you to take a look at my recommendations for determining a home for this software and other configurations:

- First and foremost, *bigger, faster, stronger* still applies here—sorry. The faster the computer, the better FileMaker Pro Server will operate. FileMaker Pro Server uses PowerMac accelerated code.
- If you can't dedicate a separate machine for FileMaker, can you install the FileMaker Pro Server software on your AppleShare server? The answer here is not really a performance-related issue, especially if you have one of the newer machines available from Apple. Rather, the answer involves a limitation in the ability of users to log on to a machine that is offering file sharing and database serving services. If you have read our book *Managing AppleTalk Networks*, you know that computers on a network are addressed by a network number and node address, and that all Network Visible Entities (NVEs) on that computer must obtain what is called a socket address. NVEs could be any piece of software running on a computer that needs network services. Now, since each computer can have a maximum of 254 socket addresses, two addresses would be taken up for AppleShare and FileMaker Pro Server as separate NVEs, right? The problem goes a little deeper—both AppleShare and FileMaker Pro Server register all their connected users as individual sockets within that computer. Are you following me here? If we had 200 people in our company needing access to AppleShare and 100 people needing access to FileMaker Pro Server, keeping both applications on one machine wouldn't work, since there aren't enough available socket addresses. You will know if this is happening if FileMaker Pro guests receive a message telling them they have "exceeded the host capacity" of that server.
- Use third-party software such as Santorini Consulting's Server Manager to regulate the time allowed for idle AppleShare users on this computer. This will free up additional sockets for FileMaker Pro users, keeping at bay the need for a second computer.
- If you will be running AppleShare and FileMaker Pro Server on the same computer, put the shared information on separate hard drives. This allows for no disruption in AppleShare service when the FileMaker Pro users need data.

- With the preceding consideration in mind, make sure that you are using hardware that supports asynchronous I/O commands. SCSI Manager 4.3 enables all SCSI devices attached to this machine to perform asynchronous commands simultaneously. If your hardware does not support SCSI Manager 4.3, SCSI commands must be stacked in logical order based on ID numbers. SCSI Manager is part of every PowerMac's native code. For older 68040 and '030 machines, System 7.5.x installs a software implementation of SCSI Manager 4.3. This way, even though you have separate hard drives for AppleShare and FileMaker Pro Server, the computer can access data for each simultaneously. For more information on SCSI Manager 4.3 and this type of configuration, see our book *Managing AppleShare and Workgroup Servers*.
- The faster the network, the faster the performance of the server's services. If you have users accessing the server from both Ethernet and LocalTalk, the LocalTalk users are going to affect the overall performance of the server for all users, even those on Ethernet. This is because it takes FileMaker a longer time to release a transaction to those users on LocalTalk than it does to those users on Ethernet. In single-process simultaneous transactions, one transaction must be finished before it can move on to the next. Remember, though, that these transactions take place in microseconds, and are therefore hard to clock.

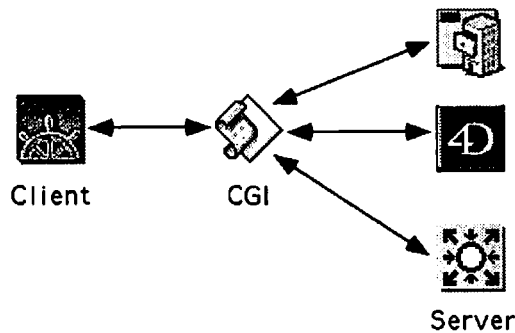
CGIs or ACGIs

The *Common Gateway Interface* (CGI) or *Asynchronous Common Gateway Interface* (ACGI) are ways of interfacing external applications with information servers, such as HTTP or World Wide Web (WWW) servers. A plain HTML document that is served to the world via the Web is a static type of document. This means it exists in a constant state that rarely changes—kind of like my Web page. Since I am always on the road, I never have time to update it, so it never changes. These documents are available for browsing but not for interaction or changes. Since a CGI program, on the other hand, is executed in real time by end user requests, it can output dynamic information. I guess you could say that CGIs respond to user input from your Web pages and perform custom tasks based on that input. CGIs enable you to serve up more than just pretty pictures and text on the Web. These applications can be relatively simple, such as a counting application, mortgage calculator, or clickable image map server, or more complex, like interfacing people from the Web with the database design of your choice. The list is as long as your imagination is vast, assuming you have a vast imagination. Pretty cool, huh?

Here is a small list of additional uses for CGI applications:

- Converting a technical manual's pages into HTML on the fly and sending the HTML results back to the client.
- Interfacing with Wide Area Information Servers (WAIS) and archive databases, then converting the results to HTML and sending the result to the client.
- Allowing user feedback about your server or organization through an HTML form, signing up for a service or class, or adding your name to a mailing list.
- Providing document fax-back or e-mail-back capabilities, and using CGIs to decode this type of feedback.
- Allowing users to add, delete, search, and sort records to a custom-designed database. These internal databases can be designed with many applications, including FileMaker Pro, 4th Dimension, and FoxPro.
- Providing users from the Web with the ability to use the same databases your company may use locally. The need for duplicate applications isn't necessary.

These are the most widespread applications of CGIs, though there are obviously many more uses.



Where Gateways Reside

Our CGIs or gateways are actually programs that can handle information requests and return the appropriate documents or generate a document on the fly. Your information servers can return information to your clients that is not inherently readable by them (such as SQL databases). The CGIs will act as mediators between your clients and servers, producing readable information for your clients. Since a CGI program is an executable application by outside users, I would say that it is basically the equivalent of letting the world run an interactive program on the servers on your network through the Web.

Does this sound like the safest thing to do? How can you protect yourself from hackers and crackers? I am glad I decided to ask that question, although you are supposed to be asking the questions. I am talking to myself again. Sorry. There are some security precautions that need to be implemented when you start looking into using CGIs as a means of transferring user requests to internal applications.

I would say the one that will affect the typical Web weaver (user) the most is the fact that CGI applications sometimes need to be in a special directory, depending on the individual who wrote the CGI. This is so that when the Web server sees the CGI in its directory, it knows to execute the program rather than to just display it to the user's browser. This directory is usually controlled by the Webmaster directly, prohibiting average users from creating CGI programs any old time they feel like it. There are other ways to allow access to CGI scripts, but it is up to your Webmaster to set these up for you. At this point, you may want to contact your

Webmaster about the feasibility of allowing CGI access. If this is a concern, please see the chapter called “Serving the World Wide Web with WebSTAR” beginning on page 303 in this book.

If you have a version of WebSTAR by StarNine Technologies and are using it as an HTML information server, you will see a directory called `/Scripts` within the WebSTAR directory. This is the special directory I just mentioned where all your CGI programs should reside.

When writing your CGI applications, you have the convenience of choosing your favorite programming language. CGIs can be written in any language that allows them to be executed by users on the system, such as the following:

- Think C/C++
- Fortran (yuck!)
- PERL (Practical Extraction and Report Language)
- TCL (Tool Control Language)
- Any UNIX shell (with the exception of AUX)
- AppleScript
- Visual Basic

It all depends on what you have available at your facility, your level of knowledge, and which language would be proper for the task at hand. It may take less time for you to write a CGI using Apple’s AppleScript language than if you were to write the same CGI using the C development environment. Programming more complex CGIs using AppleScript may cause the performance of that CGI to decrease as opposed to coding the same CGI using C. The fact that AppleScript isn’t a PPC-native development application has nothing to do with it, right? If you use programming languages such as C or Fortran, I hope you know that you must compile the program before it will run. If you don’t know what I’m talking about, than this may not be the language of choice for your CGI. AppleScript must also have the source script checked for syntax before the script can be run or saved. This is similar to compiling, as source code is checked for syntax, only the code is not reduced, unless it is saved as a stand-alone application.

A CGI's main job is to become active when a user requests its services. For example: Users enter data into an HTML form that is on a Web server. This becomes the data entry source. Once a particular button on the Web page is clicked, such as Submit, a CGI is called upon to take the data that was entered in the HTML form and move it to a particular area such as a database or an e-mail account. The CGI can then be used to produce an HTML-driven response to the users, letting them know that their actions were executed with or without error. If an error has occurred, the users receive an error message with its type. In the following picture, take a look at an example of an HTML form page we designed that provides interaction with our internal ACSE Support Database.

The screenshot shows a web form with the following elements:

- Title:** Submit Questions and Suggestions to the ACSE Database
- Subtitle:** Brought to you by Network Frontiers, Inc.
- Instructions:** Please enter all information below and click the Post Message button to send us your Questions and comments. You should receive a response to your inquiry within 24 hours.
- Test Selection:** Please select the appropriate test. (If you are unsure, this is a pop-up menu). A dropdown menu is currently set to "Server and Backup Test".
- Text Fields:** Name, Email, and Phone.
- Category:** Radio buttons for "Question" (selected) and "Suggestion".
- Text Area:** A large text area for the "Question or Suggestion".
- Buttons:** "Post Message" and "Clear this Form".

Interactive HTML Form Used on Our Web Page

What I developed here using HTML is an interactive Web page where users can submit questions pertaining to the ACSE certification program to the consultants at Network Frontiers. Once users enter all the information in the form, they then click the Post Message button. This triggers an ACGI and performs a specific function: adding a record to a FileMaker Pro database that I designed. The follow-

ing shows the HTML command that tells this HTML page to perform the add record action using a CGI named *FileMaker.acgi*.

```
<FORM ACTION="FileMaker.acgi$ADD" METHOD="post">
```

With external applications such as CGIs, we have the ability to produce *hidden* variables, or *arguments*, that return preformatted HTML data to the users' browsers once the Post Message button is clicked. Take a look at the code below and let me explain:

```
<INPUT TYPE="hidden" NAME="_databaseName" VALUE="ACSE.dbase">
```

```
<INPUT TYPE="hidden" NAME="_layoutName" VALUE="HTMLRecord">
```

When the Post Message button is clicked, our CGI grabs the information from these two hidden variables. These variables tell our CGI the name of the data file to which it will add records. An error returns if the data file called "ACSE.dbase" cannot be located. The "_layoutName" determines which layout our FileMaker database uses to send preformatted HTML data back to our Web users.

Okay, you have waited long enough—what is the difference between a CGI and an ACGI? Good question! It is about time you started asking some. Since a CGI is an external application that is called upon by user interaction, the application does not necessarily need to be active all the time. When we click the Post Message button, the server looks for the application called FileMaker.acgi. Once it has been found, the computer checks to see if it is an active application. If not, then the application must be launched before it can be used. This can decrease performance, since the CGI must be launched each time it is needed and shut down when it is finished being used. An ACGI remains active all the time, never shutting down. This gives the ACGI the capability to process requests asynchronously without opening and closing all the time. The CGI can take data entered by the users from the interactive HTML form at the same time as it is sending back a response from our database to other users. This gives multiple users the capability to use one CGI to add and search records in the same database file, simultaneously—as long as the CGI is written for this type of operation. If not, then separate CGIs could be written to add and search records. Since ACGIs need to remain active at all times, listening for calls from both sides of the court so to speak, they tend to take up a lot of the process time of the CPU on which they are residing. Take a look at the following picture:

The screenshot shows the 'Process Spy' window with a table of open applications. The table has three columns: 'Process Memory Usage', 'CPU Usage', and 'Launch Time'. Each row includes an application icon, the application name, memory usage (Assigned and Free), CPU usage (seconds and percentage of total CPU time), and the launch date and time.

Process Memory Usage		CPU Usage	Launch Time
	Finder 201K Assigned 23K Free	840 Seconds 1% of Total CPU Time	Mon, May 6, 1996 6:34:32 PM
	File Sharing Extension 210K Assigned 19K Free	7 Seconds less than 1%	Mon, May 6, 1996 6:34:38 PM
	FileMaker.acgi 316K Assigned 210K Free	35420 Seconds 48% of Total CPU Time	Mon, May 6, 1996 6:35:01 PM
	NetPresenz 1056K Assigned 512K Free	35407 Seconds 48% of Total CPU Time	Mon, May 6, 1996 6:35:05 PM
	WebSTAR 2843K Assigned 950K Free	101 Seconds less than 1%	Mon, May 6, 1996 6:35:08 PM
	FileMaker Pro 5347K Assigned 710K Free	835 Seconds 1% of Total CPU Time	Mon, May 6, 1996 6:35:15 PM
	Process Spy 416K Assigned 195K Free	24 Seconds less than 1%	Tue, May 7, 1996 2:44:43 PM

CPU Usage of Open Applications

This window is from a shareware application called Process Spy, written by Rick Christianson. You can obtain a copy of this application from most Apple archives or by contacting Rick directly at richrist@cs.nmsu.edu. This nifty little utility shows you CPU usage for all your open applications. In our window, there are two ACGIs running: FileMaker.acgi, written by me, and NetPresenz, written by Peter Lewis. Both of these ACGIs are taking up 48% of the CPU's process time while in idle mode. Because they are constantly listening to both sides of the court, they take up this substantial block of CPU time.

Does this mean that since these two applications are taking up 96% of the overall CPU time on this machine, all other applications must run in the remaining 4%? No, because you are seeing the CPU time taken up by each application for this particular second, which just happens to be while everything is in idle mode. As a certain application becomes more active, all the other applications become less active to provide better performance to the application being used. Be careful here though: There is a break-off point when an application or group of applications cannot decrease the amount of CPU time that they take up. I do not know that break-off point specifically, so I suggest you get yourself a nice utility such as Process Spy to help you reveal your CPU hogs. If you start running a bunch of applications like ACGIs on your servers, you may decrease the overall performance of every other application running on that machine.

If the ACGI can function asynchronously, what happens at the database level? Most database environments can handle the processes of reading and writing data simultaneously. While searches and reads from the data file are being pulled from the hard drive, database writes are put into application RAM. When there is inactivity to the database for a period of time determined by the database programmer, the added files residing in RAM are written to the hard drive. This improves performance for simultaneous reads and writes to one database.

Getting back to the CGIs, how does the CGI know what fields within the HTML form to grab and move over to your database or other applications? When creating your interactive HTML pages, you must have some form of record entry area where people enter custom data for the CGI to pick up. Referring back to the picture on page 204, you can see that we have given the users the ability to select a type of test category upon which their question or suggestion will be based. Here is an example of the HTML code for this:

```
<SELECT NAME="subject">
<option>Server and Backup Test
<option>Designing AppleTalk Networks Test
<option>Managing AppleTalk Networks Test
</SELECT><P>
```

“Select Name” is the variable that will receive a particular option selected by the user. The variable can have an option of any one of the three described above. Next, the users must enter some additional information—name, e-mail address, phone number, a category of question or suggestion, and finally their comment. Here is an example of the HTML code for the remaining fields:

```
Name:<INPUT NAME="name" size=45><P>
Email:<INPUT NAME="email" size=45><P>
Phone:<INPUT NAME="phone" size=14><P>
Category: <INPUT TYPE="radio" NAME="category" VALUE="Question"
CHECKED>Question <INPUT TYPE="radio" NAME="category" VALUE="Suggestion"
>Suggestion<P>
<TEXTAREA NAME="message" ROWS=6 COLS=55></TEXTAREA><P>
```

The following lines tell our interactive form what to do when the **Post Message** button or the **Clear this Form** button is clicked:

```
<INPUT TYPE="submit" NAME="submit" VALUE="Post Message">
<INPUT TYPE="reset" VALUE="Clear this Form">
```

The HTML form is predefining some variables for us—name, e-mail, phone, category, and message. Let's take a look now at how the CGI picks up these variables and puts them into our FileMaker Pro database.

I have chosen Apple's AppleScript technology as the development language for our CGI. This was because of ease of use and time constraints. Let's take a look at some of this AppleScript code:

```
-- perform the desired action

if theAction is "ADD" then set returnPage to addRecord(databaseName, layoutName, emailLayoutName, emailTo, nameAndDataList)

else if theAction is "FIND" then set returnPage to findRecords(scriptName, databaseName, layoutName, linkLayoutName, maxRecords, nameAndDataList, false)

else if theAction is "FINDLINKS" then set returnPage to findRecords(scriptName, databaseName, layoutName, linkLayoutName, maxRecords, nameAndDataList, true)

else if theAction is "GET" then set returnPage to HTTPHeader & return ↵

& "<HTML><HEAD><TITLE>Retrieved Record</TITLE></HEAD>" & return ↵

& "<BODY>" & return ↵

& getRecord(databaseName, layoutName, recID) ↵
```

As displayed below, FileMaker is sending back to our Web server some preformatted HTML from a record in our database. It knows from what layout to grab the data based on the definition of our hidden variable from our HTML form called “_layoutName.” Remember that in our HTML form we defined this variable as “HTMLRecord.” This is a preformatted field with HTML calculations in our database called “HTMLRecord” that the CGI grabs and then sends out to the users' browsers.

```
"<b>Name: </b>" & Name & "<br>¶" &
"<b>EMail: </b>" & Email & "<br>¶" &
```

```

"<b>Category: </b>" & Category & "<br>¶" &
"<b>Subject: </b>" & Subject & "<br>¶" &
"<b>MsgID: </b>" & MsgID & "<br>¶" &
"<b>Phone: </b>" & phone & "<br>¶" &
"<b>Consultant: </b>" & Consultant & "<br><br>¶" &
"<b>Message:</b><br>¶" & Message & "<br><br>" &
"<b>Answer:</b><br>¶" & Answer & "<br><br>¶" &
"<b>Reference:</b><br>¶" & Reference & "¶</pre>"

```

This AppleScript was designed so that it could perform multiple functions within one single application. Without this we would have to write separate CGIs to perform separate functions. The CGI is looking for the functions ADD, FIND, FINDLINKS, and GET, with each command executing additional operations specific to that function. The GET function is using the CGI to send preformatted HTML back to our Web server.

After the users type their data into our interactive HTML form and click the Post Message button, the CGI is called upon and the ADD function is performed.

Since a CGI is designed to take data from one application and place it in another application, it first needs to know what to take, and where to put it. I have designed this CGI in such a way that it will take whatever variables have been defined in the HTML interactive form and put them into a database with matching variable names. Pretty convenient, huh? If they don't match, then nothing is added to the database.

Take a look at the FileMaker database into which our CGI is putting data. This is the database that we use here at Network Frontiers to store all our ACSE support questions. We use FileMaker to add to and search the database here in our office, while people from the Web are using the CGI to perform these functions. Make sense? If you would like additional examples of how our CGI is being used, please take a look at the add and search capabilities found on our ACSE Support Page. It can be found at www.netfrontiers.com.

ACSE Support Database v0.48
© Copyright Network Frontiers, Inc. 1996
by Thomas R. Heesel

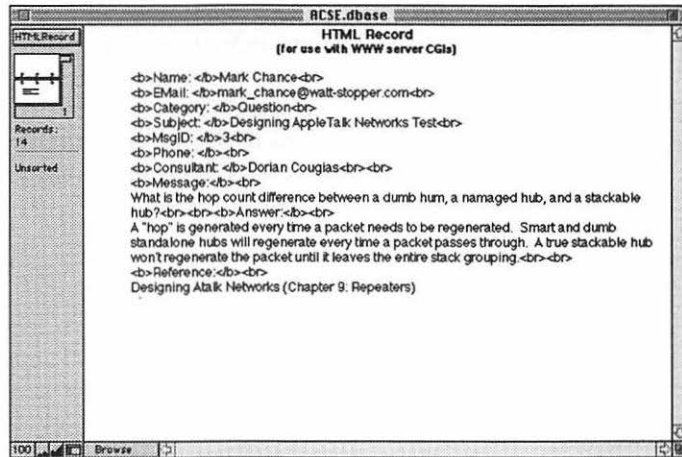
Name: Mark Chance
Date: Feb 12, 1996 Category: Question MagID: 3
Subject: Designing AppleTalk Networks Test
Message: What is the hop count difference between a dumb hub, a managed hub, and a stackable hub?
Answer: A "hop" is generated every time a packet needs to be regenerated. Smart and dumb standalone hubs will regenerate every time a packet passes through. A true stackable hub won't regenerate the packet until it leaves the entire stack grouping.
Reference: Designing Atalk Networks (Chapter 9: Repeaters)
Email: mark_chance@watt-stopper.com Consultant: Donan Cougas
Phone:

NEW DELETE NEXT PREV FIND SORT GET MAKE ALL PRINT

Sample Database into which the CGI Is Putting Data

In the preceding screen shot, you can see that there are records named exactly the same as the fields in our HTML interactive form—**Name**, **Subject**, **Message**, **Email**, **Phone**, and **Category**. The additional fields are used for the SEARCH and GET functions of the CGI. Since the fields match between our HTML interactive form and database, the CGI can easily take the data from one and put it into the other. If there is an exact match between fields, the data is added to the database; if not, you lose. Simple enough?

I had mentioned earlier that the CGI has the capability of sending data from the database back to the user—kind of the whole point behind interaction, huh? When users perform a search request from the HTML form, the CGI goes out to the database and processes the request, returning to the user all the matching items. With FileMaker Pro v3.0, aside from the CGI, you have the ability to define some HTML data at the database level before it is sent back to the user. This is accomplished by creating additional layouts within FileMaker. When search results are processed, the CGI takes the contents of these additional layouts, places them in one variable, and sends them back to the user. Take a look at one of these additional types of layouts in the following screen shot:



Additional Layouts of Preformatted HTML

When a user processes a search, the request results are sent back to the user in preformatted layouts. The picture above shows a layout called HTMLRecord. The data in this layout is placed into one variable by our CGI (also called HTMLRecord), and is then sent back to the user, producing the result of their search.

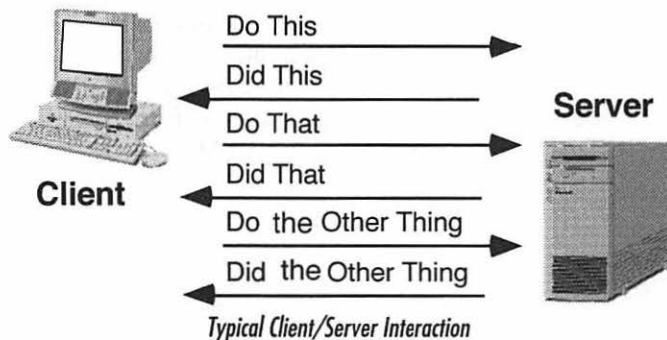
I hope that by now you have a better understanding of some of the uses of CGIs and what they are actually doing. When you plan on writing custom CGIs or using existing CGIs, please take these things into consideration:

- Remember that placing multiple CGIs on one machine can decrease the overall performance of every application running on that machine. This is because of the asynchronous capabilities of CGIs.
- Since a CGI is an interactive application that is used by people outside of your company, please consider the security problems that may arise.
- CGIs are external programs in order to allow them to run under various (possibly very different) information servers, interchangeably. CGIs are gateways that should be used to take data from one application and move it to another. CGIs should not be responsible for making calculations and decisions that will slow down the performance of the server and the CGIs. These types of operations should be performed by the application that your CGI is moving data to and from, before or after the data has been moved.

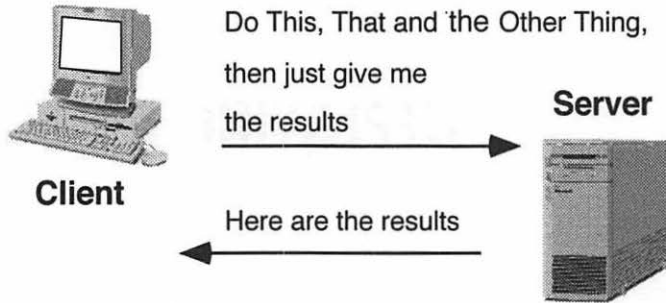
- If freeing up server processes becomes an issue, you may want to look into performing the very same tasks using Java applets. The difference between the two methods is that Java applets are downloaded and run on the user's local browser system, freeing up the server's CPU resources. Java applets tend to produce communications overhead while the applet is being downloaded to the user's browser. However, further communications and processing may be reduced by a properly written Java applet.

CHAPTER 8: DEPLOYING AN SQL DATABASE WITH BUTLER

EveryWare Development Corporation's Butler SQL is a Macintosh-native client/server database that performs its functions using the *Structured Query Language (SQL)*, commonly pronounced like "sequel." Created by IBM and first sold commercially in 1981, SQL is used to define, query, modify, and control data in a relational database. The language has evolved through numerous standards-based and proprietary versions and continues to be one of the most powerful and reliable database systems available for network deployment. Of all the various databases we speak of in this book, we consider Butler to be the most "industrial strength." In the client/server model used by Butler, search functions and data collection are handled primarily at the server end, significantly lightening the burden placed on the client and network compared to other database systems.



The client sends a small *query* to the server, which responds only when it has searched for and collected the requested results. This makes Butler a good network citizen, especially on slow networks inhabited by older and less-than-PowerPC-fast Macintoshes.



Client/Server Interaction Used by Butler SQL

Another worthwhile feature of Butler is that it was built to incorporate open standards such as *Data Access Manager (DAM)* and *Open Database Connectivity (ODBC)*. DAM is a part of the MacOS that is used to communicate with databases. ODBC is a standard developed by Microsoft to permit communication between applications and a variety of databases from different vendors. ODBC drivers are like translators, converting standard ODBC requests made by an application into a format understandable to a given database system.

Butler SQL's acceptance of these standards makes it a particularly versatile engine that can support numerous front-end applications such as AppleScript, HyperCard, Microsoft Excel, FirstClass, Informed, Web server ACGIs, and anything that can "speak" Apple's *Database Access Language (DAL)*. It can also be used in a heterogeneous network setting, supporting both Macintosh and Windows users and both AppleTalk and TCP/IP.

In this section, we will show you how to get this relational multi-user server up and running, with special emphasis on serving users of the WWW via EveryWare's Tango ACGI. Tango is a full-blown development tool that integrates ODBC-compliant SQL databases—such as Butler SQL, Oracle, Informix, Sybase, and FoxPro—with StarNine's WebSTAR.

WORKING WITH THE BUTLER SQL SERVER


Your first consideration in deploying the Butler database system should be where it will reside. The Butler SQL Server application does not have to run on a dedicated Macintosh, but of course, this is preferable. You could deploy Butler on a machine as modest as a 68030-based model, like a IIci, with 4 MB of free RAM on LocalTalk, according to EveryWare. A more realistic minimum for supporting more than a couple of users would be a 68040 or PowerPC-based Macintosh with 8 MB of free RAM on Ethernet.

Once you have designated the destination machine, we suggest that you install both the Butler server and client software on it. You will find the client handy for administrative purposes.

Butler SQL sells in two-user, five-user, ten-user, and unlimited-user packages. Be sure to buy enough licenses for the number of users you will be supporting.

Butler SQL Components

Here is a quick rundown of the system's parts.



Butler SQL Server

This is the primary engine powering the Butler SQL client/server relational database. It uses between about 3–5 MB of RAM.



Butler Databases

These are the SQL data files used by Butler SQL that you can create from scratch or into which you can import data using ButlerTools.



ButlerTools

This vital application is used to create data files and import data into them. It is also used to define users, groups, and their associated privileges.




Butler Client

This application is a rudimentary DAL client that can be used to query a Butler SQL database and to add and remove data in it. It is pretty ugly and is usually used only for testing and troubleshooting.



Butler Hosts

This application is used to configure settings for Macintosh client connections to the Macintosh running the Butler SQL server. These settings are stored in the important "ButlerHosts Data" file in the Preferences folder of the System Folder.



Butler Database Extensions

These system extensions permit Macintoshes to access Butler SQL servers from different applications such as those using DAL or HyperCard (XCMD) commands. Any application that can execute AppleScript commands is also capable of accessing Butler SQL databases. The whole pile of extensions only uses up 100K of RAM.

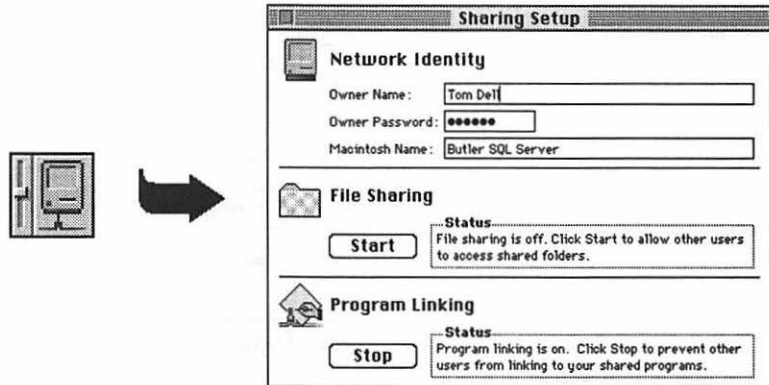


Tango

This is an add-on development tool that permits you to make Butler SQL data accessible through WebSTAR and the Internet. You don't need no stinkin' C. You don't need no stinkin' Perl. You don't need no stinkin' AppleScript.

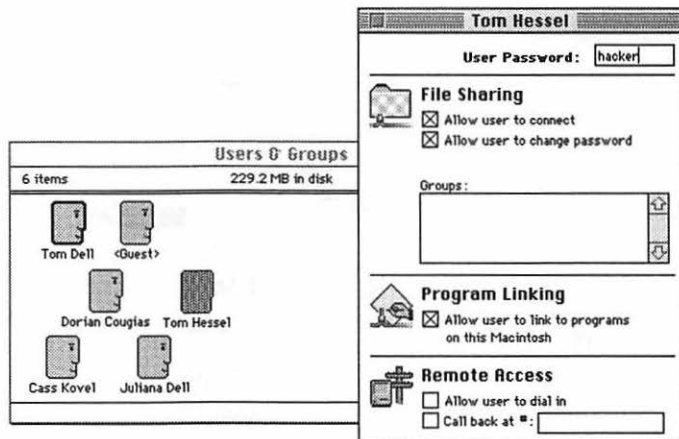
Setting up Access Privileges

Let's imagine for our purposes here that you have chosen to deploy Butler on your Macintosh Web server, that this machine is running StarNine's WebSTAR, and that this machine is not running AppleShare File Server. This makes a pretty good setup. Your first configuration step will be to enable Program Linking in the Macintosh's Sharing Setup control panel by clicking the **Start** button.



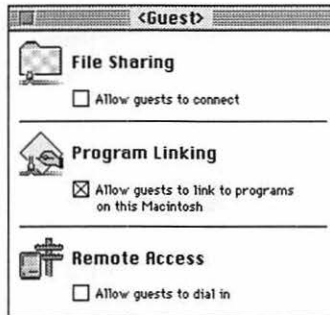
Turning on Program Linking

Once you have made Program Linking generally available to network users, designate which network users you want to permit access to it by using the Users & Groups feature of System 7 Personal File Sharing. The only thing that differs from the usual steps associated with establishing Personal File Sharing is that you will need to make sure the **Allow user to link...** checkbox is selected.



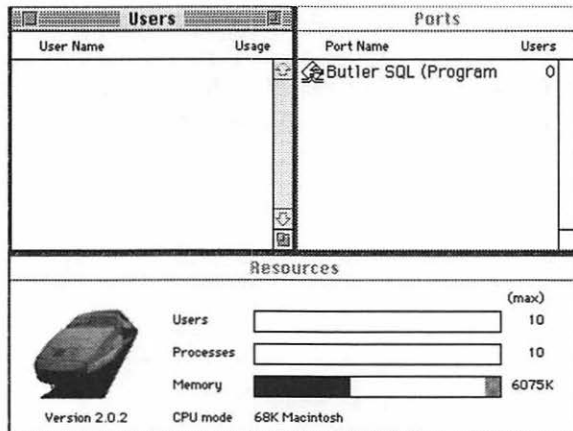
Enabling Program Linking in a User Account

If you wish to make the database available to everyone, do not bother with the process of setting up users' accounts. You can just enable Program Linking in the "Guest" account.



Enabling Program Linking in a Guest Account

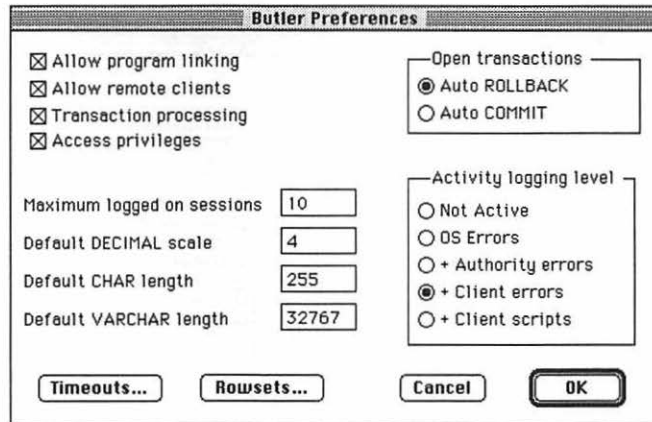
At this point, launch the Butler SQL Server application to continue. You are greeted by three windows, as illustrated in the following set of screen shots.



Butler SQL Server's Windows

The Users window displays the names of users to whom you have permitted access to Butler, their connection methods, and their usage statistics, measured in the number of low-level send and execute DAL commands they have issued.

While you can allow or deny overall access to the Butler SQL server by enabling or not enabling System 7's Program Linking feature, you need to set up additional accounts within Butler to allow and deny access to specific data files.



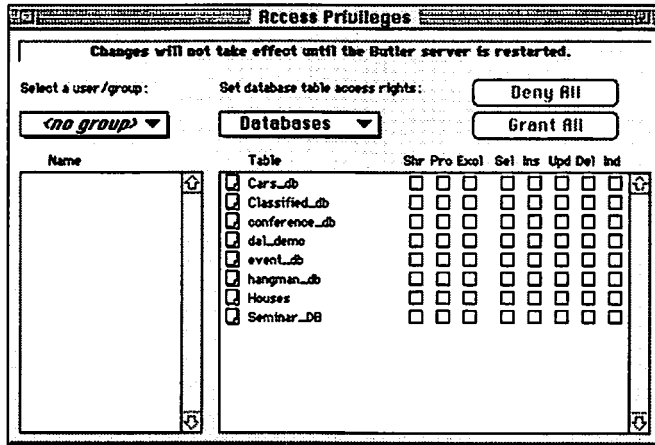
The Butler Preferences Window

To set up Butler SQL's specific users and groups privileges, open the Butler Preferences window by choosing **Preferences** from the **Edit** menu. Here select the **Access privileges** checkbox. Do not do this if you wish to leave Butler SQL open to "Guest" access.

Note that the other three checkboxes are already selected. **Allow program linking** ensures that Butler SQL can use this MacOS feature as a connection method, as we have previously planned. If **Allow remote clients** is not selected, you will be able to access Butler SQL server only from the server machine, which is not what we want to do here.

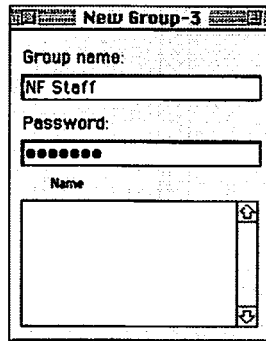
Transaction processing is a failsafe. It ensures that any modification of the database is carried out completely or not at all. Without this, a client-to-server communication that is interrupted could result in fragmentary data being submitted to the data file, causing an inaccurate or incomplete transaction. Have you ever moved money from your savings to checking accounts, and then found out at the end of the month when you receive your bank statement that the withdrawal was recorded but not the deposit? That is the kind of thing an interrupted transaction causes. We recommend that you leave this checkbox selected, but be aware of the fact that it slows server performance by as much as 15%.

Click **OK** when you are finished here, and then launch ButlerTools. Within this application, choose **Show Access** from the **Access** menu to bring up the Access Privileges window.



ButlerTools Access Privileges Window

With ButlerTools you can set up user and group privileges for each database supported by the Butler SQL server. For instance, suppose I wish to limit access to “conference_db” to my immediate co-workers. I choose **New Group** from the **Access** menu to establish my co-workers as the group “NF Staff.”



Creating a Group

I can now see my newly created group in the Access Privileges window and can select it in the pop-up menu in the upper left. To add users to the group, I choose **New User** from the **Access** menu and type in the appropriate information before selecting “NF Staff” in the **Group name** pop-up menu.

New User-1

User name:
Tom Dell

Password:
••••••

Group name:
NF Staff

Can create databases
 Can create tables

Creating a User

Unless I change it, the password defaults to the one I chose previously for the entire group.

To finish the job, I first select the group “NF Staff” in the pane on the left side of the Access Privileges window, and then I pick the database over which I want that group to have exclusive control, “conference_db,” in the pane on the right.

Access Privileges

Changes will not take effect until the Butler server is restarted.

Select a user/group: NF Staff

Set database table access rights: Databases

Deny All
Grant All

Name	Table	Shr	Pro	Excl	Sel	Ins	Upd	Del	Ind
<input checked="" type="checkbox"/>	Cars_db	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	Classified_db	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	conference_db	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	dal_demo	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	event_db	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	hangman_db	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	Houses	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	Seminar_DB	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Privileges Enabled for Access Privileges

At this point, clicking the **Grant All** button gives the group complete access to the database. If I wanted to, however, I could reduce these privileges to any combination of the following DAL statements:

Shr (SHARED)	Open database or table with SHARED permission.
Pro (PROTECTED)	Open database or table with PROTECTED permission.
Exd (EXCLUSIVE)	Open database or table with EXCLUSIVE permission.
Sel (SELECT)	Use the SELECT statement to retrieve rows.
Ins (INSERT)	Use the INSERT statement to add rows.
Upd (UPDATE)	Use the UPDATE statement to modify existing data.
Del (DELETE)	Use the DELETE statement to remove rows.
Ind (INDEX)	Use the CREATE statement to create indexes and the DROP statement to delete them.

Privileges for Groups

Changes made here do not take effect until Butler SQL Server is restarted.

Setting up Ports

Next to the Butler SQL Server application's Users window is the Ports window. Here are displayed the *port types* over which Butler SQL can communicate.

Program Linking Port

This port permits any Macintosh running System 7 or later to access a Butler SQL server via Program Linking.


Communications Toolbox Port

This port allows System 6 and System 7-equipped Macintoshes to communicate with Butler SQL using the *Apple Data Streaming Protocol (ADSP)* tool.

TCP/IP Port

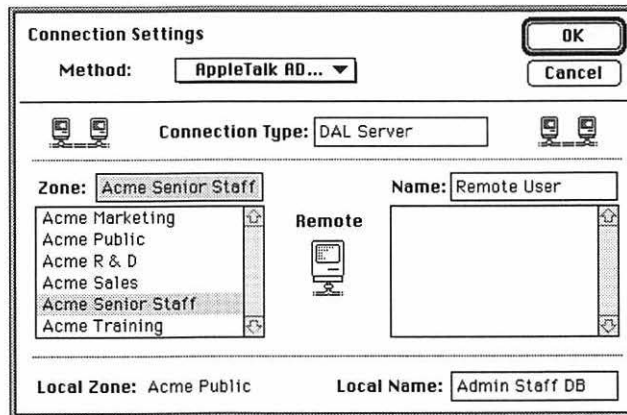
These ports permit Macintosh, Windows, and other remote client applications to connect to Butler SQL servers via the networking protocol used on the Internet, even via the Internet itself if so desired. When creating a port, select the one that uses the appropriate protocol for a given application. For example, for System 6 users, use the ADSP tool, and for Windows 95 clients, use the platform-independent TCP/IP.

You can make as many ports as you want—one for each database perhaps—but be aware that each takes up additional RAM. Butler SQL creates one Program Linking port automatically as a default. To create additional ports, choose **New Port** from the **File** menu.



Choosing a Port Type


Here you can choose the type of port you wish to create and, in the case of communications tool and TCP/IP ports, configure it.



Setting up a Communications Tool Port

In the case of the communications tool, this involves making sure **AppleTalk ADSP Tool** is selected in the window that is generated by clicking **Configure...**, calling the **Connection Type** “DAL Server,” and making the **Local Name** something meaningful like “Admin Staff DB.” You can ignore the **Zone** and **Name** fields.

In the case of TCP/IP, this means changing the Butler SQL-designated port number of “8682” to something else if you have another IP service running on the port already (which is exceedingly unlikely). Remember this number when setting up your firewalls.



Setting up a TCP/IP Port

Your ports should be saved to the Ports folder within the Butler Preferences folder. To bring them into play, choose **Install Port** from the **File** menu.

CREATING A BUTLER SQL DATABASE

There are several steps to creating databases for Butler SQL to serve. These *normalized* databases store information in *tables*, which contain *columns* and *rows*. By way of example, consider what an inventory database might look like for our fictitious company, Acme Rubber Chicken.


<i>Inventory</i>		
<i>product name</i>	<i>product code</i>	<i>number in stock</i>
<i>Economy Rubber Chicken</i>	<i>100</i>	<i>26,345</i>
<i>Classic Rubber Chicken</i>	<i>200</i>	<i>19,234</i>
<i>Electronic Rubber Chicken</i>	<i>300</i>	<i>6,123</i>
<i>Water-filled Rubber Chicken</i>	<i>400</i>	<i>13</i>

<i>Price</i>		
<i>product code</i>	<i>cost</i>	<i>price</i>
<i>100</i>	<i>.21</i>	<i>1.99</i>
<i>200</i>	<i>.43</i>	<i>2.99</i>
<i>300</i>	<i>1.26</i>	<i>4.99</i>
<i>400</i>	<i>16.84</i>	<i>89.99</i>

Here we have a single database with two tables, one for product inventory and one for merchandise price. There are columns for such things as product name, product code and number in stock, and rows—really *records*—for each kind of rubber chicken available.


Using ButlerTools

To create such a database using ButlerTools, choose **New Database** from the ButlerTools **File** menu. Name the database something logical, such as “inventory_db,” before saving it to the Public Databases folder in the Butler Preferences folder, where Butler SQL looks for databases by default. After doing so, you are greeted by an empty window.



Initial Database Window

To add tables to this Database window, choose **New Table** from the **Schema** menu. You are asked for the table's name, in this case "inventory," and are then presented with a Table window.



Initial Inventory Table Window

Add columns to this table by choosing **New Column** from the **Schema** menu.

The screenshot shows a dialog box titled "UntitledColumn3" Column. It contains the following fields and values:

- Column name: product_name
- Column title: Rubber Chicken Product Name
- Data type: Varchar
- Initial value: Null
- Size: 40
- Decimals: 0
- Remarks: These are the current product names as established by the strange little people in marketing.

At the bottom of the dialog are three buttons: "Cancel", "OK & New", and "OK".

An Initial Column

In the **Column name** field, type a brief “product_name.” You can put 31 characters of any type in this field, but no spaces. In the **Column title** field, add a less abbreviated name, such as “Rubber Chicken Product Name.”

In our example, there are three choices I could make in the **Data type** pop-up menu to designate our rubber chickens. The **Char** data type is a fixed-length character string of no greater than 255 characters. The **Varchar** data type is a variable-length character string with a maximum of 32,767 characters. A **LongChar** data type also is a variable-length character string, but it displays only the current number of characters in a data item. Because the names of our rubber chickens vary in length but are relatively short, I choose **Varchar** in the **Data type** pop-up menu. I type **40** in the **Size** field, which should be plenty of characters for these names. I also add a comment in the **Remarks** field for my own purposes. A selection in the **Initial value** field is not relevant here.

Clicking the **OK & New** button adds the column to the database. Next, add the rest of the columns. . .

The screenshot shows a window titled "Inventory" Table. At the top, it displays "Owner: TomDell", "Number Of Columns: 3", and "Row Size: 13 bytes (max)". Below this is a table with the following columns: Column Name, Data Type, Size, Dec, Title, Initial Value, and Remarks.

Column Name	Data Type	Size	Dec	Title	Initial Value	Remarks
product_name	Varchar	...40	0	Rubber Chicke...		These are the ...
product_code	Char	5	0	Rubber Chicke...		Here are the ...
stock	Integer	4	0	Number of Ru...	0	Here is the nu...

Complete Inventory Table Window

... and the price table, in turn.

The screenshot shows a window titled "inventory_db". It contains a table with two columns: "Table" and "Row".

Table	Row
inventory	0
price	0

Complete Database Window

Of course, there is a great deal more to know about setting these things up, but for that I refer you to EveryWare's (very, very long) manuals. These, plus some great tutorials and sample databases, are part of EveryWare's Butler SQL Test Drive included on this book's accompanying CD-ROM. This example should have at least given you a taste for it.

INTERFACING WITH A BUTLER SQL SERVER

Now that the Butler SQL server is available, let's look at a few ways to hit it from the client side.

Those who wish to access Butler SQL from the MacOS should have at least a Macintosh Plus with System 7 or later and a hard disk drive. One MB of RAM each is needed for the ButlerClient and ButlerHosts applications. Butler SQL can also be accessed by Macintoshes and PCs running Windows 3.1, Windows 95, or Windows NT if they have ODBC support. If you have other platforms with TCP/IP connectivity and a WebSTAR server, that is another option.

Client Connections

Here are the parts each Macintosh client needs:

Port Types	DAL/DAM	ODBC
Program Linking	Yes	Yes
Communications Toolbox	Yes	No
TCP/IP	No	Yes
Used for Client Configuration	DAL/DAM	ODBC
ButlerHosts	Yes	No
ODBC Setup	No	Yes
Control Panel/Extensions Needed	DAL/DAM	ODBC
ButlerAccess	Yes	No
ButlerLink Access	Yes	No
Butler DAL	Yes	No
ButlerFC Access	Yes	No

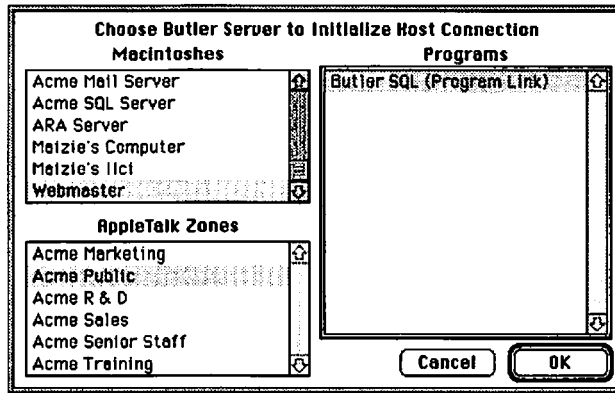
Butler SQL ODBC Driver	No	Yes
ODBC Setup Control Panel	No	Yes
Shared Library Manager	No	Yes
ODBC Configuration Manager	No	Yes
ODBC Cursor Manager	No	Yes
ODBC Driver Manager	No	Yes
Preferences Stored	DAL/DAM	ODBC
ButlerHosts Data	Yes	No
ODBC Preferences	No	Yes

We will go through each in turn.

DAL/DAM Connections

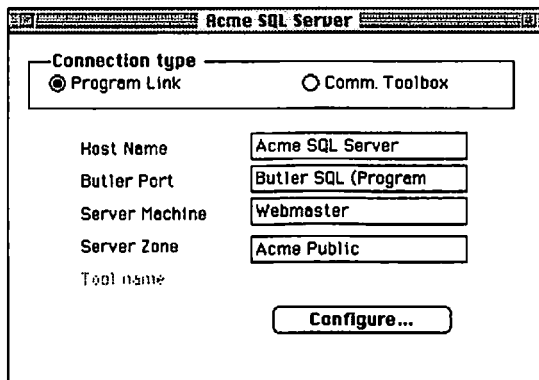
The ButlerHosts application is used to create configuration settings for client DAL/DAM connections. These configuration settings are called *hosts*, which can be a bit confusing since host is such a generic term. A ButlerHosts data file is created in the Preferences folder in the System Folder on the client Macintosh to hold these settings.

To establish a host, launch the ButlerHosts application and choose **Create Host** from the **Edit** menu. Next, choose **Program Link** in the Host Configuration window and click the **Configure...** button. Here you can select the Butler SQL server to which you wish to connect.




Connecting to a Butler SQL Server

Now name the host something logical in the Host Name field.




Naming the Host



This name will identify the settings in the ButlerHosts Data window. Close the Host Configuration window and choose **Save** from the **File** menu.


ODBC Connections

The ODBC control panel is used to configure these connection settings. In this case they are called *data sources*. Click **Setup Data Sources** in the control panel to reach the ODBC Data Sources window. Click the **Add...** button.



ODBC Data Sources Window

First you are asked to choose the ODBC driver with which you wish to work, in this case *Butler SQL Driver*.



Choosing the ODBC Driver

Choose the **Link Type** in the next pop-up menu, **Program Linking** or **TCP/IP**, and then click **Configure...** to set it up in the same manner as we configured the DAL/DAMS connection.

The screenshot shows a dialog box titled "Configure Data Source Connection". It contains the following fields and options:

- Data Source Name:** Acme SQL Server
- Description:** (empty)
- Database:** conference.db
- Link Type:** Program Linking
- Zone:** Acme Public
- Machine:** NFI Portable Web Server
- Port:** Butler SQL (Program Link)
- Allow Asynchronous Operation**
 - Full Asynchronous Mode
 - Competible Mode

At the bottom of the dialog are three buttons: "Configure...", "Cancel", and "OK".

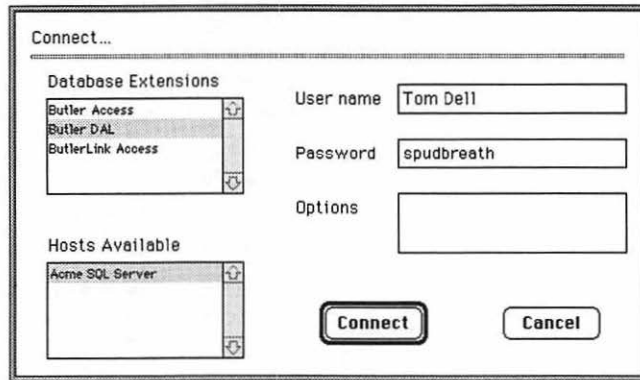
Configuring the Data Source Connection

Clicking the **Configure...** button here allows you to fill in the **Port**, **Machine**, and **Zone** fields if you are connecting via Program Linking. If you are connecting via TCP/IP, type the IP address or host name of the machine on which Butler SQL is running in the **Machine** field. The **Port** field should match the port number we set up earlier, probably **8682**.

Using ButlerClient

ButlerClient is a basic SQL client application that provides you with one ugly way to access the data in a Butler SQL database. Upon launch, it presents you with an empty document window in which you type up SQL programs for the server to execute. The data returned by Butler SQL after the program is run is also displayed here.


To connect to Butler SQL, choose **Connect** from the **Query** menu.



ButlerClient Connect Window

In the Connect window that appears, you see the names of the database extensions installed in the client Macintosh's System Folder. Select one along with a Butler SQL server in the **Hosts Available** pane. Finish by typing in your user name and password and click the **Connect** button.

Now to send, execute, and retrieve the results of a SQL program, just type a little SQL program into the window.



Sample SQL Program

Next, activate **Auto Results** in the **Options** menu, position the cursor at the end of the text, and click **Enter**. If all goes well, ButlerClient sends the selected SQL text to the server, executes the program at the server, and retrieves the results to your window for viewing. In the next illustration, the results were "5010 Acme Manufacturing Co."

```

OPEN DATABASE 'daLdemo';

SELECT cust_num, cust_name
FROM customer
WHERE cust_num < 5100
ORDER BY cust_name;

PRINTALL;

CLOSE DATABASE;

5010 Acme Manufacturing Co.

```

Results Returned

At this point I am going to make a couple of daring assumptions:

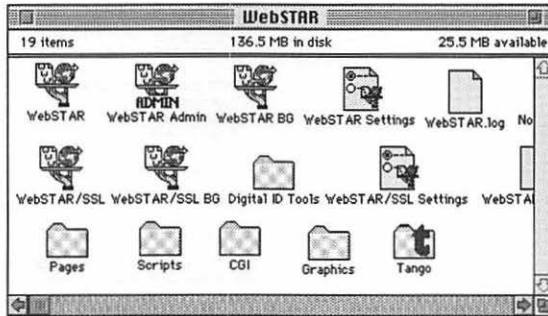
1. You are not a database programmer and are not writing a front-end in C++ so that your users can more easily access this data.
2. Your users are not adept at whipping up SQL programs.

If I am wrong, you can stop reading now. If I am right, let me show you how to create a super HTML front-end to Butler SQL and other databases using EveryWare's Tango.

Using Tango


The better way to access Butler SQL databases, especially if they use the World Wide Web as a front-end, is with EveryWare's add-on product Tango. This development tool lets you integrate your WebSTAR server with Butler SQL and other ODBC-compliant SQL database systems in much less time than it would take to program custom CGIs in AppleScript, C, or Perl. The ODBC support is particularly nice, allowing you to create Web pages that access different types of databases without you needing to know the details of how each DBMS works. I, for one, feel good about that.

Configuring Tango's interaction with WebSTAR is easy—just throw the Tango folder in the WebSTAR folder.



Tango Installed to Work with WebSTAR

In the Tango folder are two important applications: Tango Editor and Tango ACGI. The Tango ACGI accepts requests from Web browsers, queries the Butler SQL server, and returns the subsequent results.



A browser interacts with WebSTAR, which passes requests to Tango ACGI, which then queries the Butler SQL server.

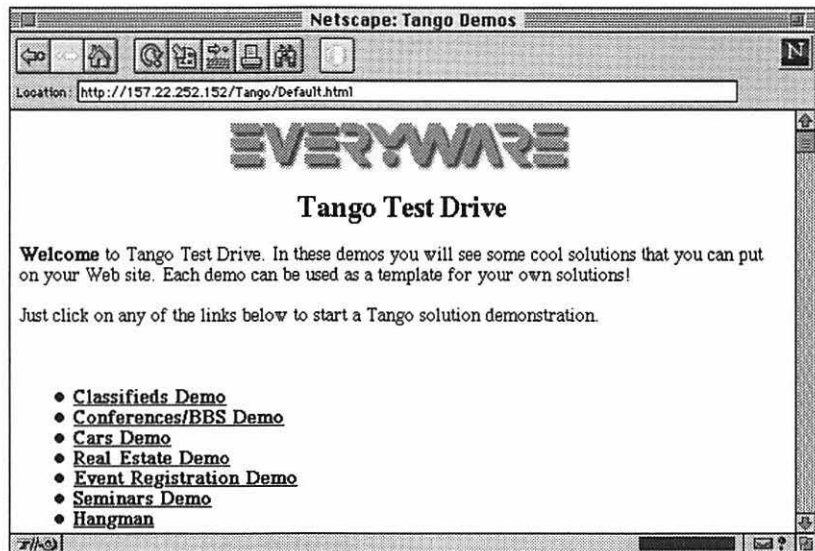
It is invoked with the following type of HTML container:

```
<A HREF="http://www.netfrontiers.com/Tango/Default.html">Tango</A>
```

Tango Editor is an application used to create *query documents*. Query documents contain instructions used by the ACGLI to perform the correct actions in communicating with Butler SQL and for formatting the results in HTML.

Using Tango Editor, query documents can be created through the use of *query builder documents*, which perform basic database functions. That way, you need write neither SQL nor HTML in making query documents!

For great examples of how this all works, go through the EveryWare Tango Test Drive included with the CD-ROM that comes with this book.



Example from the Tango Test Drive

CHAPTER 9: MAKING YOUR FORMS SMART WITH INFORMED

Are you reading this at the office? If you are, take a look at all the accumulated paper documents that you have. You have memos, letters, forms, computer print-outs, facsimiles, a “to do” list or two, and a bunch of Post-its™. You are not alone, either. Your co-workers have a bunch of paper documents, too. Your company has file cabinets full of them. Some companies have warehouses full of them. No company is without them.

All these paper documents, and of course the things that are scratched upon their surfaces, take up a lot of room. They take up a lot of people’s time, too. They have to be filed, categorized, searched for, archived, authenticated, and sometimes carefully destroyed. We humans spend much time and energy on our paper documents, and what do they do for us?

They hold and distribute information. That is all.

A paper document cannot vouch for the information that is put on it. You tell an invoice that \$200 an hour times eight hours is \$160 and it will not disagree with you. You tell a check that you are Sally Spudbreath when you are really Joe Crook and it will not challenge you to prove it. Even if you are assured that the information on a paper document is correct as recorded, the document cannot ensure you will read it back correctly. Handwriting can be messy, grammar can be unclear, and what is obvious to one human can be cryptic to another. Although they try,

paper documents cannot make sure that you encode all the necessary information on them. Despite all the bold type, many Americans still forget to sign their tax returns every year, don't they? Finally, paper documents cannot account for their whereabouts. Have you ever been told that "the check is in the mail?" A check cannot tell you when it is "in the mail," in the secretary's desk, on the boss' dashboard, or in the hedge next to your mailbox. When you finally receive it, it cannot tell you where it has been and why it has taken so long to arrive.

We have come to accept this level of stupidity in our documents because they are made of paper. However, paper is not our only choice for document encoding anymore, is it? We now have computers and their electronic media. This media has none of the same limitations. Tell a spreadsheet that 200 times eight is 160 and it will correct you. Tell an electronic check that you are someone else and then fail its authentication routine and you will be denied the money. Type with an electronic font and it will be legible, no matter how bad your penmanship is. Forget to fill in a field on an electronic form and it will not let you close the document until it is completed. Send an electronic document over a network and you can determine when it was sent, when it was received, and what path it took.

Electronic documents are as smart or smarter than we are, and they do not require nearly as much time to maintain as paper documents. Because of their inherent superiority over paper, electronic documents can greatly improve the circumstances in the office in which you are probably working right now. In several sections of this book we show you how. In this particular chapter, we show you how forms can be automated on your AppleTalk network using an impressive collection of products from Shana Corporation aptly named Informed. More than 30% of all business documents are forms, according to Shana. Informed permits you to make new electronic forms, or digitize existing paper ones, to give them the intelligence only computers can make possible. Once a form is created, it can be filled out via the keyboard, routed via an e-mail system, and authenticated using digital signatures. Data in the form can be verified, amended, and incorporated into a database system. It all takes less time, incurs fewer mistakes, and takes up a fraction of the usual storage space.

INFORMED SYSTEM COMPONENTS

The best way to visualize how electronic forms could revolutionize the work flow on your AppleTalk network is to see Informed in action. Before we show you that, however, we want to introduce you to some of its components.



Informed Forms

Informed digital forms have all the same features as paper forms. Print them out and they become paper forms just like any others. In their electronic environment, however, these forms are smart. They can look up information for you, correct you when you are wrong, offer suggestions about what answers you should make, and tell you where they are going and where they have been.



Informed Designer

This application can be used to create electronic forms. You could do this with any decent page layout program, such as Aldus PageMaker or Quark XPress, but Informed Designer also lets you add features such as calculations, defaults, choice menu lists, online help, and cross referencing to other forms. It can format *cells* for text, character, name, date, number, time, Boolean, picture, and digital signatures. AppleScript can also be used to program new functionality.



Informed Manager

This application is used to fill out and route electronic forms. It is the brains behind most of the system. It performs calculations with its built-in *if/then/else* logic, providing over 160 functions including mathematical, statistical, trigonometric, logarithmic, Boolean, text, date, time, name, spell, number, amortization, bond, depreciation, and cash flow functions. It checks formulas and verifies data. It can do a search on any cell using exact match, greater than, greater than or equal to, less than, less than or equal to, not equal to, starts/ends with, contains, and range options. It permits sending form enclosures through numerous e-mail systems. It slices. It dices. It is almost as good as the Sledge-O-Matic.

Informed Number Server

This handy server application uses Informed Manager as a client to assign unique numbers to multiple users over a network for such things as purchase orders and job requests, eliminating the possibility of duplicate transactions.

Informed Tracker

This application eliminates the time-wasting and annoying process of tracking forms through their distribution processes. It provides detailed information about the location of a specific form at any given time.

Informed Database Extensions

This software provides integration with many database systems, including Sybase, Oracle, FileMaker Pro, Office Tracker Pro, Mac P&L, and TouchBase Pro. Forms can look up information in your existing database and submit new information to that database automatically.

Informed E-Mail Extensions

This software provides integration with many e-mail systems, including MS Mail (now Quarterdeck Mail), QuickMail, cc:Mail, Eudora, and PowerTalk.

INFORMED IN ACTION

Explaining all the possible functions and deployments of Informed is beyond the scope of this book and is well documented by Shana already. What we will do here is demonstrate a few of Informed's abilities in a hypothetical situation. To do this, we will create a form with Informed Designer and fill out the form with Informed Manager. Version control will be assured by our use of Informed Number Server, and we will maintain watch on the document's whereabouts with Informed Tracker. We will use an Informed DAL database extension to access a Butler database. Also used here will be Apple's PowerTalk, digital signatures, and AppleScript.

The following was accomplished using a demonstration created by Shana. We have put this demo on the CD-ROM that accompanies this book, so that you can try it yourself.

Shana's demonstration starts with a user, in this case me, requesting a software purchase by filling out a requisition form using Informed Manager.

Purchase Requisition Form When First Opened

Notice that this application fills in some basic information for me, like the date.


Informed™ Number Server		
Form Name	Next Form Number	Increment By
Purchase Requisitions	PR-1278	1
Purchase Orders	PO-1014	1

Informed Number Server

It also communicates with Informed Number Server to obtain the next available unique purchase requisition form number, placed in the upper right of the form.

To direct Informed Manager to fill in more information for me, I enter my employee number “E100” in the **Employee No** cell and press the **tab** key. Using Program Linking, Informed looks up the information associated with this number in a Butler DAL database running elsewhere on the network.

Users		Ports	
User Name	Usage	Port Name	Users
		Informed Tracker Se...	0
		World_Corp Port	0

Resources	
	Users <input type="text"/> (max) 10
	Processes <input type="text"/> 10
	Memory <input type="text"/> 2000K

Version 1.2.7

Butler Employee Database

This number retrieves an employee’s name, mail stop, department, office name, and office number, and then Informed enters them on the form.

Purchase Requisitions

WORLD CORPORATION 12345 - 123 Street
New York, NY 12345

Purchase Req. #
PR-1277

Ship To:

Employee No	Recipient Name	Mail Stop	Order Date
E100	Tom Dell	227	4/25/96
Department to Charge		Office Name	Office Number
National Accounts		World Corp HQ	255

Orders normally take 7 to 10 days for delivery from date of order

Part #	Description	Qty	Price	Total
ORDER TOTAL				

5 Forms 5 of 5

Database Lookup for Employee Information

Although Shana's demonstration uses Butler to show off this ability, Informed database extensions support a variety of databases, including Oracle, Sybase, DAL-supported databases, and SequeLink-supported databases (such as DB2, Ingres, Informix, OS/2 Data Manager, and AS/400). Do not worry about it fitting in with your existing systems unless you have a really unusual database.

Errors and omissions on forms, as we all know, can be costly. Depending on how many people a form must pass by, the first mistake can spawn other mistakes as well, adding to the cost. Even if there is a person dedicated to finding such mistakes before the cost becomes big bucks, there is still a cost associated with that service. Arguably Informed's most important feature is its errors and omissions checking that certainly reduces user mistakes and, in some cases, eliminates them.

To demonstrate this error checking, I will delete the information in the **Department to Charge** field. When setting up this form, its author established that this field had to contain a value. Thus, when I press the **tab** key, I receive an error message. I type in "Sales" and the application permits me to continue.

The screenshot shows a 'Purchase Requisitions' window for 'WORLD CORPORATION'. The form includes fields for 'Employee No' (E100), 'Department to Charge', and 'Orders no'. A pop-up error message with a hand icon states: 'You MUST enter the department number to process this form.' The window title is 'Purchase Requisitions' and the form title is 'PURCHASE REQUISITION'. The address is '12345 - 123 Street, New York, NY 12345'. The 'Purchase Req. #' is 'PR-1277'. At the bottom, it says '5 Forms', '5 of 5', and 'ORDER TOTAL'.

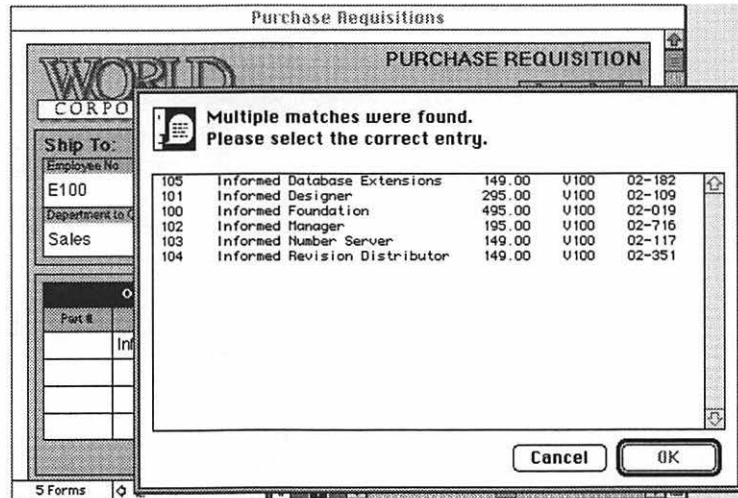
Warning of an Omission

Another way to eliminate errors is to put suitable information in choice list menus. This also saves on the time users would spend typing. When I pressed the **tab** key, it moved me to the **Office Name** field. Here a pop-up menu appeared in which I could choose a properly spelled designation for my fictitious home office. I selected the choice directly with the mouse, but I could also have just typed a character or two of the value to select it.

The screenshot shows the same 'Purchase Requisitions' window, but now with a choice list menu open over the 'Office Name' field. The menu lists: Boston, Chicago, Hawaii, New York, San Francisco, and San Jose. The 'Office Name' field is filled with 'World Corp HQ'. The 'Department to Charge' is 'Sales' and the 'Office Number' is '255'. The 'Employee No' is 'E100' and the 'Recipient Name' is 'Tom Dell'. The 'Orders no' field contains '227' and '4/25/96'. A note above the table states: 'Orders normally take 7 to 10 days for delivery from date of order'. The table has columns for 'Part #', 'Description', 'Qty', 'Price', and 'Total'. The window title is 'Purchase Requisitions' and the form title is 'PURCHASE REQUISITION'. The address is '12345 - 123 Str, New York, NY'. The 'Purchase Req. #' is 'PR-1277'. At the bottom, it says '5 Forms', '5 of 5', and 'ORDER TOTAL'.

Using a Choice List

In the same way that we can direct Informed to look up employee information, we can have it look up product information. Suppose I want to purchase Informed Foundation, a suite of applications. I simply type “Informed” into the **Description** field and press the **tab** key.



Database Lookup for Product Information

This field is designed to look up the part number and price of the item as entered in an order table. It was not necessary to enter the full description of the item I wanted to order, which is fortunate since I could not remember the exact name. In the order table, I select the item I want and click **OK**. The full description, part number, and price are all entered in the appropriate columns of the table without mistakes. Next, I enter the quantity I want and press **tab**. The order total is figured for me automatically, again without errors.

Scrolling down in the form takes me to the signature area.

Purchase Requisitions

Orders normally take 7 to 10 days for delivery from date of order

Part #	Description	Qty	Price	Total
100	Informed Foundation	1	495.00	495.00
ORDER TOTAL				\$495.00

Special Instructions:

For Management Use Only

Reject Purchase

Requestor's Signature _____ Approval Manager's Signature _____


5 Forms 5 of 5

The Signature Area

Here is where many electronic form systems have fallen short in the past: verification. A signature is legally binding, it was argued, and since there is no digital equivalent to a person's signature, electronic forms are not legal. Recent technologies, such as the digital signature ability of PowerTalk, counter this argument, and Informed makes use of this technology.

There are two main reasons why a person signs a form. First, it is to identify the originator of the information on the form. Second, it is to verify that a person is aware of and approves of the form's contents. In this demonstration, I put my digital John Hancock on the form in such a way that its integrity can be verified. I then forward it to Dorian for him to "sign off" on it.

To enter my signature, I select the **Requestor's Signature** field and then choose **Sign** from the **Signatures** submenu in the **Edit** menu. PowerTalk then asks me to select a *signer*.

*PowerTalk Signer*

After I fill in my password, the digital signature is created and attached to the form and those cells designated as *signed* are locked to prevent others from changing them. Informed forms can have multiple signatures, and each signature cell can sign different information on the form, supporting multiple fillers and multiple levels of approval. In this demonstration, my signature signs all cells except the **Reject Purchase** check box, the reject reason, and the approval signature. Dorian's signature will sign all the cells on the form. With the signature cell selected, I choose **Show Signed Cells** from the **Signatures** submenu in the **Edit** menu to see which cells a particular signature signs. All cells that are signed by that signature appear framed in red.

Part #	Description	Qty	Price	Total
1100	Informed Foundation	1	495.00	495.00
ORDER TOTAL:				\$495.00

Showing Signed Cells

Once the PowerTalk signer has been invoked, my name and the current date appear next to the signature icon in the **Requestor's Signature** cell. The signature

icon has three variations. A check mark on the icon means the signature has just been created or verified and is valid. A question mark signifies that the signature has not been verified and its validity cannot be assured. An X indicates that the signature failed verification. Either the signed data or the digital signature itself has changed since the form was signed.

Purchase Requisitions

Orders normally take 7 to 10 days for delivery from date of order

Part #	Description	Qty	Price	Total
100	Informed Foundation	1	495.00	495.00
ORDER TOTAL				\$495.00

Special Instructions:

For Management Use Only

Reject Purchase

Requestor's Signature: Tom Dell 4/25/96

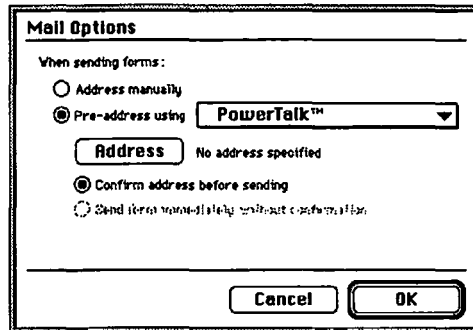
Approval Manager's Signature

5 Forms 5 of 5

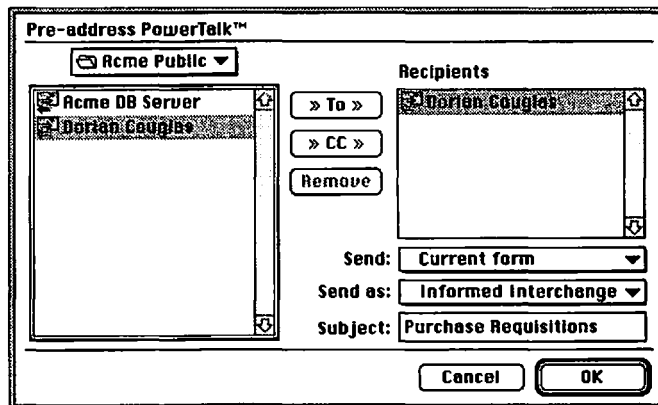
A Signed Form

Needless to say, digital signatures are more secure than handwritten ones.

Next, we call upon PowerTalk again, this time for its messaging abilities, as I forward my purchase requisition to Dorian. Since I send most purchase requisitions to the same person most of the time, I can use Informed's pre-address feature to automate the addressing and mailing of these forms. To do this, I choose **Mail Options** from the **Mail** submenu in the **File** menu.

*Mail Options Window*

In the Mail Options window, I select the **Pre-address using** radio button and **PowerTalk** in the pop-up menu. Clicking the **Address** button brings up the PowerTalk addressing window. I navigate through PowerTalk *catalogs* to find Dorian.


*Finding the Recipient in PowerTalk*

Other controls in the Pre-address PowerTalk window let me choose which forms to send—current form or collected forms—and their data format. Informed Manager lets you send forms with the layout through **Informed Interchange** or without the layout.

With the form now pre-addressed, I send the form by choosing **Send** from the **Mail** submenu in the **File** menu. Informed generates a dialog box telling me that the form is being sent and that the form is being *tracked*. A transaction has been

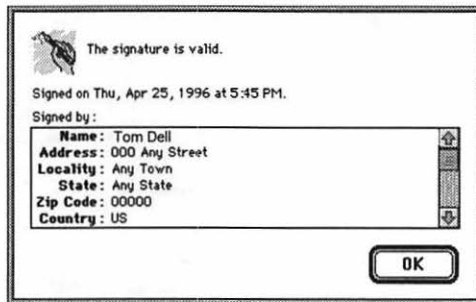
entered in the Butler DAL database automatically noting the date and time along with the names of the sender and recipient of the form.

Before long, the purchase requisition appears in Dorian's PowerTalk mailbox. Now it is his turn.



The Form Reaching the Recipient

Dorian double-clicks on the form in the mailbox and Informed Manager launches to open it. Before approving the form, Dorian ensures that I have not tried to sneak in another copy of Marathon; then he verifies my signature to make sure no one has tampered with the order. To do this, he selects the signature and chooses **Verify** in the **Signatures** submenu in the **Edit** menu. If Tom Hessel has not been hacking at it, a dialog box appears to tell Dorian that my signature is valid.



Valid Digital Signature

This dialog box also tells Dorian the date and time the signature was created and, unfortunately, where I live so he can come looking for me when the bill comes in.

To approve the form, Dorian uses the same procedure I used when signing the form and adds his digital signature to the bottom.

Dorian is now responsible for creating and submitting a purchase order to fulfill the request. That task can be automated using AppleScript. When the purchase

requisition form was designed, the author created a script and attached it to the form. Its name, **Create PO**, appears under Informed Manager's **Scripts** menu. When Dorian chooses this command, the script tells Informed Manager to open a purchase order form, create a new form document, and then copy the appropriate information from the purchase requisition into it. As with the purchase requisition, the next available purchase order number is obtained from Informed Number Server and vendor information is retrieved by looking up the vendor number in Butler.

After the purchase order has been created and Dorian fills it out, the script asks him if he is ready to submit the form. When he clicks **Yes**, its data is submitted to the Butler database, and his job is done. He can now simply close the form.

The screenshot shows a 'PURCHASE ORDER' form for World Corporation. The form is titled 'PURCHASE ORDER' and includes the following information:

- Vendor:** Shana Corp, 9650-20 Avenue, Edmonton, AB T6N 1G1
- Date:** 4/25/96
- Delivery:** 60 Days, DHL - prepaid
- Table:**

Vendor Part	Description	Qty	Price	Total
02-019	Informed Foundation	1	495.00	495.00
- Dialog Box:** "Would you like to submit this purchase order?" with "Yes" and "No" buttons.

Submitting the Purchase Order Form

I, on the other hand, still have to worry about whether my order was placed and, if not, where it is—or do I? Each time a form is sent from one person to another, Informed Tracker enters a transaction in the Butler database to make note of the date and time. I can query the database to find out where a given form has been using Informed Manager. I just choose **Tracking Status** from the **View** menu.

Form Tracking Status

Form ID: PRX
Form Number: PR-1277

Date	Time	Sent From	To	Comments
04/25/96	09:48 AM	Tom Dell	Dortan Cougias	

Buttons: Save Results..., Details, OK

Tracking a Form

In our office, locating a paper form when something does not get done involves following a bizarre path through various desks, file cabinets, in/out boxes, piles of documents, and tattered manila folders. This results in success about 50% of the time. We call this “Lynn’s Special Filing Technique,” because she is the only one who understands it. Upon finding the form, there is the additional task of finding out why the paper form is in pile “X” and not in the hands of vendor “Y.” This process results in success maybe 2% of time. It automatically triggers what I like to call the “Unnecessary but Mandatory Recriminations Cycle” that culminates in my apologizing and regretting ever having initiated the circulation of the paper document to begin with. What about the forms in the unsuccessful 50%? I do not know. I do not remember them. Hope they were not that important . . .

Electronic document tracking makes it possible to know where a digital form is 100% of the time. You can see who received it when and what they did with it, when it is in-house, and when it is out the door. This often eliminates the Unnecessary but Mandatory Recriminations Cycle, too. What are you to say when you drop the ball? “Yeah, I know it says I looked at it last Thursday, but, uh, my computer’s clock was off.” Being enabled to say “Sure it was” with a smug look on my face makes me like this feature of Informed the best.

Shana thinks “the most significant benefit of tracking the processing of forms is that the tracking information can be analyzed to monitor and improve the efficiency of the process itself.” This information can be used to measure approval cycles, identify approval bottlenecks, or locate unnecessary approvals, for example. Good point, too, I guess. Efficiency is a good thing.

To help maintain the form tracking database, Informed Tracker includes the administrative application called Informed Tracker Admin. With it you can browse, remove, and export the tracking data.

Informed Tracker

Server Administration Tool

Number of rows in database: 21 Number of rows in list: 21

Form ID	Form No.	Sent From	To	Date	Time
PRX	PR-1239	Ken Nachtigal	John Murphy	10/22/94	12:10 AM
PRX	PR-1239	John Murphy	Paul Chamberland	11/18/94	12:02 AM
PRX	PR-1239	Paul Chamberland	Glenn Smith	11/19/94	12:11 AM
PRX	PR-1240	Paul Chamberland	John Murphy	10/15/94	12:09 AM
PRX	PR-1240	John Murphy	Ken Nachtigal	11/18/94	12:02 AM
PRX	PR-1240	Ken Nachtigal	Glenn Smith	11/20/94	12:11 AM
PRX	PR-1241	John Murphy	Ken Nachtigal	10/07/94	12:09 AM

Show all
 Show only form ID
 Show only form number
 Show only on or after
 Show only on or before

View of Informed Tracker Database

What we have shown you here just touches on Informed's uses and capabilities. It should have provided you with enough information so that you can see that an electronic document system could be the most important service you deploy on your AppleTalk network, second only to backup.

PLANNING FOR ELECTRONIC FORMS

Before you go forward with deploying a system of smart electronic forms, here are a few pieces of advice we would like to offer.

- Do not migrate a paper-only organization to a digital workflow system by yourself. While you might be in charge of technology at your company, this service does not have as much to do with computers as it does with how your users work together. Enlist their support and seek their input. Ask them how the paper-based systems are supposed to work as well as how they *really* work, and then work with them to find compromises. Demand support from upper management. There must be leadership at the top if you are to make a transition work throughout the organization.
- Bear in mind that this is an opportunity for reengineering. Why waste time recreating a form nobody uses because the procedures surrounding it are flawed? If there is a better way to do something, figure that out first, then determine how networking it can make it better still.
- Eschew surplusage, as Mark Twain once advised. If there are fields in the form that do not really need to be there, remove them. If they are used only occasionally, remove them also. Occasional information can go in digital “notes.” Be particularly aware of fields needed only on paper documents, like check-boxes for how many copies and fields for the signatures of proofreaders. Inform has made these unnecessary.
- Consider this: Forms follow functions. Try to identify functions that could really use the services of a good form. In our office, for instance, we kept losing serial numbers for all the software we buy. Sometimes we would buy several copies of an application because one of us did not know another had done so. I created a simple e-mail-based form that allows us to jot in the product name and serial number when we buy something, then send it off to a database and forget about it. When we need a serial number or want to know if we have a given piece of software, we query the database. Nobody likes finding and filling out registration cards, but everyone except Dorian is willing to type up a quick e-mail.

Another idea for a form came to me when we were writing *Managing Apple-Share & Workgroup Servers*. I thought that the cardboard RAID setup sheet Apple supplied with its Apple Workgroup Servers was a little cryptic, and it required manual math calculations. I made a smarter one. When we were writing the second edition of *The Complete Guide to Mac Backup Management*, I saw that it would be handy to have a way to keep track of ARCserve's workstation INITs. I made a form. Take a look at it if you want. It is on the CD-ROM. I bet you can think of lots of new forms that would increase your organization's efficiency but that no one has created yet, perhaps for fear of making more paperwork.

- Avoid forms that are officious. While officiousness is necessary in many instances, it is off-putting in others. Try to create designs that are user-friendly through the use of simple language and friendly graphics. People are more likely to use forms that they like the looks of. It is a cheap trick, but it works.
- Design digital forms to be digital. Take advantage of every capability of the software to make the form do as much of the work and the users as little of the typing as possible. Also, do not stick to the standard 8.5" × 11" document size for forms that will never be printed! Design for the screen so that users do not have to use scroll bars.
- Avoid the "five-pound bag." A great guy I worked with early in my career when I was a typesetter coined what I call *Jules' Law*. It states that every time you present a professional ad design with a good amount of white space to an advertiser they will increase the amount of display type they want in it. This goes on through each proofing until you have what is akin to "20 pounds in a five-pound bag." This type of form will be cluttered, hard to use, and ugly. If your layout starts becoming cramped, add a page or create a second form.
- Know that sometimes a paper form should remain just a paper form. Paper forms have their place in some instances. Do not try to apply an electronic form to one of these instances in an attempt to create a paperless office. That is going too far.

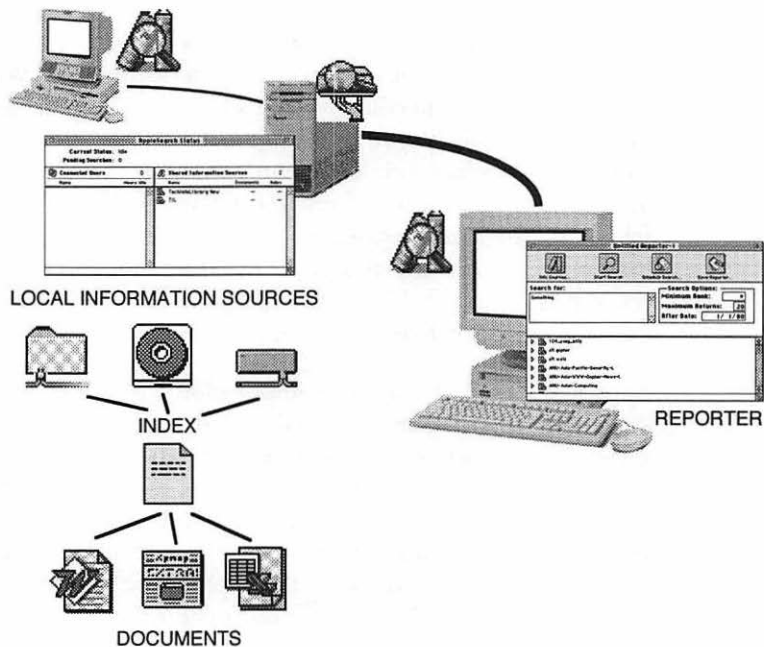
CHAPTER 10: FINDING INFORMATION WITH APPLESEARCH

You are no doubt aware of the massive volume of digital information that can be stored on your network's servers. If you have an Apple Workgroup Server with a 1-GB hard drive, you could have nearly half a million pages of ASCII text on the thing! Yikes! How do you find a particular needle in that kind of haystack? There comes a point when you have simply too much information on a server for your users to reasonably search through using the Finder—no matter how helpful your folder and file naming schemes. At this point, users either do without the information they need, or they recreate it, rather than try to find the data on the servers. When you reach this point, it is a good time to deploy a document search and retrieval system such as Apple's AppleSearch.

AppleSearch is a client/server application that helps the user find server-based data. The AppleSearch server takes that data and organizes it into *information sources* that include a given set of documents and an index of the significant text within those documents. At the client end, users submit the terms they are searching for to a *reporter*, which then queries the server immediately or at some user-specified time. Through the reporter, the server returns a list of finds ranked by how well they matched the search criteria. Users browse through the *text* that is returned, and when they find the documents they need, they can copy them down to their workstations from within AppleSearch.

Some of AppleSearch's other abilities include:

- It indexes new or modified documents automatically added to information sources at timed intervals.
- It uses DataViz' MacLink Plus XTND translators to extract text from various documents—such as reports, letters, e-mail, or spreadsheets—and presents the text to users for viewing and manipulation without them needing the creator application.
- It relies on AppleShare or Personal File Sharing access privileges for security.
- Its reporters can automatically gather information on a recurring basis, such as every week.
- It includes Windows and MacOS client software. Using a special ACgi with MacHTTP or WebSTAR, it also can be accessed through a Web browser.
- If your network has Internet connectivity, its search capabilities can be extended to Wide Area Information Server (WAIS) sites.



AppleSearch Server and Client Reporters

SETTING UP THE APPLESEARCH SERVER

If you are going to use AppleSearch solely for local information sources, the logical place to set it up is on the AppleShare server, where most of your important company documents probably are located already. Designate which *shared server folders* will become information sources. AppleSearch indexes the documents contained in these information sources to make them available for text searches. These documents can reside on any local volumes, including CD-ROMs.

AppleSearch requires a Macintosh equipped with at least a 68040 processor. It works with AppleShare 3.0.3 and later, and there is a patch available from Apple to bring AppleShare 3.0.x up to that version. If you have the old Workgroup Server 95, use the A/UX version 3.1.1 or later.

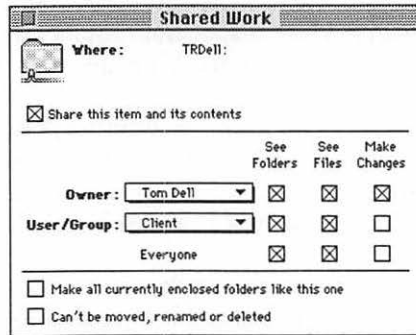
If you are going to use AppleSearch for remote information sources via the WAIS gateway only, you can add it to your Web server to work with AppleShare or Personal File Sharing. Do this only if your Web server has bandwidth to spare. If users are receiving busy messages when they log in to your home page, you should put AppleSearch up on a different machine. (See our chapter on WebSTAR beginning on page 303.)

If you wish to use AppleSearch to access both local and remote information sources, you most likely should set it up on your main AppleShare server. You then need to ensure that the server has TCP/IP installed and properly configured, so that it can “see” the Internet. Don’t worry about having TCP/IP on the client machines, as these will communicate with AppleSearch via good old AppleTalk.

Finally, you can put up more than one AppleSearch server. For instance, you can have one for local information sources only and one for remote information sources only. Each AppleSearch server can accommodate 50 concurrent users, so if you have many more active users than that, this deployment of multiple servers becomes a must.

From experience, we cannot recommend running an AppleSearch server that serves both local information sources and WAIS sources alongside a Web server. It is just too processor intensive.

AppleSearch uses the Apple Installer Script, so you should not run into any difficulties during installation. This also installs the WAIS gateway for you. Once you restart the Macintosh, you are all set to go. The only immediate configuration concerns are that your users have valid AppleShare accounts with Program Linking enabled and read privileges to those folders that contain documents that will be part of AppleSearch's information sources. Do this via Personal File Sharing or AppleShare, whichever you are running on the server machine.




Setting Access Privileges through Personal File Sharing

Give all users **See Files** and **See Folders** privileges to the folders with the information you are sharing so that they can search the documents in the designated folders. You do not have to give them **Make Changes** privileges, however, unless you want them to be able to add documents to the folders or modify documents already there. For a very complete treatment of this topic, see our book *Managing AppleShare and Workgroup Servers*.

Before you go any further, create a folder on the server called something like "AppleSearch Client" and place the contents of the Macintosh client installer diskette there. Do the same for Windows users if you have them and if they have AppleTalk connectivity. This way, your users can easily download the client software they need later.

Designating Local Information Sources

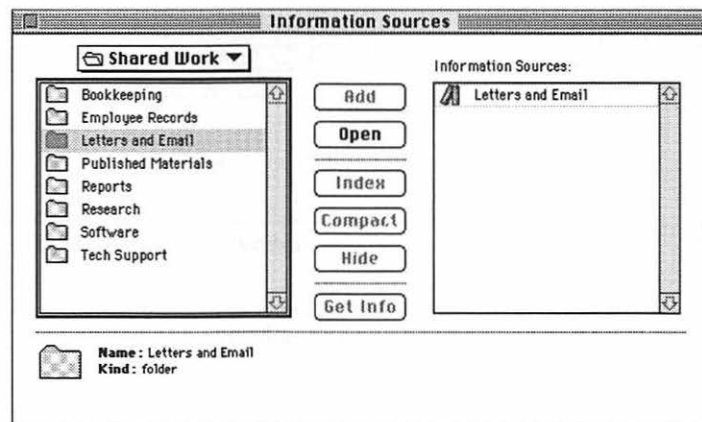
If your server's shared volumes are well organized, this process is simple. If not, you have some preliminary work to do. Look at the following illustrated example.



Example of an Organized Server Folder Hierarchy


Here we have five folders that we want to be able to search via AppleSearch and three that we do not. The folders Published Materials, Reports, Research, and Letters and Email contain the majority of this company's working texts. Bookkeeping and Employee Records are to be excluded for security reasons—which also could be accomplished with properly established AppleShare privileges, of course. Software is excluded, as it contains mostly applications and not documents. Tech Support is of special interest to you. Here is the folder in which you have placed your user how-to guides and instructions, as well as all the important “read me” files from the applications stored in the Software folder. You are very clever.

To put this foresight to work, launch AppleSearch server. Select **Information Sources** from the **Server** menu when the empty AppleSearch Status window appears. You are shown a window pane in which to select folders to be local information sources. You will not see mounted shared volumes or aliases. Highlight the appropriate folders and click the **Add** button.



Selecting a Folder to Be an Information Source


When you click the **Add** button, AppleSearch asks on which local volume it should create its index of that folder's text contents.



Designating Destination for Indexes

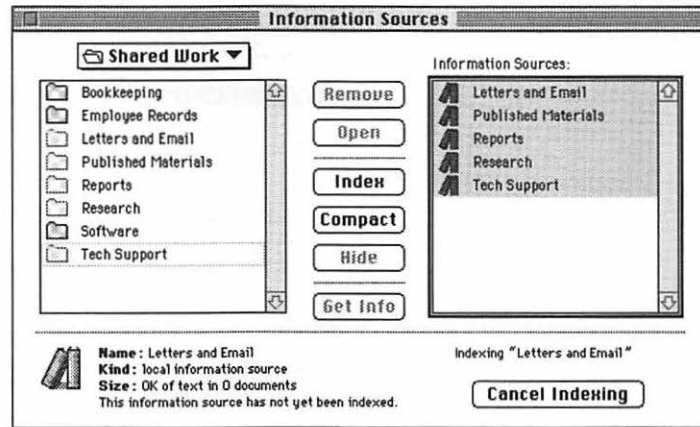
Once you have chosen a drive and clicked **OK**, AppleSearch creates a folder called AppleSearch Indexes at its root level. In this folder, it will create files that contain an index of every distinctive word in every document in that folder and in that folder's subfolders. It also indexes variations of search words. This all happens automatically. You don't have to spend your precious time designating keywords.

Make sure that the hard drive on your server has plenty of free space before you do this. You need space equal to 2.5 times the size of the information source for the index and temporary files, according to Apple. That means if you have 10 MB of assorted files to be indexed, you need 25 MB of free space: 10 MB for temporary files and 15 MB for the index. We recommend that you always have at least 20% of the hard drive space free after that. Fortunately, the space taken up by the temporary files is released when the indexing process is completed.



Automatically Created Index Files

Once you have selected your new information sources, highlight them in the pane at the left, and click the **Index** button.



Manually Indexing Information Sources

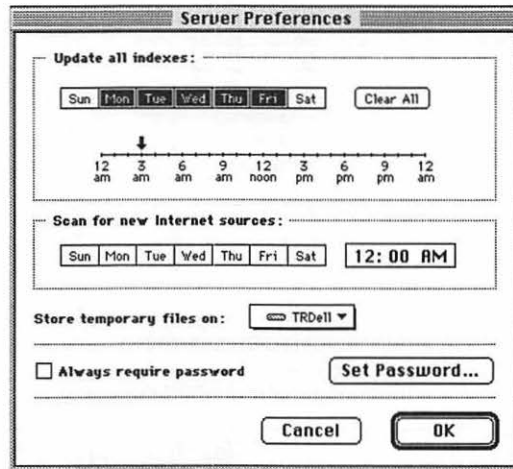
Be patient. This will take some time—up to an hour for 10–20 MB of documents. When AppleSearch is finished, you can click on any information source to see how big it is, and when it was last updated, in the lower left of the window.



Name: Letters and Email
Kind: local information source
Size: 85K of text in 20 documents
Last Updated: Tue, Dec 26, 1995, 5:38 PM

Status of an Information Source

That **Last Updated** field should concern you. Your indexes are only as good as the last time they were generated. You can repeat this process from time to time—yawn—or direct AppleSearch to do so automatically via the **Preferences** from the **Server** menu.



Choosing Days and Times for Index Updating

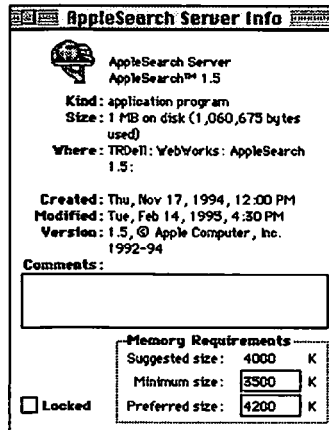
In the **Update all indexes** field, tell AppleSearch on which days and at what times to do its thing. We recommend daily, early in the morning when users are not around, and before or after the network backup takes place.

Close the Information Sources window when you are finished. Your newly designated and indexed local information sources appear for your review in the AppleSearch Status window.

AppleSearch Status				
Current Status: Idle				
Pending Searches: 0				
Connected Users		Shared Information Sources		
Name	Hours Idle	Name	Documents	Index
<Local User>	0:00	Letters and Email	20	152K
		Published Materials	0	4K
		Reports	0	4K
		Research	7	114K
		Tech Support	0	7K

Newly Indexed Local Information Sources

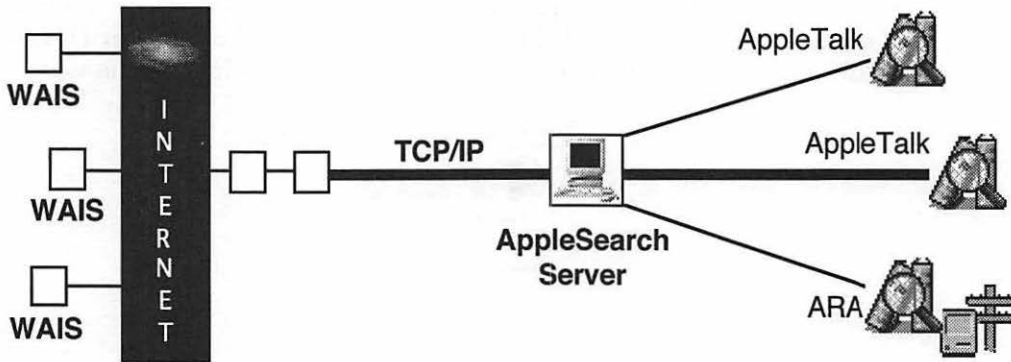
Be sure to give AppleSearch plenty of memory if you add many local information sources. The recommended allocation of 4000K is good enough for 10 sources. After that, give it another 100K for each additional local information source.



Enough Memory for a Dozen Information Sources

Designating Remote Information Sources

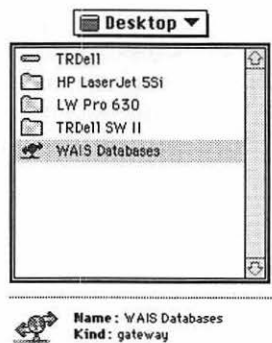
In addition to local information sources, AppleSearch gives you the power to search other types of remote information sources: *Wide Area Information Servers (WAIS)*. AppleSearch and this Internet denizen use inverted index files as their indexing method and thus work well together. They should. WAIS was originally developed by Apple in conjunction with Thinking Machines, KPMG Peat Marwick, and Dow Jones and Company. There are now about 500 WAIS servers on the Internet, put up mostly by government and educational institutions and a few large companies. A couple of the more notable sites include the Library of Congress and the CIA World Fact book. You can select up to 100 of these sites at a time to be available to your local user.



AppleSearch WAIS Gateway

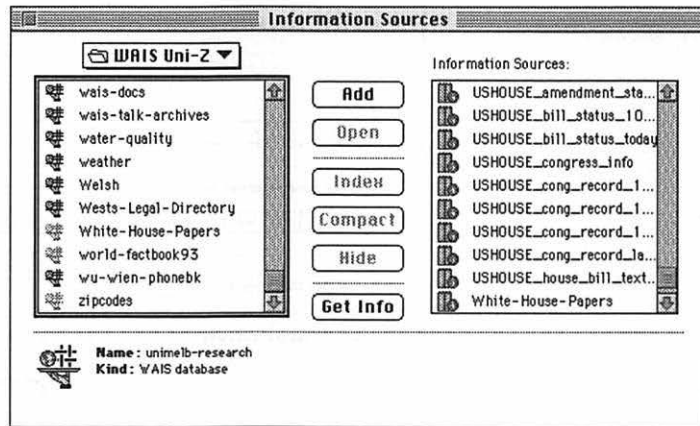
Although WAIS is an Internet service, Macintoshes on the user side do not need TCP/IP, since they access the AppleSearch server over AppleTalk. WAIS servers you select show up in the users' AppleSearch Sources windows, just like information sources on the LAN. Queries directed to the Internet go through the gateway to a WAIS directory set up by Apple, which points to all currently registered WAIS servers on the Internet. WAIS servers return a relevance-ranked list of documents matching the search criteria to the AppleSearch server, which in turn passes it on to the user. To see one of the documents, the user double-clicks in the list and the AppleSearch server retrieves it from the WAIS server. Pretty slick, no?

After you install the gateway and launch AppleSearch, it sends a query to the Internet to receive the newest list of available WAIS sources. This usually takes 10–20 minutes.



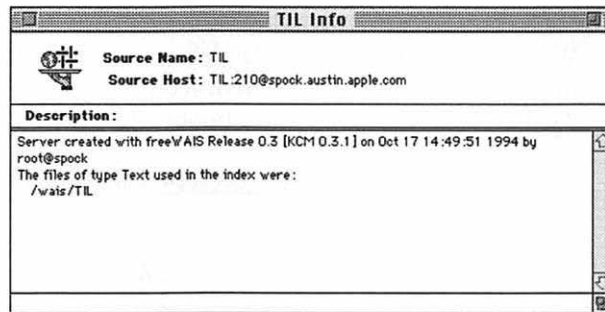
List of WAIS Sources Supplied at Launch

To use the window generated, select **Information Sources** from the **Server** menu again. Next, double-click **WAIS Databases** in the left window pane to view the libraries of data that might interest your users.



Choosing Relevant WAIS Sources

You may use the window's **Get Info** button to see more of what is available from each source.

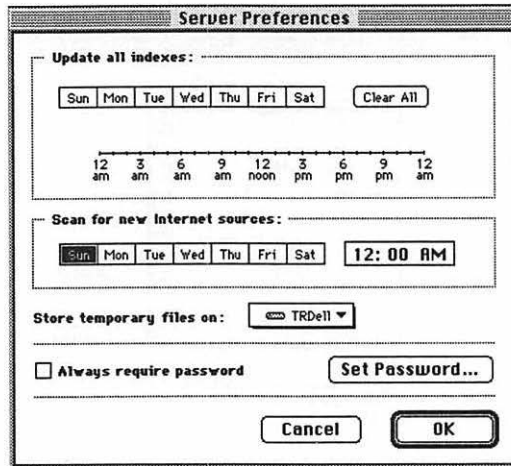


Getting Details about a WAIS Source

The one in the last screen shot should be a real help to you. It is Apple's *Technical Information Library*.

You do not have to generate a local index for these sites, but you can specify a time when AppleSearch will look around the Internet to find any new WAIS servers

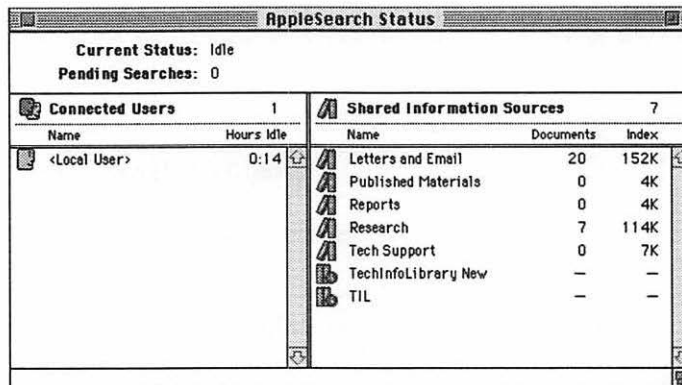
that have become available. This is accomplished by first selecting **Preferences** from the **Server** menu.



Choosing a Time to Scan for New WAIS Sites

I recommend you do this weekly and early in the morning on the weekend, when your Internet feed is relatively unused. Set some time aside on Monday morning to browse the list for any new “must have” sites.

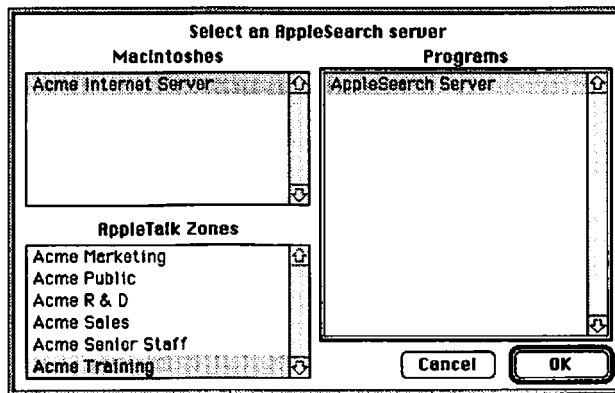
WAIS servers appear in the AppleSearch Status window much like local information sources but with a different icon. Instead of two leaning books, the WAIS database is represented by two books and a globe.



Local and Remote Information Sources

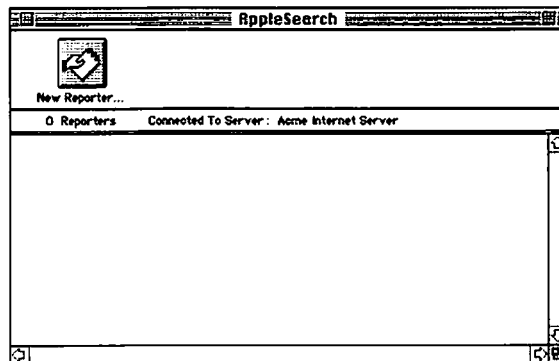
USING APPLESEARCH VIA APPTALK

On the user end, the workstation uses Program Linking to connect to the AppleSearch server. Make sure that users have Program Linking enabled in the Sharing Setup control panels of their workstations and in their AppleShare server accounts. A login screen appears when the AppleSearch client is first launched. In addition to choosing the appropriate zone and Macintosh, the user must enter an AppleShare account name and password during this process.



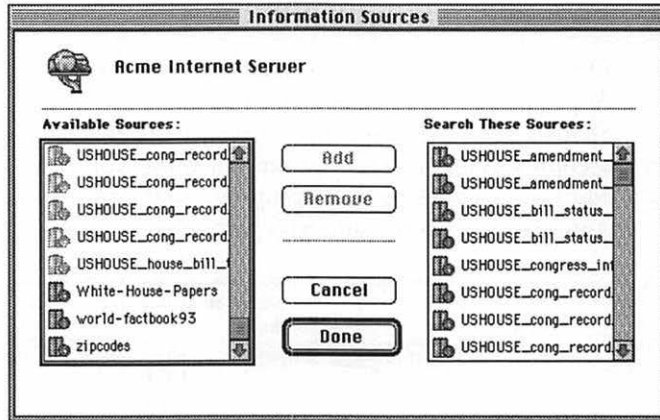
Selecting the Server via Program Linking

In the AppleSearch window that appears, users can ferret out the information they want by clicking the **New Reporter...** button.



AppleSearch Client Window

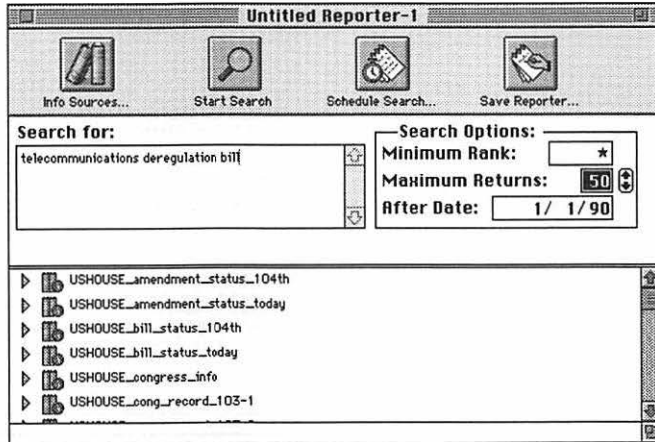
Once in the New Reporter window, users click the **Info Sources** button to choose from the databases made available through the server.



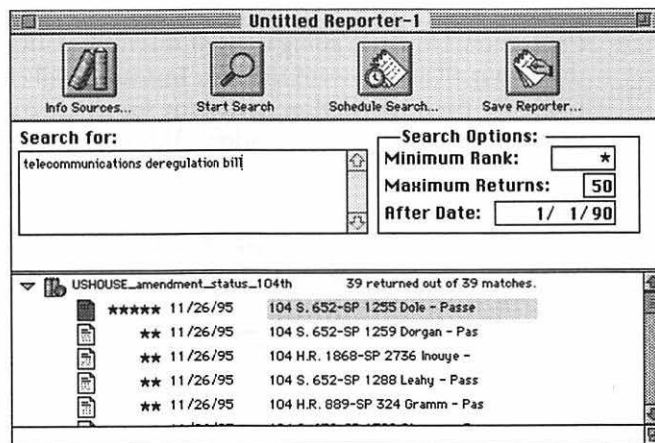
Selecting Info Sources

When your users use remote information sources, it is important that you teach them to be selective and not just shift-click all sources. That is a big waste of time, bandwidth, and processor power. For instance, in the examples given here we search for information relating to telecommunications industry deregulation proposals. Choosing the WAIS sources for the U.S. House of Representatives is a good idea. Choosing the other 80 sites that have nothing to do with it is not!

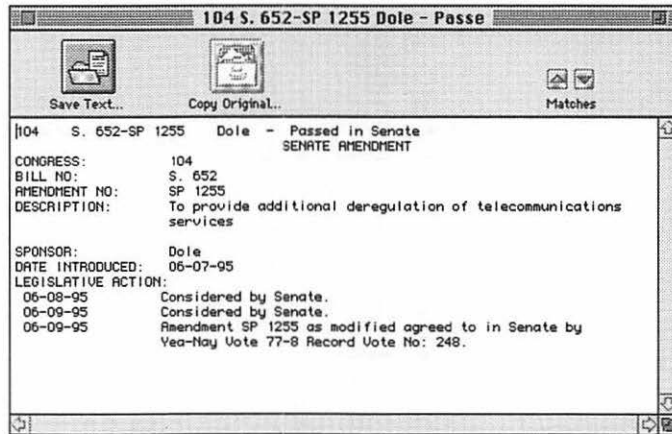
With the search sources selected, the user then enters terms in the **Search for** field and uses the **Start Search** button to retrieve data from the appropriate sources. The beauty in this is that users may enter their queries in a form of computer nomenclature lacking in many database systems—plain English (God forbid)!

*Entering a Search Request*

Users can narrow or expand the search by changing the returns and date in the **Search Options** field. Eventually, with the time depending on the complexity and scope of the search, the reporter returns a ranked list of source documents.

*Happy Returns*

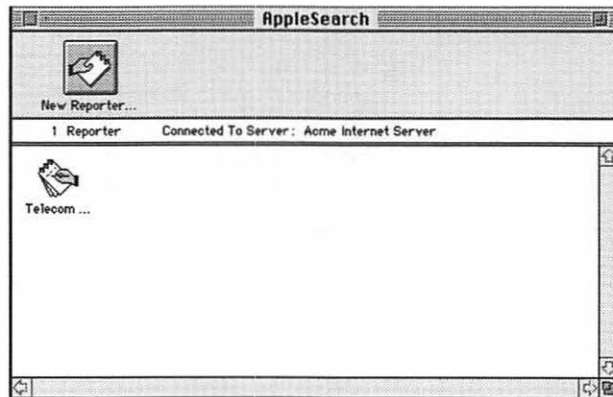
The user can double-click any one of the listed documents to read it.



Reading Retrieved Text

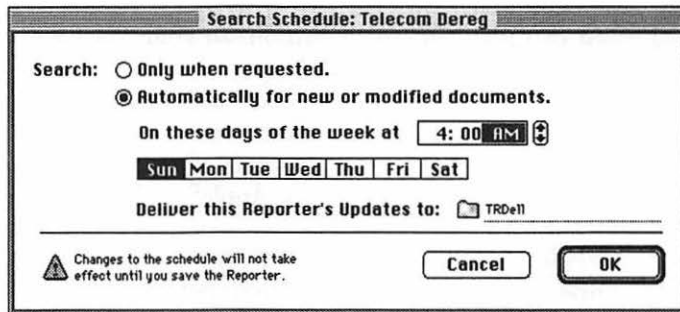
Data retrieved from a remote information source can be saved in a text file via the **Save Text...** button. With local information sources, you have two choices: You can have just the text, or you can save the original from which the text was extracted—art, formatting, and all—by clicking the **Copy Original...** button.

I just know Dorian is going to want this information, but I am far too lazy to read through it all and forward the text. Instead, I will click the **Save Reporter...** button in the previous window so that my search criteria are preserved. It now appears in the AppleSearch Client window. Although I see it here, it has in fact been saved to the AppleSearch server.



Saved Reporter

Of course, the telecommunications deregulation issue is an ongoing one, so I need to perform this search regularly to stay up to date. I can do that by selecting the **Schedule Search...** button in the Report window.



Scheduling a Search



*Telecom Dereg:
Exported Reporter*

My server is not too busy at 4:00 A.M. on Sundays, nor is my Internet feed, so I will have it gather new information on this topic then. A brisk perusal of congressional legislation will be just what I need to get my Mondays off to a good start! That and a poke in the eye with a sharp stick. But wait! I don't need this stuff—Dorian does! I can send this reporter to him by selecting its icon in the AppleSearch Client window and selecting **Export Reporter** from the **File** menu.

Now I can e-mail it to him or, better yet (hint, hint), put it in the AppleSearch Reporters folder I created on the server so everyone can have it. Dorian then uses the **Import Reporter** option from the **File** menu to make this reporter his own.

This is a very superficial explanation of how things work at the client end, but my only purpose here is to familiarize you with it. To learn about powerful stuff like Boolean search operators, stemming, ranking, and keywords, read the well-written *AppleSearch User's Guide*.

USING APPLESEARCH VIA THE WEB

At this point in our discussion, all your Macintosh and Windows users have a search mechanism for finding information on your servers. That's nice, but it leaves out your UNIX and DOS users. Furthermore, you want to make some of the information on your servers available to *everyone* via the Internet, but you do not want to deploy a WAIS server of your own. No problem. If you have a WebSTAR server, users can access AppleSearch via the World Wide Web.

Apple's AppleSearch CGI, an asynchronous Common Gateway Interface (CGI), makes this possible. You are given permission to use this with your licensed copy of WebSTAR (discussed in the chapter "Serving the World Wide Web with WebSTAR," starting on page 303). MacHTTP 1.3.1b1 or later can be used as well.

At the client end, the AppleSearch CGI supports two kinds of WWW browser interfaces. One uses <isindex> search arguments passed from a browser to perform searches on AppleSearch information sources. The other calls for a more modern browser that supports forms, like Netscape Navigator. With the latter, a nice interface is generated that uses a form to collect search arguments from the user.

Here is how the server works. The AppleSearch CGI is invoked by WebSTAR when a browser requests a particular URL. WebSTAR passes arguments, contained in the URLs, to the AppleSearch CGI via AppleEvents. The AppleSearch CGI uses these arguments to perform search operations through the AppleSearch server to which it is connected on the Web server or out on the LAN. It then automatically generates an HTML document containing the results, which it passes back to WebSTAR and, hence, the browser.

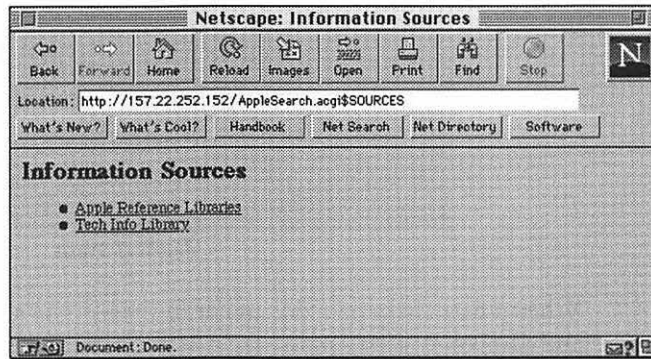
To install the AppleSearch CGI, copy it into the same folder that contains WebSTAR. If you put it anywhere else, such as in the CGI folder, you need to put an alias of it in the same folder that contains the WebSTAR application and erase the word "alias" at the end of its file name when you do. Do not rename "AppleSearch.cgi." That exact name is important.

Once you have done this, make sure AppleSearch and WebSTAR are both running. Double-click the AppleSearch CGI application. Next, choose **Connect** from the **File** menu to tell the AppleSearch CGI with which server to work in the Program Linking dialog box.

With the AppleSearch ACgi up and running, give users a way to access it by adding the proper HTML code to one of your Web pages. Again, this can be in one of two forms. The simple interface works with any WWW client, no matter how lowly. To execute this command, anchor this line to a Web page somewhere:

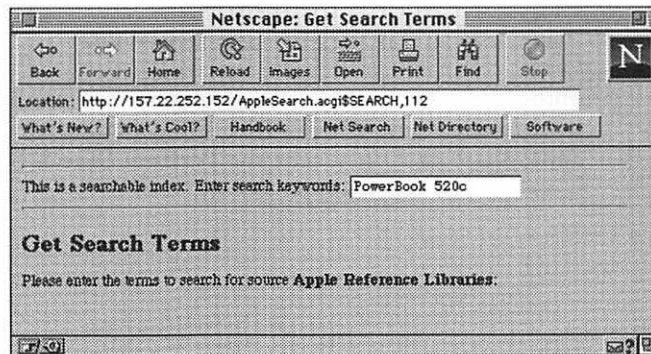
```
http://192.20.200.2/AppleSearch.acgi$SOURCES
```

Substitute your Web server's address for the one illustrated. When users click on this link, they are presented with a list of local information sources available on the AppleSearch server to which the AppleSearch ACgi is connected.



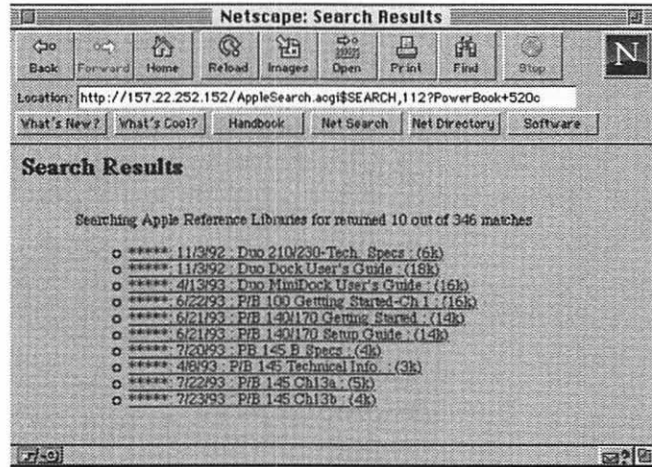
Available Local Information Sources

The user selects an item in the list and is then prompted for terms to look for in the selected information source.



Entering a Search Term

When the inquiry is complete, the gateway returns a list of documents that match the search criteria.



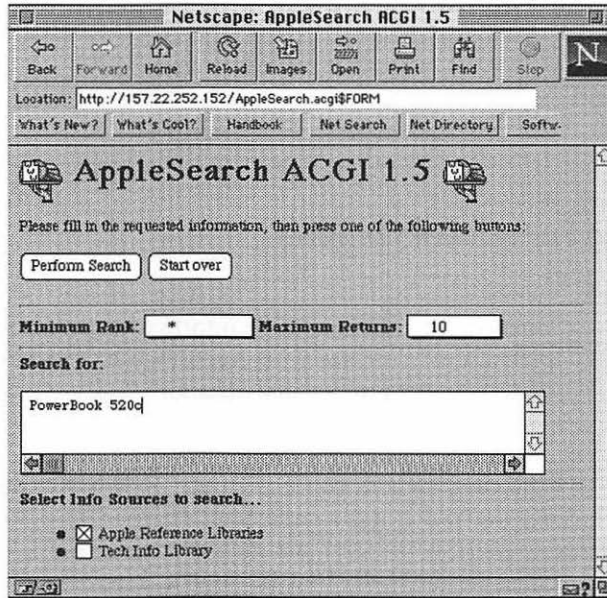
Results of Search

When a user clicks on a document in the list, the gateway retrieves the document and displays it. The AppleSearch ACgi can only return 30K segments of data at a time. Large documents are divided into multiple pages with a “click here for more” link at the bottom.

The forms interface requires a browser with forms capabilities, which most new ones have. To execute this command, anchor this line to a Web page somewhere:

`http://192.20.200.2/AppleSearch.acgi$FORM`

Substitute your Web server’s address for the one illustrated. The form that this URL invokes closely follows the AppleSearch client interface. After users enter the appropriate data, they click the **Perform Search** button. Data is returned in the same manner described previously.



AppleSearch ACGI Forms Interface


OTHER GOOD THINGS TO KNOW ABOUT APPLESEARCH

Here are a few other things you should know, in no coherent order whatsoever!

Using Pre-Indexed CD-ROMs

Some companies are making information available on CD-ROMs that have been “pre-indexed” for use with AppleSearch. The best example we have seen is the Apple Support Professional Solutions series. (Call 800-745-2775 for more information.) To get the most speed out of them, here is what you should do.

First, drag the files you find in the AppleSearch Indexes folder on the CD-ROM into the folder of the same name on your server’s hard drive. If you don’t want all the available sources, just drag over the indexes to select the ones you do want. There are three files for each source. They have the extensions *.def*, *.pls*, and *.sif*.



Selecting Index Files

Next, designate the desired folders on the CD-ROM as local information sources. When the dialog box appears asking you where its indexes should be created, be sure to select the server’s hard drive and not the CD-ROM.

You can make your own indexed CD-ROMs, too, which is a great way to deal with archival material.

What the Filters Filter

As I previously mentioned, AppleSearch uses MacLink Plus filters and the Claris XTND system to extract text from various document formats. The exact filters that came with your AppleSearch server, at the bare minimum, are these:

- AppleWorks (remember this?)
- FrameMaker MIF 2.0 and 3.0
- Microsoft Word for Macintosh 3.0, 4.0, 5.0, and 5.1
- Microsoft Word for Windows 1.0 and 2.0
- Microsoft Works (the word processor documents only) 2.0 and 3.0
- Microsoft Excel (tab text only)
- MacWrite 4.5, 5.0, and II
- Nisus 3.0
- PageMaker 4.0
- PICT (the text part only)
- RagTime 3.1
- WordPerfect for Macintosh 1.0, 2.0, and 2.1
- WordPerfect for Windows 5.1
- WriteNow 2.0 and 3.0

If you know that you have documents created by something other than the applications and versions listed, better give DataViz a call at 800-733-0030, and ask them if they have a new XTND translator for it. If they do not, call the application's developer.

Before you install your own XTND translators, there are a couple of potential problems, according to Apple. For one thing, some public domain XTNDs have

bugs that might not show up under casual use but can rear their ugly heads when they are subjected to the heavy use AppleSearch makes of them. For instance, an XTND used by a word processor to import a file that leaves a memory pointer allocated does not usually cause a problem. However, when the same translator is used to import 1,500 files, it can fragment memory and ultimately crash the server. Watch your server carefully for any such problems.

Another thing you should know: The XTND protocol allows one XTND to define itself as the default. The default XTND attempts to translate any unknown formats. The XTND shipped with AppleSearch is one of these. If you install your own default XTND, AppleSearch still gives priority to its own and your XTND will never be utilized. Sad, isn't it?


Overall, be aware that not all XTNDs work under AppleSearch. One of the most frustrating examples of this, in my opinion, is the ClarisWorks XTND. ClarisWorks is on the hard drives of thousands of shipping Macintoshes, but its files cannot be accessed through AppleSearch!

Key to the Firewall

If you have the AppleSearch server set up to use the WAIS gateway, it must be able to make a WAIS connection through whatever network firewall you set up. Firewalls are often no more than filters that look for particular types of network traffic and allow or disallow access to their packets. Internet services identify themselves by a port number attached to each of their packets, giving a firewall one easy way of filtering certain services out. WAIS operates on socket 210, so be sure to allow port 210 traffic to pass when you configure a firewall.

Accessing Unpublished WAIS Sites

The WAIS gateway is not limited to just the public WAIS servers. If your company has WAIS servers or needs access to sites that are not publicized, you can access these simply by telling AppleSearch where they are. To do this, open the "WAIS Custom Databases" file located in the AppleSearch folder within the Preferences folder of the System Folder using a text editor, such as SimpleText.



Contents of WAIS Custom Database File

Here you can add one database domain name per line, like this:

```
database:port@domain
```

For instance, if Network Frontiers had a WAIS server (we don't), it would look like the last line in the illustration.

For more information about managing the AppleSearch server once you have it up and running, read Apple's *AppleSearch Administrator's Guide* and our own *Managing AppleTalk Networks*.

CHAPTER 11: WORKING WITH THE DOMAIN NAMING SYSTEM AND QUICKDNS PRO

Hi. My name is 157.22.252.100.

If you were a computer, that would be a perfectly reasonable way to begin a conversation. Because you are a human, however, you would probably prefer a plain English designation. How about, “I am dell.netfrontiers.com?” A little better? Well, in terms of the Internet, they are one and the same. Both “157.22.252.100” and “dell.netfrontiers.com” refer to my trusty PowerBook 520c. They are its *IP address* and *host name*, respectively. The machine magic that makes this duality possible is called the *Domain Naming System (DNS)*. One implementation of this service for the MacOS is QuickDNS Pro—the subject of this chapter.

It is the job of DNS to match those dotted decimal IP addresses with host names when TCP/IP-equipped computers wish to communicate with each other. This is useful in two ways. It gives humans a more immediate recognition of the machines with which they are dealing. More important, it provides continuity by maintaining a host-name-to-IP-address mapping that can be altered without affecting network communications. For instance, suppose I get rid of the PowerBook (highly unlikely) and put another machine with a different IP address in its place. You need only remember my machine as “dell.netfrontiers.com.” If I make

a small change in the DNS, that host name will be *resolved* to the new machine's IP address. You will not know the difference.

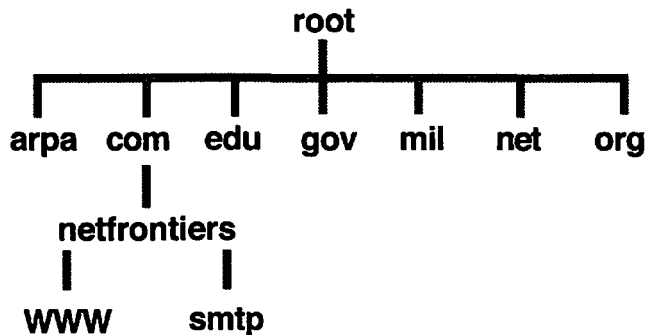
For the most part, DNS is not something you need to learn how to manipulate, unless you are a large enough organization to warrant running your own name servers. Your Internet service provider (ISP) will usually take care of this for you. You just tell the ISP which machine (as in "157.22.252.100") is who (as in "dell.netfrontiers.com"). If your ISP does not provide DNS, however, you can do it yourself with some Internet routers or software such as Apple's MacDNS or Men & Mice, Inc.'s QuickDNS Pro. Either way, here is what you should know.

NAVIGATING THE DNS DATABASE

DNS adheres to the same client/server model we have discussed throughout this book. The clients in this case are called *domain name resolvers* and the servers are *domain name servers*. Both Open Transport and MacTCP have domain name resolvers built into them to which most TCP/IP applications under the MacOS defer. The Apple Internet Mail Server, discussed in our chapter “Using the Apple Internet Mail Server as Mail Host” beginning on page 155, uses its own domain name resolver. StarNine’s products, such as Mail*Link and ListSTAR, have their own resolvers, too, but these are often used to augment the MacOS resolvers rather than to replace them.

These resolvers are actually clients to a distributed database that resides in what is called the *domain name space*. This architecture permits the local administration of domain names and relieves the responsibility from resting on the shoulders of one central authority.

The domain name space is hierarchical and can be likened to the branch of a bush, with the root directory being the thick part and various subdomains hanging off it like bristles.



Where We Fit in the DNS Hierarchy

Domain name resolvers have a pointer to the root domain, which is supported by several domain name servers owned by certain big government, educational, and ISP sites. The root domain has pointers to its top-level domains, such as *.com*, *.edu*, *.net*, and *.org*, which in turn have pointers to their subdomains, and so on.

Again, domain name *queries* are sent by domain name *resolvers* to domain name *servers*. These in turn act as name servers to domains for which they are *authoritative* (have records for) or act as resolvers when queried about domains for which they are not authoritative. Because of the distributed nature of DNS, each domain records only information about its immediate subdomains, in most cases. Domain name servers must query other servers for information about addresses outside their immediate domains.

Look again at the illustration on the last page. My network's domain name server is authoritative for any host in the "netfrontiers.com" domain. When my PowerBook's resolver asks the domain name server about the IP address for "www.netfrontiers.com," it knows to return "157.22.252.152." No problem. When I ask the domain name server for information about the Web host of some of my former Webmaster Workshop students, "www.zephyr.net," it becomes a little more complicated and, of course, takes a bit longer. My domain name server is not authoritative for "zephyr.net," so it must query the root name server. That root domain name server refers mine to a domain name server in the next-level domain, *.net*, which in turn refers it to the authoritative domain name server for "zephyr.net." Now the "zephyr.net" domain name server takes over and resolves the IP address—information that my local domain name server passes back to me. What team work!

This might appear cumbersome at first glance. If domain name servers are only good as resolvers for their local networks, for which they are authoritative, and if the TCP/IP (formerly MacTCP) control panel has its own resolvers built in anyway, why not cut out the middle man? As it turns out, in this case the middle man, the local domain name server, is your friend. He, and in most cases *only* he, knows about the root domain name servers. If the TCP/IP control panel had to keep track of this substantial chunk of information, your host configuration chores would increase greatly. Also—and this should be no surprise to you, gentle reader—the MacOS is a good deal smarter than many others out there. The resolvers that come with other machines' TCP/IP stacks are often unable to follow a path as long as the one we just traced.

The path we just traced, by the way, illustrates the process of *recursion*. My domain name server followed successive pointers until it found a machine that knew the answer to my PowerBook's *recursive query*. In so doing, my local domain name server used *iteration*. With iteration, the queried domain name server returns the requested information or the name of another domain name server that might have the requested information. My PowerBook's resolver issued a recursive query

to my local domain name server. My domain name server then began acting as a resolver, too, issuing iterative queries as it passed from one domain name server to another, until it located “zephyr.net” and the IP address I wanted.

The point of all this is that if you want to be independent of your ISP, you need a domain name server that provides full *recursive name service*. These talk directly to the root domain name servers and do not need to pass their queries off to the ISP’s domain name servers. QuickDNS Pro is one such product.



A bit of an aside here: Both my PowerBook and the Web servers it was querying are examples of Internet *hosts*. On the Internet, a host is just any old machine with a network connection and is not necessarily a machine that provides services to other computers, as you have surely seen the term used before.

Our domain name server middle man has a performance trick up his sleeve in the form of caching. When the domain name server receives a query for information it has already looked up, it recalls the requested information from a cache it keeps of recent data. Since the whole query process is not duplicated, this provides a fast response. The information in the cache does not have to be a complete match, either. Suppose you ask for the address of our “www.netfrontiers.com” and that information has been saved to the cache. Next you ask for “ftp.netfrontiers.com.” The domain name server remembers that it can query the domain name server at “netfrontiers.com” directly and thus bypasses the intervening servers.

Since everything changes on the Internet regularly, the information in this cache should change, too. After a given amount of time a record is assumed to be too old to be of value anymore and is flushed from the cache.

My local domain name server is not alone on my network. Every domain on the Internet should have at least two such servers so that if one goes down the other can maintain DNS. Logically, then, each domain will have a *primary* and a *secondary* domain name server. You or your ISP configures the primary domain name server’s database. The secondary domain name servers then read this information periodically from the primary name server. That way the information need only be entered in one place.

Now that you know what the DNS database is and how it is used, let’s take a look at what is in it.

DNS RESOURCE RECORDS

When you set up your primary domain name server you will find yourself typing in DNS *resource records*. The information in these resource records is what the domain name server returns to a resolver when queried and may be of several different types. The most important of these are:

Domain Information (SOA) Record

The Domain Information, or *Start of Authority (SOA)* record, tells other domain name servers that yours is the authoritative domain name server for your domain. This record contains the following:

Domain Name

This is the domain name to which the record pertains, as in “netfrontiers.com.”

MName

This is the name of the primary domain name server for this *zone*. The zone is that part of the domain name space for which the domain name server is authoritative, meaning the space for which it has complete records. Depending on how your network is set up, you could have one domain name server responsible for everything in a domain, or you might have multiple domain name servers, each responsible for different subdomains.

RName

This is the e-mail address of the human being responsible for this zone. Go after this person if the records are wrong!

Serial

This is very important. This is a unique serial number for the current version of the *Hosts file* that contains your resource records. Secondary name servers use this

number to determine whether or not the data in the primary's Hosts file has been updated. You need to increase this number whenever you make changes to the primary domain name server's resource records.

Refresh

This is the interval of time that elapses between a secondary domain name server checking in with the primary domain name server to see if the Hosts file has changed. If the serial value has changed, the secondary domain name server will update its records from the primary domain name server's Host file, in what is called a *zone transfer*. A common setting here is 36,000 seconds (10 hours).

Retry

If a secondary domain name server fails to communicate with a primary domain name server after the refresh value is reached, it will try again after the amount of time entered here has passed. This is usually set to 7,200 seconds (2 hours).

Expire

What if the retry does not work and the secondary server just cannot find the primary domain name server? In that event, the information supplied by the secondary domain name server becomes suspect. If the two do not communicate within the period of time designated here, the secondary server is no longer considered authoritative. The usual setting here is 604,800 seconds (1 week).

Minimum Time to Live (TTL)

As we mentioned previously, this is how long a resource record is kept in a domain name server's cache. This usually is set to 86,400 seconds (1 day).

Address (A) Record

As the name implies, this record provides a domain-name-to-IP-address mapping. Its parts are simply:

Domain Name

This is the domain name for the host to which this record refers, such as “www.netfrontiers.com.”

Host Address

This is the IP address of the host to which this record refers, such as “157.22.252.152.”

Name Server (NS) Record

This record lists the domain name servers in a given domain. Its parts are:

Domain Name

This is the domain name to which the record applies, such as “netfrontiers.com.”

Name Server Name

This is the domain name server that is authoritative for the domain above, such as “primary.netfrontiers.com.”

Mail Exchanger (MX) Record

This record points to the hosts that can accept mail for this domain. For example, the MX record for “netfrontiers.com” points to “smtp.netfrontiers.com.” When you send me mail at “tom_dell@netfrontiers.com,” it is routed to the mail host at “smtp.netfrontiers.com.” These records contain the following parts:

Domain Name

This is the domain name to which the record applies, such as “netfrontiers.com.”

Preference

This is a number. There may be multiple MX records for the same domain, such as “smtp.netfrontiers.com” and “smtp2.netfrontiers.com.” When this is the case, the preference number is used to determine the priority of the records, with lower numbers equating with higher priority. Say “smtp.netfrontiers.com” has a preference number of 10 and “smtp2.netfrontiers.com” has a preference number of 20. When you send mail to my domain, it will be routed to the lower numbered mail host first, then the next highest, and so on, until some machine accepts it. It is common to use multiples of ten when you set these up, so that you have room to slip other numbers in between later, if needed.

Mail Exchange

This is the host name of a mail host, such as “smtp.netfrontiers.com.”

Host Info (HINFO) Record

This is an optional record that contains information about the domain name server’s machine. Its parts are:

Hardware Type

This is the type of computer on which the domain name server is running, such as “Macintosh.”

Operating System

This is whatever operating system the host is running, such as “MacOS.” You can enter the words “rules man!” here, too, if you want.

Canonical Name (CNAME) Record

Here is a nifty one: This record lets you define an alias for an actual (*canonical*) host name stored elsewhere in an A record of the Hosts file. This can be really useful when you wish to associate more than one domain name with a given host and IP address. For instance, our Web server and Anonymous FTP site are on the same

machine. That machine is “www.netfrontiers.com.” Using this record, I can also associate “ftp.netfrontiers.com” with this same machine. A Canonical Name record contains these parts:

Domain Name

This is the alias domain name of the host to which this record refers, such as “ftp.netfrontiers.com.”

Canonical Name

This is the real name of the machine, such as “www.netfrontiers.com.”

Text (TXT) Record

This is another optional record like HINFO. Here you can type whatever you want, up to 255 characters. Something like “Server 3rd Floor” would be appropriate, but “Fred Loves Sue Ellen” would probably be acceptable, too, if your significant other were geeky enough to ever see it.

Pointer (PTR) Record

This record is used for *reverse* domain name mapping. Instead of trying to associate a host name with an IP address, the idea here is to associate an IP address with a host name. This is handy for such things as logging statistics and excluding people you do not like from domains.

These records map host names to IP addresses using the domain “in-addr.arpa.” This special domain was created so that you do not need to search the entire domain name space to relate a host name to a specific IP address. Instead, they can be recorded in this single domain. The “in-addr.arpa” address of a host that has an IP address of “157.22.252.1” would be “1.252.22.157.in-addr.arpa.” The IP address is reversed because host names are ordered left to right, with the most specific domain at the beginning and the least specific domain at the end. IP addresses are ordered in the opposite direction, with the least specific number on the left and the most specific on the right. A Pointer record contains these parts:

Pointer Name

This is the “in-addr.arpa” domain name to which this record refers, such as “152.252.22.157.in-addr.arpa.”

Domain Name

This is the domain name for the host to which this record refers, such as “www.netfrontiers.com.”

PUTTING UP QUICKDNS PRO SERVER

Now that you have some idea of what it is you would type into a domain name server's Hosts file, let's take a look at how you would do it using QuickDNS Pro.

This software will run on something as lowly as a Mac Plus, but it is PowerPC native. It requires System 7 or later and MacTCP 2.0.6 or Open Transport. You will need at least a megabyte of RAM for the server application and another megabyte for its administration program. You will also need a properly configured permanent connection to the Internet.

Domains are created and edited using the QuickDNS Admin application. To apply what we have been discussing, launch this application and choose **New** from the **File** menu. A blank window is generated initially with the title "untitled.com."


Before you do anything, select **Get Info** from the **Domain** menu. The fields in this window should look familiar. They are the ones we just talked about in relation to SOA records.

Domain Information	
Name	netfrontiers.com.
Primary	primary
Hostmaster	Postmaster
Serial	1995092801
Expire	604800
Refresh	36000
Minimum	66400
Retry	7200
Time-to-live	
[Cancel] [OK]	

SOA Values in QuickDNS Pro


Type your network's domain name, as in "netfrontiers.com," into the **Name** field. Type the host name for the machine on which you are setting up QuickDNS Pro into the **Primary** field. You do not need the whole name, so if you want to designate "primary.netfrontiers.com," you just type in "primary." This is also true of the e-mail address you type into the **Hostmaster** field. If you want the address "postmaster@netfrontiers.com," just type in "postmaster." The rest of the fields contain defaults you can leave alone. Remember: The numbers are all in seconds.

Click **OK** when you are finished. Note that the title that was “untitled.com” is now the name of your domain. Next select **Save** from the **File** menu and put the new file in the Primary Data folder within the QuickDNS Pro folder. This is the first time you have worked with this file, so you need not update the **Serial** field. The next time you edit this file, and each time thereafter, you will! At this point you are ready to set up your own resource records. Choose **Create Record** from the **Domain** menu.



Creating a Resource Record

QuickDNS Pro dumps some preliminary information in there for you that you need to amend. The information gives you a clue as to what should go there, just in case you did not get around to memorizing the previous pages! Click the **Type** field for a pop-up menu of the various resource record types that will change the information in the other fields accordingly. A good place to start is with a record for your network’s authoritative domain name server. In my case, it looks like the following screen shot.



Creating an NS Resource Record

Here, too, I need only put in the relative host name. QuickDNS Pro knows well enough to tack on “netfrontiers.com.” I also do not bother to add a number in the TTL (Time to Live) field. The number I set previously in the Domain Information window will be used. Once I have a few more lines filled in, I add my all-important MX records.

Name	TTL	Type	Parameters	Comment
netfrontiers.com.		NS	primary_netfrontie...	
netfrontiers.com.		NS	secondary_netfront...	
primary_netfrontie...		HINFO	Macintosh	MacOS
primary_netfrontie...		TXT	Macintosh Forever!	
netfrontiers.com.		MX	10	smtp.netfrontiers.c...
netfrontiers.com.		MX	20	smtp2.netfrontiers...

Creating an MX Resource Record

By default QuickDNS Pro gives me a 10 for the Parameters. I change the number for my secondary mail host to 20 to make it next in line after the primary mail host should the primary be unavailable or busy. I could make it "11," but I want to leave a little space in case I decide to put another mail host in between what are now my primary and secondary servers.

Now for my favorite trick, the "round robin." Here I map out the A record for my Web server, which is "www.netfrontiers.com" and "157.22.252.152." However, maybe I also have two more identical Web servers with which I want to share the work load. I map them in exactly the same way but for the IP addresses.

Name	TTL	Type	Parameters	Comment
netfrontiers.com.		NS	primary_netfrontie...	
netfrontiers.com.		NS	secondary_netfront...	
primary_netfrontie...		HINFO	Macintosh	MacOS
primary_netfrontie...		TXT	Macintosh Forever!	
netfrontiers.com.		MX	10	smtp.netfrontiers.c...
netfrontiers.com.		MX	20	smtp2.netfrontiers...
www.netfrontiers...		A	157.22.252.152	
www.netfrontiers...		A	157.22.252.153	
www.netfrontiers...		A	157.22.252.154	

Creating Round Robin A Resource Records

Now incoming connections will be passed off to each in turn, balancing the HTTP load.

For my next trick, I also have FTP server software running on each of those machines, and I use the CNAME record to direct requests coming from FTP clients to those computers as well.

Domain netfrontiers.com.				
Name	TTL	Type	Parameters	Comment
netfrontiers.com.		NS	primary_netfrontie...	
netfrontiers.com.		NS	secondary_netfront...	
primary_netfrontie...		HINFO	Macintosh	MacOS
primary_netfrontie...		TXT	Macintosh Forever!	
netfrontiers.com.		MX	10	smtp.netfrontiers.c...
netfrontiers.com.		MX	20	smtp2.netfrontiers...
www.netfrontiers...		A	157.22.252.152	
www.netfrontiers...		A	157.22.252.153	
www.netfrontiers...		A	157.22.252.154	
ftp.netfrontiers.com.		CNAME	www.netfrontiers...	

Creating a CNAME Record

At this point, I am sure you have the gist of it. I create A records for every Macintosh I want to advertise on the Internet: the domain name server, mail hosts, Web servers, FTP servers, Gopher servers, whatever. I may or may not add HINFO records for each machine. Before you know it, I am finished.

Domain netfrontiers.com.				
Name	TTL	Type	Parameters	Comment
netfrontiers.com.		NS	primary_netfrontie...	
primary_netfrontie...		A	157.22.252.1	
primary_netfrontie...		HINFO	Macintosh	MacOS
netfrontiers.com.		NS	secondary_netfront...	
secondary_netfront...		A	157.22.252.2	
secondary_netfront...		HINFO	Macintosh	MacOS
netfrontiers.com.		MX	10	smtp.netfrontiers.c...
smtp.netfrontiers.c...		A	157.22.252.160	
smtp.netfrontiers.c...		HINFO	Macintosh	MacOS
netfrontiers.com.		MX	20	smtp2.netfrontiers...
smtp2.netfrontiers...		A	157.22.252.161	
smtp2.netfrontiers...		HINFO	Macintosh	MacOS
www.netfrontiers...		A	157.22.252.152	
www.netfrontiers...		A	157.22.252.153	
www.netfrontiers...		A	157.22.252.154	
ftp.netfrontiers.com.		CNAME	www.netfrontiers...	
gopher.netfrontiers...		CNAME	www.netfrontiers...	
dell.netfrontiers.co...		A	157.22.252.100	
dell.netfrontiers.co...		HINFO	PowerBook 520	MacOS Rules man!


Completed Resource Records?

Almost. . . For every host that has an A record there should be a PTR record as well for reverse domain mapping. This requires that you create a file just like the first one, except that it will contain only PTR and NS records, and its contents will be reversed in the manner we described earlier.

Domain 252.22.157.in-addr.arpa.				
Name	TTL	Type	Parameters	Comment
252.22.157.in-addr...		NS	primary.netfrontie...	
252.22.157.in-addr...		NS	secondary.netfront...	
1.252.22.157.in-a...		PTR	primary.netfrontie...	
2.252.22.157.in-a...		PTR	secondary.netfront...	
160.252.22.157.in...		PTR	smtp.netfrontiers.c...	
161.252.22.157.in...		PTR	smtp2.252.22.157...	
152.252.22.157.in...		PTR	www.netfrontiers...	
153.252.22.157.in...		PTR	www.netfrontiers...	
154.252.22.157.in...		PTR	www.netfrontiers...	
100.252.22.157.in...		PTR	dell.netfrontiers.co...	

Reverse Domain Mappings

There is not much to installing QuickDNS Pro Server. Just drag it onto the hard drive and double-click the QuickDNS Pro Server icon. You will be asked for the activation key. Once that is typed in, the server boots up and loads whatever is in the Primary Data folder—most likely the two files you just created. QuickDNS Pro will check back with this folder every minute for new information. Needless to say, do not move this folder from the folder that contains QuickDNS Server!

QuickDNS	
 QuickDNS - Caching only. Version 1.0b © 1995 Men & Mice Free memory: 145780	
Packets in: 116	Packets out: 122
Requests in: 43	Requests out: 79
Replies in: 67	Replies out: 43
Replies from cache: 12	
Entries in cache: 198	purged: 0
Records in cache: 69	purged: 324

QuickDNS Pro at Work

When running, QuickDNS Pro gives you totals showing how active it is. This information is also saved to a log file.

HOST NAME DOS AND DON'TS

Before we leave the topic of setting up DNS, let me offer a few dos and don'ts when naming hosts:

- *Don't* use special (nonalphanumeric) characters—even underscores. Many types of software have trouble with them.
- *Do* keep the names short and simple. The whole idea behind domain names is that they are supposed to be easier for humans to remember than the IP addresses. A name like “www.sf.king.205.netfrontiers.com” is not easier to remember than “157.22.252.152,” now is it?
- *Don't* use words that can be spelled two ways. One of my favorite Web sites is at “www.lynda.com,” but when I refer to it in my Webmaster Workshop class, a few people type “www.linda.com” into their browsers. Naturally!
- *Do* use names that are descriptive of the host's function. Something like “mailhost.netfrontiers.com” says all you need to know, but “mailhost.tech.netfrontiers.com” or “mailhost.classes.netfrontiers.com” further delineates the respective hosts' duties.
- *Do* remember that these host names are relative to your domain so that conflicts are not much of an issue. For instance, “apple.support.netfrontiers.com” is not going to conflict with “apple.support.apple.com”—at least not as far as the machines are concerned.
- *Do* send me \$5 every time you name a computer. I would like to visit Bali.

CHAPTER 12: SERVING THE WORLD WIDE WEB WITH WEBSTAR

In this chapter, we discuss taking advantage of running MacTCP concurrently over your AppleTalk network by putting up a World Wide Web server. Although the obvious use for a Web server is to provide a digital billboard from your organization to the Internet-surfing world, it has applications as an internal service as well. This is especially true in heterogeneous network environments in which the Web server provides the same face to Macintosh, PC, or UNIX workstation users.

In the example in this chapter, we will set up services for both internal and external users on an Apple Workgroup Server running WebSTAR from StarNine Technologies. In doing so, we briefly discuss HyperText Markup Language (HTML), HyperText Transfer Protocol (HTTP), File Transfer Protocol (FTP), Universal Resource Locators (URLs), Common Gateway Interfaces (CGIs), and other frightening acronyms. We also discuss implementation strategies and security.

Before we go on, I want to impress upon you an important lesson. Those of you who have already read *Managing AppleShare & Workgroup Servers* know it. The rest of you, repeat after me: *Guest access can be very, very bad. Guest access can be scary. I will not use guest access carelessly.*

Good. We'll talk more about that later.

HTML FOR FAST TYPISTS

Arguably the most enjoyable aspect of publishing on the WWW is creating your HTML documents. Before beginning this creation process, think about this: HTML documents generally contain a large number of *links*, or alias-like references, to other documents. When the locations of referenced files change, your HTML links can be broken. Thus, it is a good idea to establish your Web materials in a relatively static place. Another important reason for setting aside this special location is security. Think of this “webspaces” as your network’s lobby. You do not mind if people drop in and roam about a public area, but you want to be able to restrict their movements into the rest of the offices.

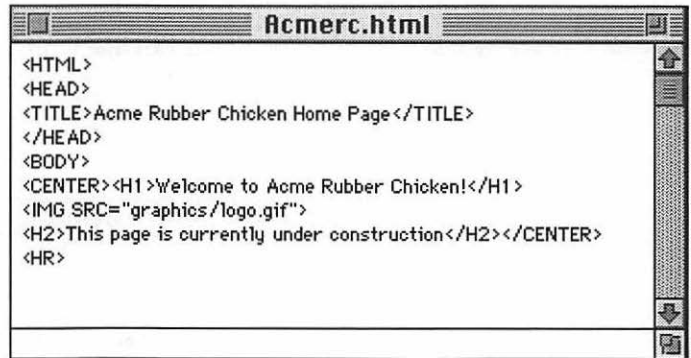
Set aside a special area on your Apple server’s hard drive for the sole use of Internet documents and software. You could make this an entire hard drive, a hard drive partition, or simply a folder named something obvious like “WWW” or “Web-space.” Many people just use the WebSTAR folder. This is where we will store your first foray into HTML.

There are many HTML editors available today, and more are arriving all the time. This makes the simple process of writing HTML even easier. We are not going to use one of those here, however. Noooo. If this is your first experience with HTML, you will garner a better understanding of it if you create a document from scratch. To do this, use Apple’s tried-and-true SimpleText, which comes with the Apple system software. What better way to see how easy this really is?

If you have any art that should be included in the home page, like a cool company logo, save it in a non-platform-dependent format (GIF or JPEG). Put it in the same folder that is to contain the SimpleText document on which you will be working. It is best to keep these art files at or below 50K so that they load quickly for viewers, even over slow modem connections.

Begin the creation process by giving the document a meaningful title. For instance, our fictitious company, Acme Rubber Chicken, will use the title “Acme Rubber Chicken Home Page.” Titling the document is accomplished in the following manner.

Looking at the illustration to the right, you should notice that whatever is between the < and the > characters is code. These characters and the code within them are the key to HTML. The first reference

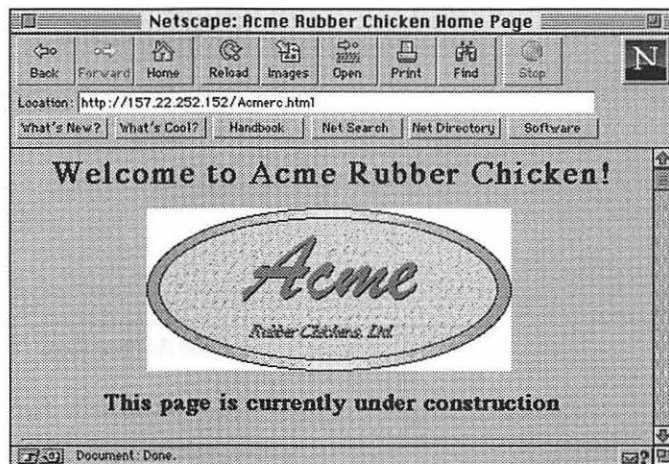


```
<HTML>
<HEAD>
<TITLE>Acme Rubber Chicken Home Page</TITLE>
</HEAD>
<BODY>
<CENTER><H1>Welcome to Acme Rubber Chicken!</H1>
<IMG SRC="graphics/logo.gif">
<H2>This page is currently under construction</H2></CENTER>
<HR>
```

Home Page Title

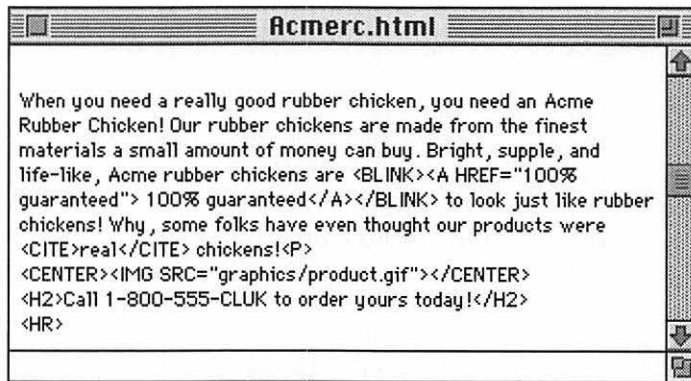
turns a given attribute on. The second reference, preceded by a /, turns the attribute off. These commands are called *tags*.

We have told any interested browser that this is the <HEAD> of an <HTML> document with a given <TITLE>. We then told it to start off the <BODY> with a heading of a certain display style, <H1>, in the <CENTER>. Beneath this is an image that can be found in the same directory as the HTML document. Beneath the image is another heading, <H2>. Beneath that is a horizontal rule, <HR>. HTML is evolving, so not all browsers handle all the same tags, but the basics are uniformly used. To see what a browser would make of this, use the **Open File** command from the **File** menu in Netscape Navigator.



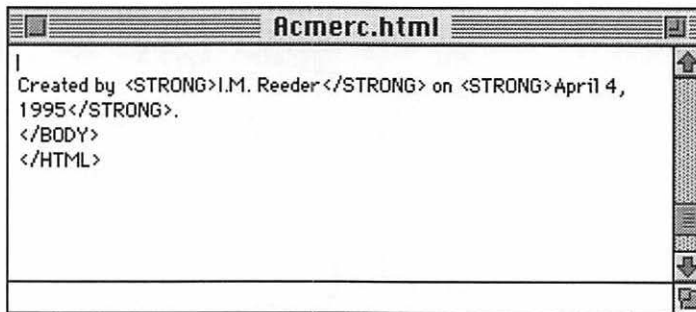
Viewing Title Area with Netscape Navigator

Resize Netscape Navigator's window. Notice that the text reflows to fit the new window size, ignoring your carriage returns in the SimpleText file. Hard returns are accomplished with a `
` tag or a `<P>` tag. The latter adds a blank line between text blocks. Try some of these as you add a marketing message for visiting Internet surfers:



Home Page Body

Do you notice a few new elements here? The `<BLINK>` tag will cause the enclosed text to flash. The `<A HREF>` anchor tag creates a link to another document or another section of the same document. You can use the `<CITE>` or `<I>` tags to create italics.



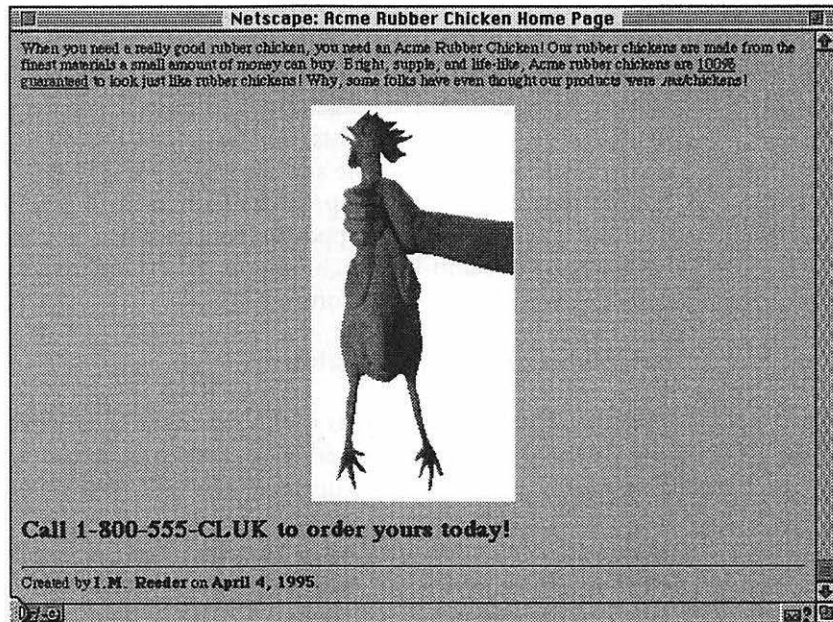
End of the Home Page

End your task with a modest credit to yourself, if you want.

The page we have created here is essentially a dead end to surfers, since we have not yet added links that go anywhere. At least it does let visitors know that you

have a Web presence and that you are working on the site. It also gives you something to test.

If you want to see how this page really looks as we have just written it, log into Tom Dell's Home Page at the Network Frontiers WWW site (www.netfrontiers.com). I put it there for you.



Viewing the Body with Netscape Navigator

PUTTING UP A WEBSTAR: THE BASICS

Now that you have created an initial home page, you can use WebSTAR to advertise it to the rest of the Internet-connected world. WebSTAR uses the Apple Thread Manager for multitasking. This makes it three to four times faster than the shareware application on which it is built, MacHTTP. In addition, WebSTAR is Apple Events- and AppleScript-aware, works with Macintosh database CGIs (such as AppleSearch, FileMaker Pro, and Butler), supports pre- and postprocessing functions on URL requests, and has extensive logging options for keeping track of who is using your site and for what. WebSTAR add-on products include the Security Toolkit, which provides authentication and encryption, and the Commerce Toolkit, which supports secure commercial transactions. It also comes with several free add-ons, including an FTP/Gopher server and an evaluation copy of ListSTAR list server software.

But let's not get ahead of ourselves . . .


Installing WebSTAR is easy, as it uses the traditional Apple installer script. Once installed on your Internet server's hard drive, it is ready to be launched and tested. To perform a test, first boot up the WebSTAR application. Although you don't have to do anything else with it at this point, I recommend enabling **Verbose Messages** in the **Options** menu. This way, you can gather as much operational information as possible for troubleshooting. Additional information is displayed only, not saved to the server's log.

Next, go to another TCP/IP-capable Macintosh on your network and launch a browser such as Netscape Navigator. Use the **Open Location** command to log into the WebSTAR server by typing in the server machine's IP address or domain name (e.g., "http://157.22.252.152"). If all is well, Netscape Navigator will pull up an HTML page displaying the colorful WebSTAR logo and information on where to obtain further documentation. At the server end, you should see WebSTAR faithfully logging the transaction.

If this does not occur, troubleshoot your network software and connections. At least you can still access the useful HTML-based documentation by using a browser locally, which is what we did when proofing your home page.

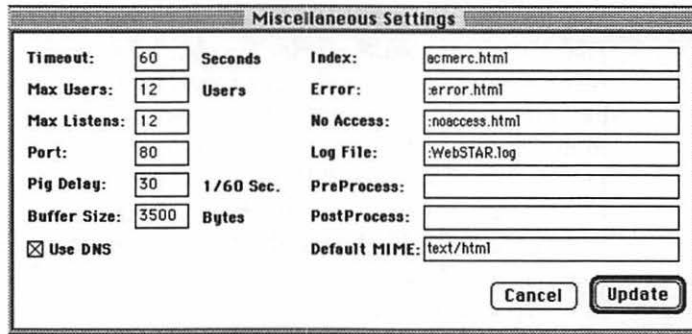
For security reasons, *WebSTAR does not allow URLs to access files outside of its own folder hierarchy*. Before putting up your home page, think about the folder structure with which you will want to work. For instance, you will probably have graphics in all your HTML documents, so you might as well make a Graphics folder for them. If you are going to have a collection of Acrobat files for visitors to pull down, why not a PDF folder? If you created a special storage area on the server already, you can merge this with the material in the WebSTAR folder. Unlike MacHTTP, which only allows you to alias files, WebSTAR allows you to alias folders. Thus, you could also put aliases in the WebSTAR folder to the folders in the special storage area.

Create all the folders you can think of (such as CGI, Graphics, Movies, and Sounds) and drag the appropriate files into them. Next, drag your home page into the WebSTAR folder. Because you have changed the folder structure, update any links that may have been affected in the document. For instance, the Acme Rubber Chicken Home Page had to be changed to reflect a new path for the art files in the illustrated folder structure, as in “graphics/logo.gif.”



Setting up Folders

This home page will be the server's *index* document, the default front door for anyone accessing this site. To designate an HTML document as WebSTAR's index file, launch the WebSTAR Admin application and choose **Misc. Settings** from its **Configure** menu. Type the name of your home page in the **Index** field. Click **Update** when you are finished, and go back out to that workstation to launch the browser again.



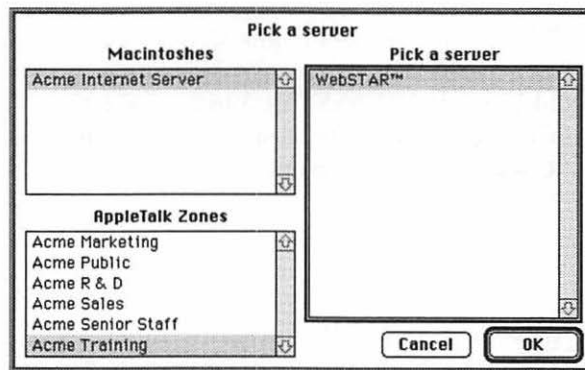
Setting the Index Document

If you see your home page and not the WebSTAR documentation page, you have been successful!

Tom's Time-Saving Tips

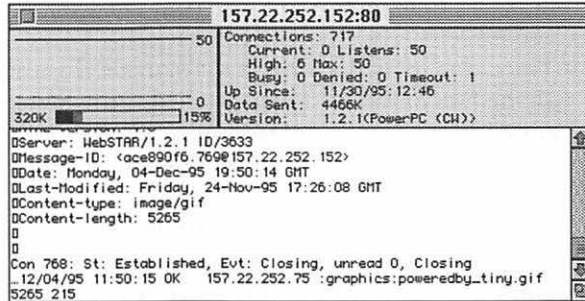
Ever since Apple introduced the PowerBook, I have been unable to drag myself into the cramped server room without trepidation. I love products like Farallon's Timbuktu Pro and Santorini's Server Manager (for AppleShare) because they let me accomplish whatever I need to do from anywhere, even over a modem connection. Here is my recipe for managing a WebSTAR remotely.

First, put a copy of WebSTAR Admin on your PowerBook or network management machine. This program uses Program Linking, so make sure this is enabled on both the server and the remote computers. Launch WebSTAR Admin.



Finding the WebSTAR to Link Programs

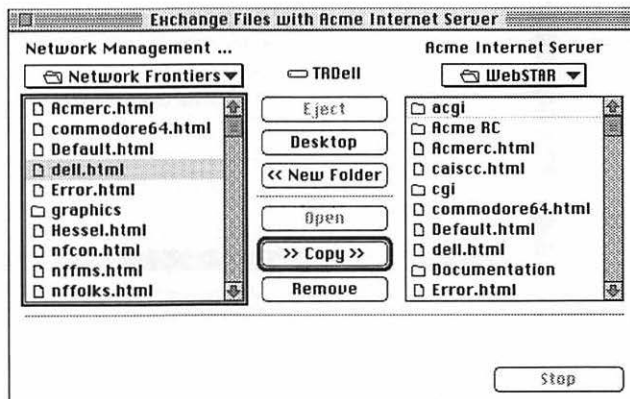
After you have finished the usual password-protected Program Linking procedure, you are presented with the statistics for the currently running server. You have full configuration powers over the WebSTAR. Remember, take that Program Linking password seriously and do not permit guest access to it!



WebSTAR Admin Running Remotely

That takes care of WebSTAR itself. Now, consider its documents. Like all similar publications in an organization, these are going to change quite a bit. People add things. People notice things. Do you want to have to go to the server and boot up an HTML editor every time you are told to change the “e” at the end of someone’s name to a “y”? Naaaaaa. Instead, create a mirror Webpace on your remote Macintosh that contains all the folders, art, and HTML of the server. When changes need to be made, perform them here, and proof your work with a browser.

There are many ways to move these updated files to the server, but my favorite method is Timbuktu Pro.



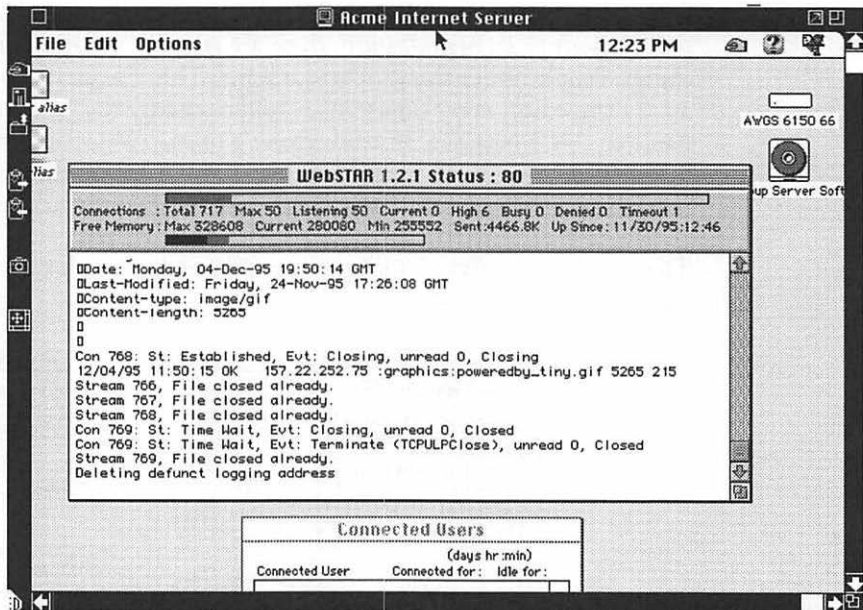
Exchanging Files with Timbuktu Pro

With this application, you can use the **Exchange Files** function to place the remote and server folder structures side by side for comparison, and then easily copy over the new documents.

WebSTAR doesn't file lock! That means you can update these files "live." Use your browser's **Reload** button to see your changes.

There are a lot of documents to reconcile. To ensure that I do not forget to update all the files on which I have worked, I use System 7's **Label** menu to change the edited files on my workstation to another color. After they are copied over I change them back to the label **None** until the next time I edit them.

Of course, you will probably have more applications running on the Web server than just WebSTAR, and only a few will have their own management programs. For these, use Timbuktu Pro's screen-sharing **Control** function to take over the server's keyboard and to do the job remotely.




Controlling the Server Remotely

One final reason to love Timbuktu Pro: It works over AppleTalk *and* TCP/IP.



Safety Tip: Timbuktu Pro has passwords, too. Use them! Also, make sure that access is not enabled for TCP/IP unless you really need it to be. Otherwise, you may find yourself playing mouse wars with some hacker in Bolivia!



Set only the access privileges you need!

Sharing Files with FTPd

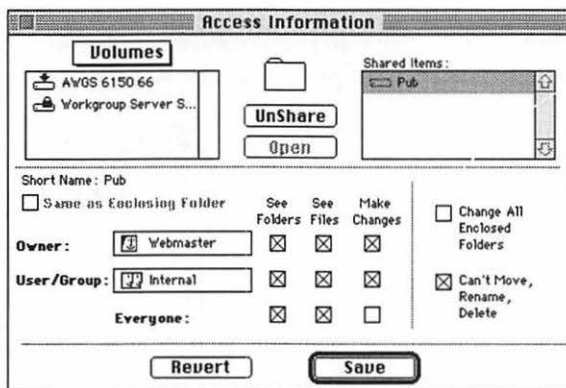
Before we leave the realm of the basic, we should add one additional service to your Web server: FTP. This makes it possible for the public to download designated files from anywhere in the world. To do this, use the shareware application FTPd, by Peter Lewis, which comes with WebSTAR. It supports both FTP and Gopher client requests.



Warning: If you don't have a real need to provide anonymous FTP access to the Internet, don't use this service. It is not worth the security risks "just to have it." If you do use it, make sure that you have a secure router firewall set up. Otherwise, you risk opening up every shared volume on your LAN to the whole world, and some of those folks aren't very nice.

To begin, create a folder at the root level of your hard disk's hierarchy—not in your Web stuff area—and name it Pub. Copy into this folder any folders and files you wish to share with the world. Avoid folder names with spaces in them. FTP doesn't like those very much.

Next, launch AppleShare File Server and AppleShare Admin. Share the Pub folder, giving the administrator ownership privileges, a group composed of your internal users read and write privileges, and everyone else read-only privileges.



Sharing the Pub Folder


We are using AppleShare here, but System 7 Personal File Sharing works fine, too. Either way, make sure that the administrator and internal users have secure passwords! I just have to quote from FTPd's author, Peter Lewis, here:

Keep your password secure! Anyone on the Internet with your user name, machine address, and password will likely be able to delete every file on your hard disk. This is a scary thought. You should be scared. Don't give your password out and don't use an obvious password. Obvious passwords include, but are not limited to, any of the following patterns (in decreasing obviousness) . . .

- your user name
- your real name
- your initials
- any of the above backwards
- your husband's /wife's /girlfriend's /boyfriend's /dog's /frog's / machine's /etc. name
- your car license plate, make, model, etc.
- your birthday
- your student/MediCare/social security/tax file/etc. number
- any of the above backwards
- any word from a dictionary (especially an electronic dictionary)

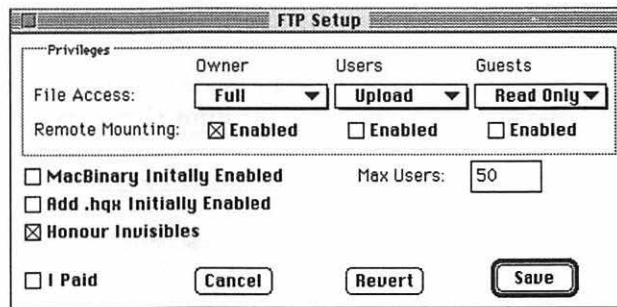
Good passwords are found by making up nonsense words or using the first letters from a common saying and by including non-alphanumeric ASCII characters.

Next, launch the FTPd Setup application and dismiss the About... window when it appears.



FTPd Menu Bar

Click the **FTP Setup** button in FTPd's menu bar.



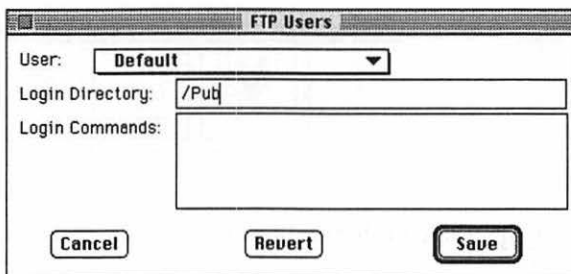
FTP Setup Window

In the FTP Setup window, give guests **Read Only** privileges in the pop-up menu. This allows those who log in anonymously to read and download files but not to upload files or change file names. Give users the same or **Upload** privileges, if you prefer. Upload access allows users to read, download, and upload files, but they still cannot change or overwrite what is already there.

You could give the owner **Full** privileges, as shown in the picture above. That gives the owner, presumably you, the power to both upload and download files, as well as the ability to delete and rename them. You could do that, but frankly, *I wouldn't*. If your user name and password combination isn't clever enough, you could open up not just the Pub folder but all the server's drives to malicious intent. Leave server administrator privileges to AppleTalk or ARA, not FTP. I would set owner privileges to **None**.

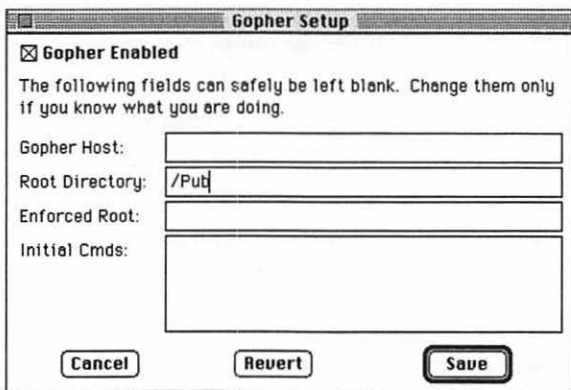
Notice the **Remote Mounting** checkboxes. These permit FTP users to mount your AppleShare server volumes. *We strongly recommend that you do not enable those checkboxes for users and guests—very strongly.* Click **Save** when you are finished.

Next, click the **FTP Users** button in FTPd's menu bar. Type **/Pub** into the default **Login Directory** field, then click **Save**.



FTP Users Window

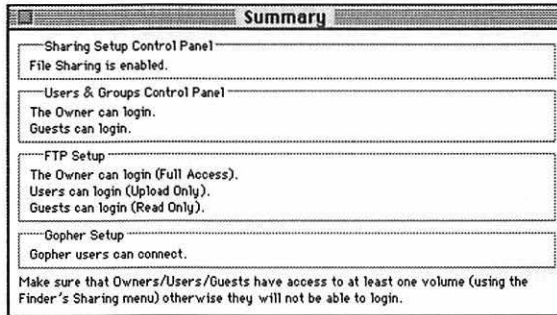
Next, click **Gopher Setup** in FTPd's menu bar. Select the **Gopher Enabled** checkbox, type **/Pub** into the **Root Directory** field, and click **Save**.



Gopher Setup Window

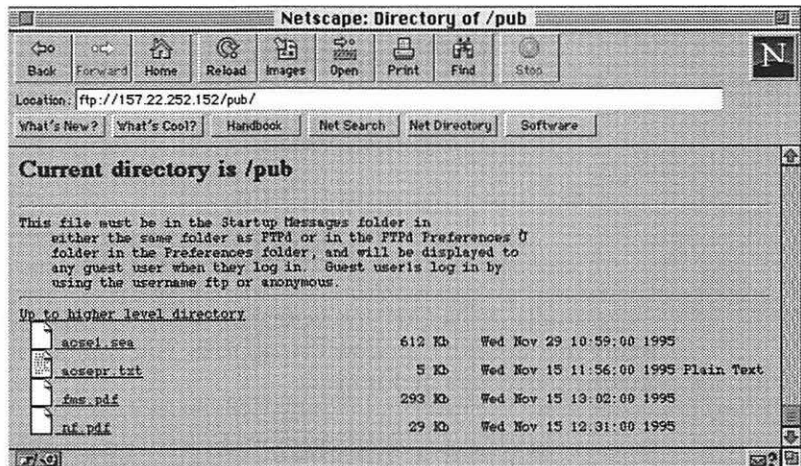
Notice that we skipped the **Security** button in the FTPd menu bar. I have just one thing to say here: Don't you dare enable the **Allow Clear Text Passwords** checkbox! That would be bad.

Finally, click **Summary** in FTPd's menu bar to make sure there are no problems.



Setup Summary

Assuming that there are no problems, quit FTPd Setup and launch FTPd. As a final test, use your Web browser, a shareware copy of TurboGopher, or Fetch to try it out.



Pub Folder as Seen by Netscape Navigator

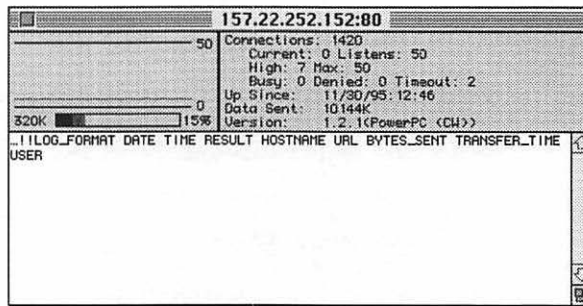
At this point, you have a perfectly workable Internet server solution. Before you move on, however, consider the suggestions in the next section, “Putting up a WebSTAR: Additional Considerations.”

PUTTING UP A WEBSTAR: ADDITIONAL CONSIDERATIONS

Now that you have a feel for WebSTAR, it is time to talk about more advanced issues that affect how it fits into your AppleTalk network design.

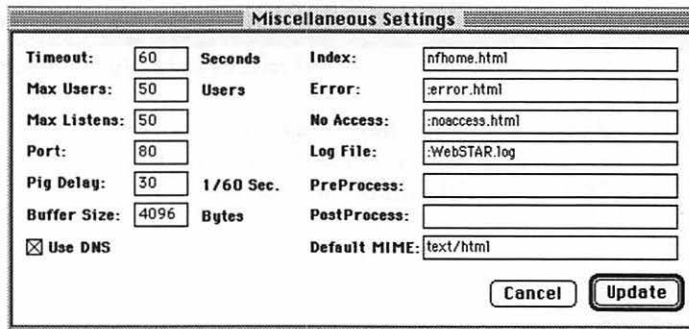
Nuts and Bolts

Most of WebSTAR's internal functions are summarized in the upper right corner of WebSTAR Admin's main window.



WebSTAR Admin Main Window

To get under WebSTAR's proverbial hood, choose **Misc. Settings** from the **Configure** menu. In the window that appears, you will find a few simple options for altering the server's performance.



WebSTAR's Miscellaneous Settings Window

The number in the **Timeout** field determines how long WebSTAR will keep a connection with a client before giving it up as lost. This is also how long WebSTAR will wait for an external application to reply to an AppleEvent query it has sent. The number can be set between 15–600 seconds. Figuring out what number to use is accomplished through trial and error. You can try different numbers and see what happens to the **Timeout** counter in the WebSTAR Admin window. Increase the number until it does not reduce the number of timeouts. You can then adjust the number backward until it starts to affect the number of timeouts again. A number in between will be your answer. Overall, you probably will not need to pay attention to this until you start adding CGIs.

My favorite field in this window is **Max Users**. In theory, your WebSTAR can handle 1,000 incoming connections. Cool. In reality, MacTCP can handle 64 simultaneous connections. Whoa. The maximum number WebSTAR currently supports is 50, as it saves the other 14 for such things as FTP, e-mail, and Telnet. Ouch. Open Transport changes all that. Here the limitation is imposed by RAM and processor bandwidth rather than software.

How many users your WebSTAR supports is also directly related to how much memory you give the application. You should have at least 1 MB for 12 users and 3 MB for 25 users. Memory also affects what you can put in the **Max Listens** field. This is the number of processes WebSTAR sets aside to listen for new connections. StarNine's rule of thumb is 100K for each. Fortunately, WebSTAR will let you know at startup if you have sufficient memory for the choices selected here.

```

WebSTAR 1.2.1 Status : 80
-----
Connections : Total 3049 Max 50 Listening 50 Current 0 High 10 Busy 0 Denied 0 Timeout 16
Free Memory : Max 329608 Current 245920 Min 206560 Sent:21608.7K Up Since:11/30/95:12:46
-----
12/18/95 11:07:21 OK 149.netfrontiers.com : graphics:construction.gif 4921 76
Mozilla/2.0b3 (Macintosh; I; 68K) CONDITIONAL_GET
http://157.22.252.152/hessel.html
12/18/95 11:07:22 OK 149.netfrontiers.com : graphics:lights.gif 2882 79
Mozilla/2.0b3 (Macintosh; I; 68K) CONDITIONAL_GET
http://157.22.252.152/hessel.html
12/18/95 11:07:29 OK 149.netfrontiers.com : tips&tricks.html 3742 43
Mozilla/2.0b3 (Macintosh; I; 68K) CONDITIONAL_GET
http://157.22.252.152/hessel.html
12/18/95 11:07:31 OK 149.netfrontiers.com : graphics:const2.gif 1159 53
Mozilla/2.0b3 (Macintosh; I; 68K) CONDITIONAL_GET
http://157.22.252.152/tips&tricks.html
12/18/95 11:07:42 OK 149.netfrontiers.com : Hessel.html 5125 68
Mozilla/2.0b3 (Macintosh; I; 68K) CONDITIONAL_GET
http://157.22.252.152/tips&tricks.html

```

Free Memory in WebSTAR Status Window

Despite this, you should still take a look at the memory usage displayed in the WebSTAR Status window from time to time. If the **Min** number falls below 150K, some scripts will not execute.

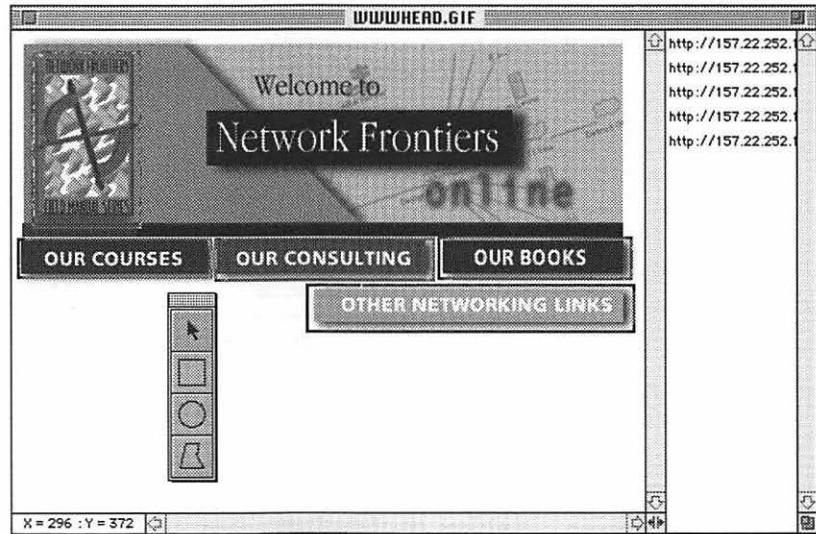
Digressing back to the Miscellaneous Settings window, the last thing I want to mention is the **Buffer Size** field. This reflects the number of bytes WebSTAR sends in a single write to the client, between 256–10,240. Like the **Timeout** field, this one requires a certain amount of experimentation. If you have the majority of your users or visitors logging in over slow connections, keep this number low, around 2048K or less. The smaller this buffer, the faster WebSTAR sends data to a slow client while not slowing down its reaction to other connections. Set it too low, however, and WebSTAR spends more time servicing the connection than it does sending data over it! If you are using the WebSTAR internally on a fast network, try a value of 4096K or greater.

What's in a File Name?

How the WebSTAR server handles information when contacted by a WWW browser is determined by *suffix mappings*. Simply put, these are the file name extensions that tell WebSTAR to associate a given file with a given application. To put it more accurately, this is a process in which the WebSTAR server looks at the file name extension or the Macintosh file type and creator codes of a retrieved document, and then processes that data based on a predetermined MIME type and action. It is the ability to use this method of launching various “helper” applications and scripts—both at the server and the client ends—that makes the Web so powerful. As you have seen, HTML by itself is very basic, but an HTML document that acts as a front end to, say, a database, can be very powerful indeed.

To see how this works, let's take a look at a typical CGI program that makes use of them. On your WebSTAR CD-ROM is a copy of MapServe, a shareware CGI from Kelly Campbell that allows the server to interpret the clicking of various “hot spots” on an image as the pressing of buttons. If you haven't already, create a folder named something like “CGI” in your server's WebSTAR folder and place the MapServe folder in it.

Next, you will need an image designed for this purpose. Here is our site's first one:



Making an Image

To create hot spots, you need an editor, such as City Net Express's shareware Web-Map shown in the illustration. This results in your producing a text file with the suffix ".map", such as, "logo.gif.map". Throw this into the MapServe folder (within the CGI folder).

Next, you need to go into your HTML document and tell it that this image has a special function. Let's say we were to do this to the logo on top of our Acme Rubber Chicken Home Page. In that case . . .

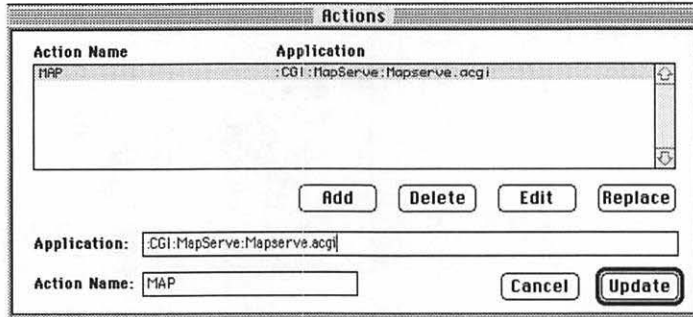
```
<IMG SRC="graphics/logo.gif">
```

could become:

```
<A HREF="CGI/MapServe/logo.gif.map"> <IMG SRC="graphics/logo.gif" ISMAP
BORDER=0> </A>
```

This tells the CGI the location of the mapping information and to which image it applies.

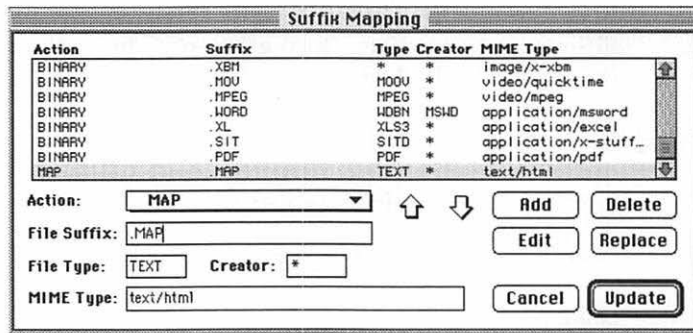
Now launch WebSTAR Admin and choose **Actions** from the **Configure** menu. In the Actions window, type ":CGI:MapServe:Mapserve.acgi" in the **Application** field and "MAP" in the **Action Name** field.



Designating an Action

Click **Add**, and then click **Update** to complete this step.

Now that we have told WebSTAR what to do, it needs to know when to do it. Choose **Suffix Mapping** from the **Configure** menu. In the Suffix Mapping window, type “.MAP” in the **File Suffix** field, “TEXT” in the **File Type** field, “* ” (wildcard) in the **Creator** field, and “text/html” in the **MIME Type** field. Next, choose the newly added **MAP** item from the **Action** pop-up menu.



Establishing a Suffix Mapping

Click **Add**, and then click **Update** to complete the process. Now when someone presses the hot area on a graphic, WebSTAR will launch MapServe, which will in turn interpret the mapping and execute the proper URL request.

We could go much deeper into all this, but that is better left to the more technical documentation that comes with WebSTAR, MapServe, and WebMap. Also, review the sections in this book in which we talk about AppleSearch, Butler, File-Maker Pro, and Internet mail for information on their ACGIs.

What's Not in a Name?

As long as we are more or less on the subject, here is a rule for file names in general: Keep them short and simple. WebSTAR is limited in how it serves files and folders with names containing special characters and spaces. Let's say you want to serve up a file called "My Resume.HTML." Well, okay, but that space in the link will have to be encoded with the proper ASCII hexadecimal code and appear to be used. The space is hexadecimal 32, which translates to "my%20resume.html." Unless you have the ASCII reference chart tattooed on your palm, you are better off chucking your extravagant Macintosh naming conventions and calling it "myres.pdf." Sure it's ugly, but it will always work. Keep this in mind when we talk about realms. If you insist on using special characters, most decent HTML editors can handle the encoding for you.

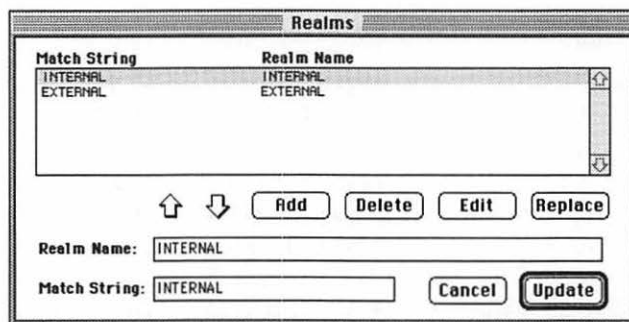
Who's In and Who's Out

I am often amazed by what I find available on corporate sites on the Web. I am not a hacker, but in many cases I don't have to be. It seems that there is a lot of potentially sensitive stuff available in the /Pub folder, or near it, at many sites. Does the Web make your server security that unreliable?

No. What makes your corporate underwear drawer fly open to the world is the failure of network administrators to set up password protections and firewalls, or more often, to realize which data should be put behind a firewall. The latter problem can be addressed by good administrator-user communication; the former, with WebSTAR's security features.

Realms

If you are using one WebSTAR server to support both your internal and external pages, you can require user name and password authentication on files and folders by assigning them to different *realms*. To begin, let's set up two realms: Internal (users) and External (customers and clients). Do this in the Realms window presented when you choose **Realms** from the **Configure** menu.




Establishing Realms

Files and folders are grouped into specific realms by their names. For instance, the following files would be part of the Internal realm we just set up:

```
http://157.22.252.152/internal_db.html
http://157.22.252.152/internal/phonelist.html
http://157.22.252.152/employ/internal_manual.pdf
http://157.22.252.152/graphics/internal_map.gif
```

This is because their URLs contained the match string “internal.” Each such match string must be unique. If WebSTAR can match up a URL with more than one realm, it will go with the first one it finds as listed in the Realms window. Note that when WebSTAR matches a string with a folder name, that folder’s contents become part of that realm, too, no matter what their names happen to be.

To assign user names and passwords to these realms, choose **Add Password** from the **Configure** menu.



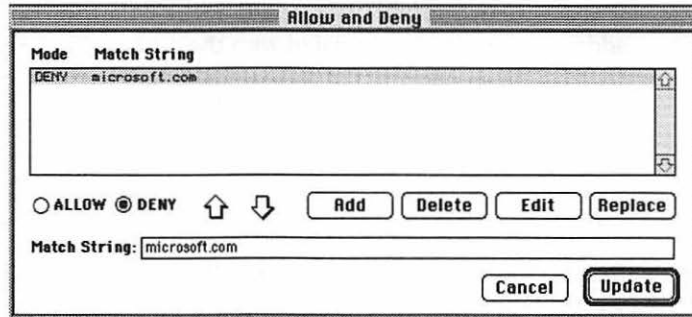
Establishing User Name and Password

When users attempt to access a URL that includes the realm match string from within their browsers, they will be prompted to input this information first. Thereafter, the browser usually remembers the information.

Keep in mind that you actually have one more group of users to think about at all times: those who have no specific realm privileges. In other words, everyone else in the world!

Allow and Deny

In addition to restricting who can see what when they log in to your Web server, you can also control who is allowed to log in at all. To do this, choose **Allow/Deny** from the **Configure** menu. In the Allow and Deny window, type a given computer's IP address or domain name into the **Match String** field, and then select the **Allow** or **Deny** radio button.



Denying Client Access

Click the **Add** button, and then click **Update** when you are finished.

When any computer using that IP address or domain name attempts to access your Web page, it will be confronted with a terse “not authorized” message or a user name/password authentication dialog box. Set up the latter the same way we set up the realms passwords—through the **Add Password** command from the **Configure** menu.

This is a great way to isolate “problem” machines or domains. It is a very powerful feature. In a way, it may be too powerful. By default, WebSTAR permits access to anyone. After you have established one Allow or Deny statement, however, the reverse becomes true. *No client can access the server without your explicit permission!*

Here is what to do if you only want to isolate a few addresses. Let's say you are sick and tired of being bothered by users from that rotten “netfrontiers.com” domain.

To solve this problem, you could configure WebSTAR's Allow and Deny window in the following way:

```
Allow 1
Allow 2
Allow 3
Allow 4
Allow 5
Allow 6
Allow 7
Allow 8
Allow 9
Deny netfrontiers.com
```

It is a bit cumbersome, but it tells WebSTAR to permit access to every possible address except those associated with the domain "netfrontiers.com." What if there are a couple of machines in this same domain to which you do not mind allowing access? Then you could set up the Allow and Deny fields like this:

```
Allow 157.22.
Deny 157.22.2.
Allow 157.22.252.50.
```

First, WebSTAR allows access to all machines with addresses beginning with 157.22. It denies access, however, to all machines with addresses beginning with 157.22.2. That means 157.22.2, 157.22.20, 157.22.200, and so on. Any subnet under that first address is also excluded. If you just want to deny the 157.22.2 subnet, *you must put a period after the address*. In this case, you do want to exclude all those subnets, but there is one machine you want to include: 157.22.252.50. Note that the order in which the statements are listed *does* matter!

None Shall Pass

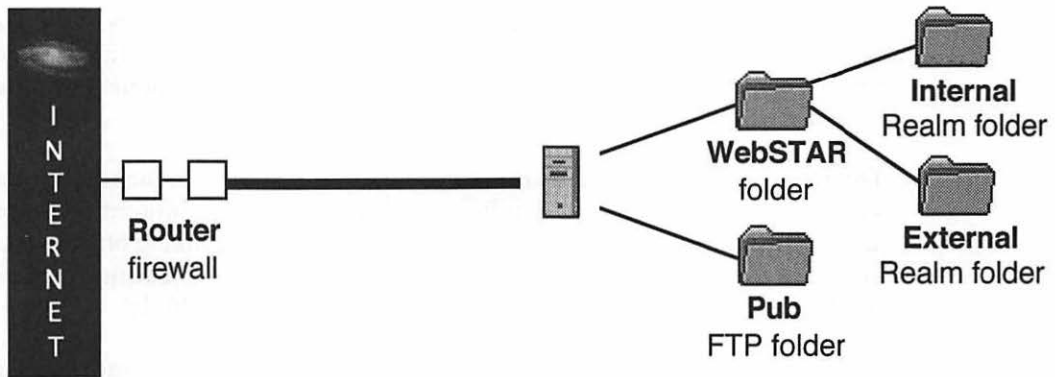
The most powerful security of WebSTAR is that it will not allow access to anyone using HTTP outside of its own folder hierarchy. Regardless of where it is on the server's hard drive, it handles URLs only relative to its own location. As far as browsers are concerned, the root directory is WebSTAR's folder. There is no backing out to "parent directories," as with UNIX machines.

You can use aliases to permit access to files and folders that reside outside WebSTAR's folder, even if those files reside on different volumes or different networked machines. For any URLs in those aliased files to be able to work, however,

they must contain paths relative to WebSTAR's folder. No staggering around the hard drive like a bull in a China shop will be permitted.

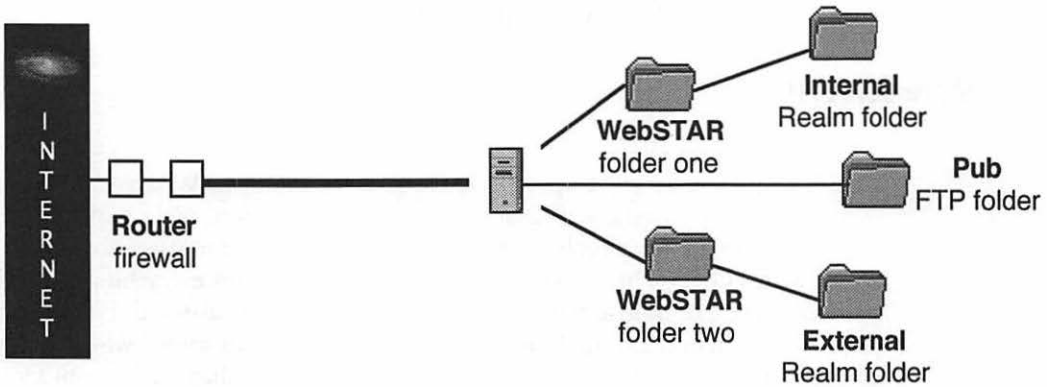
More Servers!

At this point we have set up a good Web server solution. We have one Apple Workgroup Server dedicated to this purpose, and it has four main parts. There is the WebSTAR folder, which contains everything available to users accessing the server with HTTP. There is the Pub folder, which contains everything that FTP users can access or upload with the proper user name and password. Finally, there are the Internal and External folders, within the WebSTAR folder, which contain everything that only those in the employee and customer/client realms can access, respectively.



One Server, One WebSTAR

Suppose you wanted the same result we have here with some pages private (Internal) and some pages public (External), but you did not want to mess around with security realms? In that case, you could give each group their own WebSTAR server on the same Macintosh. Because all configuration information is stored in settings files in the same folder that contains the WebSTAR application, and not only in the Preferences folder of the System Folder, you can run multiple copies of WebSTAR on the same Macintosh with completely different HTML pages and associated access restrictions.




One Server, Multiple WebSTARs

You do not need to buy additional copies of WebSTAR, either. One serial number per Macintosh. You do, however, need to enter unique port numbers in the **Port** field of the Miscellaneous Settings window. The default is 80. All other numbers under 1024 are reserved for other services. It is best to use additional port numbers in the 8000 range.

For users to access a WebSTAR running on a port other than 80, their URLs must include this port designation. For instance, if you direct your browser to “http://www.netfrontiers.com,” you will access the server we have running on port 80. If you want to access another WebSTAR running on that same machine, you will need to know its port, such as “http://www.netfrontiers.com:8001.”

The basic rules of throughput still apply. You have simply changed the procedure for administration. The machine’s ability to handle traffic has not improved. For that you need multiple servers and multiple licensed copies of WebSTAR.



One Domain, Multiple Servers and WebSTARs

To set up a Web site capable of sustaining many thousands of hits per day, you need a Web RAIC (Redundant Array of Independent Computers). Typically this involves adding a custom “BIND” code on your site’s DNS server so that it associates several Macintosh mirrors with a single domain address. Enlist the aid of your ISP for this.

Another way to pull this off is with the freeware Round Robin CGI, written by Jim Hall and available through the StarNine Web site. This application redirects traffic to any number of different URLs that you designate in a text file. You configure the WebSTAR on your first Macintosh—the one with your publicized address—to use “Round_Robin.cgi” as the index page. Then you create the file “robin_addresses.txt” in which you list URLs that point to copies of your home page on the array, such as “acmerc.html,” separating URLs with carriage returns or commas like this:

```
http://157.22.252.151/acmerc.html
http://157.22.252.152/acmerc.html
http://157.22.252.153/acmerc.html
```

It is quick and dirty, but it will do the job.

Most organizations don’t need to worry about Web arrays. You will know that you do if you pay attention to how many times your server turns down a connection because it was too busy servicing others, as displayed in both the WebSTAR Status and WebSTAR Admin windows. If you crank WebSTAR all the way up to its 50-user limit and still have busy messages, it might be time to consider expanding your server pool.

Do not get too excited about the number you see after **Connections** in these windows. It can be misleading because each element on a given Web page requires a separate connection. For example, let's say we have put up the Acme Rubber Chicken Home Page, and three people have accessed it. That would be nine connections! There is one connection for the text page and two for graphics files, multiplied by three hits. You can see how this number can become unduly impressive!

For a more complete discussion of how to analyze Web server activity, see our book *Managing AppleTalk Networks*. For a more complete discussion of putting up a Web server in general, see *Webmaster Macintosh* by Bob LeVitus and Jeff Evans, also published by AP PROFESSIONAL. If you like the casual tone Dorian and I use in our writing, you will like this book, too.

CHAPTER 13: ERASING BARRIERS IN DIGITAL PUBLISHING WITH ADOBE ACROBAT

Futurists have predicted the end of paper-based communications for some time. We will soon live in a “Paperless Society,” they exclaim, in which communication takes place over electricity, air waves, and light rather than compressed and molded flora.

If you consider the technologies available today and the ever-growing numbers of people who are buying and using computers, you have to admit that the futurists are probably right. Up to now, however, the path to a tree lover’s Utopia has run into some substantial barriers, not the least of which has been the diverse and incompatible document formats used by our computers.

This problem is bad enough on one platform. How many times have you double-clicked a Macintosh file and been told that the application that created it could not be found? How many times have you opened such a document, perhaps with the aid of Macintosh Easy Open or DataViz’s MacLink Plus, only to find it full of text in the Courier font with no coherent page layout? I should think at least once.

One of the things we advise all network administrators to do is to standardize on the applications used across the network, as well as on their versions. We also recommend that fonts and printer drivers be standardized across the network. This helps to fend off the incompatibility problems on one platform. Many networks are heterogeneous, however, containing Macintosh, Windows, DOS, and UNIX

computers, if not others. This compounds the problems of document incompatibility and has made ASCII text the only, albeit the ugliest, universal format.

Until, perhaps, today, when we have several vendors vying to give the world the true building blocks for the Paperless Society. The most successful so far is Adobe, the fourth largest software company in the world and the company responsible for so many of the beautiful fonts we have been spraying onto our slices of dead tree since the beginning of the desktop publishing revolution. Now Adobe is bidding to lead the digital publishing revolution.

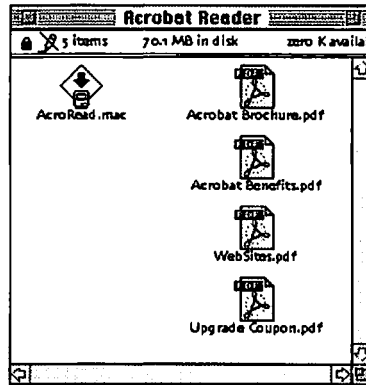
Most digital publishing products share two characteristics. First, they translate a computer document into a format that can be read and printed *without significant alteration* on more than one platform and without the recipient needing the application that originally created it. All that is required at the receiving end is their proprietary translator, called a *reader* or *viewer*. Second, they make use of the technology inherent in a digital environment to give a computer-created document powers that no mere paper facsimile could have, such as keyword search engines, embedded sounds and movies, vibrant graphics, and security encryption. These documents have become more than just a means to a printed end and are generally referred to as *portable* or *digital documents*.

Adobe calls their technology *Portable Document Format (PDF)*. If you have a heterogeneous networking environment, it can provide you with just what you need to help your various computer camps work together without reducing quality to a lowest-common-denominator format such as HTML or straight ASCII. Here I will continue a discussion I began in *Managing AppleShare & Workgroup Servers* and look at using PDF to *repurpose* your documentation for use across platforms, internally and externally, and even on the World Wide Web.

COMPONENTS OF ADOBE ADOBE ADOBE PUBLISHING SOFTWARE

Acrobat Reader

This is the application that permits users to open, read, navigate through, and print PDF documents. You may legally distribute versions of Acrobat Reader for Macintosh and Windows computers, as long as all who receive it accept a special licensing agreement. There are also versions for DOS and UNIX computers. You may distribute the Acrobat Reader from your AppleShare or FTP server. Just be sure to make available the installer (AcroRead.mac) and not just the stand-alone Reader application. The installer includes a special version of Adobe Type Manager (ATM) that must be included, as well as the licensing agreement.



Adobe Reader Installer in a Server Volume

Acrobat Reader gives you convenient controls for such things as navigating through the document via bookmarks, thumbnails, scroll buttons, and keywords. You are also permitted to view the document in a variety of sizes and to cut and paste its text into other documents (unless the author has prohibited this). You cannot alter anything in a PDF document, however.



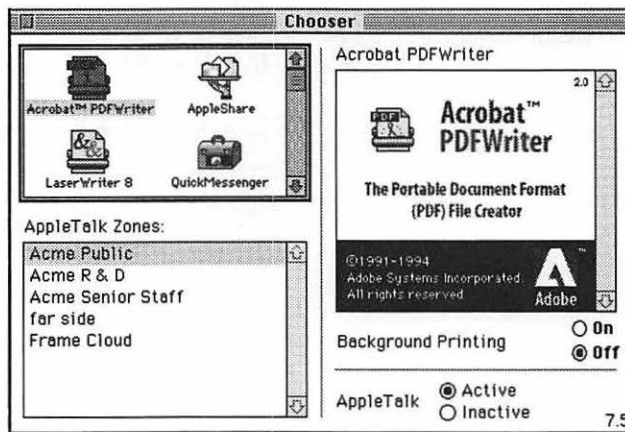
Reader Navigation Controls

Acrobat Exchange

Exchange is Acrobat's document editor. In addition to providing you with all the functions of the Reader, it gives you the ability to set up navigation links within documents, create annotations, create indexes, and secure documents so that they cannot be easily plagiarized.

Acrobat PDF Writer

Included with Acrobat Exchange, this Chooser driver permits users to generate digital documents in the same way they print documents.



Selecting Acrobat PDFWriter in the Chooser

To produce a basic digital document, you simply create it in whatever application you wish, then "print" it after selecting Acrobat PDFWriter in the Chooser.



Some applications do not behave well when creating PDF files in the background. Turn off **Background Printing** if you encounter this problem.

Acrobat Distiller

This application permits you to convert PostScript files into PDF format. This is useful when you are dealing with high-resolution art images and graphics files that

have been saved as “PS” (PostScript) or “EPS” (Encapsulated PostScript). All the above items are bundled together as Adobe Acrobat Pro.

Acrobat Catalog

This application allows you to create indexes containing networked collections of PDF documents. An additional application, Acrobat Search, performs full-text searches on indexes created with Acrobat Catalog. These items are bundled together as part of Adobe Acrobat for Workgroups.

Acrobat Plug-Ins

Like many of its other products, most notably Photoshop, Adobe has incorporated the ability to support “plug-ins.” These are special-purpose Adobe applications and third-party products that add functionality to Acrobat. One of them, Acrobat Search, performs full-text searches on indexes created with Acrobat Catalog. Another, WebLink, permits PDF files to contain hypertext links just like HTML. A third-party plug-in I will talk about here is Remark from Software Partners, Inc. It makes PDF a real workgroup tool by permitting readers to add text annotations, editing marks, file attachments, and voice notes to these otherwise read-only documents.



Acrobat Plug-Ins

Now that you know what tools you have, let's apply them to a problem most network administrators face.

AN EXERCISE IN DIGITAL PUBLISHING

One of the first things you should do with this software is create digital documents for all your important employee-related company publications, like the *Employee Manual*, *Internet Usage Policy*, *E-Mail Privacy Policy*, *Network Orientation Manual*—you know, those. What? *You don't have any of those?* Shame on you!

Examples of Digital Publications on Network-Related Tasks

I want you to make a series of illustrated brochures for your users on various network-related tasks. Topics that might be appropriate for your network are:

Who to Call for Technical Support (and in What Order)

Tell your users where to go when they have trouble. I do not mean the place with all the fire and brimstone. I mean you, your subordinates, your consultants, and the vendor community. Tell them who is the appropriate contact for which types of trouble, too. For instance, “If you have trouble with your network connection, call the MIS,” and “If you need to know how to print letters from FileMaker Pro, first look in the manual, then call Claris at 1-800-735-7393.”

Do this and you *will* thank me later. (All thank-you notes may be sent to tom_dell@netfrontiers.com, along with any good jokes you have heard lately.)

How to Connect to the File Server

Tell users about such things as logging in, logging off, passwords, the issues inherent in copying a file from the server and running it locally versus opening a file from the server, what volumes contain what, privacy issues, and usage policies.

How to Print across the Network

Tell users which printers should be used when, as in “Do not print your e-mail on the Color LazoMatic.” Tell them which printers should be used where, as in “Do not print to devices in the Japan zone if you are in the New York zone.” Tell them

which Chooser extensions to use with what printers and how they should be set up. Tell them your company's standard fonts and where they can find them on the server. Tell them about "direct" versus "background" printing and print spooling.

Here's a good one: Tell users how long a print job should take. Go out and test it. Then they will know if something is really wrong or if printing just seems "slow," and they will complain to you less often (in theory).

How to Obtain Your E-Mail

Tell users about such things as logging in, logging off, passwords, enclosures, LAN-based e-mail conventions versus e-mail destined for gateways (as in Internet mail), privacy, and usage policies.

How to Search Your Network Databases

Tell users where they can find database servers, what different types of databases are available and what they are for (as in calendaring, contact management, and inventory), logging in, logging off, passwords, adding data, removing data, and searching for data.

How to Dial out Using the Network Modem

Tell users how they can access this shared resource and what the common setup parameters are (as in stop bits, data bits, and parity), as well as your privacy and usage policies.

How to Connect from Home or the Road

Tell users how they can access their LAN-based resources from outside the office and when it is acceptable to do so. Stress security here.

How to Connect to the Web Server

The whole purpose of having a Web server is to make information navigation simple, so you should need to concern yourself only with telling users how to reach the home page (URLs) and what passwords they might need. After that, if you have designed your intranet well, they will not get lost.

How to Use the Internet Feed

Tell users what the Internet is, how they are connecting to it, the services it makes available to them (as in WWW, FTP, Gopher, and WAIS), what they should be concerned about in the area of security, and what your usage policies are.

The Works

Get the idea? Do not make these very long—say, five pages maximum. Think of these as “cheat sheets” or “quick start” guides. Compile more lengthy material into full-blown policy manuals with titles such as:

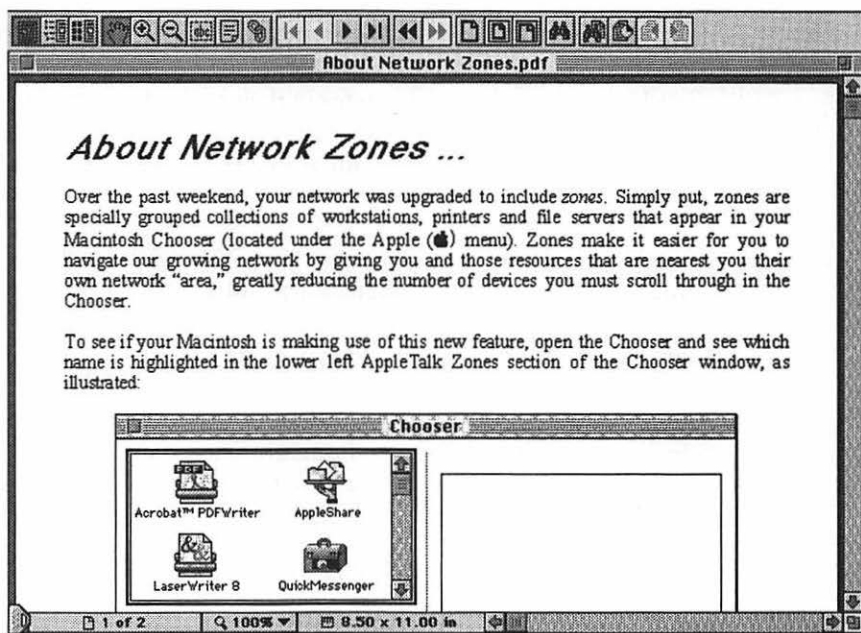
- *Our Company E-Mail Policy*
- *Our Company Data Backup Policy*
- *Our Company Data Privacy Policy*
- *How to Use Your Network Resources*
- *Our Company Security Policy*
- *Our Company Internet Usage Policy*
- *Don't Let Me Catch You Playing Marathon* (optional)

These documents should be the mainstay of your user training library, but there are a couple more types of documents you should also have.

The first are *FAQs* (*Frequently Asked Questions*). You know there are questions you hear a lot from your network users. Put the explanations in a PDF document. Store them in a prominent place on the file server in the unlikely event your users attempt to work it out themselves before calling you. FAQs can also be incorporated into any Web site or e-mail-on-demand list services you might have.

The second are network event *bulletins*. We have created many of these when consulting and they have always made our lives and the lives of our clients much easier. These are documents that you create before you make a major change to your network, introduce a new service, or move the company. They are designed to get users up and running as soon as possible after you, or other circumstances, have disturbed their status quo.

In the following illustration, we show an example of a PDF document similar to one that we created for a client who had just had its first router installed and, with it, *zones*.



Special User Training Document Made in Preparation for a Network Event

The trick here is to make the training documents well in advance of the event. After the event, you won't have a chance to and nobody will read them anyway!

In the following pages, I assume that you have at least some of this material ready. To this material we will apply the Acrobat digital publishing system.

CONVERTING EXISTING DOCUMENTS INTO PDF MATERIAL

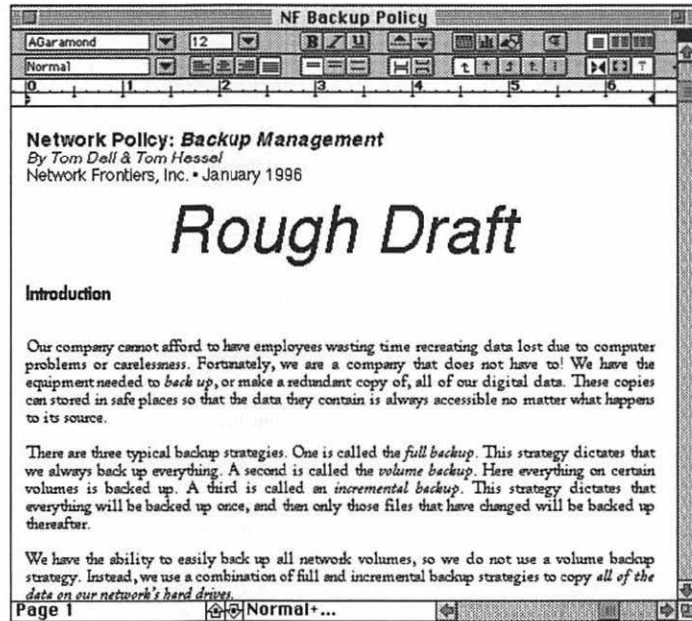
Acrobat has, in my opinion, two really important uses for the average network user. The first is to add functionality to documents such as those listed above. You could distribute these documents in their native formats, say Microsoft Word, but then you would have to deal with operating system and font incompatibility issues. You could send them out as PICTs, but then you would lose the ability to navigate and search for key text.

You could also put documents up on a Web server in HTML, which would give you both cross-platform readability and enhanced functionality. This will not work with all documents, however. In moving your data to HTML you also confine yourself to the design and topography limits not only of HTML but also of the browsers that are reading it. Suppose you have a beautiful full-color brochure that was created in an expensive desktop publishing application. Do you want to lose your graphics, layout, and carefully chosen fonts just to put it on the Web? If not, then you will be interested in Acrobat's second useful function—creating digital document replicas without significant alteration to the original.

With these two uses in mind, let's take a look at how an Acrobat document can be created and shared.


Publishing an Internal Document

As an example, I will use one of the most important network-related documents that you should have, the *Data Backup Policy*. Here is one in Microsoft Word:



Original Document in Word

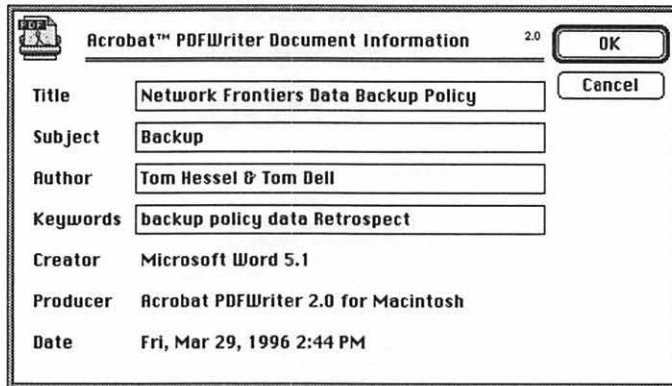
I want to send this out for comments, but I do not want the text altered. Thus, I “print” it as an Acrobat document by first selecting the PDFWriter extension in the Chooser and then using Word’s **Print** command.



Acrobat Print Dialog Box

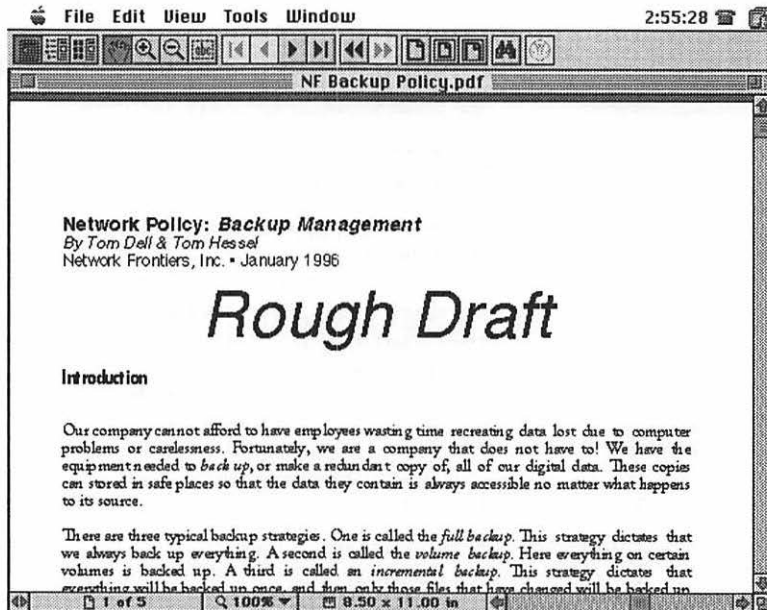
This generates the usual Print dialog box, with a few new details. Selecting the **View PDF File** checkbox will trigger the launch of Acrobat Reader so that I may proof the PDF file after it is created. Selecting the **Short (DOS) File Names** checkbox will cause whatever name you gave the document to be truncated to a DOS-readable eight-character name plus the extension “.pdf.” This is useful in

heterogeneous network environments. Selecting the **Prompt for Document Info** checkbox opens another window in which I can add authorship information and search terms (for use with Adobe Catalog), as illustrated.



Acrobat Document Information Dialog Box


After I click **OK**, I am prompted to tell Acrobat where I want the document saved on my hard drive. After this, the PDF file can be displayed using the Reader.



Read-Only Replica Document in PDF

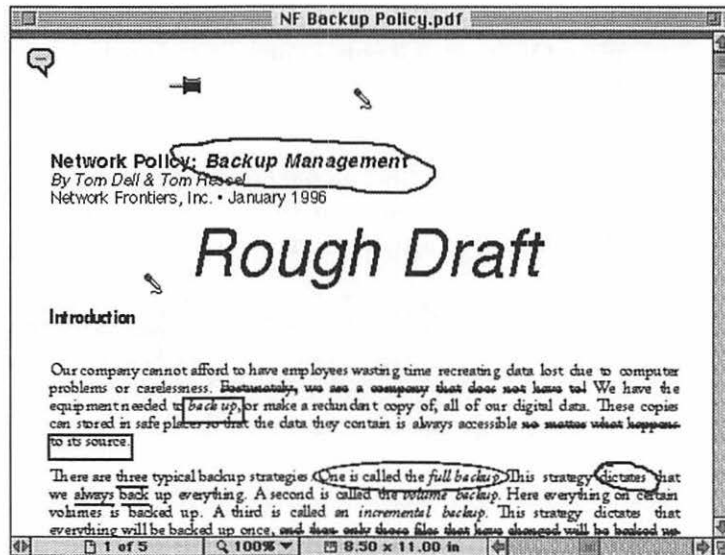
A very close approximation of my original file in a read-only format is now complete. Of course, it is only a rough draft. I could never implement a policy such as this without getting my co-workers to “buy in” to it, so I forward this read-only document to a few of them for comments.

Dorian likes to comment on most things, and his comments are usually . . . forceful. Although he cannot modify this read-only text, he can use Exchange and the Remark plug-in to add comments before forwarding it on for other comments.



Annotating the PDF File

I sent the file out as an e-mail enclosure, and in due course it returns full of helpful and supportive annotations, file attachments, and voice messages from my helpful co-workers—none of which I can print here for fear of offending those with sensitive natures!



Annotated PDF Document

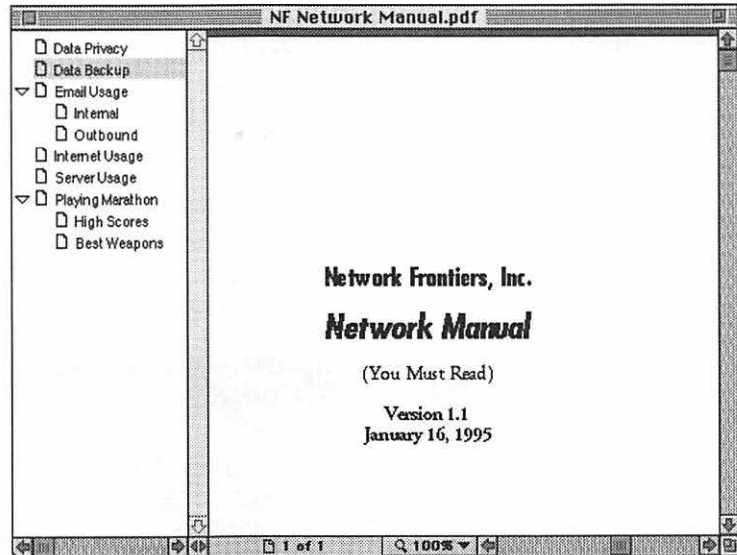
Giving this commentary due consideration, I amend the original document, reprint it in PDF, and distribute it to my network users as a short guide. It is a mere five pages in length and 36K in size.



Word 48K and PDF 36K

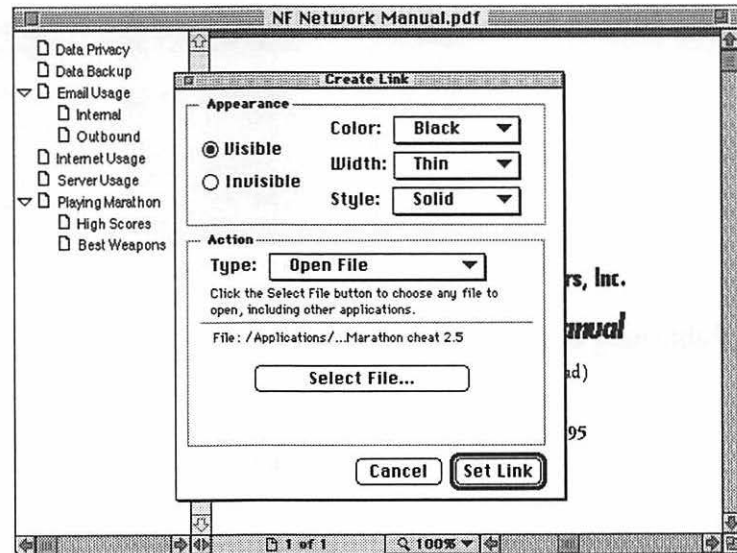
The Word file was 48K. PDF files are usually smaller than the originals, but compression ratios vary.

I put together a longer version with other documents in a more comprehensive manual. Here, too, I use Exchange to add functionality, such as a column of book-marks that provide a point-and-click link to each topic.



PDF Document with Bookmarks


Bookmarks can also be used to link the PDF file to other important documents.



Linking a PDF File to Another File

The link feature also works to give users a point-and-click function for changing the viewing size of the document (as in clicking on a picture displayed at 50% to see it at 100%), launching another file (such as a helper application or a self-running presentation), or following a URL by launching a Web browser and taking the users to a home page.


If you have not yet looked at the CD-ROM that comes with this book, do so now. It contains the Acrobat Reader, and its PDF interface makes use of most of the features we are discussing here. We have also made Acrobat our interface of choice for our quarterly digital newsletter, *Frontiers in Networking*.



PDF Interface to the *Frontiers in Networking* CD-ROM

Publishing an External Document

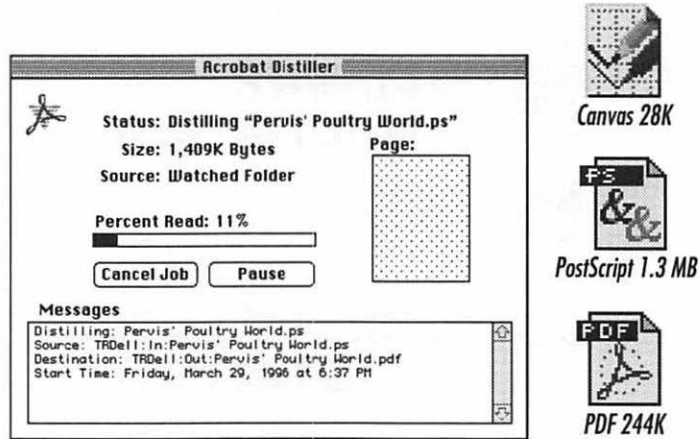
Now imagine yourself at Acme Rubber Chicken. Your ever-talented art department has created an advertisement with Deneba's Canvas that appears in *Pervis' Poultry World* each month.



Ad Created in Deneba's Canvas


This ad, with its multiple art elements and numerous fonts, is output to an EPS file so that only one file—containing both fonts and linked art—is sent to the printer. Your company's ad manager has this file, and he wants to be able to distribute it electronically, too.

For this you use Acrobat Distiller. The job of this application is simple: It makes high-resolution PostScript files into PDF files. When it is first launched, Distiller creates two folders at the root level of the hard drive named "In" and "Out." Throw the PostScript file into the In folder. Distiller will notice it and produce a PDF file for you in the Out folder.



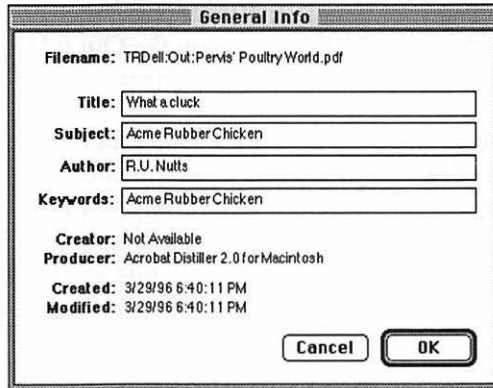
Distiller Processing PostScript File (left); Relative File Sizes (right)

This gives you a file ready for distribution or for further modification. The file is larger than the original Canvas file because the fonts and art have been embedded. However, it is much smaller than the PostScript file.



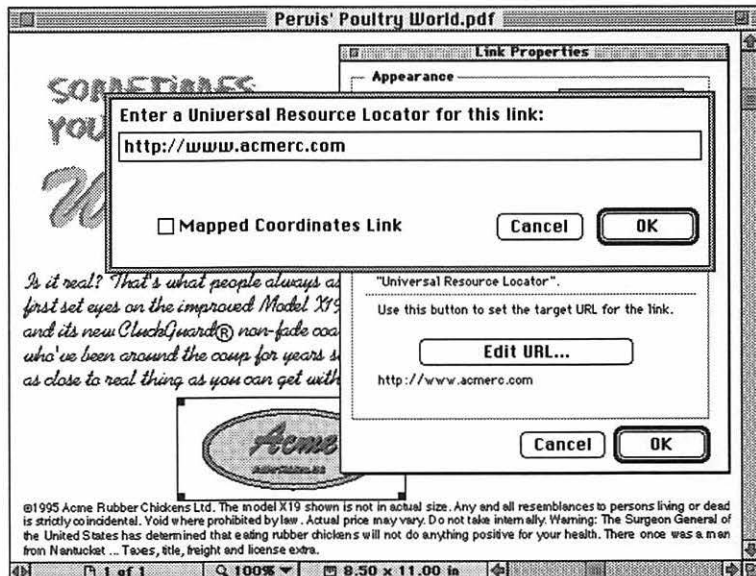
Ad in PDF Format

Before you put the file up on the Web or upload it to a BBS, make some modifications in Exchange. First, add document information by choosing **Document Info** and then **General** under the **File** menu.



Adding Document Comments

Next, why not create a hyperlink to the corporate Web site using the **Link** button?



Adding a Hyperlink

Now the ad can be distributed.

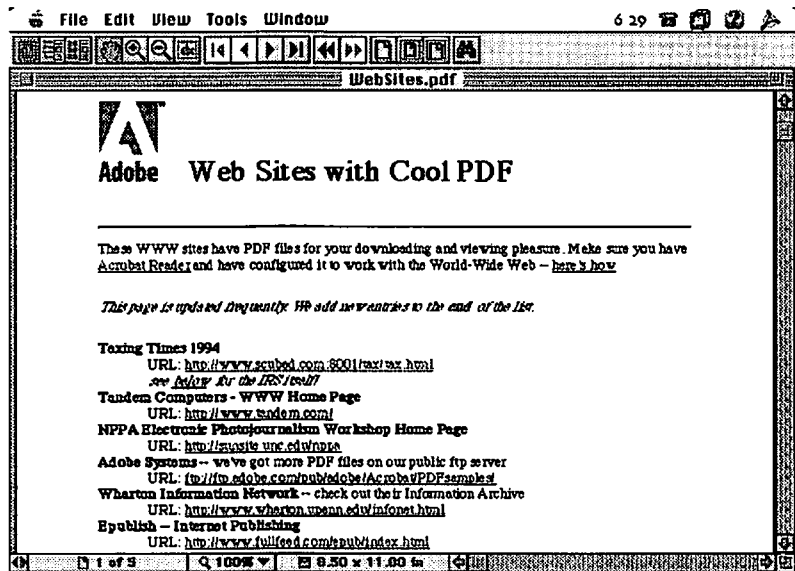
Here is a handy trick for graphic arts forms. Install and launch Distiller on your AppleShare server. Then, make both the In and Out folders shared volumes. When your designers finish a piece, they dump it into the In folder. The proofers pick up the PDF results from the Out folder, edit it with Exchange and Remark, and then send it back to the artists.

By now you should have a good idea of how a technology such as Acrobat can be applied in your network environment. Going into greater detail about the use and features of Acrobat is beyond the scope of this book. I think I would be remiss, however, if I did not tell you a little bit about what you need to do to support those who will be reading your PDF files on the WWW.

USING ADOBE ON THE WWW

I believe Acrobat is one of the most useful applications created for the Internet because it makes possible the distribution of documents that are a true replica of the originals, not just some HTML or text reduction. I advise students in my Webmaster Workshop class to use HTML to provide an interface with links to their documentation and to use Acrobat to publish the actual documentation.

Check out the Adobe Web site at www.adobe.com for some great examples of PDF publishing on the Web.



Acrobat—Great for Publishing on the Web

The one problem with Acrobat regarding publishing on the Web is that, at the time of this writing, most browsers cannot inherently read PDF. I will show you how to overcome this obstacle in this section.

Making WWW Browsers Cooperate with Acrobat

Once you have posted Acrobat files on your intranet Web server, there is one more step. Sit back and wait for this phone call:

You: *Hello, happy-go-lucky network administrator here.*

User: *Uh, yeah, uh . . . I just tried to read that . . . uh . . . 'Company Policy on Spitting' and all I got was a bunch of gibberish. You better post the file again.*

You: *That is an Adobe Acrobat document. Make sure you have installed the Adobe Acrobat Reader application found on the server in the 'Net Installs' folder.*

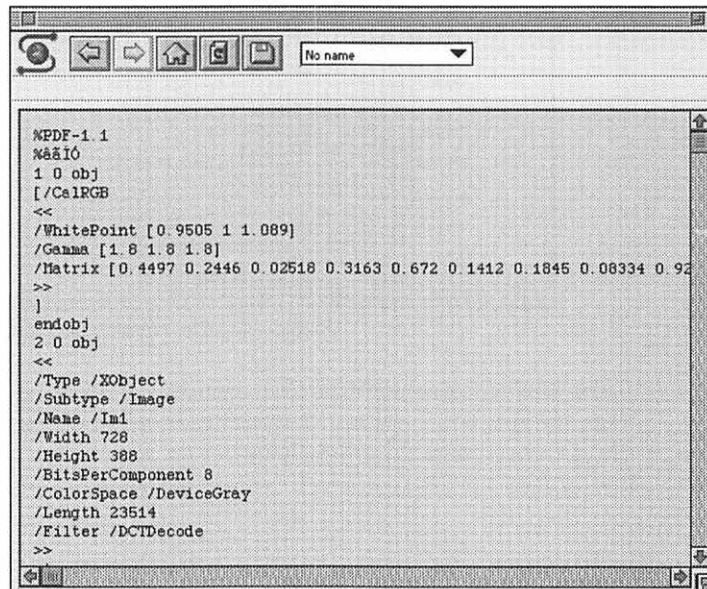
User: *Yeah, I did that. It boots up fine. The problem is your file.*

You: *I opened that file—just now—without a hitch!*

User: *Well, it doesn't work here.*

You: *Okay, try reinstalling the software and call me back at (checking your watch to make sure you won't be there) three o'clock.*

Individual (*sic*)? Not necessarily. In this case *you* are at fault. For shame! This is what the user was seeing:



PDF Document as Seen from Unconfigured Mosaic

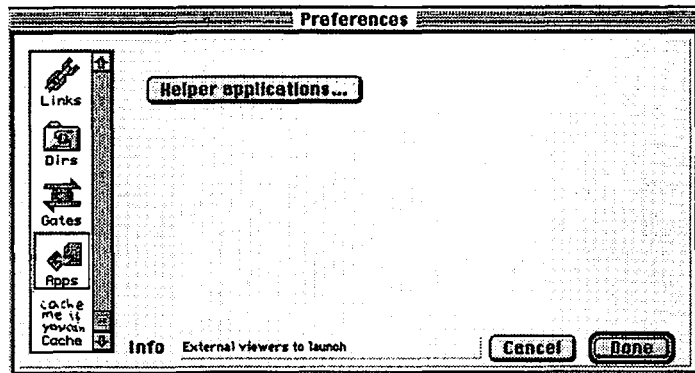
It is not enough to have the Reader sitting there on a hard drive. The browser must also be told that it is there and what to do with it. Depending on the browsers your users and public have out there, this information might have to be manually entered. Otherwise, the browser may try to open the Acrobat document directly as text, with less than pleasing results.

Before you publish Acrobat documents on the Web or an intranet, be sure to create an HTML page describing to your users how they should configure their browsers to use the documents—unless you love phone calls. Adobe has a great page for this, so if you are on the Web, just create a hyperlink to it (<http://www.adobe.com/Acrobat/acroweb.html>). If you are only running an intranet, create your own page that includes the following sorts of explanations.



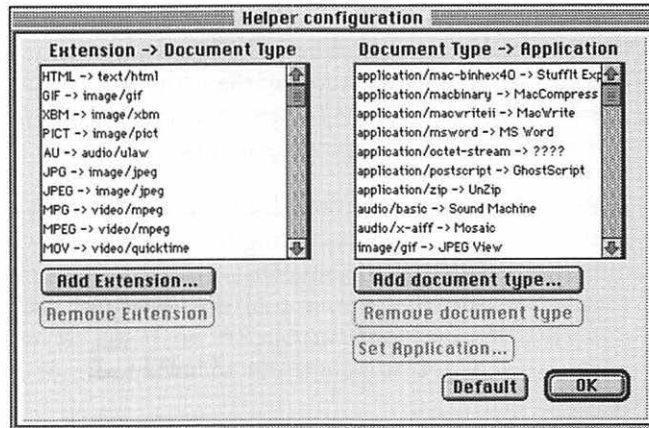
NCSA Mosaic

To configure the National Center for Supercomputer Applications' (NCSA's) Mosaic browser, first choose **Preferences** from its **Options** menu.




Preferences Window

Here select the **Apps** pane and click the **Helper applications...** button.




Helper Applications Configuration Window

Look at the pane on the right. Click the **Add document type...** button and type “application/pdf” in the dialog box that is generated.



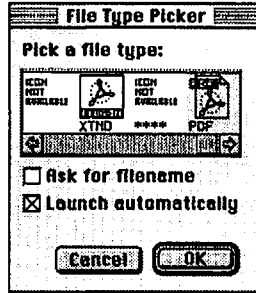
Typing in Document Type

When you click **OK**, a new line appears in the Helper configuration window. Highlight this and click the **Set Application...** button. This generates the usual Open File dialog box in which you can seek out and select Adobe Acrobat Reader.




Finding Adobe Acrobat Reader

Selecting the application generates a File Type Picker dialog box. Choose the icon for PDF and make sure the **Launch automatically** checkbox is selected.



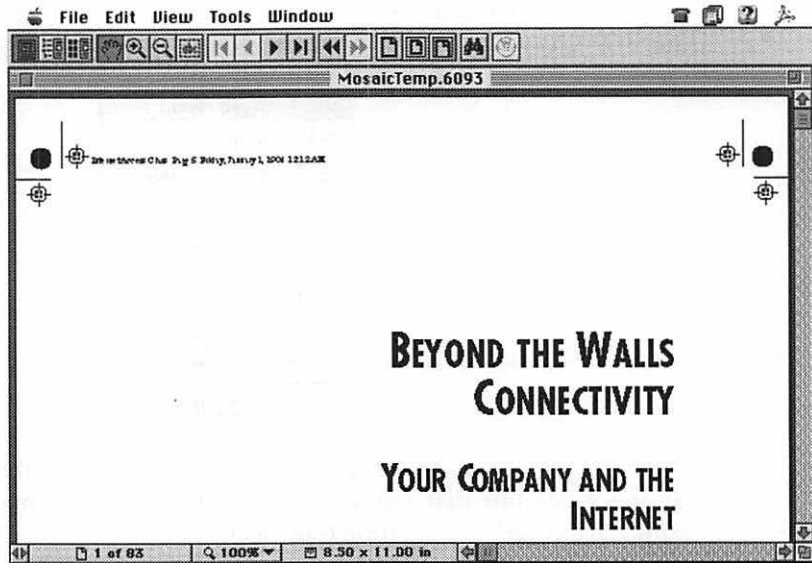
Selecting PDF Icon

When you return to the Helpers window, choose **Add Extension...** on the left. In the dialog box that appears, type “PDF” in the **Extension** field and choose **application/pdf** in the **Mime type** field.



Typing in Extension and Choosing Mime

Restart Mosaic. The PDF document is now a bit more readable than in the previous illustration.

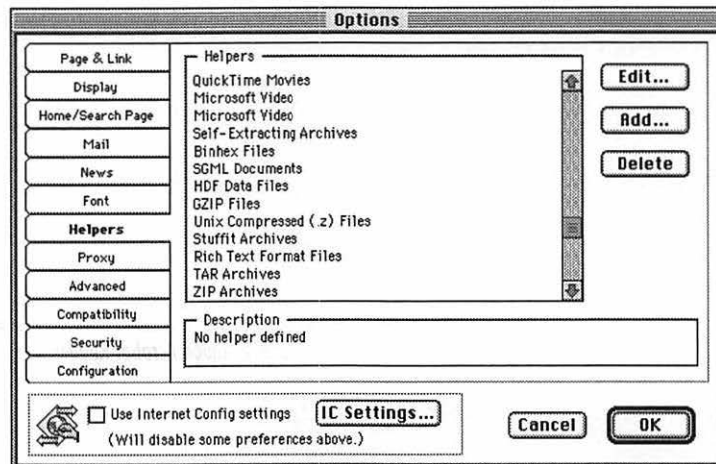


Properly Configured Mosaic Downloading PDF and Launching Acrobat Reader



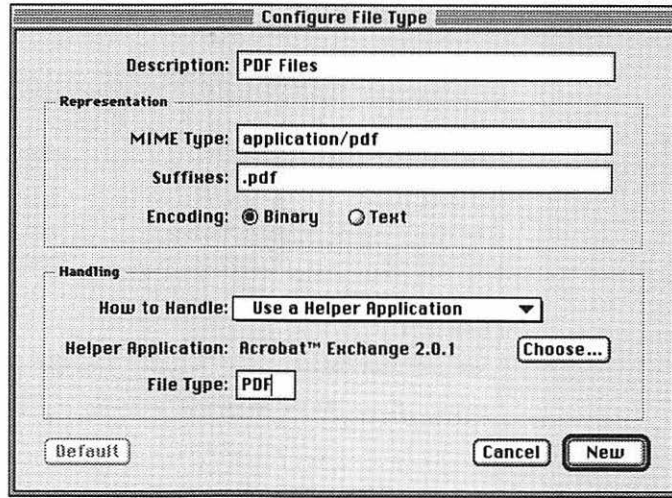
Internet Explorer

To configure Microsoft's Internet Explorer, choose **Options** from the **Edit** menu. The Options window will be displayed.




Options Window

Here select the **Helpers** tab. You may see that “PDF Files” is already listed here. Click the **Add...** button.



Configure File Type Window

A designation such as “PDF Files” should be in the **Description** field. The **MIME Type** field should contain “application/pdf,” and the **Suffixes** field should contain “.pdf.” The **Binary** radio button for **Encoding** should be selected. The last user-definable field in the window, **File Type**, should be—you guessed it—“PDF.” Click the **Choose...** button to generate the usual Open File dialog box in which you can seek out and select Adobe Acrobat Reader.



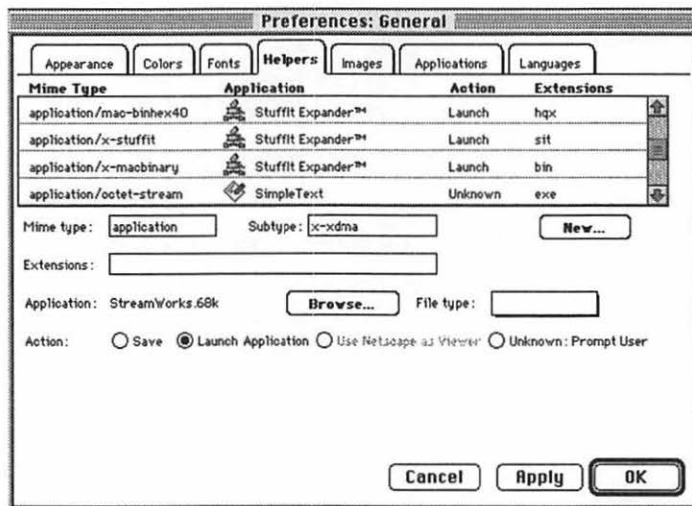
Finding Adobe Acrobat Reader

Finally, verify that **Use a Helper Application** is selected in the **How to Handle** pop-up menu. Restart Internet Explorer when you are finished.




Netscape Navigator

To configure Netscape Navigator browser, choose **General Preferences...** from its **Options** menu. The following window will appear.



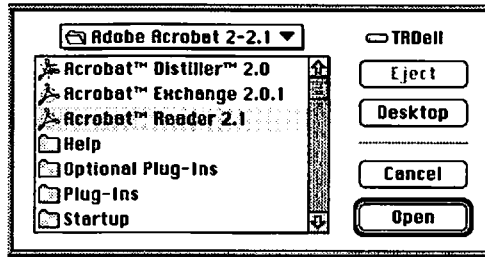
General Preferences Window

Here select the **Helpers** tab and click the **New...** button. In the dialog box that appears, type “application” in the **Mime type** field and “pdf” in the **Mime subtype** field.



Typing in Mime Type and Subtype

After you click **OK**, you should note a new line entry for “unknown” in the General Preferences window. Make sure this line is selected, and then type “pdf” in the **Extensions** field. Also choose **PDF** in the **File Type** pop-up menu. Finally, click the **Browse...** button. This generates the Open File dialog box, in which you can find and select Adobe Acrobat Reader.



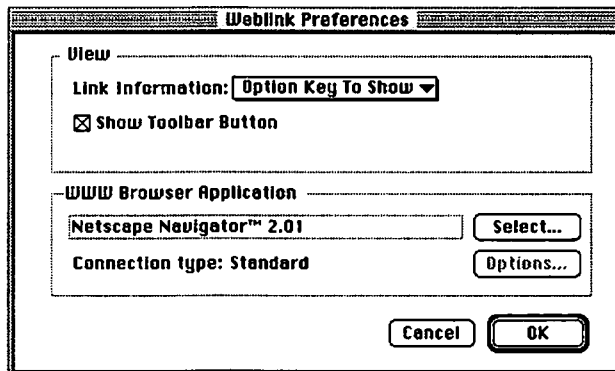
Finding Adobe Acrobat Reader

Back in the General Preferences window, make sure the **Launch Application** radio button is chosen. Then, click the **Apply** button followed by **OK**.



Acrobat Reader

Newer browsers are smart enough to look for Acrobat Reader when it is installed, so most users will not have to deal with this type of thing. Nevertheless, it is good for you to know how it is accomplished. We have done all of this to tell browsers what to do when they encounter a PDF file. If you embed hyperlinks in your PDF files using the WebLink plug-in, an additional step is to configure Acrobat Reader so that it knows what to do when it encounters a URL. To configure Acrobat Reader to work with WebLink, choose **WebLink** from the **Preferences** menu within the application's **Edit** menu.



Selecting the Browser

Here use the **Select...** button to choose the browser.

CHAPTER 14: SELECTING PRINTERS

If we had written this book when AppleTalk was first introduced we would have titled it *AppleTalk Network Services: Printing*. Back then, shared printing was all you thought about when you heard mention of AppleTalk—some Macintoshes, some cable, and a LaserWriter. We didn't need no stinkin' routers.

In those days, it made a lot of sense to share one printer on a network because to put one at each workstation was very expensive. Today, although printers are much less expensive, it still makes sense to put them on a network because they come in so many different flavors, such as inkjet, laserprint, high-resolution, color, large-format, label, and fax. You might find it affordable to put a basic printer at each workstation now, but you will probably still fall back on the network to share specialized printers.

Even though AppleTalk is now used to support a vast number of network services, what was true at AppleTalk's inception is still true today: Printing is what people do most on a network. In this chapter, we pose the questions you need to answer to determine what type of printer you should purchase.

What Kind of Network Access Do the Printers Need?

Many computers come with serial, parallel, and LocalTalk connections. Many more come with Ethernet installed or installable as an option. Sometimes this option can run as much as \$450, not including the Ethernet transceiver.

There aren't many printers that are TCP/IP-based, LAT-based, or Token Ring-based. Check your network access type for the printer before you hook up to it.

Some of the printers can access their Ethernet, LocalTalk, and parallel ports *simultaneously*. Therefore, they can switch between jobs coming into the different ports. If the printers are hooked up to each of the different ports at once, how do you handle priority levels for each of the ports? Ethernet would undoubtedly run over the LocalTalk network with its speed and hunger for printing.

How Many Users Can Have Physical Access to the Printers at a Time?

Knowing *where* to put your printers, and who will have physical access to them, goes a long way in planning your printing needs. It won't do you any good to buy an industrial-strength printer for your art department to share with your sales department if the two of them are on different floors. People want printers to be nearby. They don't want to have to walk halfway across the building to pick up a print job. *You* don't want your users to walk halfway across the building to pick up their print jobs. Have you ever seen a cartoon in the Sunday funnies called "Family Circus?" It always shows the path taken by this little kid named Billie who has to traverse the entire neighborhood to get from one point in the yard to next.

What Type of Resolution Is Needed?

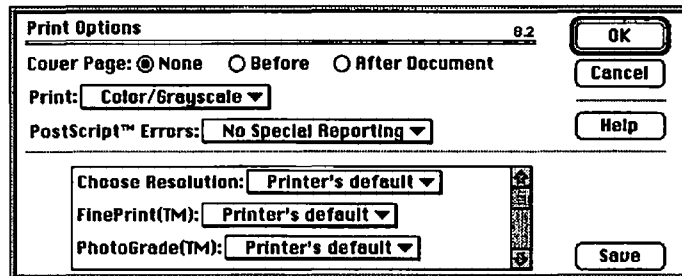
You need to determine the crispness, or *resolution*, of what you are printing—not just for text but also for line drawings, scanned information, photographic input (PhotoCD), slide presentations, charts, or maybe all of these within a single piece, as in a book like this one.

Resolution Isn't Always What You Think It Is

Most printers print at 300 dpi, or *dots per inch*. Many printers, though, surpass that setting and print at 400 dpi, 600 dpi, or 1200 dpi. There is a catch, however. Some printers print at 300 × 600 dpi and call themselves 600 dpi printers. That's because they use a software technique called interpolation that *emulates* 600 dpi printing, or so they say. When examining these interpolated printers, the loupe I use makes their printing look closer to 300 × 300 dpi than to 600 × 600 dpi.

Resolution Helpers

There are a great many printers on the market these days that have what I've come to call *resolution helpers*. These make the printing of pictures or text more clear. However, as with the LaserWriter Pro 630 in our office, if you enable FinePrint or PhotoGrade printing, you automatically *disable* printing at 600 dpi. The opposite is true as well. If you enable 600 dpi printing, you lose the PhotoGrade and FinePrint options.



Apple's LaserWriter 630 Print Options Dialog Box

RAM Takes You There

No, I'm not talking about the Dodge Ram. I'm talking about those expensive little chips we gobble up like grandma's cookies. The answer to higher resolution is somewhat simpler for most printers. The more RAM you put in your printer, the higher the resolution, up to the total resolution specified for your printer. Many color printers run at 300 dpi with 8–12 MB of RAM. At 12–24 MB, they run at around 600 × 600 dpi. With 24 MB of RAM, one of them ran at 1200 × 300 dpi.

Do the Users Require Color Output?

This is the question of the day. If the users do need color output, and they can only afford a single printer, look at a printer that provides *both color and* black and white output *without* breaking the bank on such consumables as paper and color inks or waxes.

Color Wax Printers

Color wax printers use a technology that melts rolls or bits of colored wax onto specially coated paper. If you are using *rolled* wax printing technology, you have a single roll of colored paper, usually with a semi-equal amount of colors, *including black*. This means that when you print a black ink print job, you are using just as much of the other colors as you are the black, even though you aren't printing the other colors.

At the same time, most of these color wax printers use a special coated paper, which is as expensive as the color ink you are printing on the paper or more so. Some of the newer color wax printers come with a special "pre-coat" wax so you can use "plain paper" with the color printer instead of having to purchase the special coated paper.

Color Toner

Another type of color printer uses a color toner and places as many as four layers of color toner—cyan, magenta, yellow, and black—on the printer's laser transfer drum, heats it, and imprints it on plain paper. These printers use a separate black toner cartridge. Thus, printing regular type doesn't cost as much in consumables.

Inkjet

The ever-popular inkjet printers use colored inks sprayed onto a plain paper surface. These printers sometimes have up to four different inkwells for printing. Sometimes they have a "color inkwell" with three colors and a separate "black inkwell" for printing in black and white mode.

Continuous Tone

Finally, there are those color printers that should really be called color proofers. They use photographic technology, called *continuous tone printing*, and photographic paper. These printers are *not* for normal usage, as each page can cost anywhere from \$8 to \$45, depending on the size of the page and the printer.

Do the Users Require Larger Paper Sizes?

Not all printers are 8.5" × 11". Most of them also accept European sizes, such as A4, and a great many of them accept 8.5" × 14" with or without special-purpose feeder trays. Some printers allow users to print up to 11" × 17", or what is known as full-bleed 11" × 17"—paper slightly larger than 11" × 17". Of course, the 11" × 17" full-bleed paper is *much* more expensive than the regular 11" × 17" paper because manufacturers don't sell as much of it.

When selecting a printer that takes different sizes of paper, realize that some of them have to be configured for multiple sized trays. Others can only print the larger-format papers and don't handle the 8.5" × 11" sized papers at all.

Just Because You Can Load It Doesn't Mean You Can Print to It

There is something that the sales folks won't tell you: Some color printers accept 11" × 17" paper, but won't print on any more than 8.5" × 11" of it. I don't know why. They just don't. Some of them cover a legal-sized area when printing to a legal-sized page, but only when printing black and white. I loved dragging that one out of the sales goons.

Will the Users Be Printing Saturated Pages?

Saturated pages are pages with more ink than white space on them. That's a big deal. Really. When you create a slide, overhead, or AV presentation, you usually create a colored background for it. Thus, when you print the presentation, you *print* the colored background as well. That's called saturation printing.

When you print saturated pages, the cost of toner or wax goes up significantly, as does the time it takes to print the page. How well the page is printed with other colors is affected also. Sometimes the background color on saturated pages "pools"

behind text and creates shadow effects. Some printers are not able to handle the color gradations of large, saturated areas as well as they handle them on pages that are less saturated.

While thinking about saturated pages and presentations, if you want a color printer for printing transparencies, make sure you buy one with transparent color inks. Solid color inks show as black on a transparency.

Do the Users Need Special Papers?

There are some printers, many of them color printers, that need certain weights of papers to work properly. Many printers choke on heavy stock (36–42-pound paper), transparencies, labels, and envelopes. If you need to print to these types of media, you may want a printer with special bins to handle them. Swapping paper in and out of the trays, especially 250-page feeder trays, can be time-consuming and irritating, especially if someone forgets to take the paper *out* of a tray. I hate printing reports or charts on our letterhead. Therefore, it is best to look for printers that support multiple bins.

While thinking about this, also check the paper-curl factor of the printers. Some printer's pages curl more than others while printing. If you are printing a single report that you are going to bind and send to someone, that might not be much of a factor. However, if you are printing a report to fax to someone else, or a report of which you need multiple copies, you'll encounter problems. A copy shop may make a photocopy of your original so that the paper comes out straight, and will then use *that* photocopy to make more copies. This means that you would be running third-generation copies instead of second-generation copies. Each time you move through a generation during the photocopy process, your text and pictures become muddier.

Finally, if you need to print pages on both sides, check the printer to see how well it handles your refeeding the paper through for the second side of printing. Some printers handle two-page printing better than others.

What Type of Font Technology Do You Want to Use?

The fonts used in your printers are more important than you think. There are two types of font for your printers: Type 1 and TrueType. If you don't know what these

are, we cover them in great depth in *Managing AppleTalk Networks*, and therefore don't cover them at all here. If you don't understand this font technology, don't try and specify printers. Seriously.

If you are using a printer needing Type 1 or TrueType fonts, define the font list for each printer, and determine whether you need to add a hard drive to the printer to support extra fonts. Sending fonts over the network each time a user wants to print is not a good idea. This takes time and a large number of packets—around 328K worth of data per font is sent when a font needs to be downloaded. Therefore, if you need more fonts than are on the printer, look for a printer with a SCSI port so that you can attach an external hard drive to hold the fonts.

Will the Users Be Printing from Macs Alone or from Macs and PCs?


If all the printer's users are Macintosh users, you don't have the same problems as people who have Macintoshes *and* PCs. Many of the laser printers today, such as Apple's 16/600, come with both Macintosh *and* Windows-based printer drivers. Some printers come with *both* printer drivers, but have setup manuals designed for one platform or the other. The Brother HJ-400 printer, for example, can be printed to by Mac, Windows, and DOS computers. However, the setup manual included with the printer is heavily DOS-oriented. The new Virtual Printer Technology (VPT) printers from DataProducts (the LZR 1580 and 2080) can be set up to be viewed by different computer types as different printers, and even though they support TCP/IP, AppleTalk, DECNet, LAT, and IPX, their manuals for IPX really stink, and the printers are extremely difficult to set up for Novell users.

What Do You Expect from Your Printer's Performance?

The first thing you need to know about a printer's page-per-minute rating, as in 8 ppm or 16 ppm, is that the rating is merely indicative of how many pages can be moved from the paper tray to the output tray per minute. That doesn't have a darned thing to do with rasterizing the page or putting type on it. The more complex the print job, the slower your printer will print the pages you send to it. If I sent the printer a page that contained Courier type on 5% of its area, it would print quickly—probably close to its ppm rating. However, if I sent our training brochure or our logo, it would take *much* longer to print. Therefore, gauge your printer's speed by the type of printing you are doing and by the applications from which you will be printing.

More RAM Means More Abilities

The RAM you give to your printer, to a certain extent, determines what you can print and how fast. I've worked with and tested some printers over the last few years that couldn't print certain pictures at all . . . until, of course, we added more and more RAM. Just because your printer comes with 8 MB of RAM, don't think you have all of that for your print jobs. Take a look at the screen shot of our LaserWriter 630 in the following picture. It is in its standard configuration with no fonts downloaded. Even though it comes with 8 MB of RAM, there are only 1.6 MB available. Adding more RAM to this printer would greatly speed the processing of certain high-density graphics files or complex pages.



Printer Info from the LaserWriter 8.2 Driver Setup

Will the Printer Support Your Applications?

Some printers do some very strange things with graphics and EPS files. One printer that I tested printed all TIFF files as negatives. Other types of printers, those that print only QuickDraw and not PostScript, print all EPS files at 72 dpi instead of 300 dpi. Many color printers are QuickDraw instead of PostScript. Check out the type of work you are printing.

How Are the Consumables to Be Handled?

One of the things you need to know about your printer is how it handles its consumables. Consumables are things like paper, toner or some type of ink (which might be a roll of wax or a bottle of ink in an inkjet), and, with some printers, the laser drum.

When printer marketing agents refer to a “consumables” cost for the printer, what they are usually referring to is a 5% coverage of the paper. In other words, a page, much like this one, with a lot of white space. Once you begin printing graphics on the pages, you move toward higher and higher coverage and your consumables price goes up, along with the longevity of your consumables going down. Printing transparencies and presentation slides, with the page mostly covered with art or background fills, drives your price through the roof when using color printers, and costs you significantly more in black and white mode.

Some of the black and white printers ship with a toner-printing drum single cartridge, and some of them ship with a toner cartridge separate from the printing drum. One thing is for sure, you are bound to run out of toner, and you are bound to need another printing drum. However, what I’ve found is that you often run out of toner *way* before you need a new printing drum. Therefore, be careful when purchasing these types of printers, because each time you change toner, you are paying for a new printing drum as well.

Does the Printer Use Standard AppleTalk Printer Drivers?

Many printers today can be accessed and printed to via the standard AppleTalk printer driver. However, some printers, even ones from Apple, need special drivers. Do you really want to manage different printer drivers for different printers on your network? If you do, that means each and every computer printing to these printers will need the drivers loaded properly and will need them updated as AppleTalk or systems themselves are updated.

A prime example of this is many of the HP printers. They work rather well with their drivers, but when we installed QuickDraw GX along with System 7.5, we couldn’t print to those printers. They used their own, nonstandard drivers, and those drivers hadn’t been converted to QuickDraw GX yet. Sometimes it takes up to six months or longer before a printer manufacturer updates its drivers. The key thing for you to check is that you have the correct PDS files for the printers on your LAN.

CHAPTER 15: ADDING FILE SHARING AND FILE SERVERS

After your printer, the most important aspect of your network will be your file server. This is your data repository. If you ask me, most file servers become *crap* repositories after the first year, as most network administrators don't clean them out often enough—but hey, you probably wouldn't ask me. This chapter is broken into one small section about what a file server is and does, followed by a quick introduction to Apple's Workgroup Server and Network Server lines. After that, we run through some questions you need to consider about file servers. Finally, we end up with a miniature planning guide for building a file server.

WHAT DOES IT TAKE TO BE A FILE SERVER?

This is a review of information in our book *Managing AppleShare and Workgroup Servers*, published by AP PROFESSIONAL and available at most bookstores (or any bookstore where my father has called and ordered it).


All AppleTalk file servers, whether they are workstations running System 7 Personal File Sharing or AppleShare Workgroup Servers, have three basic elements: a computer; a hard drive or other information storage device, such as a CD-ROM or WORM drive; and file server software allowing users on a network to have access to the computer and the information stored on the hard drive. When we talk about AppleShare, we mean the file server software—either Macintosh System 7 Personal File Sharing or AppleShare dedicated software version 3 and later—that runs on Macintosh computers on AppleTalk networks.

The file server software controls the storage systems on the computer, and through the software, the entire storage system (hard drive or CD) or one of the directories (folders) is published as a *shared volume*. The file server controls how users access these volumes and subvolumes through user and group access controls to each of the volumes and folders within the volumes. The file server also controls who has access to the server's volumes by creating user names, or *Users*, and collections of users, or *Groups*. AppleShare security is maintained through login passwords and access levels to folders on shared volumes. Finally, what makes AppleShare work is the *AppleTalk Filing Protocol (AFP)*.

This might sound overly simplistic, but this is generally the way file sharing works. Within this framework, there are two ways that file servers can be set up and files can be shared across the network: *dedicated* file sharing on one large file server and *non-dedicated* file sharing on many smaller file servers. Within the family of Apple-published file sharing software, Macintosh System 7 Personal File Sharing falls into the non-dedicated category, and AppleShare 3 and later fall into the dedicated category. If you need to know more about how the servers themselves work, see *Managing AppleShare and Workgroup Servers*. See our Open Transport primer for more information on the protocols and how AFP works. For now, let's take a look at each of these types of servers in more depth.

NON-DEDICATED FILE SHARING

Non-dedicated access allows any user on the network to act as a file sharing host. The user's files and disk space can be made available to other appropriate users on the network. Apple itself, through System 7, now provides a form of non-dedicated file sharing.



One of the benefits of non-dedicated file sharing is that a dedicated workstation and special file sharing software are not needed. File sharing is a built-in part of System 7, and because of that, doesn't cost the user any extra money to operate. There are, however, some performance issues and other drawbacks related to distributed file sharing.

Completing file service requests from users logged in to a non-dedicated file server requires a certain amount of CPU cycles, RAM, and disk I/O. Depending on the nature of the file service requests and the file sizes, this could place a burden on the server node. With a non-dedicated server, these CPU cycles and other processing power requests must be shared with other foreground and background operations. This means interruptions to the user of the non-dedicated server and longer access times for the those requesting files from the non-dedicated server.


Performance degradation on a workstation sharing files is *really* noticeable. Once when I was teaching a class, I made my hard drive accessible for a few minutes so that I could copy some files from my machine. I went back to teaching, but soon noticed that the Aldus Persuasion file I was running on my PowerBook was bogged down and moving slowly. I couldn't understand why until I opened the File Sharing Monitor control panel. Lo and behold, I found one of my students attempting to pirate software from my computer! I guess if file sharing hadn't

degraded my computer's performance so badly, I never would have noticed that my personal database was being lifted. You know, I still get mad about that, but I don't complain *that* much about degradation these days.

Even worse than performance degradation is what happens to the other users' files when the non-dedicated server is not powered on. With a dedicated server, the only time files are not accessible is when the server is down for repairs. But what happens, say, if in the previous diagram, Sam saves a file on Scooter's machine, and Scooter doesn't come in the next day to turn on his machine? Sam can't access the file on Scooter's machine that day. Worse, if Scooter has a password on his machine, his machine can't be turned on by anyone but him.

DEDICATED FILE SHARING—HARDWARE THAT DOESN'T MOVE

A dedicated file sharing method uses one computer committed to providing access to its files to any other computers on the network that request access. In other words, this file server acts as a host to the other computers. The other computers become the clients of the host.



This client-host method of file sharing requires a dedicated machine with enough storage space to accommodate the needs of everyone accessing it for storage and retrieval. This means that a user cannot actively conduct work, such as word processing, on the machine that is the dedicated file server. However, the same machine can be used for other shared services, such as e-mail or print servers.

The benefits to having a dedicated file server are that all files can be stored in one location for easy lookup, retrieval, and archiving. Also, whatever processes the server must perform to support the network will not drain the existing computing power of any of the individual computers. However, one of the disadvantages of a dedicated system is that when it fails, every active client of that server loses the ability to gain access to the information stored in it.

There are three types of dedicated systems open to Macintosh users: Macintosh-based servers, non-Macintosh-based servers, and minicomputers.


Dedicated Macintosh-based Servers

Even though this approach requires expensive Macintosh hardware, there are some hidden benefits. A dedicated AppleShare server makes a great platform for simultaneously running the Apple Internet Router, as well as other services, like

the mail server or ARA server. Also, a dedicated Macintosh server (or several of them) running the AppleShare software can easily be administered in the areas of backups and security.

AppleShare 3

AppleShare 3.x runs on all Macintosh models, from the Classic through the PowerMacs. However, there won't be any significant speed gains using AppleShare 3 on a Macintosh equipped with a 68040 or PowerPC processor. We recommend that AppleShare 3 be run on 68000, 68020, and 68030 Macintoshes. This software needs at least System 7.0x to run, but we recommend that you use the most current version.



AppleShare 3 can run on anything, but is recommended for '030 machines, such as this LCIII.


AppleShare 4

AppleShare 4.x runs on all 68040-equipped and Power Macintoshes, such as the Quadra and PowerPC series computers. It is three to four times faster than AppleShare 3 software. It *will not run* on anything less than a 68040-based computer. This software needs a minimum of System 7.1 to run. If you are running AppleShare 4.2.1 or later, you need at least System 7.5.3 on the server *and* the end-user workstations. This is because Apple has decided to bundle hardware and software together. The Workgroup Server 7250 is essentially a PowerMac 7200 with a 256K cache card and AppleShare 4.2.1 bundled as an option. The Workgroup Server 8550 is essentially a PowerMac 8500 with a 512K cache card and AppleShare 4.2.1 bundled as an option.

Meet the Apple Server Family

Here is a semi-detailed listing of the file servers from Apple. The low-end products are appropriately dubbed the “Workgroup Server” family, instead of something like “enterprise servers.” While they aren’t as robust and all-powerful as a Novell tower-built server, they are more than ample for most workgroup needs. The high-end “Network Server” family, however, can take on the other server “big boys” in its ability to handle the needs of the most demanding enterprise networks.

Apple Workgroup Server 7250



Front of a 7250

The Workgroup Server 7250 uses the same form-factor as the PowerPC 7200 or 7500. There is room internally for two hard drives and a CD player. This is not the server you want if you can have only one of them. This is the server you pick *after* you have your first server and want to extend file services to a smaller workgroup that needs its own server, but won’t need backup services or other concurrent services running on the same machine, such as e-mail or an AppleSearch Server. If pushed, though, I’d be willing to put a small AppleShare server and a small e-mail group of about 20 people on one of these servers. This would be fine, as long as the users don’t exchange huge files over e-mail while running large documents from the server.

Power, Speed, and Memory

- 120-MHz PowerPC 601 RISC microprocessor with floating-point processor and 32K on-chip cache
- 256K level 2 cache

- 16 MB of RAM, expandable to 256 MB


Networking

- Built-in Ethernet with AAUI-15 and 10BaseT connectors
- Two DMA serial (RS-232/RS-422) ports compatible with LocalTalk

Expansion and Storage

- Three internal Peripheral Component Interconnect (PCI) expansion slots
- Internal 1.2 GB hard drive
- One AppleCD 600i Plus CD-ROM drive (up to 600K-per-second data transfer rate)
- One Apple SuperDrive 1.4 MB floppy disk drive that uses high-density floppy disks and reads, writes, and formats Macintosh, Windows, MS-DOS, and OS/2 disks
- High-speed asynchronous SCSI interface that connects up to seven SCSI devices—three internal, four external
- Three software bundle options: *Applications Server Solution* includes Apple RAID and workgroup software; *AppleShare Server Solution* includes such products as AppleShare 4.2.1 and ARA Multiport Server; *Internet Server Solution for the WWW* includes such products as WebSTAR and PageMill

Apple Workgroup Server 8550



Front of an 8550 with CD and Tape Drive

When I think of a moderate to medium workgroup server, I think of the Workgroup Server 8550. It's not huge like the Network Server 700, and it's not small like the 7250. I guess, then, like the little bear said about his porridge, "it's just right." I love this machine. It comes with both a CD player *and* an internal DDS2 tape drive. It comes with three PCI expansion slots and *just* enough memory: 24

MB of RAM. This puppy should be able to run most of your workgroup's needs without hiccuping. The one thing I wouldn't do is load the ports with ARA, routers, and the like, *and* expect your core users to run their file sharing needs from the same computer. Even though I've seen some of my clients try to do that, I wouldn't recommend it.

***Power, Speed,
and Memory***

- 132-MHz PowerPC 604 RISC microprocessor with floating-point processor and 32K on-chip cache
- 512K level 2 cache
- 24 MB of RAM, expandable to 512 MB


Networking

- Built-in Ethernet AAUI-15 and 10BaseT connectors
- Two DMA serial (RS-232/RS-422) ports compatible with LocalTalk

***Expansion and
Storage***

- Three PCI slots for expansion cards
- Internal 2 GB hard disk drive
- Internal DDS-2 Digital Audio Tape (DAT) drive with Retrospect Remote server and client backup software (optional)
- Internal quadruple-speed CD-ROM drive
- Two SCSI DMA buses supporting as many as eight SCSI devices with data transfer rates as high as 10 MB/s
- Three software bundle options: *Applications Server Solution* includes Apple RAID and workgroup software; *AppleShare Server Solution* includes such products as AppleShare 4.2.1 and ARA Multiport Server; *Internet Server Solution for the WWW* includes such products as WebSTAR and PageMill

Apple Network Server 500 and 700



Front of Network Server 700 with both CD and Internal DAT Drive

These AIX-based (as in UNIX) servers are unlike anything Apple has ever done before and are as far removed from the SE/30 running AppleShare 3.0 as the space shuttle is from the V2. Packed into its wheeled cabinet is room for hot-swappable and/or redundant power supplies, hard drives, optical drives, and tape drives. Its PowerPC processors are mounted on upgradeable cards. It has PCI support for up to seven separate Ethernet segments. Packed, this baby can hold up to 512 MB of RAM and 256 GB of storage space. Good God! All this, and Apple claims its design permits all major components—logic board included—to be replaced in 60 seconds or less!

- Power, Speed, and Memory***
 - Upgradeable 132-MHz (model 500) or 150-Mhz (model 700) PowerPC 604 RISC microprocessor with floating-point processor and 32K on-chip cache
 - 512K level 2 cache
 - 32 MB (model 500) or 48 MB (model 700) of error-detecting parity DRAM, expandable to 512 MB
- Networking***
 - Built-in Ethernet AAUI-15 Ethernet connector
 - Two DMA serial (RS-232/RS-422) ports compatible with LocalTalk
- Expansion and Storage***
 - Six PCI slots for expansion cards
 - Two Fast/Wide SCSI-2 channels supporting seven internal devices, such as 4 GB hard drives; two fixed-mount drives in the model 700
 - SCSI-1 channel supporting seven external devices
 - Quadruple-speed CD-ROM drive; DAT drive optional

SHORT VERSION OF A NEEDS ANALYSIS

Okay, here are some questions to ask yourself or your client when you are trying to decide what type of server you should buy. If any of you wants to make a snappy form for this or wants to add more information, please feel free to do so. If we publish it in *Frontiers in Networking*, our quarterly CD-ROM journal, you get something cool in the mail. I don't know what yet, but we'll think of something.

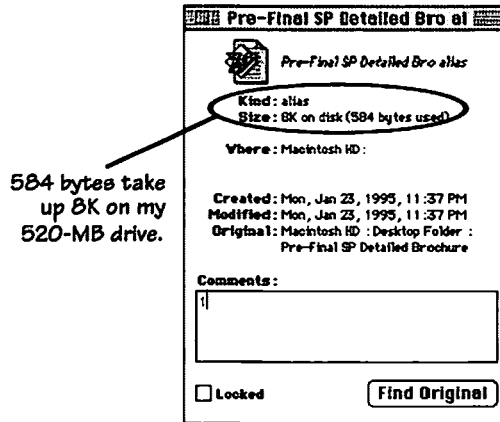
Are Occasional Point-to-Point File Transfers Sufficient for the Users' Needs?

This is the best place to start. Don't be lured by the fact that the AppleShare Workgroup Servers are cheap and easy to manage. If you only have six computers, you probably don't need a server—especially if all six are PowerBooks.

How Large Are the Files to Be Stored on the File Servers?

Simply put, smaller files require a smaller drive, and larger files require a larger drive. That's because of the block sizes of the drives. Remember, a drive only has 65,536 *total* blocks, which means that a 64-MB hard drive has block sizes equivalent to 1K per block ($64 \times 1024 = 65,536$). A 2-GB hard drive (2,048 MB) would have blocks the size of 32K each. That means the smallest of files, like an alias file, would *still* take up a total of 32K, even though the actual size of the file is less than a single kilobyte.

To prove this, I created an alias on my 520-MB hard drive and selected **Get Info...** from the Finder's **File** menu. The file is 584 bytes, or approximately 0.5K, but it takes up 8K on my hard drive.



Drive sizes determine how much room files take up.

Some people make the mistake of obtaining large format drives and putting small files on them, and then they wonder where all the drive space went. For example, you wouldn't want to put Great Plains or MediMac accounting software on a large format drive. Each of these applications creates hundreds of small 1K files and pointer files. Moving these files to a larger drive would only take up more room, without buying you any space.

Will Users Be Storing Personal Information on the File Servers?

If users will be allowed to store personal information on the file servers, allow a certain amount of space for their personal files and folders. Many companies that *do* allow personal folders on the file servers often elect to put these folders on separate hard drives that aren't backed up and are not the main drives. You don't want personal folders on the main drives, as they can grow quite large. For example, one of my former clients allows personal folders on the file server as long as they are kept on the "personal" volume. Of course, this is the volume that is continually running out of hard drive space. Go figure, huh?

Will the Company's Critical Information Be on the File Servers?

When I'm talking about critical information, I'm talking about information that could lose the company business if the information is not backed up in such a way

as to be restorable within a given time frame. If you don't know what type of information is critical to your company, read our book on backup management, *The Complete Guide to Mac Backup Management*. It'll explain that to you. Determine how you will protect critical information. This means something *more* than, not *instead of*, your normal backup plan. Luckily for you, the Workgroup Server 8550 is well equipped to handle data mirroring. If you need more than the software mirroring that comes with the servers, MicroNet has a great RAID setup you can employ. RAID drives, coupled with locking folders in place and labeling them for fast restoration, all go toward creating a safer system for your critical information.

Will the Company's Confidential Information Be on the File Servers?

This question and the previous question are actually quite different. With confidential information, you not only have to protect the data from being destroyed, you also have to protect the data from being seen and copied from the file server. Copying confidential data from the file server doesn't just mean copying files across the network through AppleShare. It also means copying the files across the network through Timbuktu (which copies even faster than AppleShare), or copying files directly from the file server to a floppy disk or a removable drive.

I was at a client site once, doing a security audit. They asked me what I thought I could take from their file server. So, I walked over to the server, unplugged one of their hard drives, and walked out the front door past security with it. I had it in the car for 20 minutes before anyone asked me where the server drive went. I could have been long gone before anyone would have done a thing. The moral of the story is that if you have confidential information, you need more than protocol security. You need some type of physical security, as well.

BUILDING A SERVER PLAN FROM SCRATCH

We have tried to make setting up user accounts on a server as painless as possible. We have divided this section into two parts. This first part is for people who are just beginning to manage their file servers. The second part is for those who have been assigned the task of managing an existing server situation or have been managing one and want to hear *our* two cents.

If you are starting from scratch, think about which users you want on the file servers, how you want to organize them into workgroups, and how best to set up their folders so that you can administer them and the users can get their work done. You won't have any past information, but you won't be inheriting a messy situation, either.

Buying Your Machine

We aren't going to go into all the decision trees about which machine and hard drive system is right for you. We give you our own personal opinion. You can take it or leave it.

CPU

Buy a Workgroup Server 8550 if you want a single, dedicated server. Load it with as much RAM as you can get your hands on. You might think this is overkill, but by the time you think you have grown into the server, you will be growing out of it—trust me on this one. *Most new purchasers* (of course we aren't talking about *you*) buy too cheap and end up with a server that doesn't have the horsepower they need. Don't be one of those people.

Hard Drive

Buy the biggest, fastest hard drive you can afford, and buy it from APS or MicroNet. They have the best warranties in the business. You can never have too much hard drive space, and the faster the better. APS's 1.2-GB hard drive can handle *anything* we throw at it. We also like MicroNet's fast and wide Raven Pro line, as well as their dedicated RAID arrays.

UPS

Buy an Uninterruptible Power Supply (UPS) and attach it to your server. Buy it from APS. A UPS is a big battery that sits between the power your building supplies and your server. If the power goes kablooie, so does your UPS, but *not* your server. Which would *you* rather replace?

Network Backup and Archiving System

Read our book on backup management, *The Complete Guide to Mac Backup Management*, and read the chapter in this book entitled “Planning for Backup Management.” We give you an update on running a backup server on your network, which is something extra from the backup book. Remember, nobody’s ever been fired for spending too much time or money backing up, only for spending too much time and money trying to restore what they *didn’t* back up properly.

WORKGROUP SERVER USER ACCESS PLANNING

Many network administrators often fail to plan for and document the users' connecting to and then leaving the file server. I can't count how many of my clients and students have admitted that they have users in their Users & Groups lists for their file servers who haven't been with the organization for years. I even had one student tell me that he found a user account to be active, even though the user had been gone from the company for over five years! Knowing that you probably hate forms, I created another one just to drive you crazy. Seriously though, you might find this one useful. It will help you plan out the process of putting users on the server, and then removing them when the time comes.

AppleShare User Account Planner

Add User to Server Remove User from Server Change User's Group

User's name _____

Authorized by _____ Date Authorized _____

If New, Trained by _____ Date Trained _____

If Removing User, Assign Privileges to _____

Authorized by _____ Date Authorized _____

Name	Computer	Connection Type	Extension	Location

Groups User Currently Belongs To	Groups User Should Belong To

Actions Performed By _____ Date Performed _____

User Account Planner Form

Broken down into its constituent parts, the form has five major sections. Before you say that nobody would ever fill one of these things out, think about this for a minute: Wouldn't you want that person's boss or whoever is responsible for their part of the data on the server to know this information? Wouldn't you want somebody in the organization to have a slightly better understanding of how the system operates, so that they can take ownership of it? This kind of simple form gives the end users a glimpse into the server's operations and how it affects users who work with them. Hopefully it will give them some insight into the organization of data on the network. If they simply ignore it, that will tell you something about the level of involvement they want to have with the system.

Name, Action, and Authorization

This sets down the user's name as it should appear in the Users & Groups list, and tells the administrator what actions should be performed, from adding the user to the server to removing the user from the server. It also gives an authorization so that the administrator knows to whom the user reports.

Training

If there are new users, *somebody* should train them about the system, show them where to find information, and give them some kind of clue about the information privacy policy the company has created.

Removal

If you are removing users, reassign their folders to someone within the organization—*not* the "Administrator" account. Again, this should be covered somewhere in the corporate data system policy. If you need more information about it, see *Managing AppleShare and Workgroup Servers*, pages 359–384.

Network Information

This is basic information about the name of the user's computer on the network, the computer type, the network connection type, the user's phone extension, and

physical location. Why do you need all that when you are simply assigning a user to the server? You need this information because there are many users out there who don't understand network speed. They sit next to a user with a faster computer and wonder why the "network is slower at my cube." Thus, you want to know some particulars about the user's connectivity and computer system.

Group Information

This should be used to assign or deassign users to groups. Most users don't know to which groups they belong, or what groups exist on the server. Educating them is important if you want the server maintained correctly. Educating the person responsible for the creation and management of the group's information for that particular workgroup is pretty important. That person should know who belongs to what group and who doesn't.

Design	1/6/96
Proposed WorkGroup Name Ali Beckman	Date of Request ext 492
Contact	Contact's Phone


*Basic Information***System Information**

Provide information for each user in the workgroup. Include the users' names as they will appear in the file server, the types of computers on which they are working, the network connection, telephone numbers to contact them if problems arise, and where they are located within the building.

User Names and Information					
Name	Computer	Connection Type	Extension	Location	
Robert Blueberry	PCI Open Transport	Ethernet	173	Room 200	
Veronica Clearwater	Perfoma Series	Ethernet	734	Room 204	
Cynthia Sesspöhl	Duo	Ethernet	676	Room 200	
Jack Ambuster	Quadra Series	Ethernet	222	Room 202	
Louis Loll Sr.	Power Mac Series	Ethernet	152	Room 204	
John P. Putsibckuhm	Power Mac Series	Ethernet	111	Room 204	

System Information

In our electronic form, we have provided choices for the **Computer** type and the **Connection** type. As you are filling out the form, these lists pop up onto the screen as necessary.

*Computer and Network Type Assistant Windows*

AppleShare WorkGroup Planner

Sample Folder Setup

Name & Position of Folder	See Folders	See Files	Make Changes	Owner & User/Group Assignment
Group Folder	X	X	X	Owner will be Administrator
	X	X	X	WorkGroup Name for Assignment
Projects (indent 4 spaces to show sub-folders)	X	X	X	Owner will be WorkGroup Leader
	X	X	X	WorkGroup Name for Assignment
Joe's folder (indent 8 spaces for sub-sub-folder)	X	X	X	Joe Jones (user Name as owner)
	X	X	X	WorkGroup Name (drop folder)
Read Me (indent 4 spaces to show sub-folder)	X	X	X	User Name as owner
	X	X		WorkGroup Name (bulletin board)

**Use 4 spaces for each sub-folder relationship you would like to set

Name & Position of Folder	See Folders	See Files	Make Changes	Owner & User/Group Assignment
Design	X	X	X	All Beckman
	X			Design
- Finished Art	X	X	X	John Putsbokuhm
	X	X		Design
- In Progress Art	X	X	X	Design
- ClipArt	X	X	X	All Beckman
	X			Design
-- EPS	X	X	X	All Beckman
	X	X	X	Design
-- TIFF	X	X	X	All Beckman
	X	X	X	Design
- To Editing	X	X	X	John Putsbokuhm
			X	Design

Page 2 of the Planner

Folder Hierarchy and Access Privileges

Once you know which users you want assigned to the workgroup, have the person requesting this information give you a few hints about folder setup. If this is the first time this person has ever asked for a workgroup setup, walk the person through the steps to filling out this part of the form. This part really is what we call a mechanical “thinker,” or something to start a person thinking the right way

about network management practices. If they feel they had a hand in the design of the folder, they will go that extra mile in keeping the folder organized.

This section is as close as we can get to providing folder setup and security information in a form format. The left column provides an area where you can enter the folders' names and what should be stored in them, if necessary. When filling out this area, space four times for each subfolder indentation. Doing this will show the hierarchical relationship between the folders in the server.

Assign folder access privileges and ownership among the members of the group. If you are filling this out electronically, all you need to do is double-click each of the privilege fields to enter a check mark.

Name & Position of Folder	See Folders	See Files	Make Changes	Owner & User/Group Assignment
Design	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Ali Beckman
	<input checked="" type="checkbox"/>			Design
--Finished Art	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	John Putsibokuhm
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		Design
--In Progress Art	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Design
--ClipArt	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Ali Beckman
	<input checked="" type="checkbox"/>			Design
---EPS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Ali Beckman
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Design
---TIFF	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Ali Beckman
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Design
--To Editing	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	John Putsibokuhm
			<input checked="" type="checkbox"/>	Design

Assigning Privileges

We've filled out the form above as an example for you to follow. Notice that the network administrator owns the Acme Design folder itself. Ensure that the network administrator owns all the root-level folders on the servers. All other folders are owned by Ali Beckman, the group's leader, or in the cases of the individual users' folders, the users themselves are the owners. Notice also that most of the folders in this list do *not* give change access to the users of the folder. Only the owner has that ability. This ensures that the users don't add folders where they shouldn't or put files in places where those files don't belong.

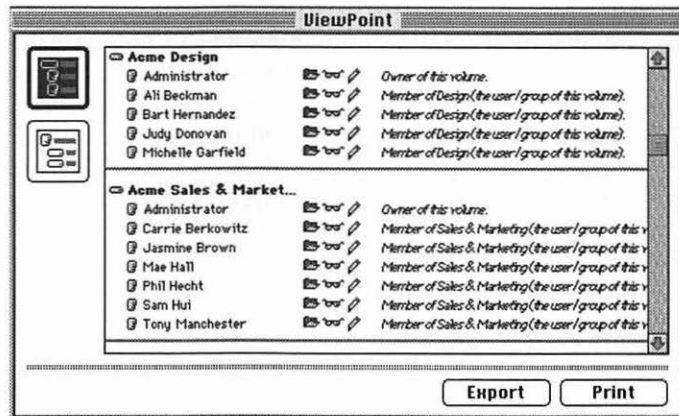
BUILDING A SERVER PLAN FOR AN INHERITED NETWORK

If you are like most of us, you were probably assigned the job as network or server administrator well after the first decisions to computerize were made. You probably inherited a server with a great many Users & Groups and tons of folders and files that neither you nor anyone else understand.

Before you begin reorganizing what you have, spend some time discerning *who* has *what* capabilities on your server. We *highly* recommend that you acquire a package published by Santorini Design and Consulting in San Francisco called Server Tools. This package is a group of small utilities that will make you wonder what you ever did without them.

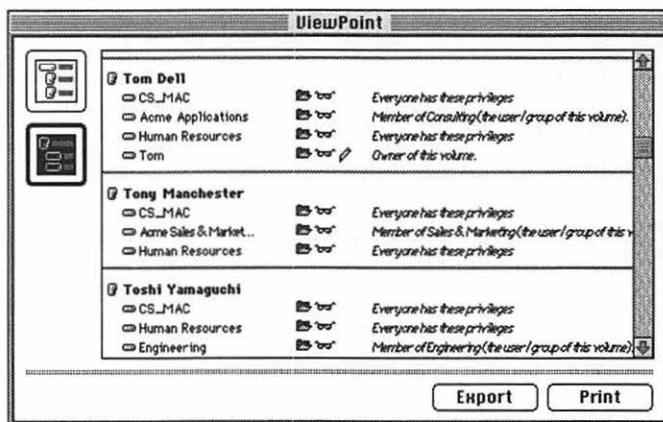
User Account Review

Review what accounts have privileges on which volumes, and what those privileges are. Another of Santorini's tools is called ViewPoint, and it does just that. ViewPoint lets you look at each of the folders on a particular hard drive and gives you privileges for each folder.



ViewPoint Software Showing Privileges

You can also view privileges from a user's perspective, showing for which folders the user has privileges, and what those privileges are.

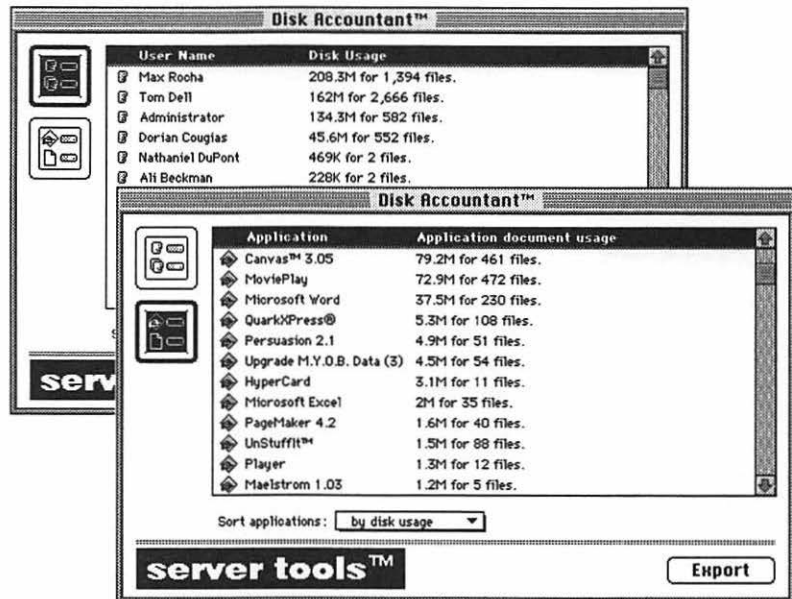


ViewPoint Showing User Privileges

Look for folders that have *all privileges* enabled for Guest access. Also note which groups have privileges to sensitive folders.

Storage Evaluation

You should soon become interested in how much information each user and group is taking up on your file servers. There's an old saying that goes, "There's *never* enough storage." I think it was originally about cupboards and closets in a house, but it applies equally well to disk space and file servers. To gain this information, we use Santorini's Disk Accountant. This utility application lets us view the disk storage usage in three ways: Users, Groups, and Applications.



Users Seen by Disk Accountant (Groups look the same) and Applications Seen by Disk Accountant

You are looking for any one user or group who owns a disproportionate amount of storage space. You can probably bet that this user or group is also using the server's processing time more than anyone else. Once you find that person or group, have a meeting to find out what they are doing and how it affects the rest of the users on the network.

CHAPTER 16: PLANNING FOR BACKUP MANAGEMENT

The first book I ever wrote was *The Complete Guide to Mac Backup Management*, and it is now in its second edition. When I wanted to publish it, nobody but Doc McGraw at APS believed it would sell—not even my own staff. Nobody thought there was a market for a backup book. One of my friends at Dantz asked me “Why on earth would anyone want to read a book about backup?” My answer was plain and simple: Most people do it wrong, and most people suffer at their jobs because of it. Retrospect, the software from Dantz Development, and ARCserve, the software from Cheyenne, back up both PCs and Macs from either a Macintosh backup platform or a PC backup platform. It is our firm belief at Network Frontiers that you should back up each and every one of your computing devices. On a small network, that is pretty easy. Purchase an AppleShare server and it comes with Retrospect. Install the Remote control panels and away you go. Set your backup to start after people go home for the day; it will run during the night, and in the morning everything is backed up and hunky dory.

What if you have portable computers with nomad users like myself? I’m writing this portion on a plane to Idaho while the backup is running in my office. What if you have so much data to back up that the backup spills over into your working hours? Sometimes it takes more time to back up than your network has downtime at night. Obviously, you have to do some further planning. You can’t have your backups running during the day when your users are trying to work. Thus, you have three choices:

- Don't back up everything every night.
- Buy another computer and run it as the backup platform, dividing the backups into two and thus buying yourself 50% more time.
- Run Retrospect 3.x in the Backup Server mode so that it runs constantly and picks up machines as it can.

Obviously the first choice isn't one we recommend. If you want to know why, stop in during one of my classes on backups and watch me do what one student has called "the crazy chicken dance" while I scream at the top of my lungs about the importance of backups. This, then, leaves you with the second or third choice. We have covered the implementation and reasoning behind the second choice *very* thoroughly in our book *The Complete Guide to Mac Backup Management, Second Edition*, published by AP PROFESSIONAL. What does that leave us with for this section? Yup, door number three. I'll explain how to implement a Retrospect Backup Server and the planning you have to do for it. For actual script writing or anything else, you'll have to refer to our backup book.

PLANNING FOR AND IMPLEMENTING RETROSPECT BACKUP SERVER

All planning for backups should begin with a backup capacity and throughput analysis. When planning for your backups, your goals should be simple. You want to be able to:

- Finish a full or new backup within 48 hours.
- Finish daily normal or incremental backups within 10 hours.
- Fit all backups onto a maximum of six tapes, since the best DAT changers, like the MicroNet 48-6, hold up to six tapes.

If, for some reason, any *one* of these goals cannot be met using *one* backup station with conventional backup scripts, you need to use two backup stations to allow for more time, add another DAT changer, or add another drive to your backup computer. Another option, and the one that we discuss here, is to use Retrospect 3.x's ability to create a Backup Server script. This script runs *continuously* during the time periods that you set for it. A combination of full backup scripts and a Backup Server script work very well to back up your computers within the time limits you set daily.

Creating a Backup Planner

In our book *The Complete Guide to Mac Backup Management, Second Edition*, "Chapter 2: Defining Critical Systems," pages 37–43, we discuss in detail how to fill out the backup planner form. Because of page constraints in this book, I can't go as in-depth here. However, here is the first page of the backup planner filled out for our network. Because of the number of systems that we have, our network spills into the second page of the planner as well. On this page, the critical information is shown for the computers, especially concerning how much the backups should be storing on a daily basis.

Capacity & Throughput Planning										
Macintosh Backup ThroughPut Analysis										
Computer Name	AppleTalk Connection	Zone	Volume Name	Vol. Cap.	Free Space	% Change	Daily Backup	Volume Backup Date	Type	Physical Location
Dorian's Computer	EtherTalk 2.0	Acme Public	Macintosh HD	500	50	0.10	50.00		Standard Fixed Disk	118 King, Room 200
Tom's Traveler	EtherTalk 2.0	Senior Staff	Internal Drive	324	100	0.15	48.60		Standard Fixed Disk	118 King, Room 206
Tom H's 520	EtherTalk 2.0	Senior Staff	HD	500	200	0.12	60.00		Standard Fixed Disk	118 King, Room 200
Lynn's PowerMac	EtherTalk 2.0	R&D	HD	240	100	0.50	120.00		Standard Fixed Disk	118 King, Room 200
Brando's Cl	EtherTalk 2.0	R&D	My Drive	80	20	0.05	4.00		Standard Fixed Disk	118 King, Room 200
Maizie's Cl	EtherTalk 2.0	R&D	MaizieDrive	80	40	0.05	4.00		Standard Fixed Disk	118 King, Room 200
Cass's Computer	EtherTalk 2.0	R&D	TheDrive	240	60	0.05	12.00		Standard Fixed Disk	118 King, Room 200
WebServer	EtherTalk 2.0	R&D	Internal HD	1000	200	0.20	200.00		Mirrored	118 King, Room 204
Calendar/ etc. Server	EtherTalk 2.0	R&D	Internal 1	4000	2000	0.10	400.00		RAIDS	118 King, Room 200
	EtherTalk 2.0	R&D	Internal 2	4000	1500	0.05	200.00		RAIDS	118 King, Room 200
	EtherTalk 2.0	R&D	Internal 3	4000	2500	0.05	200.00		RAIDS	118 King, Room 200
	EtherTalk 2.0	R&D	Internal 4	4000	2500	0.03	120.00		RAIDS	118 King, Room 200
Student 1	EtherTalk 2.0	Acme Students	Internal HD	240	120	0.01	2.40		standard	118 King, Room 204
Student 2	EtherTalk 2.0	Acme Students	Internal HD	240	120	0.01	2.40		Standard Fixed Disk	118 King, Room 204
Student 3	EtherTalk 2.0	Acme Students	Internal HD	240	120	0.01	2.40		Standard Fixed Disk	118 King, Room 204
Student 4	EtherTalk 2.0	Acme Students	Internal HD	240	120	0.01	2.40		Standard Fixed Disk	118 King, Room 204

Page 1 of the Backup Planning Chart

Remember when planning for your backups that you can plan for the amount of storage space you are using today, *or* you can plan for the *total* storage space you have available today. We plan for the total amount of storage space. This way, we guarantee that our plans are good for about a year. If you plan for the total capacity you are using today, your plan will be outdated as soon as somebody creates a new file, downloads something from the Internet, or God forbid, installs some Microsoft product at a bizillion megabytes.

The last page of our backup planner shows that it will take approximately 85 hours for a *full* backup, around nine hours for *daily* backups, and a total of seven tapes to capture the information on our network if the drives are fully utilized. Again, since we are planning for full utilization—which gives us about a year's worth of realistic usage before we outgrow the plan—we have surpassed our design guidelines for full backups and requisite numbers of tapes. Something needs to be done about this. We decided to create a Backup Server that will run continuously on our network.

Capacity & Throughput Planning	
LocalTalk Planning	
Add the total of all LocalTalk server capacities (Size in Meg column) and enter that figure into the Capacity in Mbytes cell to the right.	Full Backup Capacity in Mbytes: 0
Multiply the Capacity in Mbytes cell by 1.16 to allow for verification (we highly recommend that you use verification).	Total Transmission with Verification: 0
Divide the Total Transmission amount by the 51 Mbytes per hour throughput rate to obtain the estimated transmission time.	Total LocalTalk Backup Transmission Time (Hrs): 0
Ethernet Planning	
Add the total of all Ethernet server capacities (Size in Meg column) and enter that figure into the Capacity in Mbytes cell to the right.	Full Backup Capacity in Mbytes: 23904
Multiply the Capacity in Mbytes cell by 1.16 to allow for verification (we highly recommend that you use verification).	Total Transmission with Verification: 27728.64
Divide the Total Transmission amount by the 326 Mbytes per hour throughput rate to obtain the estimated transmission time.	Total Ethernet Backup Transmission Time (Hrs): 85.06
Capacity and Throughput Estimate Totals	
Add the Total LocalTalk backup transmission time (along with the LocalTalk to Ethernet Router additional time), and the Ethernet transmission time (along with the Backbone Router additional time) for a total backup time estimate.	Total Full Backup Transmission Time Estimate: 85.06
This calculation will be done for you by adding the change rate for each computer together, calculating both the Ethernet and LocalTalk transmission times, and then creating a summary in this field.	Total Daily Backup Transmission Time Estimate: 9.09
Add the total LocalTalk and Ethernet capacity figures together, along with the daily backup calculations, then divide that number by 6.4 (DAT Compressed tapes offer 6.4 Gig of storage per tape) Gigabytes to obtain the number of tapes you'll need to backup this info. All fractions should be rounded up to the next highest integer.	Total Weekly Backup Capacity Estimate: 7

Notes:

Backup Planner, Last Page

Setting up the Backup Server

Your Backup Server doesn't have to be the fastest machine on your network, or the biggest, but it shouldn't be one of the slowest. Since our network comprises 660 AVs, 6100s, 8100s, 500 series PowerBooks, and two IICi's, we decided to make our lone Quadra 610 the Backup Server. Nobody wanted to use it for anything else, and as it has 16 MB of RAM, we decided it would be the perfect machine. It is *highly* recommended that your Backup Server have at least 8 MB of RAM dedicated to the Retrospect application.

Since we knew immediately that we were going to overrun a single DDS-2 backup tape, our options were to use a bunch of DAT drives stacked on top of each other, or to use a 6-DAT changer. We decided that the best thing for us would be to use a DAT changer. The DAT drive in our main server is used for a different type of backup. We decided that not only do we want to have backups for everything, we also want to have RAID5 storage for the server, mirrored systems for the e-mail and DB server, nearline backups for our Projects folder containing information we need restored within four hours, and an ongoing archive of information about

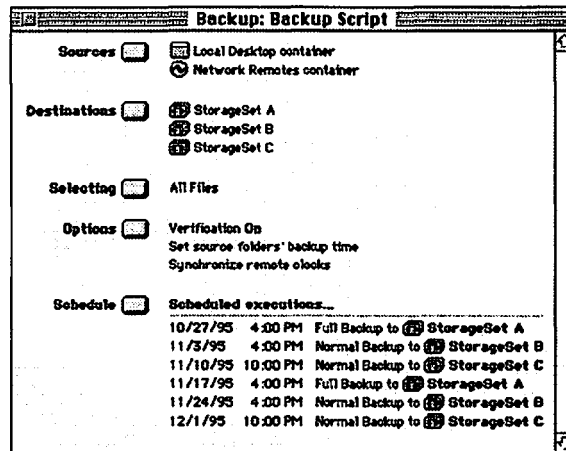
everything we do so that users can retrieve files they might delete accidentally. This will require some of the following setup.

RAID 0 and 5

These are set up according to pages 74–97 of one of our other books, *Managing AppleShare and Workgroup Servers*, using a secondary standard drive for the mirroring, and the MicroNet RAIDbank for the RAID5 system. As they don't have much to do with the network's design, we won't cover them here.

Backup Server

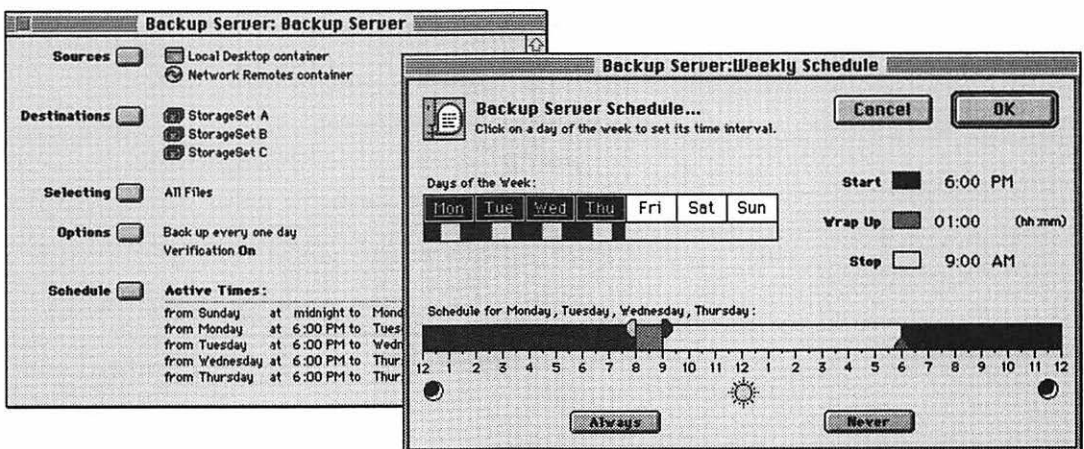
The first thing to script is your full backups. Make sure that you include all your remotes, your groups, and your Local Desktop. If you need to know what those are and their significance when creating a backup script, see our backup book, pages 255–278. The following screen shot shows a completed rotational backup script for three weekly tapes—StorageSets A, B, and C. They are set to run on Fridays at 4:00 P.M. Setting the script for 4:00 P.M. gives you a chance to see Retro-spect winking and blinking at you when you forget to put the correct tape in the tape drive. There is nothing worse than coming in on Monday and finding out that the backup didn't work because you forgot to put in a tape.



Full Backup Script

That's all you need to script for normal backup tapes. Script this separately from the Backup Server script because the Backup Server handles tapes differently. When creating your Backup Server script, make sure that you set your schedule like the one in the following picture, which backs up Monday through Thursday, and not Friday through Sunday. Friday through Sunday is when your full backups are running. You don't want the Backup Server running at the same time as your scripted backups. If the two overlap, Retrospect will try to resolve the conflict by rescheduling one or the other. Don't make life any harder than it needs to be.

The only other point I should mention for your Backup Server script is to put all your StorageSets (which represent the tapes you are using) into it. This allows Retrospect to choose the set into which to place files. It searches your drive chain and backs up to whichever sets it finds. If it finds multiples—as in having both A and C in the tape drive—it backs up to whichever has the earliest backup dates. When finished, your Backup Server script should have all the remote desktops as well as your Local Desktop as sources; it should have however many tapes you are using in its Destinations field; and it should have a set schedule for when you want the backups to start and stop, along with any wrap-up times you need. In the following pictures, our script is shown on the left and the Schedule window is open on the right showing when the backups should and should not run. Our script shows that it should run from 6 P.M. until 9 A.M. Monday through Thursday, with a wrap-up time of one hour. Remember that the wrap-up time is based upon volumes, and the final “stop” time will halt whatever it is doing.



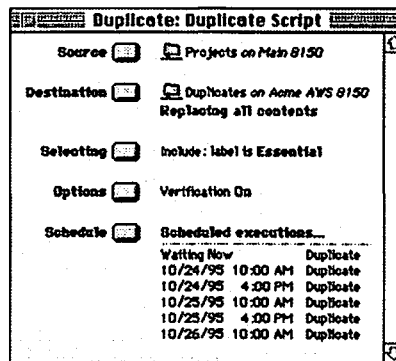
Backup Server Script

This nicely coincides with our regular full backup scripts, erasing the tapes when necessary so that the catalogs don't become too large. That isn't good. Right, Steve?

Now you are ready to set up your archive and duplicate scripts. Run them from one of your file servers with a DAT drive built into it. Since you don't want to confuse departmental backup scripts with departmental archives (remember—backups are for restoration and archives are for file retrieval), you may want to run them from separate computers as well.

Setting up Duplicate Scripts

The duplicate script won't need any special tapes to run and, theoretically, could be run from the same computer from which you are running your backup scripts. First, single out a volume on one of your servers for duplication. In our case, we singled out the active Projects folder of the main server. We set it up so that users can label folders they think are important and need to be duplicated. Each of the folders labeled Essential is duplicated twice a day: at 10:00 A.M. and at 4:00 P.M. Any changes made to the files in these folders are replicated twice daily onto a reserved drive on another server. If a user trashes or corrupts a file, we have a duplicate that has been saved twice daily. RAID systems don't prevent stupid user errors, but this does.



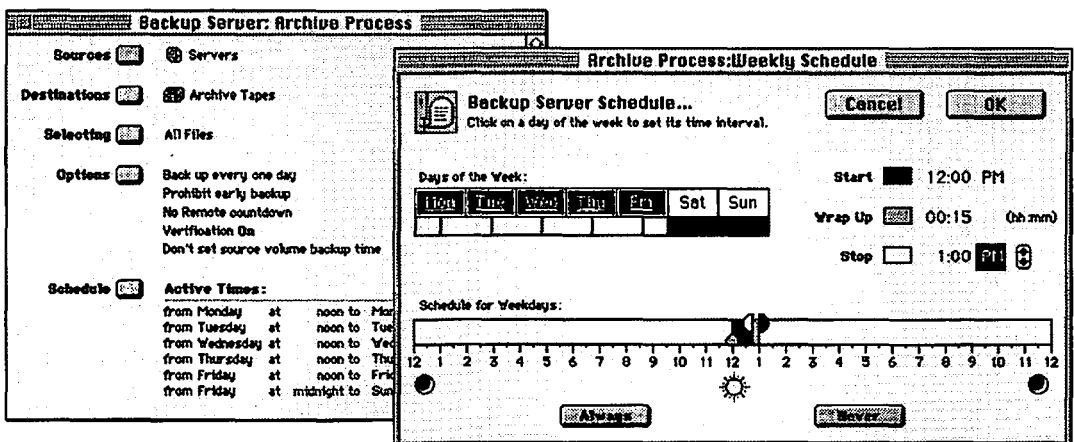
Duplicate Script

Setting up Archive Scripts

Now we are ready for setting up the archive script. As we explained in the backup book, you aren't really creating an archive since you aren't taking files or folders off the server. For creating *real* archives, see our backup book, pages 397–438.

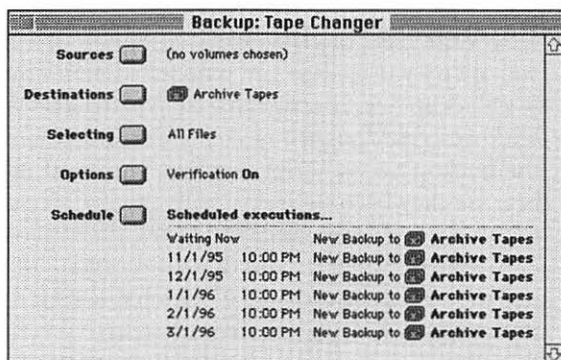
Give your users the lazy man's guide to file management—this is what I believe you are providing, but hey, who asked me?—by setting up two scripts. Script one is a Backup Server script that runs daily and script two is essentially a “no files” backup script designed to create a monthly change in file catalogs and tapes that the Backup Server script requests. You must change your catalogs once a month because they become really, really large. Once they are about 4 MB in size, equal to hundreds of thousands of entries, they become unwieldy. One of my students once asked, “What do you do with a catalog that's 32 MB in size?” All I could say was to change the catalog. Wow, what insanity! Anyway, here is the low-down on the scripts for implementing an archive server, assuming that you have a DAT drive for your tapes.

Set the Backup Server script to run from 12:00–1:00 P.M. daily, with a 15-minute wrap-up time. This backs up computers every lunch and then cycles through your available computers on a catch-as-catch-can basis. Don't worry about whether it can finish all the computers each day. If it doesn't get to one of them today, it will get to it tomorrow or the next day. Don't worry; the files will be in there.



Archive Backup Server Script

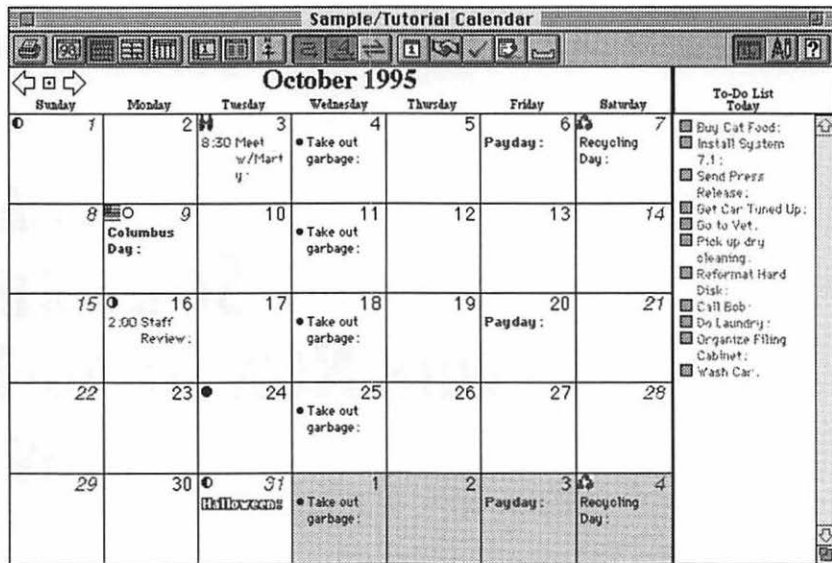
To change your catalogs, run another backup script with no sources that only activates itself to change the catalog and ask for a new tape. We set ours to do its thing once a month. This way, we have January, February, and each consecutive month as an archive. It makes searching for lost files easier and keeps the catalogs to a manageable size.



Consecutive Month Archives

CHAPTER 17: SCHEDULING USERS WITH NOW UP-TO-DATE AND NOW CONTACT

A useful addition to most AppleTalk networks is a workgroup calendar/scheduler. Here users record appointments and other events that can be seen by other users on the network, permitting everyone to know at a glance who is and is not available at any given time. The best of these applications also provide a scheduler. Users can tell the application the names of those they wish to see at a given meeting. The scheduler reviews the calendars and selects a time when all other users are available to attend. With that meeting time established, users can direct the scheduler to record the meeting and inform co-workers of the plan, or to look for the next possible time.




Now Up-to-Date's Calendar Window

Two of the nicest applications made for workgroup scheduling are in the Up-to-Date and Contact bundle from Now Software, Inc. Up-to-Date provides both calendaring and scheduling and can be cross-linked to a Contact database to bring up a wealth of information about individuals or companies mentioned in calendar event fields.

SETTING UP THE UP-TO-DATE AND CONTACT SERVERS

Now makes an easy-to-use installer for Up-to-Date and Contact. You may install their respective server software, client software, or both by choosing the self-explanatory menu items.



Now Installer Screen

General Considerations

Before you use the installer, however, consider which Macintosh will be home to the *Public Event Server*. Up-to-Date is designed to run on both 68K and PowerPC-based Macintoshes with a minimum 450K of free RAM, 2 MB of free hard drive space, and System 7.0.1 (with System 7 Tuner) or later. Busy servers use approximately 1 MB of hard drive space as a work area.

Next, decide whether or not you wish to use Contact along with Up-to-Date and, if so, where the *Public Contact Server* should be located. We recommend that you install Contact to extend the features of Up-to-Date. For example, let's say that you see in your calendar file that you and some colleagues have been scheduled for classes at Network Frontiers, Inc. "What is a Network Frontiers?" a subordinate asks. A mouse click links you to a Contact file and you know what it is, where it is, and why you are going there!

We use Contact the most in the consulting end of our business. It is great for PowerBooks and Newtons, as it remembers the records you can access on the network, regardless of whether or not you are on the network. That means your business address data can move around with you as easily as your personal addresses. You can even make changes to the data off-line. When you reconnect to the network, Contact's server component will synchronize your address book and keep updating both with the newest entries.

Contact has similar hardware requirements, needing at least 1 MB of free RAM and 3 MB of free hard drive space for 68K installations or 4 MB of free hard drive space for Power Macintosh use. It also requires System 7.0.1 (with System 7 Tuner) or later.

You do not have to install the client portion of the application on the machine on which the server elements reside, unless you plan to do all configurations locally.

Now suggests that you use these guidelines for the destination Macintosh:

If you have up to . . .	Generating up to . . .	Use at least a . . .
50 users	1,000 events or contacts	Performa 200
250 users	10,000 events or contacts	Quadra 630
500 users	10,000 or more events or contacts	Power Macintosh 7100

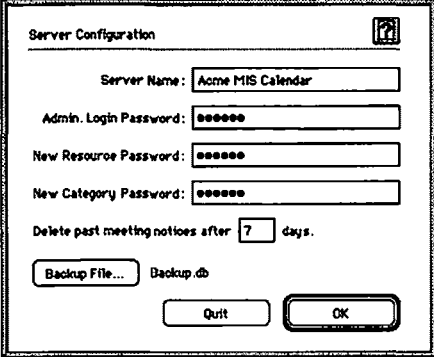
Don't take those numbers as gospel, however. I cannot imagine trying to give 500 people a shared calendar that fits on a 14-inch screen! For more than 10,000 events or contacts, remember that when a workstation is off the network, or the server is—perish the thought—down, server-based data must be saved locally at least in part. If you have a really big Contact database, that could eat up an inordinate amount of workstation hard drive space. We tried one with 64,000 contacts. I do not know how that would have turned out because I became tired of listening to my PowerBook trying to synchronize after about a day! In my opinion, databases of that size belong in FileMaker Pro or Butler. Which brings us to the most important thing in deploying this service: These are truly *workgroup* applications. Apply them to small groups of, say, 25 users or less. You can do this by department, by project, or by whatever mechanism impels the users to work together. Also keep the Contact databases you wish to link with Up-to-Date to a maximum of, say, 1,000 records.

You might decide to have one or many networked calendar/contact combinations. We will discuss servicing multiple workgroups from one combination server. You could also give each workgroup their own server on separate machines. You could even run multiple separate server applications on a single machine, giving each workgroup their own calendar/contact database in that way, if this makes more sense to you administratively.

Next, we set up one Up-to-Date server and an associated Contact server for the Information Services department of our fictitious company. You can expand upon this in deploying it throughout your own organization.

Initial Configurations

When the installer has finished, you are prompted to restart the Macintosh. When it comes back up, you can begin your server configuration.




Public Event Server Configuration Screen

In the **Server Name** field, enter a name indicative of the company or department the calendar is to serve. We have chosen “Acme MIS Calendar.” In the **Admin. Login Password** field, enter a code that will prevent others from messing around with the Public Event Server. In the **New Resource Password** field, enter a code that must be entered for anyone to add users, conference rooms, or resources to the group scheduling system. In the **New Category Password** field, enter a code that must be presented before anyone is allowed to create, delete, or move public event categories. Finally, click the **Backup File...** button. You are asked to specify a place and name for a backup of the Public Event Server’s database. If you change

the name from the default of “backup.db,” be sure to include the “.db” extension in the new file name.

The backup feature is important. The Public Event Server database being backed up contains all the calendar categories, events, users, and resources. This is not data that you want to have to recreate from memory every so often. All information added to the Public Event Server database is saved every five minutes. The backup feature makes sure that, in addition, the entire database file is duplicated at midnight, assuming that you always have it running. Otherwise, it backs up the next time you turn on the server. This redundancy protects you should the main database become damaged or corrupted. Your daily network backups should provide even further protection.

Click **OK** when you are finished here. What remains is the unobtrusive Status window from which you will configure and maintain some features of your network calendar.



Public Event Server Status Window

Repeat these steps in the Public Contact Server configuration screen.

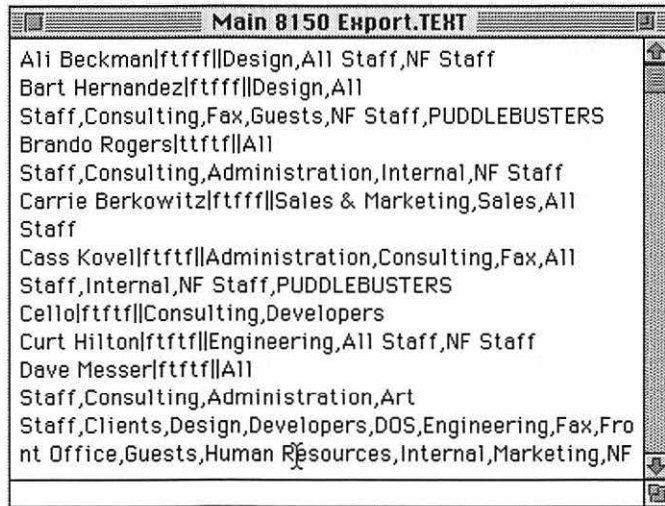
Configuring the Public Event Server

You are ready to add the people, places, and things for Up-to-Date to track.

Installing Users, Rooms, and Resources

The first step in setting up the Public Event Server is to add the names of the people who will be using it. You could enter the name and account information of each individual remotely from the Up-to-Date client application. If you have many users, however, a more efficient way is to export the information from an AppleShare or e-mail server and then import that database file into the Public Event Server.

First, edit the list of users you have gathered from an application such as AppleShare. Delete all the information that does not pertain to users, and then delete all those users who will not be a part of the network calendar. Be sure to count the number of remaining user names to make sure that you have enough legal copies of Up-to-Date to support them.

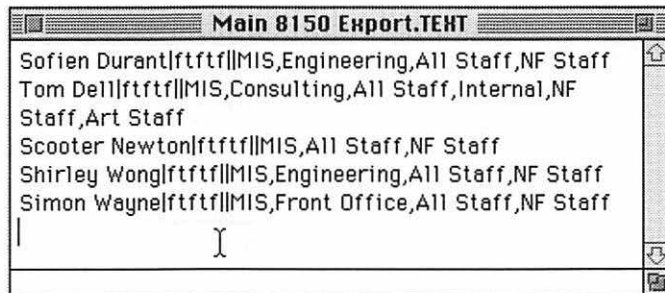


```

Main 8150 Export.TEXT
Ali Beckman|ftfff||Design,All Staff,NF Staff
Bart Hernandez|ftfff||Design,All
Staff,Consulting,Fax,Guests,NF Staff,PUDDLEBUSTERS
Brando Rogers|tftft||All
Staff,Consulting,Administration,Internal,NF Staff
Carrie Berkowitz|ftfff||Sales & Marketing,Sales,All
Staff
Cass Kovel|ftftf||Administration,Consulting,Fax,All
Staff,Internal,NF Staff,PUDDLEBUSTERS
Cello|ftftf||Consulting,Developers
Curt Hilton|ftftf||Engineering,All Staff,NF Staff
Dave Messer|ftftf||All
Staff,Consulting,Administration,Art
Staff,Clients,Design,Developers,DOS,Engineering,Fax,Fro
nt Office,Guests,Human Resources,Internal,Marketing,NF
  
```

Text List of Users Generated by AppleShare

In our example, we are making a calendar that will serve only a small IS department, so we have whittled the list down quite a bit.



```

Main 8150 Export.TEXT
Sofien Durant|ftftf||MIS,Engineering,All Staff,NF Staff
Tom Dell|ftftf||MIS,Consulting,All Staff,Internal,NF
Staff,Art Staff
Scooter Newton|ftftf||MIS,All Staff,NF Staff
Shirley Wong|ftftf||MIS,Engineering,All Staff,NF Staff
Simon Wayne|ftftf||MIS,Front Office,All Staff,NF Staff
|
|
  
```

Edited List of Network Users


The next step is to reformat the list. Delete everything but the user's name and the group designation that is the one you wish to use in the calendar. In our example,

all the users are in the AppleShare “MIS” group, but that’s still not specific enough for the calendar. Thus, we have deleted the group designation as well.

The reformatting process works like this: Type some data, press **tab**, type some more data, press **tab**, and when you reach the last piece of data to enter, press **return**. Everything before a tab is a data *field*. Everything before a return is a data *record*. When you are finished, you have a *tab-delimited text file*.

Here is the data you need to enter in each field:

1. **Type:** In this case, it is *user*.
2. **Name:** You have the user’s name from the list of users we just discussed.
3. **Password:** This could be a new password, or the password used for AppleShare or e-mail accounts. If you don’t care about passwords, just press **tab**.
4. **Description:** This could be the group designation, as in “MIS,” or it could be something more descriptive, like a job title, such as “Network Manager.”
5. **Capacity:** Press **tab** here. We’ll talk about it later.
6. **Extension:** Put in the users’ telephone numbers or extensions.
7. **Department:** Again, this could correspond to an AppleShare group designation or it could be something more descriptive, like a physical location, as in “fourth floor.”

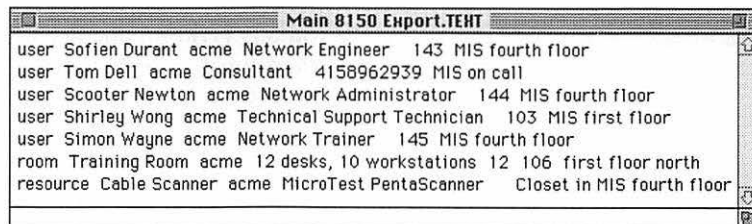


```
user Sofien Durant acme Network Engineer 143 MIS fourth floor
user Tom Dell acme Consultant 4158962939 MIS on call
user Scooter Newton acme Network Administrator 144 MIS fourth floor
user Shirley Wong acme Technical Support Technician 103 MIS first floor
user Simon Wayne acme Network Trainer 145 MIS fourth floor
```

Formatted List of Network Users

There may be a few more things you want to do before you import this file into the Public Event Server. In addition to users, Up-to-Date can also keep track of the availability of rooms and resources. Enter this data as follows:

1. **Type:** Instead of *user*, make it *room* or *resource*.
2. **Name:** This is the name of the room, as in “Training Room,” or the resource, as in “Cable Scanner.”
3. **Password:** This could be a new password, or just press **tab** to leave it blank.
4. **Description:** This could be, for example, that the room has “12 desks, 10 workstations,” or that the resource is a “MicroTest PentaScanner.”
5. **Capacity:** Enter the number of people that Up-to-Date could schedule to use this room at the same time. Leave it blank for a resource by pressing **tab**.
6. **Extension:** Enter the room’s telephone number or extension. Leave it blank for a resource by pressing **tab**.
7. **Location:** Enter the physical location of the room or resource. For instance, for a room you could enter something like “first floor north” and for a resource you could enter something like “closet in MIS, fourth floor.”



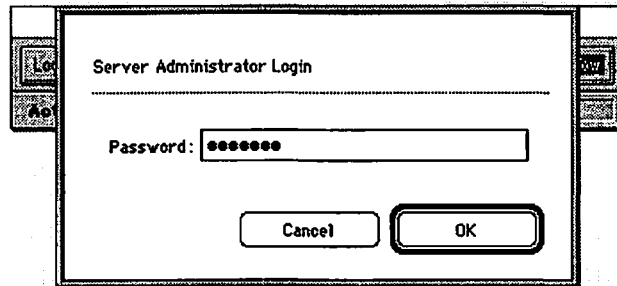
```

user Sofien Durant acme Network Engineer 143 MIS fourth floor
user Tom Dell acme Consultant 4158962939 MIS on call
user Scooter Newton acme Network Administrator 144 MIS fourth floor
user Shirley Wong acme Technical Support Technician 103 MIS first floor
user Simon Wayne acme Network Trainer 145 MIS fourth floor
room Training Room acme 12 desks, 10 workstations 12 106 first floor north
resource Cable Scanner acme MicroTest PentaScanner Closet in MIS fourth floor

```

Formatted List of Network Users, Rooms, and Resources


You are ready to import this file into the Public Event Server. If it is not there already, copy the file to the server machine’s desktop. Next, gain access to the Public Event Server’s database by clicking the **Log In** button in the Status window. You will be prompted for the password you set up earlier.



Logging in to the Public Event Server

Once you are in, the **Log In** button changes to **Log Out**. Select **Import Users** from the **Server** menu, or click the spigot button and enter the name of your tab-delimited text file in the dialog box that is generated.

If you made a mistake in the file's creation, Up-to-Date will tell you about it:



You screwed up. Try again.

Go back into the file and fix the error. The error message will tell you which records are wrong.

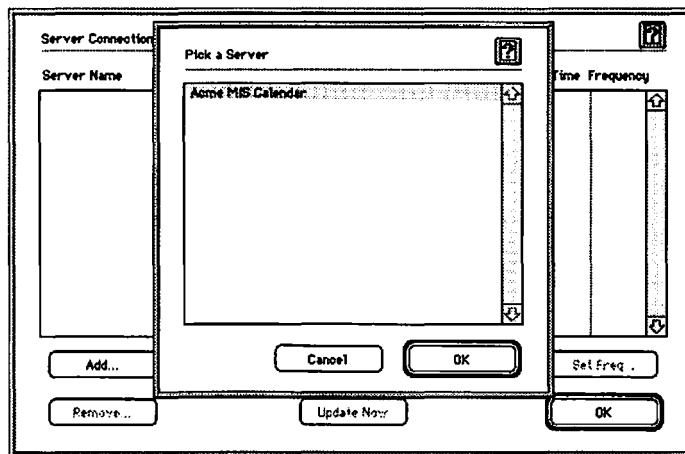


Note: Avoid all mistakes because the Public Event Server does not overwrite data added during the interrupted import process. Suppose you make ten mistakes, with the first mistake interrupting the process at line six. When you are finished, you will have ten listings of the accounts in lines one through five! You can delete these superfluous accounts, but you can only select one account at a time and you will be asked for the user's password each time! If you find yourself in this situation, I suggest you shut down the Public Event Server, trash both the "Public Event Server.db" and "Backup.db" files, and restart. You will need to reenter the

name for the Public Event Server, but after that you can perform a clean install of your user list.

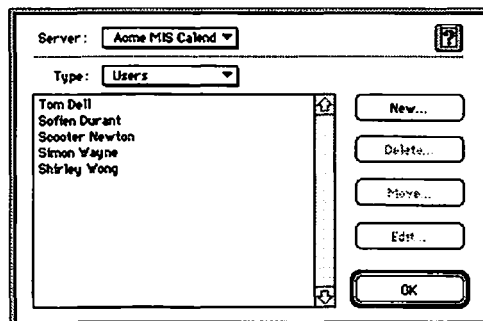
Once you have made any necessary corrections and have imported the data into the Public Event Server, verify the work by launching the Up-to-Date application.

First, tell the client which calendar server it should reference by selecting **Server Connections** from the **Define** menu. In the windows that are generated, select the Public Event Server you just set up.



Showing Up-to-Date the Location of the Server

Now select **Administration** from the **Meeting** menu.



Reviewing the List of Imported Users

In this window, you can examine your database of users, rooms, and resources. They can be modified or removed with the **Edit...** and **Delete...** buttons.

User Info

Name: Scooter Newton Password

Extension: 144 Login: ****

Department: MIS fourth floor

Description: Network Administrator

Meeting Avail... Set Proxies... Cancel OK

Modified: Main 8150 Export.TEXT, 12/15/95, 11:17 AM

Editing User Information

Finally, designate those hours when users are available for Up-to-Date to schedule meetings, and the people who will act on behalf of individual users when they are not available. This can be accomplished with the **Meeting Avail...** and **Set Proxies...** buttons available in the Edit dialog box.

Type: User Name: Scooter Newton

Select days and times available for meetings

Place checkmarks next to days available for meetings. Select a day to set times for that day. Lunch times apply to every day of the week.

Sunday
 Monday
 Tuesday
 Wednesday
 Thursday
 Friday
 Saturday

From: 8:00 AM To: 5:00 PM

Unavailable during lunch: From: 12:00 PM To: 1:00 PM


Reserve personal time

0 blocks of: 30 minutes

Cancel Save

Establishing a User's Availability

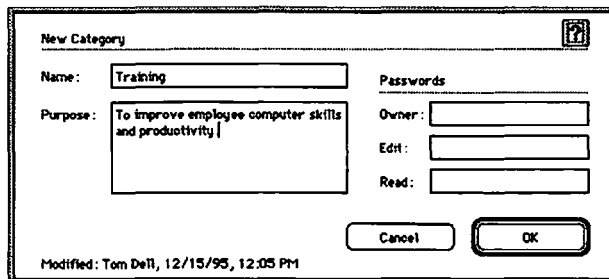
You can't set availability for rooms and resources. These are always available except when they are scheduled to be used in a meeting. You can set proxies for both, however. This is a good way to make someone responsible for a room or resource to ensure they are always ready for use, as in placing manuals in the training room or charging the batteries in the cable scanner.



Establishing a Room's Proxy

Installing Public Categories

In Up-to-Date, events are in either the *personal* or *public* category. Users can employ the personal category to keep track of events they do not need to share with the rest of the world, such as when their mortgage payments are due. Public categories are used to advertise events that affect entire workgroups, such as when a project deadline falls. You could let your users set up these public categories, but for the sake of consistency we suggest that you set them up. You are *required* by the software to set up at least one public category. To do this, first select **Administration** from the **Meeting** menu in the Up-to-Date application. In the window that appears, pick the name of your Public Event Server in the pop-up menu at the top and **Categories** in the pop-up menu beneath it. Click **New** to enter the first category.



Setting up a Public Category

As a general rule, name and describe the category in such a way that users can easily determine whether they should be interested in the events associated with it. You can create many categories. Users have the option of subscribing only to the categories that are relevant to them and ignoring the rest.

Like the Public Event Server database itself, you can assign password protection to public categories. The **Owner** password must be produced by anyone who wishes to archive, delete, move, rename, or otherwise alter a category. It is a good idea to set this. The **Edit** password is needed by those wanting to move, edit, or delete a category's events. You want to allow most users to do this so they can make corrections to calendar data. The **Read** password is needed simply to view the events associated with a category and is not necessary in most cases. In fact, you *should not* set this password if you wish to have meetings automatically scheduled under this category; it will hamper Up-to-Date's group scheduling abilities.

Configuring the Public Contact Server

Now you are ready to add the Contact data that will relate to the calendar.


Installing Public Categories

Contact is capable of handling *private* and *public* categories of address information. As the names imply, private categories are used for information that does not need to be shared with the rest of the company, such as the addresses of a user's family and friends. Public categories contain information that should be shared with the rest of the company or the workgroup. In our example, we create an interoffice list of employees and office-related service personnel, with emphasis on what a network administrator would need in such a database.

When you install the Public Contact Server, one public category is created automatically. Its name contains whatever you called your Public Contact Server with something like "Shared Contacts" tacked on. For instance, our example was created as "Acme Contacts Server Shared." Users also receive three automatically created personal categories: Business, General, and Personal.

The public category can be accessed from network workstations using the Contact application, as long as it has established a connection with the Public Contact Server. There is no need to set up a users database as you did with Up-to-Date.

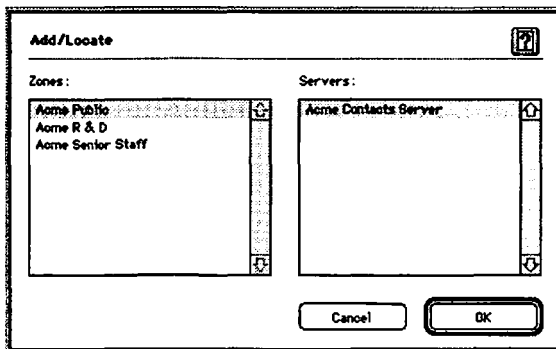
For our example, we are going to add the public category “Personnel.” To do this, launch the Contact application, and then select **New** from the **File** menu to create a new address book. Next, choose **Server Connections** from the **Define** menu.



Define Server Connections Window

You see an empty Define Server Connections window. Here, click **Add Server....**

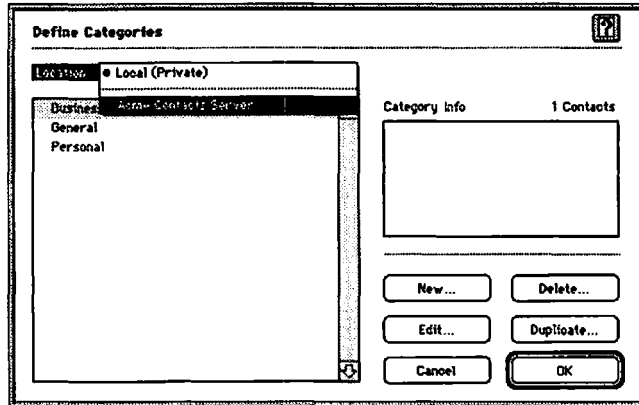
In the Add Server window, select the zone on the left in which the Public Contact Server is located, and then on the right, select the name of that server.



Selecting the Public Contact Server

Click **OK** when you are finished. You are returned to the Define Server Connections window. Click **Done**.

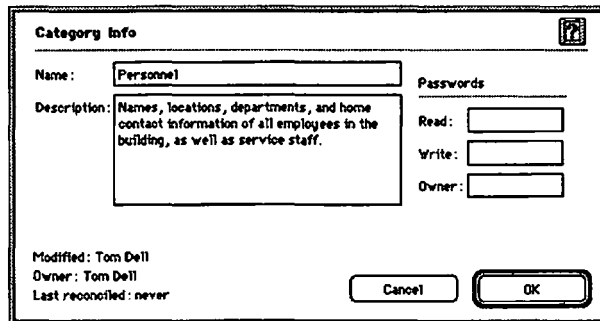
Now that you are logged in to the server, you are ready to create a category that can be accessed by all the connected users. Choose **Categories** from the **Define** menu, and switch from local to networked categories in the pop-up menu at the upper left of the Define Categories window.



Choosing Shared Categories

Click the **New...** button to create a category. Enter the password you established for the Public Contact Server to continue.

Name the category, describing it so that your users will know whether it's a category they want to use, and add passwords, if you wish.



Establishing a Shared Category

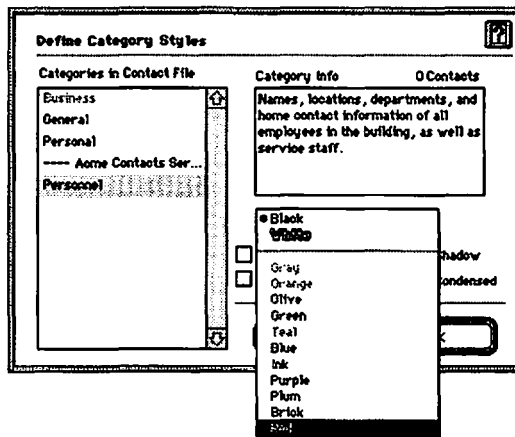
The **Read** field allows you to hide these contacts from all users except those privy to the password. The **Write** field allows users to read the contacts but not to make

changes to them unless they have the password. The **Owner** password prevents users from deleting the category unless they are authorized to do so.

For our example, we named the category “Personnel” and provided a fairly good description of its contents. We have chosen to put employees’ home contact information in here as well, which is never something you want to do until you have talked with Human Resources about company policy on the matter. They may have you add something like “not to be disseminated to non-employees” to your description.

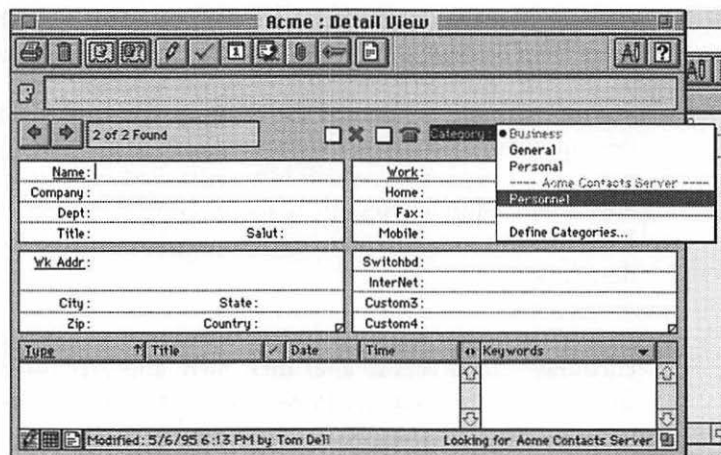
Finally, we do not feel a need to password protect anything here. We would rather encourage employees to alter their own employee information whenever something changes. Assigning a password for Owner is probably a good idea, however, to ensure that the category is not deleted by accident.

By the way, all categories can be colored-coded. If you want to make the contacts in this category stand out from other categories when users view their address books as a whole, you can set the text color by selecting **Category Styles** from the **Define** menu.



Color-Coding a Category

How you actually enter data into a database is well covered in Now’s manual and is simple enough that the loss of the manual wouldn’t hurt you much. To put it simply: Create a new contact record, choose the public category in the pop-up menu, and type in the contact information.



Adding a Contact to the Public Category

What most database programmers identify as a *record*, Now calls a *contact*, adding further that it can contain data for an individual or an entity. Whatever you call it, each record contains the data you would expect to find—name, telephone, and address—as well as a few fields that can be customized.

In our example, we have been creating a personnel address book, so we will use this customization feature. Since your company's employees are also network users, add some information to the database that can help you reconcile the information from backup, configuration management, AppleShare, and e-mail logs. These logs can tell you what your users are doing or are not able to do. Contacts can tell you where in the building they are located.

We talk more about the importance of being able to associate network resources with users in our book *Managing AppleTalk Networks*. This is essential in troubleshooting and accounting management.

	Long Label	Short	Field Type
Custom 1:	Asset Tag		Text
Custom 2:			Text
Custom 3:			Text
Custom 4:			Text
Custom 5:			Text
Custom 6:			Text
Custom 7:			Text
Custom 8:			Text

Buttons: Copy Fields From..., Cancel, OK

Creating Custom Fields

To use this feature, select **Custom Fields** from the **Define** menu. What is entered in the custom fields depends on how you have organized your network. Some of the basic fields you may wish to create are:

- **Network jack:** The numbers of the wall jacks in the users' offices, which can be handy if you have a nice big blueprint on your wall showing where all the wall jacks are and how they are labeled.
- **Asset code:** This could be the number that appears on the asset tags you have attached to each user's computer.
- **Physical location:** This could be something like an office number or a description, as in "4th floor, NW."
- **Responder number:** This could be the number you gave to the responder of a network management tool, such as SaberLAN. The trouble is that you can rely on the users to enter information for the previous categories, but you probably can't here.
- **Telephone extension:** This is applicable in most companies.
- **Some other thing:** You get the idea. The point of adding these fields is to give you a quick way to access the people associated with what you see happening on the network's wires. *How* exactly you do this is not as important as the fact that you *do* do it.

SETTING UP THE UP-TO-DATE AND CONTACT CLIENTS

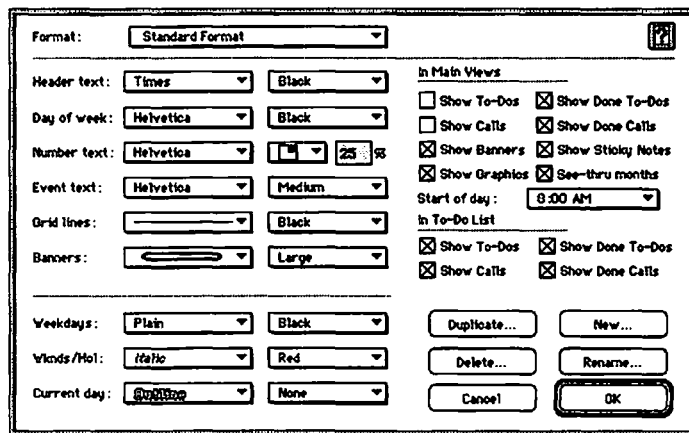
Here are a few things you can do to get your users up and running.

Installing Clients

An easy way to distribute the Up-to-Date and Contact applications to your users is from an AppleShare server volume. Create a folder with the appropriate privileges and name it Calendar/Contact Install. Then drag the files from the master diskette into the folder. Your users can install the client by double-clicking the installer on the mounted server disk. Be sure to tell your users not to install the Public Event Server or Public Contact Server. They do not need it and probably should not have it. You might even want to delete it from the server install volume.

Creating Standardized Up-to-Date and Contact Files

Although not strictly necessary, it is a good idea to start your users off on the right foot by distributing a preconfigured and standardized calendar, making use of the categories you created in setting up the Public Event Server. There are a great many things you can configure when doing this, such as category sets, formats, and layouts, but the explanation of these is better left to the Up-to-Date manual.

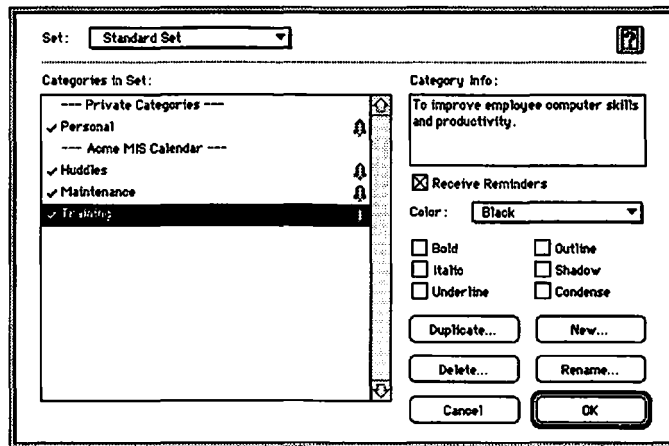


Some of the Many Options Available to Customize Up-to-Date's Look and Feel

We will go over the basics, again using the IS department as an example.

Create a fresh calendar using **New** from the **File** menu. Next, tell Up-to-Date which Public Event Server(s) to monitor using the **Server Connections** command, as we discussed earlier (p. 417).

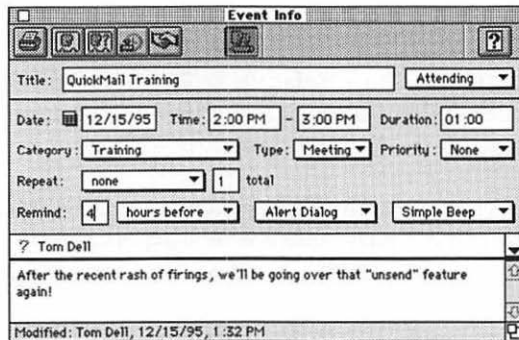
Select **Set** from the **Define** menu to create a *set* comprised of a group of categories. You can designate office-wide categories that you do not want any of your users to be without—office holidays, vacations, production deadlines, and the like.



Selecting Categories to Include in a Standard Set

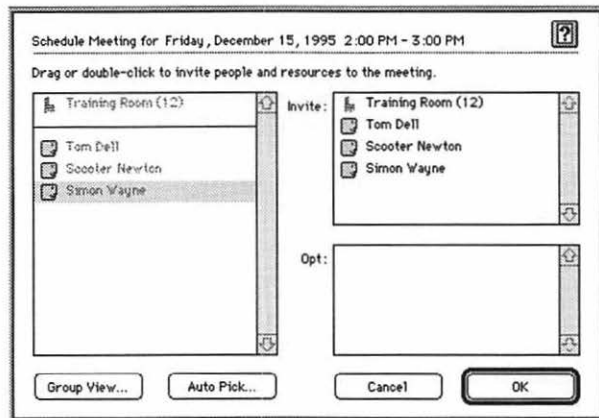
In our IS department example, we are concerned about three categories in particular. The first, “Huddles,” contains events that involve the department getting together as a group for planning purposes. The “Maintenance” category contains chores that network administrators must perform on a recurring basis. The “Training” category contains event information for both the IS department’s ongoing education of employees and the training of the IS personnel through seminars and conventions.

Users who are subsequently given this calendar file will be aware of these events, regardless of which employee scheduled them. For example, the network trainer might announce an afternoon training session in the window generated by selecting **New Meeting** from the **Event** menu.



Example of a Meeting

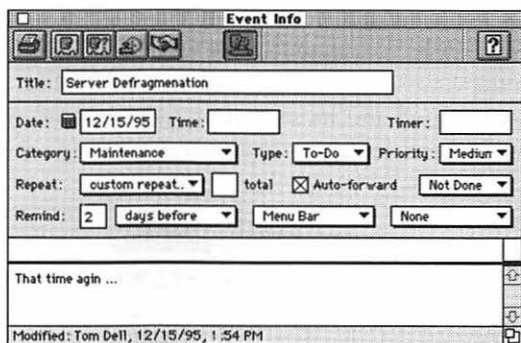
All the users who have "Training" as one of the categories in their event set will be reminded that there is a QuickMail training session at 2:00 P.M. on Friday. They will have the entry displayed on their desktop calendars and will receive an alert generated by the Reminder control panel at 10:00 A.M. These users could find themselves automatically scheduled for this event even though nobody "said" anything about it. When entering the new meeting notice, the trainer selected everyone he wanted to attend and the resources he wanted to use.



Choosing Meeting Attendees and Resources

Up-to-Date can then look over the designated users' schedules and choose a time when all can attend, or tell the trainer whether his choice will work. A user who knows something Up-to-Date does not know and cannot attend the meeting can simply *decline* it, at which time Up-to-Date can try again to find an acceptable time slot.

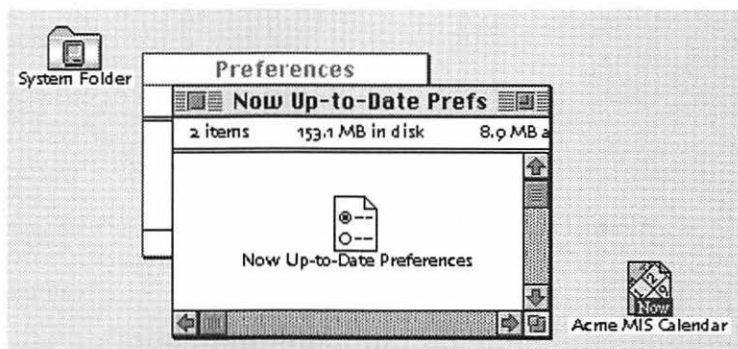
Finally, users of this calendar can be reminded of their recurring maintenance duties throughout the year even though the administrator only scheduled the event once. This is accomplished with the **Repeat** option in the Event Info dialog box. For example, by entering a **New To-Do** item from the **Events** menu, the network administrator might establish a notification to remind himself and others on a quarterly basis that the server's hard drive must be defragmented. In this way, users cannot claim that they didn't know the server would be down on that particular night, because it has been on the calendar for months!



Establishing a Recurring Maintenance Reminder

You no doubt get the idea. For more specifics on how to use Up-to-Date, see the manual. For more specifics on what should go into a calendar used by network administrators, see our book *Managing AppleTalk Networks*.

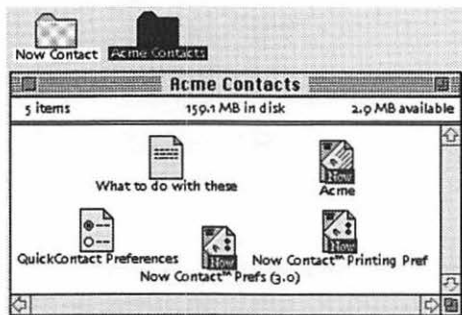
To distribute this standardized file to your users, simply save it under some appropriate name and send it to them, along with the "Now Up-to-Date Preferences" file from the Preferences folder of the System Folder.



Files Needed for Distributing the Up-to-Date Calendar

Tell your users to drag the “Now Up-to-Date Preferences” file into the Preferences folder of their System Folders, replacing the files already there, and then to double-click on the calendar file. When they are asked whether they wish to make it their primary calendar file, tell them to answer **Yes** and to select the **Don’t ask me again** checkbox. Now they will log into the standard calendar every time they launch Up-to-Date and will be informed immediately of the events you have decided are most important to them.

If you have created a Contact file with customized fields and layout options, as we did in our example, make sure that this also is distributed to your network users. In this way, everyone has the same information and works with it and records it in the same way. To do this, create a new folder and place in it the Contact file you have created and its related preferences files, which are stored in the Now Contact™ Prefs folder in the Preferences folder of the System Folder.



Contact Files to Distribute to Users

Call the newly created folder something like “Company Contacts” and place it on the server volume with the Contact installer folder, or e-mail it to your users. Create a “read me” file with SimpleText explaining how to install these files.

CHAPTER 18: SERVICE PERFORMANCE

This chapter was written by E. L. Heiberger with Tom Hessel

I have been teaching the Designing AppleTalk Networks course for over a year and I would be rich if I had a nickel for every time a student asked me, “What kind of a server should I buy if I want to run AppleShare?” The easy answer is that there is no easy answer. The better, longer answer is what this chapter is all about.

Before we dig in, let me tell you about the method to our madness. We began by selecting some of the most common services:

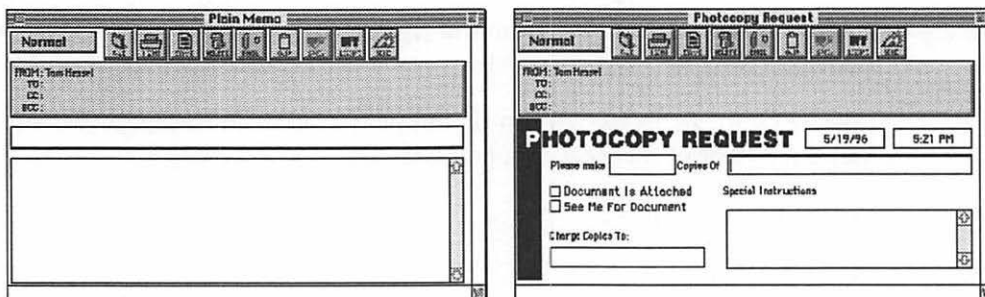
- E-mail servers
- File sharing
- Printing
- Databases
- Web servers
- Backup servers

Next, we ran a battery of tests pertaining to each service. For example, how much traffic does one person vs. many people generate for a particular service, how does the performance of these services compare to alternative types of services, and so

on. The results were analyzed using multiple tools: NetMinder Ethernet by Neon Software, Skyline Satellite and EtherPeek by AG Group, and LANTest by HELIOS Software. NetMinder was used during each test to capture the packets and their contents. This helps us determine how many bytes were transmitted over the network and how long each transaction took. We chose Skyline Satellite to examine network utilization and to determine overall packet sizes and bytes per second. We used LANTest to push sustained amounts of data to a machine in an attempt to retard its performance. Lastly, EtherPeek v.3 was used to create a ton-o-network traffic.

We ran our tests on a variety of machines. The server was always a PowerWave 604/150 by PowerComputing, and the controlled-client was a PowerComputing PowerWave 604/120. The PowerWave 604/150 is a boom-chaka-laka fast machine. Both machines are a good representation of what the market has to offer today for both clients and servers. An assortment of older machines were used to generate traffic during our testing.

We tried to keep our test parameters as constant as possible. For small files, the *Plain Memo* form from QuickMail was used. Larger files were represented by the *Photocopy Request* form, also from QuickMail and a 1.3-MB stuffed file was used.



QuickMail Forms Used in Our Test

Our results are presented using a combination of tabular format tables and more detailed discussions. We show the original size of the document in bytes before it was transmitted across the network, and then the total bytes it took to transmit the original document across a LAN and a 56K Frame Relay WAN. Both of these transactions are measured for speed and assessed by their role in overall network utilization. Finally, we show the amount of information overhead for each transaction and what percentage of that transaction was actual payload versus overhead information created by network transactions. Sometimes the results were exactly

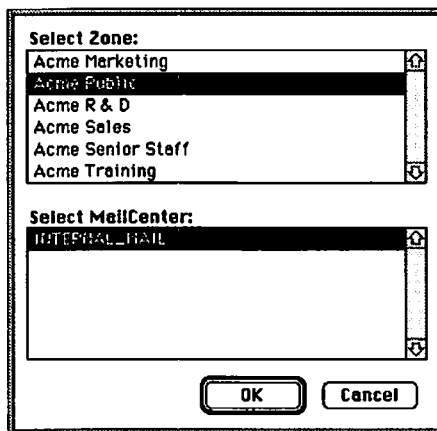
what we were expecting. At other times, we could not believe our eyes. Remember, the service is only as good as the programmer who created it. An inefficient service on a fast machine with tons of RAM and a speedy hard drive is still an inefficient service. An efficient service running on a butt-slow computer and operating system is still an efficient service. Throughout this chapter, we will be giving you tips and techniques so that your network clients can better utilize their networks and network services.

E-MAIL

E-mail services were not found to be high consumers of network bandwidth. Electronic mail is more well known for its potential to consume the processes of a server. QuickMail is designed in such a way as to allow no more than 32 users to simultaneously use its services. More important, we found that it never used more than 10% of the available Ethernet bandwidth for any one server, no matter what we did to the service.

Signing In

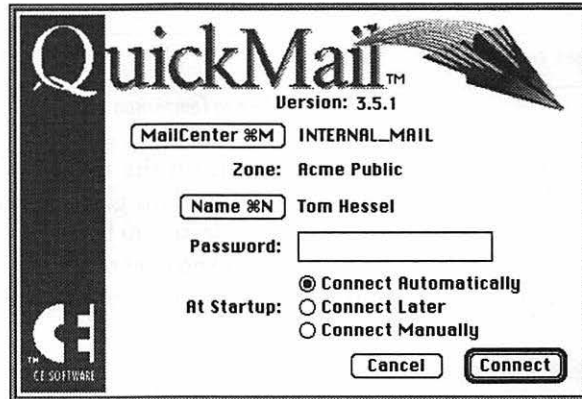
In a nutshell, when you start up the QuickMail client application a MailCenter must be selected. It is important to select your server and then get the heck out of there. As long as you have the Select MailCenter window open, NBP-LKup (lookup) packets will be transmitted onto your networks indefinitely, and all the available servers will answer back with NBPREps (replies) indefinitely.



QuickMail's MailCenter Browser Window

The NBP activity stops 30 seconds after a MailCenter server has been selected and the Select window has been closed. No matter what you do after you select a MailCenter, those NBP packets will still be generated for an additional 30 seconds. So, no fence-sitting folks—pick that server ASAP.

If an e-mail server has already been selected in the login window, the operation is just like using an alias to log into an AppleShare server. The client remembers the AppleTalk address and sends a packet out to see if the server is still there. If the server is found, that's the end of the lookup traffic.



Signing into QuickMail

If the server is not found, you must start the process all over again by searching for a MailCenter.

Log Bog

A well-known author and a good friend (who will go unnamed) and I have a pun routine we do together. Tag, you're it.

Usually, e-mail administrators ask their users to clean out their e-mail logs so that hard drive space on the e-mail server is cleared out for new e-mail. There is another good reason to clear out your log: to reduce the amount of network traffic. For every e-mail message in your log, additional bytes of information are sent across the network from the server to the client every time you log into your mail server. Here is what we found:

Signing In	Bytes over LAN	Overhead per Message	Bytes over WAN	Overhead per Message
No messages waiting	5,687	NA	6,085	NA
One message waiting	5,926	239	6,623	538
18 messages waiting	9,133	191	11,005	4,920

Sign-in Comparison



The amount of data depends on the nature of the e-mail. We used the Plain Memo form with no title at all. This gave us a generic baseline form for testing. According to our test data, at least 200 bytes were sent for every additional e-mail message in the user's log. If you go over the WAN, that amount of data nearly doubles. So, our hot tip is to clear your logs on a regular basis.

General Performance

The biggest news flash was the maximum amount of network bandwidth that QuickMail took up: 10%. No matter how many users, how big the e-mail messages were, or what kind of processor we were using, only 10% of the available Ethernet bandwidth was ever occupied by the traffic of our e-mail server. It was sometimes less, but never more.

As with any network service, there is no such thing as 100% efficiency, meaning that if you are sending e-mail that is 102 bytes total for the message and the form, the actual information transferred over the network is never 102 bytes. There are always some overhead packets involved. How many overhead packets is part of what determines network efficiency.

What Was Sent	Original Bytes	Bytes over LAN	LAN Payload	LAN Time
Plain Memo	102	1,806	6%	0:00:03
Plain Memo with Text	1,930	3,450	56%	0:00:02
Plain Memo with Enclosure	18,483	22,455	82%	0:00:03
Photocopy Request	5,604	8,142	69%	0:00:02
Photocopy Request with Text	7,432	9,691	77%	0:00:02
Photocopy Request with Enclosure	23,985	28,960	83%	0:00:03

Efficiency of QuickMail over a LAN

There were several notable occurrences in our testing of QuickMail:

- The bigger the file, the more bytes are transmitted over the network. Regardless of form size, however, the overhead packets remain the same.
- Sending an enclosure is less efficient than sending the same information within the context of the e-mail message. Ironically, e-mail etiquette suggests that if your e-mail is large you should make it an enclosure. One reason for this is that most e-mail packages don't make good text editors or browsers. However, good network practice dictates the opposite. The nature of most Internet mail servers and gateways is that if incoming mail is quite large, it is made into an enclosure automatically by the gateway or server. Either way, if you compress your large documents, this creates less traffic overall. Thus, as a suggestion, if your e-mail text grows to be over 20K, make it an enclosure. If what you are enclosing is extremely large, compress it first.
- Overall, QuickMail can achieve up to 90% efficiency when transmitting data across the network. Sending smaller e-mail messages is less efficient than sending larger e-mail, as the overhead tends to overcome the actual data in smaller e-mail. There will always be a base number of overhead packets, but this is something to be expected with any network service, as we will see in the next analysis.

WAN Performance

When sending files that were under 10K, there was no substantial performance difference between sending over an Ethernet network or over a 56K Frame Relay WAN link.

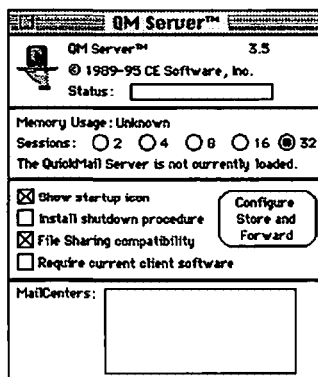
As the files became larger, the transmission time more than doubled. But considering that we were sending files over 1 MB in size, e-mail provides good service over the WAN as long as the files are not too large. Since bigger files take longer, files that would normally take 20 minutes to transmit on your LAN can take two and three times that on your WAN. It better be a pretty important e-mail message to tie up your WAN connection for an hour. While QuickMail can use only 10% of your Ethernet pipe, it can use up almost the entire pipe available on your WAN.

What Was Sent	Original Bytes	Bytes over WAN	WAN Payload	WAN Time
Plain Memo	102	1,806	6%	0:00:03
Plain Memo with Text	1,930	3,725	52%	0:00:03
Plain Memo with Enclosure	18,483	22,341	83%	0:00:03
Photocopy Request	5,604	8,053	70%	0:00:03
Photocopy Request with Text	7,432	9,690	77%	0:00:03
Photocopy Request with Enclosure	23,985	29,073	82%	0:00:02

Efficiency of QuickMail over a WAN

Fighting for the Bandwidth

QuickMail can be set for 4, 8, 16, or 32 concurrent users. If you set it for 4 users, no more than 4 users can be sending e-mail through the server at one time, and so on for 8, 16 and 32 users. When you install QuickMail Server, the default setting is 8 users. You should change this accordingly.



Setting up QM Server

As shown above, QuickMail's maximum setting for concurrent usage is 32 users. If you have less than 32 people within your organization using QuickMail's services, it does nothing for you to set this to less than 32 even though you'll never have 32 people requesting service at the same time. Remember, 32 is the maximum setting.

Finally, if there is 40% utilization on the wire, a user would have no problem sending e-mail as long as the QuickMail server is not tied up. So, in short, if you have a robust server, the service itself is most often the bottleneck. In every instance, as soon as we hit 10% utilization for QuickMail, the service stopped serving. It would not take on another process until some of that 10% available bandwidth could be used for new processes.

QuickMail Usage	Total Bytes	LAN Time	LAN Payload
1 User	1,508,396	0.00113426	90%
2 Users	1,507,959	0.00111111	90%
9 Users (8 Sessions Selected)	1,568,823	0.00223379	90%
9 Users (32 Sessions Selected)	1,512,879	0.00155093	90%
1 User (40% Net Utilization)	1,501,219	0.00112269	90%

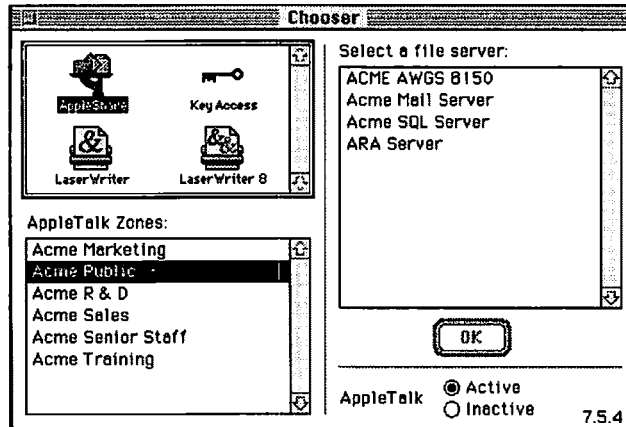
Sharing the Ethernet Bandwidth

FILE SHARING

We have dedicated an entire book to file sharing, *Managing AppleShare and Workgroup Servers*, but aw gee, let's do it one more time just for fun. There are two types of file sharing in the land of the Macintosh: System 7's Personal File Sharing and AppleShare file serving. Is there a difference? You betcha!

The Chooser

The Chooser is an application or network search engine that allows us to search for and sign into existing services, such as services for printing, faxing, and file sharing. The Chooser, not the icons we select, is the application that controls the nature of the network traffic we generate when searching for and logging into these services.



The Chooser

When you open the Chooser and select a service, NBP broadcast packets are sent out to the network, seeking out those computers publishing the selected service. For every computer or printer that is running the requested service, a response is returned back to the user's Chooser. This goes on nonstop for 35 seconds and then begins to die off. The Chooser begins to pace its NBP broadcasts by increments of 6×2^z where $z = 0, 1, 2$, and so on. However, the client never stops asking. If

you don't already know this, every computer has to listen to an AppleTalk broadcast and process the request. The process could be a response if the computer is running the service being sought, or just the process of disregarding the packet because the device is not running that service. The process time involved in this type of operation is said to degrade the performance of any shared service on a network. Well, I am here to tell you that we put 7% of the traffic on our network of type NBP without even a hiccup. With six computers opening their Choosers and producing NBP-LKups, we created a whopping 0.0% network traffic.

✦ Our hot tip here is not to allow your network administrator, or anyone else for that matter, to use an application like AG Group's EtherPeek to create 7% NBP traffic on your networks.

AppleShare vs. Personal File Sharing

This is where it gets hot and heavy. There are several major differences between AppleShare and Personal File Sharing. All of these differences contribute to the efficiency of AppleShare or the lack of efficiency of Personal File Sharing. "What?" you say, "Personal File Sharing is not as efficient?" Well, bought any Personal File Sharing file servers lately?

Packin' the Packets

AppleShare and Personal File Sharing use two different algorithms. Personal File Sharing was designed to be a one-to-one or peer-to-peer file transfer system. The assumption is that users are sharing files among themselves and that no one person is acting as the file server.

Sending Files with Personal File Sharing	Original Bytes	Bytes over LAN	% LAN Utilization	LAN Payload
Plain Memo (form, no users)	102	5,231	2%	2%
Photocopy Request (form, no users)	4,914	10,818	2%	45%
1.3-MB File (no users)	1,357,596	1,569,854	20%	86%
1.3-MB File (1 user LANTest)	1,357,596	1,569,071	25%	87%
1.3-MB File (2 users LANTest)	1,357,596	1,569,476	30%	86%

Performance of Personal File Sharing

Personal File Sharing is very efficient when one user is transferring a file to another. As a matter of fact, it is just about as efficient as AppleShare when it comes to overhead packets. According to our tests, both types of file sharing were able to deliver the same payload percentages.

Sending Files with AppleShare	Original Bytes	Bytes over LAN	% LAN Utilization	LAN Payload
Plain Memo (form, no users)	102	5,231	2%	2%
Photocopy Request (form, no users)	4,914	11,068	2%	44%
1.3-MB File (no users)	1,357,596	1,568,761	18%	87%
1.3-MB File (1 user LANTest)	1,357,596	1,569,256	20%	87%
1.3-MB File (2 users LANTest)	1,357,596	1,568,540	25%	87%

Performance of AppleShare

Transmission Time

This is where we begin to see the differences in efficiency between the two types of file sharing. AppleShare is able to deliver data more quickly to the client than Personal File Sharing. Personal File Sharing becomes inefficient when multiple users are transferring files at the same time.

Sending Files with Personal File Sharing	Original Bytes	Bytes over LAN	LAN Time
Plain Memo (form, no users)	102	5,231	0:00:01
Photocopy Request (form, no users)	4,914	10,818	0:00:01
1.3-MB File (no users)	1,357,596	1,569,854	0:00:05
1.3-MB File (1 user LANTest)	1,357,596	1,569,071	0:00:08
1.3-MB File (2 users LANTest)	1,357,596	1,569,476	0:00:10

Transmission Times for Personal File Sharing

Sending Files with AppleShare	Original Bytes	Bytes over LAN	LAN Time
Plain Memo (form, no users)	102	5,231	0:00:00
Photocopy Request (form, no users)	4,914	11,068	0:00:00
1.3-MB File (no users)	1,357,596	1,568,761	0:00:04
1.3-MB File (1 user LANTest)	1,357,596	1,569,256	0:00:05
1.3-MB File (2 users LANTest)	1,357,596	1,568,540	0:00:07

Transmission Times for AppleShare

Although the actual network overhead is no different, the rate at which AppleShare can deliver data to users on the network is faster than that of Personal File Sharing. That's just the way Personal File Sharing was built; it was never meant to be able to serve multiple clients quickly and efficiently.

The Benefits of Caching

AppleShare 4 and later allows for you to set aside an amount of the computer's total RAM for temporarily storing frequently accessed items such as files, folders, and icons. It does this by using a technology called read-ahead and write-behind caching. If you would like additional information on how this type of caching is used, please see our book *Managing AppleShare and Workgroup Servers*. If you are already a wiz at this caching stuff, let me give you just a couple of additional hints to think about when trying to configure your cache settings within AppleShare 4.0.1–4.1. AppleShare 4.2 and 4.2.1 use dynamic caching algorithms that automatically adjust cache settings for the types of files being used on a per microsecond basis.

- Allow one buffer or one file for each active user on your servers. Try to allocate RAM for the file cache based on this formula:

$$\text{Active clients} \times \text{buffer size} - 640\text{K} = \text{additional RAM needed}$$

- Allocate your folder cache based on this formula:

$$\text{Shared files and folders} \times 0.2\text{K} = \text{folder cache}$$

- Decrease your folder cache size by a percentage of the total active users on your file servers at any given time. If 50% of your users are very active, then try this formula:

$$\text{Shared files and folders} \times 0.2K \times \text{usage percentage} = \text{folder cache}$$

- When trying to determine the amount of cache to set aside for icons not recognized by the system software, try setting the icon cache to 25% of the folder cache size. Try this formula:

$$\text{Your folder cache size} \times 0.25 = \text{icon cache}$$

The best rule I can give you to follow is that the more you can afford to cache, the better you will be, especially when using the dynamic caching capabilities within AppleShare 4.2.1.

File Sharing over the WAN

The good news is that very little overhead packets were added when transmitting files over the WAN. This is true for both types of file sharing. The bad news is what we already knew—going over the WAN is much, much slower. If you are going over a 56K WAN link like we were in our tests, transmission times are 50 times slower. I will explain the support of a 56K WAN link later on in this chapter.

Sending Files with Personal File Sharing	Original Bytes	Bytes over WAN	WAN Payload	WAN Times
Plain Memo (form, no users)	102	5,964	2%	0:00:01
Photocopy Request (form, no users)	4,914	10,389	47%	0:00:01
1.3-MB File (no users)	1,357,596	1,573,793	86%	0:00:05

WAN Transmission Speeds with Personal File Sharing

Sending Files with AppleShare	Original Bytes	Bytes over WAN	WAN Payload	WAN Times
Plain Memo (form, no users)	102	7,894	1%	0:00:02
Photocopy Request (form, no users)	4,914	12,762	39%	0:00:03
1.3-MB File (no users)	1,357,596	1,574,323	86%	0:00:04

WAN Transmission Speeds with AppleShare

Here's a riddle: What is faster over a 56K WAN link—sending a 1.5-MB file as a QuickMail enclosure or reading the similar size file from an AppleShare file server?

QuickMail Usage	Bytes over LAN	LAN Time	LAN Payload	Bytes over WAN	WAN Time	WAN Payload
1 User	1,508,396	0:01:38	90%	1,507,767	0:03:23	90%
2 Users	1,507,959	0:01:36	90%	1,507,447	0:03:22	90%

WAN Performance of QuickMail

There's your answer: File sharing is at least 35% more efficient.



Our hot tip here is to grab your files with file sharing if you can and avoid sending large enclosed files through QuickMail.

PRINTING

We spend a lot of time talking about printing in all our books. In case you missed it the first, second, and third time, *the newer LaserWriter drivers are more efficient than the older drivers*. The old drivers generate more overhead packets.

In our tests, we used a PDF file and the LaserWriter Driver 8.3.3. A PDF file is like a PostScript file in that all the information about that file—like fonts and formatting—is contained within the file. The document is not dependent on any outside sources, such as the preferences set in the application, for its performance. The speed and efficiency of the print job is based solely upon the print driver and your selected printer.

The Hewlett-Packard 5si proved to be a much faster printer than the Apple LaserWriter Pro 630. It also took up more of our network's overall utilization to do the same print job as the LaserWriter Pro 630. Of course it takes more bandwidth. It must. If the same amount of information arrives somewhere in less than a third of the time, then more information is on the wire during transmission.

Printer	Original Bytes	Bytes over LAN	LAN Time	% LAN Utilization	LAN Payload
Laser Writer Pro 630	1,210,816	1,780,037	0:14:49	0.16%	68%
Hewlett Packard 5si MX	1,210,816	1,720,248	0:03:53	0.62%	70%

Printer Throughput Times (Printing Speed)

We would expect the HP printer to be faster since it costs more, right? There's one thing to remember here, bean-counters, the HP printer costs about \$3,000 for the base model, while an Apple LaserWriter 16/600 costs approximately \$2,300. So, if you are doing a lot of printing and the users' time is worth money, your goal should be to buy a faster printer.

Here's another riddle: How many users could print a 1.2-MB document using the Hewlett Packard printer at the same time? Give up? Just one. Only one person can send a print job at a time. This holds true when printing to any printer, unless the

printer has a built-in spooler or you're printing to a virtual print queue. The print job just is processed faster than if users were printing to a slower Laser printer.

Here's another: How much network traffic would 100 users printing to the Apple LaserWriter Pro 630 create? Give up? The answer is about 0.16%. The reasoning is the same as before—the printer controls any and all print jobs. Since the printer can only accept one job at a time and one user produces 0.16% network traffic, we would see a sustained rate of 0.16% until all 100 jobs were printed. Of course, there may be slightly more traffic since the printer has to keep sending occasional packets back to those users waiting to print, but this overhead is minimal.

Just like all the other services we have seen, there are a certain number of overhead packets in any network transaction. The payload efficiency maxes out at about 80% for printing. There is little change to payload efficiency for printing over a LAN versus a WAN, but the time it takes to print over the WAN versus the LAN is monumental.

Printer	Original Bytes	Bytes over LAN	LAN Payload	LAN Time	Bytes over WAN	WAN Payload	WAN Time
Laser Writer Pro 630	1,210,816	1,780,037	68%	0:14:49	1,789,613	68%	0:15:30

WAN Throughput Times (Printing Speed)



Our really big, hot tip here is that even though the traffic may be minimal when printing to a printer on a LAN, depending on your WAN link you may decrease the efficiency of your WAN when printing. Try not to print over a WAN.

WEB SERVERS

A Web server is a static means of distributing information, whether it is locally across a LAN or remotely across a WAN or the Internet. That information can be browsed, downloaded, or designed for dynamic user interaction. We are going to examine the three types of traffic and utilization that stem from Internet activity: from your local users who are looking at their organization's Web services internally, from local users going out to surf the Net, and from users coming into your page to browse around. The applications selected for testing were Netscape v2.0.1 on the client and WebSTAR v1.2.4 on our server running on a PowerMac 6150/66 with 40 MB of RAM.

Hitting Our Home Page Locally

We started our testing by letting one user who has been our controlled client, the PPC 604/120 by PowerComputing, hit our Web page for the first time, with nothing cached. Next, we tried the same test with eight simultaneous users.

Local Hits to WWW Page	LAN Bytes	LAN Time	% LAN Traffic
1 User	46,923	0:00:06	0.02%
8 Users	376,872	0:00:28	0.90%

LAN Throughput

It took six seconds for one person to log into the server locally. The network utilization was only 0.2% overall for the time of this transaction. It peaked to a whole 3% for a split second then dropped back off again. Web servers are built to take multiple hits simultaneously, as most servers are designed. The only bottleneck is the speed at which our Web server can push data to the users requesting its services. This is because we are accessing this server over our LAN.

It took 28 seconds for eight people to log into the server. The network utilization was 0.89%. The total number of bytes went from 46,923 to 376,872 or around eight times as much. But the utilization was only four times as much with eight times as many users. As with most servers, when a server is dedicating its time to

just one user, it has the ability to send or receive data at x rate producing $x\%$ network utilization. When multiple people request data from a server, network utilization is never that number of users times $x\%$ network utilization. Because the Macintosh platform is a single-processing, multi-threading platform, when multiple people request data from a server, the server no longer dedicates its time to just one user. It has to split its time and give equal time to everyone needing its services. When giving equal time, you decrease the per-person performance by the total number of users requesting services at that time. When working on a LAN, the amount of network utilization used by this is solely determined by the type of hardware being used and the efficiency of the software in processing multiple requests. Let's add another variable here: logging in from the outside.

Coming in over the WAN

Can people from the outside world ever slow down your network by hitting your Web server? This is not a trick question. The answer is absolutely not, unless of course you have a T3 WAN link and the rest of your users are on LocalTalk. Hell, if it's a T3 service, your net could be Ethernet for that matter and you might have problems. But that would never happen . . . please tell me that would never happen . . . What is more likely to happen is that users from the outside world could decrease the performance of the *machine* they are accessing. Our bottleneck for this testing is our 56K WAN link. It has the ability to transmit data at 56 kilobytes per second, and with that ability users can only produce $x\%$ network traffic no matter what they are doing over the WAN link. Wait a second, let's take a moment and determine just what our WAN link can support. A 56K line is 56 kilobits per second, with each character consisting of 8 data bits and no start/stop bits, because 56K links are almost always *synchronous* (meaning just 8 data bits, no start/stop bits, OK?). Also, to the telephone company, 56K means 56,000 bits/sec. If we did the math, this would give you 7000 bytes/sec maximum throughput. Adding in all of the overhead packets, approximately 40 bytes of TCP/IP headers and around 5 bytes for framing headers, I guess it is safe to say that a 56K link will handle a maximum of approximately 5.6K characters per second. With this as your bottleneck, it is almost impossible to create enough network traffic on your LANs to become at all concerned.

What happens when our eight people begin to surf? The chart below shows eight people hitting Apple's home page (www.apple.com).

Going out to the Internet	Total Bytes	WAN Time	% WAN Traffic
1 User	89,389	0:00:23	0.35%
8 Users	661,176	0:01:47	0.60%

WAN Throughput

While eight times as many users moved approximately eight times as much data, it took them a little less than eight times the number of seconds and only drove up the network's utilization from 0.65% to 1%. No surprises here: Our WAN link was the bottleneck that could only produce a certain amount of network traffic. It just took longer for all eight users to get their data. If we were rich, famous, and had T1 service, we would have been able to pull down the same amount of data in far less time, but our network utilization would have gone up. Also, if we had a T1 and, for discussion's sake, Apple had only a 56K link, guess what the throughput would be, as well as our overall network utilization? You got it—about what you're seeing above.

A Bonnet Full of Bees

I stole this phrase from *Second City*, but it sure sums up the nature of Web page activity. It is next to impossible to nail down exactly what may be slowing down the data transfer of your Web page. Here is a list of just a few potential problems and the impacts that they have on data transfer rates or, better put, the efficiency of your Web server:

- While CGIs are great and make the transition of your static WWW pages to dynamic user-interactive pages, their overall use has a substantial impact on the performance of your servers. The reason for this is that most CGIs are not built to handle the request for data either by lots of users at the same time or in an asynchronous fashion. Web pages may freeze on users who were timed out of the CGI's processes, and the result becomes a reload of the page.
- This one is simple—the bigger the graphic, the longer it will take to transmit that image from server to client. Try to keep your images small, simple, and 72 dpi. Remember that you are viewing images that are on the screen, not in print. Your screens support 72 dpi and images still look great.

- Also, please remember this: More RAM doesn't mean that computers will operate faster! More RAM means a computer can run more applications simultaneously. However, doing this may decrease the performance of the key services that machine is providing. If you need to offload additional services on different machines, please do so.

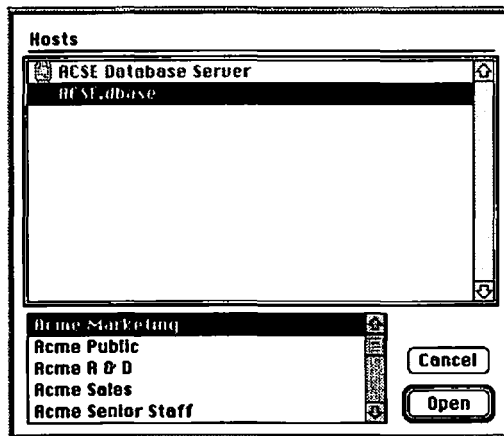
For additional helpful hints on the development of your Web pages or setup of your Web servers, visit our chapter earlier in this book entitled "Serving the World Wide Web with WebSTAR," starting on page 303.

DATABASES

As we discussed in our chapter called “Database Servers and CGIs” (page 185), there are two different types of databases: flat file and relational file. Within this realm there are two types of database serving methods: peer-to-peer hosting and true client/server. We will discuss the performance of FileMaker Pro when using its peer-to-peer hosting technology. FileMaker Pro Server 3 was unavailable at the time of this testing.

Signing In

As with most other services on your network, the first step is always signing in to those services. We discussed earlier the effect NBP has on your network when signing in to different types of services. Claris seems to have taken the correct path here in the development of their software with respect to the effects of NBP broadcast packets. When opening a database file across the network using the AppleTalk protocol, a user must select Open from the File menu, and then click on the Hosts button. From there a browser-type window is displayed.



FileMaker's Host Browsing Window

The FileMaker browser when using AppleTalk sends just five NBP lookup packets onto the network and then displays the results. There is no indefinite lookup or

additional lookups occurring after the browser is closed—just five. It gets even better: If you are hosting your databases using the TCP/IP protocol, this browser sends just two, count them *two*, IP broadcast packets looking for database hosts. If none are found, the user has the ability to type in a direct IP address, if it's known. The test we performed here was the difference in signing in to a hosted database file using AppleTalk versus IP. Here's what we found:

Operation	Original Bytes	AppleTalk over LAN	AppleTalk Time	IP over LAN	IP Time
Flat File	220,160	45,218	0:00:07	39,620	0:00:05
Relational File	319,488	54,854	0:00:08	47,709	0:00:06

Tests on Databases

When opening the file using IP, we see that it took slightly less time than it did with AppleTalk. It also took less bytes over the network to open the data file using IP. The reason it took less time and less traffic for IP over AppleTalk is due to the difference in packet sizes and robustness of the protocol. IP has a packet size that is almost three times that of AppleTalk with less handshaking at the beginning. This allows us to get more data, quicker.

The Big Mistake

Another test we performed was to show the extra overhead involved when a user makes the mistake of opening a FileMaker database file that has been shared across the network via AppleShare rather than opening the file using the `Host` command. This type of mistake is explained in more depth starting on page 191 in the aforementioned chapter, "Database Servers and CGIs."

Operation	Original Bytes	Shared via AppleShare over LAN	Using the Open Host Command
Opening Flat File	220,160	92,324	45,218

Traffic from Hosting a Database File

This extremely costly mistake has produced 47,106 additional bytes on our network. The AppleShare file server has to send the data across the network using the AFP protocol to the users who double-clicked the data file. This is much slower and holds smaller packet sizes than the ADSP protocol, which is what you are using when the data file is opened correctly through the `Host` command.

As is true for most database-type servers, we were unable to create a ton of network traffic with FileMaker Pro's hosting technology. Upon the client's request, databases send data to remote users one page at a time, much like a WWW server, or in database terminology, one record at time. When users add data to a database remotely, the data is sent to the server every time a user clicks a single field within that record. You can say that data is sent to the server in short, bursty spurts. Even a pasted image of 100K would take a split second to send to the server. Your concern here should be whether the software you have or the processor on which the software is running will efficiently handle incoming and outgoing requests to multiple users. The network traffic generated by this type of use is very minimal.

BACKUP SERVERS

A few things you should know about the efficiency of backup services:

- Backup applications are very efficient.
- Retrospect is a foreground application that needs all the attention of the CPU. It was designed to monopolize the processor and to run at night when nothing else is happening. Thus, it is a bad idea to have any other service running on your backup server that someone may be using at night, such as ARA or fax serving.
- Retrospect can reach up to 65% Ethernet utilization. This is why you don't want to run your backups over LocalTalk or during normal production hours.

There are several factors that determine the efficiency of your backup server:

- The hard drive speed, bus speed, and processor speed of the remote as well as the file size and the file fragmentation.
- The bus speed and processor speed of the backup server.
- The type of the media you are using—DAT, CD, hard disk, floppy, or other removable media.

If you want to know more about backup, there is a whole book, *The Complete Guide to Mac Backup Management*. Guess who wrote it? We did.

ALL TOGETHER NOW

There is a prevalent theme in all this discussion regarding performance of services: the beauty of Ethernet. Ethernet is velvet when it comes to handling all the traffic running across the wire. Unless you drive your network utilization up over 80%, Ethernet just keeps chugging along, with little and sometimes no error. At this level of utilization, collisions barely occur. This type of result was determined using Neon Software's NetMinder Ethernet software to capture packets and look at overall percent network utilization. If collisions begin occurring, Ethernet's Carrier Sense Multiple Access/Collision Detection (CSMA/CD) access control scheme begins to back the senders off, and we would see up and down traffic patterns. This is the nature of Ethernet. As collisions begin to occur, everyone backs off until it is safe to send more data. As long as the utilization remains below 86%, anything and anyone is treated equally for access on the wire. Fortunately, the nature of Ethernet makes understanding service utilization a little bit easier.

Service bandwidth utilization is additive. If the QuickMail Server is using up 10% of the bandwidth (the maximum bandwidth QuickMail can take) and the AppleShare File Server is using up 40% of the available bandwidth, your users are using up 50% of the overall capabilities of Ethernet. This means that if your services are eating up to 80% of your overall network traffic, it's definitely time to look into devices such as switches and routers. It's that simple. For additional information about designing networks with switches and routers, please see the companion book to this one, *Designing AppleTalk Network Architectures*.

Eenie, Meenie, Chili, Beenie

Being a bean-counter type, I always wonder what would be the cheapest server I could afford and still get the job done. In all the testing that we did, we were never able to bury an Ethernet network with any single service. It always took a combination of services with a multitude of people using those services. In the end, the bottleneck was always the server, the service itself, or our WAN link. So the question is, how do you match up your servers with your services?

First, if your network utilization has flattened off, either the service or the server has maxed out. If you can run an additional service over the network from the same machine without any trouble, the service being evaluated is the bottleneck,

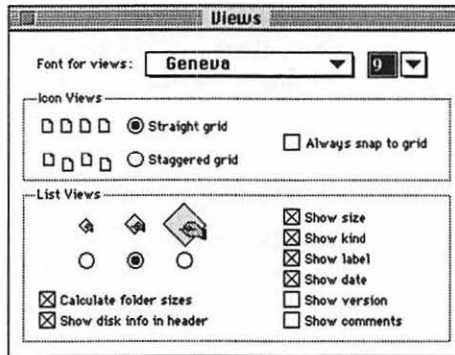
not the server. Therefore, a new machine isn't going to do you any good. You may need an upgrade or there is a configuration problem with the service itself. In either case, Fast Ethernet is not going to do you any good. Read this paragraph five times.

By running these tests, you can logically decide if it is time to buy a machine with a faster processor, better bus speed, or faster internal hard drive . . . but hey wait, I've never seen a hard drive rating. That's because there isn't such a thing. Life's not fair, folks. All you can do is test the machines yourself and make sure you are not wasting resources. Experience and asking the right questions become the keys to planning for the placement of your services. As you grow your network and become more familiar with your users, you will decide how to best allocate your services.

A Final Note

Pretty scary, huh?

Did we cover every service you ever wanted to know about? If not, you can set up your own Ziff-Davis labs in the comfort of your very own office. The first step is to choose five or more machines that are representative of your network. The drives on each computer will need to be reformatted and a clean version of the system software should be installed. As we have stated time and time again, fragmented hard drives, running applications that have extensions and control panels that use network services, and even more subtle vices such as having **Calculate folder sizes** turned on contribute to retarded performance both locally and when accessing remote volumes.



Views Control Panel with "Calculate Folder Sizes" Selected

Thus, you need to use clean machines. This may even be a good idea for your real servers . . . hint, hint.

Now, it's time to designate your machines. The best of the litter should be your test server. Otherwise, all you will achieve is slowing down your users. If your tests are for the performance of a slower machine, then have at it. Remember when using an accelerated service on a slower machine that the service is no longer the bottleneck. The processor or the SCSI bus will be the bottleneck. Run the analysis software on the two next-best machines. If you use a dumpy machine to collect data, you are apt to collect and gather improper data. The final step is to install the server software and the client software for the software you are testing on the appropriate machines. Voilà!

While these tests are not for the fainthearted, you and the kids can have some real family fun. Who knows? Maybe some publication will even pay you for your data.

INDEX

- * (wildcard character), 163
- <> HTML characters, 305

A

Access menu (ButlerTools)

- New Group command, 220
- New User command, 220
- Show Access command, 219

access privileges (Butler SQL), 216–222

- adding users for, 220
- changing, 221
- group, 222
- group creation, 220
- privileges enabled for, 221
- See also* Butler SQL

Access Privileges window (Butler SQL)

- illustrated, 220
- opening, 219

Account Information window (Apple Internet Mail Server), 159–161

- Add button, 160
- Forwarding pop-up menu, 161–162
 - illustrated, 161
 - options, 162
- illustrated, 160, 161
- Master Privileges checkbox, 161
- Password field, 160
- Require APOP field, 160–161, 165
- Size limit field, 161
- See also* Apple Internet Mail Server

ACGIs

- AppleSearch, 276–279
- CGIs vs., 205–207
- at database level, 207
- defined, 201
- FileMaker.acgi, 206

NetPresenz, 206

See also CGIs

Acrobat Catalog, 335

Acrobat Distiller, 333–335

- defined, 334
- launching, 347
- PostScript file processing, 348
- uses, 334–335, 347
- using, 347–348, 350
- See also* Adobe Acrobat

Acrobat Exchange, 334

- Acrobat PDF Writer, 334
- defined, 334
- using, 349
- See also* Adobe Acrobat

Acrobat Reader, 333, 359

- configuring, 359
- finding
 - Internet Explorer, 357
 - NCSA Mosaic, 354
 - Netscape Navigator, 359
- installer, 333
- navigation controls, 333
- versions, 333
- See also* Adobe Acrobat

Acrobat Search, 335

Actions window (WebSTAR), 322

Address (A) record, 291–292

- creating, 298
- Domain Name, 292
- Host Address, 292
- See also* resource records

address books, 39

Addressing dialog box (QuickMail), 39–42

- Address Books list, 39
- BCC section, 39
- CC section, 39

- Delay button, 42
- Find button, 40–41
- From List button, 41–42
- recipient list, 39
- Send button, 42
- Special button, 40
- TO section, 39
- See also* QuickMail
- Addressing preferences (Mail*Link Remote), 119–120
 - Alias Creation options, 119–120
 - illustrated, 119
 - Incoming Mail options, 120
 - Outgoing Mail options, 120
 - See also* Mail*Link gateways
- Add Server window (Contact), 421
- Administration service (ListSTAR), 93
- Administration window (FileMaker Pro Server), 198–199
- Administrator Address List window (ListSTAR), 95
- Adobe Acrobat, 331–359
 - Acrobat Catalog, 335
 - Acrobat Distiller, 334–335
 - Acrobat Exchange, 334
 - Acrobat Reader, 333
 - Acrobat Search, 335
 - annotations, 343
 - components, 333–335
 - documents
 - annotations, 343–344
 - with bookmarks, 345
 - browser configuration, 353
 - comments, adding, 349
 - linking, 345
 - printing as, 341
 - publishing external, 346–350
 - publishing internal, 340–346
 - save location, 342
 - from unconfigured documents, 352
 - on WWW, 351–359
 - PDFWriter, 341
 - plug-ins, 335
 - Remark plug-in, 335
 - uses, 340
 - WebLink plug-in, 359
 - WWW browser cooperation, 352–359
 - Internet Explorer, 356–357
 - NCSA Mosaic, 353–356
 - Netscape Navigator, 358–359
- Adobe Web site, 351
- AFP file server
 - AppleShare client creation on, 51
 - selecting volume, 53
- AFP File Server window, 52–53
- aliases
 - changing, 133
 - creating, 120
 - defined, 119
 - requesting, 134
 - using, 120
 - See also* Mail*Link gateways
- Allow and Deny window (WebSTAR), 325–326
 - configuring, 326
 - illustrated, 325
- America Online, 3, 4
- APOP, 160–161
 - defined, 160–161
 - enabling, 165
- Apple Data Streaming Protocol (ADSP) tool, 222
- Apple Filing Protocol (AFP), 192
- Apple Installer Script, 262
- Apple Internet Mail Server, 110, 155–165
 - account information, 159–161
 - Debug window, 156–157
 - defined, 155
 - with Eudora, 155
 - installing, 156–164
 - with Mail*Link, 155
 - mail host configuration, 157–164
 - with Netscape Navigator, 155
 - with NotifyMail, 155
 - requirements for using, 155
 - user setup, 165
- AppleLink, 4
- AppleMail, 146
 - defined, 141
 - format, 148
 - launching, 149
- Apple Modem Tool, 71
- Apple Network Server 500/700, 380

- Apple Open Collaborative Environment (AOCE), 10
- AppleScripts, selecting, 102
- AppleSearch, 259–283
 - to access local/remote information sources, 261
 - access privileges and, 262
 - Apple Installer Script, 262
 - with AppleTalk, 271–275
 - client window, 271
 - defined, 259
 - filters, 281
 - firewalls and, 282
 - functions, 259–260
 - indexes, 264
 - automatically created files, 264
 - days/times for updating, 266
 - destination designation for, 264
 - manual, 265
 - quality of, 265
 - updating, 266
 - indexes folder, 264
 - information sources, 259, 262–270, 272
 - Information Sources window, 263–265
 - launching, 263
 - local information sources designation, 262–267
 - locating, 261
 - MacLink Plus filters, 281–282
 - memory for, 267
 - multiple servers, 261
 - preferences, 266
 - pre-indexed CD-ROMs, 280
 - Program Linking, 271
 - for remote information sources, 261
 - remote information sources designation, 267–270
 - reporters, 259, 260
 - requirements for using, 261
 - search requests
 - entering, 273
 - reading, 274
 - returned, 273
 - saving, 274
 - scheduling, 275
 - server setup, 261–270
 - things to know about, 280–283
- WAIS
 - gateway, 268
 - queries, 268
 - searching, 267–270
 - unpublished sites, 282–283
 - with Web, 276–279
- AppleSearch ACgi, 276–279
 - browser interface support, 276
 - defined, 276
 - forms interface, 279
 - functioning of, 276
 - installing, 276
 - launching, 276–277
 - local information sources, 277
 - returned data, 278
 - search results, 278
 - See also* AppleSearch
 - AppleSearch Administrator's Guide*, 283
 - AppleSearch Status window, 263, 270
 - AppleSearch User's Guide*, 275
- AppleShare
 - Admin, 313
 - algorithm, 441
 - AppleShare 3, 376
 - AppleShare 4, 376
 - caching, 443–444
 - client creation, 51
 - with FileMaker Pro Server, 199
 - File Server, 313
 - file servers, 195
 - performance, 442
 - Personal File Sharing vs., 441–444
 - privileges, 263
 - transmission time, 443
 - WAN transmission speeds with, 445
- AppleTalk, 144
 - AppleSearch via, 271–275
 - file servers, 372
 - hosting data files using, 192
 - printer drivers, 369
 - Remote Access, 74
 - software drivers, 49
- AppleTalk Data Streaming Protocol (ADSP), 192
- Apple Thread Manager, 308
- Apple Workgroup Server 7250, 377–378

- expansion and storage, 378
 - form-factor, 377
 - illustrated, 377
 - networking, 378
 - power, speed, memory, 377
 - See also* Apple Workgroup Servers
 - Apple Workgroup Server 8550, 378–379
 - expansion and storage, 379
 - illustrated, 378
 - networking, 379
 - power, speed, memory, 379
 - software bundle options, 379
 - See also* Apple Workgroup Servers
 - Apple Workgroup Servers, 259, 377–380
 - user access planning, 386–388
 - users & groups planning, 389–392
 - archive scripts, 405–406
 - ARCserve, 397
 - Asynchronous Common Gateway Interfaces. *See* ACGIs
 - asynchronous I/O commands, 200
 - Authenticated Post Office Protocol. *See* APOP
 - AutoResponder (ListSTAR), 99
- B**
- backup management, 397–406
 - choices, 398
 - Retrospect, 397, 399–406
 - software, 397
 - backup planning, 399–401
 - creating, 399–401
 - illustrated, 400, 401
 - backup scripts, 402–404
 - archive, 405–406
 - duplicate, 404
 - backup servers, 399–406
 - archive scripts, 405–406
 - duplicate script, 404
 - efficiency of, 455
 - performance, 455
 - planner creation, 399–401
 - RAID 0 and 5, 402
 - RAM and, 401
 - script, 402–404
 - setting up, 401–404
 - See also* backup management
 - bandwidth
 - e-mail performance and, 436–438
 - Ethernet, sharing, 439
 - binary format, 83
 - BinHex format, 83, 85, 121, 122
 - BITNET addresses, 163
 - bridge MailCenters, 20
 - bridges (QuickMail), 20–21
 - Printer Bridge, 20, 53–58
 - QM-Direct Bridge, 20, 58–61
 - QM-QM Bridge, 21, 61–69
 - See also* QuickMail
 - browsers, Acrobat and, 352–359
 - Internet Explorer, 356–357
 - NCSA Mosaic, 353–356
 - Netscape Navigator, 358–359
 - See also* Adobe Acrobat
 - building servers, 384–385
 - CPU, 384
 - hard drive, 384
 - network backup/archiving system, 385
 - UPS, 385
 - ButlerClient, 233–235
 - Connection window, 234
 - defined, 216
 - launching, 233
 - using, 233–235
 - See also* Butler SQL
 - ButlerHosts
 - Data window, 231
 - defined, 216
 - launching, 230
 - using, 230–231
 - See also* Butler SQL
 - Butler Preferences window
 - Access privileges checkbox, 219
 - Allow program checkbox, 219
 - Allow remote clients checkbox, 219
 - illustrated, 219
 - opening, 219
 - Transaction processing checkbox, 219
 - Butler SQL, 213–237
 - accessing from MacOS, 229

- access privileges, 216–222
 - adding users for, 220
 - changing, 221
 - group, 222
 - group creation, 220
 - privileges enabled for, 221
 - ButlerClient, 216, 233–235
 - ButlerHosts, 216, 230–231
 - Butler SQL Server, 215, 218, 229–237
 - ButlerTools, 215, 219, 225–228
 - client connections, 229–233
 - control panel/extensions needed, 229–230
 - DAL/DAM connections, 230–231
 - ODBC connections, 231–233
 - port types, 229
 - preferences stored, 230
 - used for client configuration, 229
 - client/server interaction used by, 214
 - components, 215–216
 - DAM and, 214
 - database creation, 225–228
 - database extensions, 216
 - databases, 215
 - ODBC and, 214
 - port setup, 222–224
 - Programming Linking, 216–217
 - requirements, 215
 - Tango, 216, 235–237
 - Test Drive, 228
 - user packages, 215
 - WebSTAR integration with, 235
 - working with, 215–224
- Butler SQL Server, 215
 - interfacing with, 229–237
 - launching, 218
 - windows, 218
 - See also* Butler SQL
 - ButlerTools, 215
 - launching, 219
 - using, 225–228
 - See also* Butler SQL
- C**
- Canonical Name (CNAME) record, 293–294
 - Canonical Name, 294
 - creating, 299
 - defined, 293
 - Domain Name, 294
 - See also* resource records
 - Canvas, 346–347
 - cataloging (Cumulus), 173–175
 - Catalog menu (PowerTalk)
 - illustrated, 142
 - New User command, 143
 - Catalog preferences (Cumulus), 172
 - catalogs, 141–146
 - defined, 141
 - desktop, 142
 - live, 141
 - opening, 147
 - personal, 142–143
 - See also* PowerTalk
 - Catalog window (PowerTalk), 148
 - categories
 - adding contact to, 424
 - color-coded, 423
 - defining, 422
 - naming, 422
 - passwords, 423
 - personal, 419
 - private, 420
 - public, 419–425
 - selecting, 427
 - sets of, 427
 - shared, 422
 - See also* Contact; Up-to-Date
 - CGIs, 187–188, 201–212
 - ACGIs vs., 205–207
 - application languages, 203
 - application uses, 201
 - defined, 201
 - example, 209–212
 - hidden variables and, 205
 - HTML commands calling, 204–205
 - Java applets vs., 212
 - main job, 204
 - as mediators, 202
 - multiple on one machine, 211
 - performance and, 211–212

- to perform separate functions, 209
 - Round Robin, 329
 - script access, 202–203
 - security and, 211
 - using, 211–212
 - Web server performance and, 450
 - writing applications, 203
 - WWW, 188
 - See also* ACGIs
- Chooser
- defined, 440
 - illustrated, 440
 - performance and, 440–441
- color printers, 364–365
- continuous tone, 365
 - inkjet, 364
 - toner, 364
 - wax, 364
 - See also* printers
- Common Gateway Interfaces. *See* CGIs
- CompuServe, 3, 4, 144
- The Complete Guide to Macintosh Backup Management*, 257, 383, 385, 397, 399, 455
- Configure File Type window (Internet Explorer), 357
- Configure MailCenter window (QuickMail), 25
- Configure menu (WebSTAR)
- Actions command, 321
 - Add Password command, 324, 325
 - Allow/Deny command, 325
 - Misc. Settings command, 309, 318
 - Realms command, 323
 - Suffix Mapping command, 322
- Configure Online MailCenter window (QuickMail), 28–29
- checkboxes, 28–29
 - Days until removal of public messages, 29
 - Mail priorities, 29
 - Maximum entries in user's mail list, 29
 - Maximum entries in user's mail log, 29
- Connection Statistics window (Apple Internet Mail Server), 164
- Connect Times window (Mail*Link), 115–116
- Connect window (ButlerClient), 234
- consumables, 368–369
- cost, 369
 - defined, 368
 - See also* printers
- Contact, 407–430
- Add Server window, 421
 - client installation, 426
 - client setup, 426–430
 - contacts
 - adding, 424
 - defined, 424
 - fields, 425
 - stand-out, 423
 - databases, 408, 410
 - Define Categories window, 422
 - Define Server Connections window, 421
 - destination Macintosh guidelines, 410
 - distributing, 426
 - distribution files, 430
 - file creation, 426–430
 - initial configuration, 411–412
 - installing, 409
 - launching, 421
 - private category, 420
 - public category, 420–425
 - Public Contact Server, 409, 412
 - requirements for using, 410
 - server location, 409
 - setting up, 409–425
 - uses, 410
 - See also* Up-to-Date
- continuous tone printing, 365
- CPU
- choice for databases, 191–196
 - file server, 384
 - usage for open applications, 206
- Create a Special Address dialog box (QuickMail), 40
- Create Phone Book Entry window (QuickMail), 62–63
- If no messages are sent option, 63
 - illustrated, 62
 - When sending to this site... option, 63
- Create Special Address window (Mail*Link), 133, 134
- CSMA/CD (Carrier Sense Multiple Access/Collision Detection), 456

- Cumulus*Client, 168
 - artwork sort methods, 177
 - copying pictures, 179
 - defined, 168
 - previewing, 178–179
 - RAM and, 181
 - text display view, 177
 - thumbnail display view, 177
 - view choices, 179
 - working with, 176–179
 - Cumulus, 167–184
 - access speed, 181–182
 - cataloging, 173–175
 - client/server art database catalog system, 168
 - copying files in, 172
 - databases
 - ability of, 169
 - access to, 173
 - adding to from client, 173
 - dividing, 180–181
 - file creation, 170–171
 - multiple across multiple servers, 180
 - one vs. multiple decision, 180–181
 - records, 175
 - saving changes immediately, 171
 - sharing, 176
 - defined, 167
 - file formats, 183
 - filters, 183
 - graphics database machine and, 169
 - hard drive speed and, 182
 - heavy users, 181
 - high-resolution format, 183
 - implementation considerations, 180–184
 - importing to, 174
 - installing, 170
 - keywords, 173
 - logging activity, 171
 - low-resolution format, 183
 - memory usage, 182
 - on-line resolution, 184
 - password, 176–177
 - planning issues, 169
 - preferences, 171–173
 - Program Linking, 176
 - RAM and, 181
 - records, 175
 - registration, 170
 - screen resolution, 183–184
 - server, working with, 169–175
 - system planning, 180–184
 - usefulness of, 167–168
 - See also* Cumulus*Client
 - custodians, 25
 - coordination of, 26
 - QM-QM Bridge, 68–69
 - QM-Script Gateway, 73
 - QuickMail, 26
- D**
- Data Access Manager (DAM)
 - connections, 230–231
 - defined, 214
 - Data Backup Policy document, 340–346
 - Database Access Language (DAL)
 - connections, 230–231
 - defined, 214
 - database creation (Butler SQL), 225–228
 - complete database window, 228
 - complete table window, 228
 - initial column, 227
 - initial database window, 226
 - initial table window, 226
 - normalized, 225
 - tables, 226
 - using ButlerTools, 225–228
 - See also* Butler SQL
 - Database menu (Cumulus)
 - Allow Sharing command, 171
 - Database Setup command, 171
 - databases (Contact), 408
 - large, 410
 - Up-to-Date linked, 410
 - See also* Contact
 - databases (Cumulus)
 - access to, 173
 - adding to, from client, 178
 - dividing, 180–181
 - multiple across multiple servers, 180

- one vs. multiple decision, 180–181
- records, 175
- saving changes immediately, 171
- sharing, 176
- See also* Cumulus
- databases, xi–xii, 185–200
 - CPU choice for, 191–196
 - creation considerations, 189
 - current developments of, 187
 - defined, 185
 - FileMaker Pro, 188, 195
 - flat file, 185–186
 - layouts for, 190
 - multi-user, 190, 191
 - performance, 452–454
 - the big mistake and, 453–454
 - signing in and, 452–453
 - tests on, 453
 - See also* service performance
 - planning, 189–200
 - RAM and, 191
 - relational, 186–187
 - tracking, 255
- data files, 189
 - ghosting hosting, 193
 - hosting using AppleTalk protocol, 192
 - layouts and, 190
 - planning, 189
 - See also* databases
- DataViz Inc., 281
- DAT drives, 401, 404
- Date & Time control panel, 157
- Debug window (Apple Internet Mail Server), 156–157
- dedicated file sharing, 375–380
 - Apple servers, 377–380
 - benefits, 375
 - defined, 372
 - Macintosh-based servers, 375–376
 - model, 375
 - types of, 375
 - See also* file sharing
- dedicated Macintosh-based servers, 375–376
 - AppleShare 3, 376
 - AppleShare 4, 376
- default service (ListSTAR), 89–93
 - configuring, 89–90
 - defined, 89
 - Miscellaneous pane, 91
 - Preferences Account pane, 90
 - Preferences Rules pane, 90
 - Reply pane, 91
 - Rules Comment pane, 92
 - See also* ListSTAR
- Define Categories window (Contact), 422
- Define menu (Contact)
 - Categories command, 422
 - Category Styles command, 423
 - Custom Fields command, 425
 - Server Connections command, 421
 - Set command, 427
- Define menu (Up-to-Date)
 - Connections command, 417
 - Set command, 427
- Define Server Connections window (Contact), 421
- Designing AppleTalk Network Architectures*, xi, 456
- Designing AppleTalk Networks course, 431
- Destination preferences (smtp.daemon), 130
- digital documents, 332
- digital publications, 336–339
 - examples of, 336–339
 - FAQs, 338
 - how to connect from home or the road, 337
 - how to connect the file server, 336
 - how to connect to the Web server, 337
 - how to dial out using the network modem, 337
 - how to obtain your e-mail, 337
 - how to print across the network, 336–337
 - how to search your network databases, 337
 - how to use the Internet feed, 338
 - network event bulletins, 338
 - policy manuals, 338
 - user training example, 339
 - who to call for technical support, 336
- Disk Accountant, 394–395
- DNS, 156, 285–301
 - authoritative domains, 288
 - database navigation, 287–289
 - defined, 156, 285
 - domain name resolvers, 287

- domain name servers, 287, 289
 - domain name space, 287
 - dos and don'ts, 301
 - function of, 285
 - hierarchy, 287
 - host naming, 301
 - ISPs and, 286
 - queries, 288
 - recursive name service, 289
 - resource records, 290–295
 - reverse domain mapping, 294, 300
 - root domain, 287
 - servers, 130
 - See also* QuickDNS Pro
 - Document Information dialog box (Acrobat), 342
 - documents (Acrobat)
 - annotations, 343–344
 - with bookmarks, 345
 - browser configuration, 353
 - comments, adding, 349
 - linking, 345
 - printing as, 341
 - publishing external, 346–350
 - publishing internal, 340–346
 - save location, 342
 - from unconfigured browsers, 352
 - See also* Adobe Acrobat
 - Domain Information window (QuickDNS Pro), 296–297
 - Domain menu (QuickDNS Pro)
 - Create Record command, 297
 - Get Info command, 296
 - domain name resolvers, 287
 - domain name servers, 287
 - local, 289
 - primary, 289
 - secondary, 289
 - See also* DNS
 - domain name space, 287
 - Domain Naming System. *See* DNS
 - DOS client (QuickMail), 32–33
 - CE.INI, 32
 - CONFIG.EXE, 33
 - connection screen, 33
 - QC.EXE, 32
 - QCONF.EXE, 32
 - QMAILATP.EXE, 32
 - QMATP, 32
 - See also* QuickMail
 - duplicate scripts, 404
- E**
- Easy Open, 331
 - Edit menu (Butler SQL)
 - Create Host command, 230
 - Preferences command, 219
 - Edit menu (Cumulus)
 - Copy command, 179
 - Zoom In/Zoom Out command, 179
 - Edit menu (PowerTalk), 248, 249, 252
 - Edit menu (smtp.daemon), 126
 - Electronic bulletin board systems, 5–6
 - chat areas, 5
 - costs, 6
 - defined, 5
 - illustrated, 5
 - network design, 6
 - security, 6
 - Sysop, 5
 - threads, 5
 - electronic documents. *See* forms
 - Electronic Invasion: Brave New World of Business Communications*, 16
 - e-mail, xi
 - adjective stage, 16
 - defined, 1–2, 16–17
 - electronic bulletin board systems, 5–6
 - enterprise LAN/WAN e-mail systems, 12–15
 - Internet, 109–137
 - logs, 435–436
 - mainframe messaging centers, 3–4
 - messaging, 1–76
 - Network Frontiers, 14
 - noun stage, 17
 - peer-to-peer e-mail systems, 10–11
 - performance, 434–439
 - bandwidth and, 438–439
 - general, 436–437
 - logs and, 435–436

- signing in and, 434–435
 - WAN, 438
 - See also* service performance
 - publishing with ListSTAR, 77–108
 - sending
 - methods, 2
 - PowerTalk, 146–149
 - QuickMail, 37–42
 - usage scope, 1
 - verb stage, 17
 - E-Mail-on-Demand demo (ListSTAR), 96–107
 - Error Handling rule, 105–106
 - Mailer-Daemon Mail rule, 98–99
 - Send Desired Files rule, 101–102
 - Send File Index rule, 99
 - timer, 106–107
 - Track All Visitors rule, 103–105
 - Update Classes List rule, 106–107
 - Welcome/Help rule, 103
 - See also* ListSTAR
 - Enclosure preferences (Mail*Link Remote), 120–122
 - Binary enclosures, 121–122
 - illustrated, 121
 - Limit total enclosure size to, 122
 - TEXT enclosures, 121
 - See also* Mail*Link
 - Enclosures preferences (ListSTAR), 83–85
 - illustrated, 84
 - Outgoing Body Translation, 85
 - Outgoing Enclosure Translations, 84–85
 - See also* Preferences window (ListSTAR)
 - EPSF files, 55
 - Ethernet, 362
 - CSMA/CD, 456
 - performance and, 456
 - Eudora, 155, 165
 - Apple Internet Mail Server with, 155
 - enabling APOP with, 165
 - See also* e-mail
 - Event menu (Up-to-Date)
 - Event command, 427
 - New To-Do command, 429
 - Exchange Files function, 312
 - exporting, phone book lists, 66
- F**
- Fetch, 317
 - file-based MailCenters, 20, 49–53
 - password, 52
 - from QM Administrator, 50
 - as seen by QMNET, 51
 - server volume, 51
 - setting up, 51–53
 - See also* MailCenters
 - FileMaker Pro, 187
 - access privileges, 190
 - databases, 188, 195
 - ghosting hosting, 193
 - multi-user databases, 191
 - nondedicated hosting, 192
 - password security, 190
 - peer-to-peer hosting, 191
 - RAM allocation, 195–196
 - FileMaker Pro Server, 187
 - administrators, 198
 - AppleShare running with, 199
 - bigger, faster, stronger and, 199
 - capacity, 196
 - database file size, 196
 - default preferences, 197–198
 - defined, 196–197
 - with different protocol clients, 197
 - Ethernet and, 200
 - LocalTalk and, 200
 - recommendations, 199–200
 - support, 196
 - when to use, 196–200
 - See also* FileMaker Pro
 - File menu (Apple Internet Mail Server), 159
 - File menu (AppleSearch)
 - Export Reporter command, 275
 - Import Reporter command, 275
 - File menu (Butler SQL)
 - Install Port command, 224
 - New Database command, 225
 - New Port command, 223
 - File menu (Contact), 421
 - File menu (Cumulus)
 - New command, 170

- Share Catalog command, 176
- File menu (FileMaker Pro)
 - Administrator command, 198
 - Open command, 193
- File menu (Finder), 381
- File menu (Informed), 250
- File menu (QuickDNS Pro)
 - New command, 296
 - Save command, 297
- File menu (QuickMail)
 - Configure command, 28
 - New MailCenter command, 24, 53, 59, 61, 70, 113
 - Newsflash command, 30
- file servers, 371–395
 - AFP, 51, 53
 - Apple family, 377–380
 - Apple Network Server 500 and 700, 380
 - AppleTalk, 372
 - Apple Workgroup Server 7250, 377–378
 - Apple Workgroup Server 8550, 378–379
 - building, 384–385
 - confidential information on, 383
 - CPU, 384
 - critical information on, 382–383
 - elements of, 372
 - file size, 381–382
 - hard drive, 384
 - inherited network plan, 393–395
 - Macintosh-based, 375–376
 - needs analysis, 381–383
 - network backup/archiving system, 385
 - personal information on, 382
 - software, 372
 - types of, 372
 - UPS, 385
 - user's needs, 381
- file sharing, 371–395
 - Apple software, 372
 - dedicated, 372, 375–380
 - functioning of, 372
 - non-dedicated, 372, 373–374
 - performance, 440–445
 - AppleShare vs. Personal File Sharing, 441–444
 - Chooser and, 440–441
 - over WANs, 444–445
 - See also* service performance
 - See also* file servers
- File Transfer Protocol. *See* FTP
- File Type Picker dialog box (Mosaic), 355
- FinePrint printing, 363
- firewalls
 - AppleSearch and, 282
 - WebSTAR and, 323
- flat file database structure, 185–186
 - defined, 185
 - disadvantage of, 186
 - illustrated, 186
 - planning, 189–190
 - See also* databases
- fonts, 366–367
 - TrueType, 366, 367
 - Type 1, 366, 367
 - See also* printers
- forms, 37–38, 239–257
 - approving, 252
 - defined, 37
 - designing, 257
 - electronic, 240
 - errors/omissions on, 245
 - fields, nonuseful, 256
 - functions needing, 256
 - Informed, 241
 - Memo, 38, 57
 - migration to, 256
 - officious, 257
 - paper, 257
 - planning for, 256–257
 - pre-addressed, 251
 - sending, 251–252
 - signature, 248–250
 - tracking, 251, 254, 255
 - “Unnecessary but Mandatory Recriminations Cycle” and, 254
 - verification, 248
 - See also* Informed
- Frontiers in Networking*, 346
- FTP, 303, 313–317
- FTPd, 313–317

defined, 313
 menu bar
 FTP Setup button, 315
 FTP Users button, 315
 Gopher Setup button, 316
 illustrated, 315
 Security button, 316
 Summary button, 316
 Setup application, 314–315
See also WebSTAR
 FTP Users window (FTPd), 316

G

gateways (QuickMail), 20–21
 QM-AOL Gateway, 21
 QM-Link Gateway, 21
 QM-Script Gateway, 13, 21, 69–74
 General Info window (Acrobat Exchange), 349
 General preferences (Cumulus), 171
 General preferences (ListSTAR), 82–83
 illustrated, 82
 Logging Level, 83
 Outgoing Mail Queuing, 83
 Threads, 83
 See also Preferences window (ListSTAR)
 General preferences (Mail*Link Remote), 116–117
 Display button, 136
 illustrated, 116
 Logging Level, 116–117
 Outgoing Mail Queuing, 116
 QuickMail Connect Times, 117
 See also Mail*Link gateways
 General preferences (smtp.daemon), 127
 General Preferences window (Netscape Navigator), 358
 Get Info window (Mail*Link), 136
 ghosting hosting, 193
 GIF format, 184
 Gopher Setup window (FTPd), 316
 graphic letterhead, 55
 groups
 Apple Workgroup Server planning, 389–392
 defined, 372
 guest access, 303, 311

H

headers, 120
 Helper Configuration window (Mosaic), 354
 hidden variables, 205
 Host Configuration window (ButlerHosts), 230
 Host Info (HINFO) record, 293
 hosting, 191–193
 database files to multiple users, 196
 ghosting, 193
 nondedicated, 192
 peer-to-peer, 191
 Hosts dialog box (FileMaker Pro), 194
 HTML
 <> characters, 305
 commands calling CGIs, 105
 data definition, 210
 document publishing in, 340
 for fast typists, 304–307
 moving data to, 340
 preformatted layouts, 210–211
 HTML editors, 304
 HTML form, 204
 hidden variable definition, 208
 predefining variables, 208
 record entry area, 207
 HTML links, 304
 HTML tags, 305–306
 HTTP, 303
 HyperText Markup Language. *See* HTML
 HyperText Transfer Protocol. *See* HTTP

I

Import File window (QuickMail), 67–68, 135
 column title, 68
 Confirm entries before importing checkbox, 68
 illustrated, 67
 See also QuickMail
 importing
 artwork to Cumulus, 174
 Internet address lists, 135
 into Public Event Server, 415–416
 phone book lists, 67
 indexes, 264

- automatically created files, 264
 - days/times for updating, 266
 - destination designation for, 264
 - manual, 265
 - pre-indexed CD-ROMs, 280
 - quality of, 265
 - updating, 266
 - See also* AppleSearch
- information sources, 259
 - automatic indexing of, 264
 - folder selection for, 263
 - local, designating, 262–267
 - manually indexing, 265
 - memory for, 267
 - remote, designating, 267–270
 - selecting, 272
 - status of, 265
 - See also* AppleSearch
- Information Sources window (AppleSearch), 263–265
 - Add button, 263, 264
 - Get Info button, 269
 - illustrated, 263, 265
 - Index button, 264
- Informed, 239–257
 - choice list, 246
 - Database extensions, 242, 245
 - Designer, 241
 - e-mail extensions, 242
 - error/omission checking, 245
 - forms, 241
 - approving, 252
 - sending, 251–252
 - tracking, 251, 254
 - Foundation, 247
 - Manager, 241
 - multiple signatures, 249
 - Number Server, 242
 - omission warning, 246
 - Program Linking, 244
 - system components, 241–242
 - Tracker, 242
 - Tracker Admin, 255
 - using, 243–255
 - See also* forms; signatures
- Informed Designer
 - defined, 241
 - See also* Informed
- Informed Foundation, 247
- Informed Manager
 - database query, 253
 - defined, 241
 - Scripts menu, 243
 - using, 244–250
 - View menu, 253
 - See also* Informed
- Informed Number Server
 - defined, 242
 - illustrated, 244
 - See also* Informed
- Informed Tracker
 - data and time entry, 253
 - defined, 242
 - See also* Informed
- Informed Tracker Admin, 255
- inherited network server plan, 393–395
 - storage evaluation, 394–395
 - user account review, 393–394
- inkjet printers, 365
- installation
 - Apple Internet Mail Server, 156–164
 - AppleSearch ACGI, 276
 - Contact, 409
 - Contact clients, 426
 - Cumulus, 170
 - ListSTAR for SMTP, 79–88
 - PowerTalk, 141
 - public categories, 419–425
 - Up-to-Date, 409
 - Up-to-Date clients, 426
 - WebSTAR, 308
- Install PowerTalk folder, 145–146
- Internet
 - address lists, importing, 135
 - hosts, 289
 - Standard for the Format of ARPA Internet Text Messages*, 160
- Internet e-mail, 109–137
 - PowerTalk, 151–154
 - See also* Mail*Link gateways

- Internet Explorer, 356–357
 - Configure File Type window, 357
 - Options window, 356
 - See also* documents (Acrobat); PDF
- Internet Service Providers (ISPs), 155
 - DNS and, 286
 - e-mail hosts, 110
- IP addresses, 294
- iteration, 288
- J**
- Java applets, CGIs vs., 212
- JPEG format, 184
- K**
- Key Chain (PowerTalk), 151
- Keywords menu (Cumulus)
 - New Keyword command, 173
 - Show Keywords command, 173
- L**
- Label menu (System 7), 312
- LAN client/server e-mail systems, 7–9
 - administration, 8
 - costs, 8
 - defined, 7
 - illustrated, 7
 - network design, 8
 - security, 9
 - See also* QuickMail
- LANs
 - printing throughput times, 446
 - QuickMail efficiency over, 437
 - Web server throughput, 448
 - See also* WANs
- LANTest, 432
- LAN/WAN e-mail systems, 12–15
 - administration, 15
 - connect to everything/anything, 13–15
 - LAN-to-host, 12
 - LAN-to-LAN, 12–13
 - legal issues, 15
 - network design, 15
 - Network Frontiers, 14
 - security, 15
 - See also* e-mail
- letterhead
 - correctly placed, 57
 - incorrectly placed, 56
 - selecting, 54
- list servers
 - defined, 77
 - testing, 108
- ListSTAR, 77–108
 - AutoResponder, 99
 - correct time, 87–88
 - defined, 77–78
 - direct mailing lists, 78
 - E-Mail-on-Demand demo, 96–107
 - Error Handling rule, 105–106
 - Mailer-Daemon Mail rule, 98–99
 - Send Desired Files rule, 101–102
 - Send File Index rule, 99–101
 - timer, 106–107
 - Track All Visitors rule, 103–105
 - Update Classes List rule, 106–107
 - Welcome/Help rule, 103
 - functions of, 77
 - host name, 80
 - installing for SMTP, 79–88
 - launching, 80–81
 - mailing lists, 78
 - for POP, 77, 79–80
 - preferences
 - About ListSTAR, 88
 - Enclosures, 83–85
 - General, 82–83
 - Naming Hosts, 81–82
 - SMTP/TCP, 80, 86–87
 - Status Window, 85–86
 - reserved addresses, 98–99
 - rules, 91
 - serial number, 81
 - servers used with, 77
 - services, 89–95
 - Administration, 93
 - default, 89–93

- demo, 96–97
- e-mail-on-demand setup, 96–107
- Postmaster, 93–95
- for SMTP, 77, 79–88
 - folder, 80
 - installing, 79–88
- for StarNine Mail, 77, 79
- technical support numbers, 88
- testing and, 108
- Welcome Reply, adding, 103
- windows, 81
- ListSTAR Administrator's Guide*, 83, 101, 108
- ListSTAR Services window
 - Duplicate button, 96
 - illustrated, 97
 - Rename button, 99, 101
- local information sources, 262–267
- LocalTalk, 362
- Login Script window (Mail*Link), 124
- logs, 435–436

M

- MacHTTP, 308
- Macintosh client (QuickMail), 33–42
 - connection screen, 34
 - Desk Accessory, 33
 - function buttons, 36
 - mailbox, 36
 - mail log, 36
 - Main window, 35–36
 - New button, 37
 - newsflash area, 37
 - personal folders, 37
 - public section, 36–37
 - QuickConference, 33
 - sending mail with, 37–42
 - See also* QuickMail
- MacLink Plus, 331
- MacLink Plus XTND translators, 260, 281–282
 - default, 282
 - potential problems, 281–282
 - protocol, 282
 - See also* AppleSearch
- MacPPP, 152

- MacSLIP, 152
- MacTCP, 296
- Mail*Link gateways, 109–137
 - Apple Internet Mail Server with, 155
 - defined, 109
 - domain, 110
 - information gathering and, 110
 - location of, 111
 - login procedures, 110
 - MailCenter configuration, 113–130
 - message path, 111
 - passwords, 110
 - queued messages, 136–137
 - Remote, 109, 110
 - Remote preferences, 116–126
 - Addressing, 119–120
 - Enclosures, 120–122
 - General, 116–117
 - Naming Hosts, 117–118
 - Remote Receiving, 125–126
 - Remote Systems, 122–125
 - Users, 118–119
 - setting up, 112–130
 - SMTP, 15, 109, 110
 - preferences, 115–116
 - smtp.daemon, 126–130
 - TCP/IP configuration, 113
 - smtp.daemon preferences, 111, 126–130
 - Destination, 130
 - General, 127
 - Naming Hosts, 127–128
 - SMTP/TCP, 129–130
 - Status Window, 128–129
 - user accounts, establishing, 131–135
 - See also* e-mail
- Mail*Link Monitor, 153, 154
- Mailbox (PowerTalk), 141
 - checking, 149–150
 - open, 150
 - See also* PowerTalk
- MailCenters, 20–21
 - bridge, 20
 - configuration, changing, 28–29
 - custodian, 25
 - defined, 20

- file-based, 20, 49–53
- gateway, 20
- Mail*Link gateway configuration, 113–130
- Mail*Link users list, 134–135
- maintaining, 26
- names, 35
- online, 20
- passwords, 24, 48
- QM-AOL Gateway, 21
- QM-Direct Bridge, 20, 58–61
- QM-Link Gateway, 14, 21
- QM-Printer Bridge, 20, 53–58
- QM-Script Gateway, 13, 21, 69–74
- Store & Forward, 21, 46, 47
- user creation, 27
- user list, 28
- See also* QuickMail
- MailCenter Selection dialog box (QuickMail), 35
- MailCenter window (QuickMail), 24
- Mailers (PowerTalk), 147
- Mailer services (ListSTAR), 89
- Mail Exchange (MX) record, 292–293
 - creating, 298
 - Domain Name, 292
 - Mail Exchange, 293
 - Preference, 293
 - See also* resource records
- mail host configuration, 157–164
 - backing up drive and, 164
 - connection statistics, 164
 - data and time, 157
 - preferences, 157–159
 - sending via other hosts, 163–164
 - user accounts, 159–162
 - See also* Apple Internet Mail Server
- mailing lists, 78
 - creating, 104
 - direct, 78
 - See also* ListSTAR
- Mail menu (PowerTalk), Send command, 148
- Mail Options window (Informed), 251
- mail servers
 - configuring TCP/IP, 113
 - setting time zone, 112
- mainframe messaging centers, 3–4
 - costs, 4
 - defined, 3
 - illustrated, 3
 - legal issues, 4
 - network design, 3
 - See also* e-mail
 - Managing AppleShare & Workgroup Servers*, 170, 200, 262, 372, 443
 - Managing AppleTalk Networks*, 199, 283, 330
- many-to-one relationship, 186
- Map control panel, 87–88, 112, 157
- MapServe, 320, 322
- margins, 54–55
 - default, 56
 - incorrect setting of, 55
 - setting, 54
 - See also* Printer Bridge
- Meeting menu (Up-to-Date), Administration command, 417, 419
- messages (Mail*Link)
 - headers, 120
 - path steps, 111
 - queued, 136–137
 - See also* Mail*Link gateways
- messages (PowerTalk), 149–150
- messages (QuickMail)
 - creating “on the fly” user, 40
 - date/time sending of, 42
 - finding user for, 40–41
 - header information, 71
 - paths, 43–47
 - on different mail server, 45–47
 - QuickMail for Windows, 49–51
 - in same MailCenter, 44–45
 - wrong idea of, 43–44
 - sending, 42
 - See also* Macintosh client (QuickMail); Quick-Mail
- Miscellaneous Settings window (WebSTAR), 309–310
 - Buffer Size field, 320
 - illustrated, 318
 - Max Listens field, 319
 - Max Users field, 319
 - Port field, 328

Timeout field, 319, 320

N

Name Server (NS) record, 292
creating, 297

See also resource records

NameServer control panel, 19

database, 19

illustrated, 19

See also QuickMail

Naming Hosts preferences (ListSTAR), 81–82

illustrated, 82

See also Preferences window (ListSTAR)

Naming Hosts preferences (Mail*Link Remote),
117–118

Default Mail Host, 118

illustrated, 117

My Hostname, 117

Outgoing Hostnames, 118

See also Mail*Link gateways

Naming Hosts preferences (smtp.daemon),
127–128

illustrated, 127

My Hostname field, 128

Send all mail to default host checkbox, 128

See also Mail*Link gateways

NBP broadcast packets, 440

NBP-LKups, 434, 441

NBPReps, 434

NCSA Mosaic, 353–356

File Type Picker dialog box, 355

Helper Configuration window, 354

Preferences window, 353

proper PDF document configuration, 356

See also documents (Acrobat); PDF

NetMinder Ethernet, 432

Netscape Navigator, 155, 305–306

General Preferences window, 358

New MIME Type window, 358

See also documents (Acrobat); PDF

network designs

electronic bulletin board systems, 6

LAN client/server e-mail systems, 8

LAN/WAN e-mail systems, 15

mainframe message centers, 3

peer-to-peer e-mail systems, 10–11

Network Frontiers, xiii–xiv

e-mail addresses, xiii

e-mail system, 13–15

Web site, xiv

network traffic. *See* service performance

Network Visible Entities (NVEs), 199

New Form dialog box (QuickMail), 37

New MIME Type window (Netscape Navigator),
358

New Reports window (AppleSearch), 272

newsflash, 29–30

defined, 29

sending, 29–30

See also QuickMail

Newsflash window (QuickMail), 30

non-dedicated file sharing, 373–374

benefits of, 373

defined, 372

illustrated, 373

performance degradation, 373–374

See also file sharing

nondedicated hosting, 192

NotifyMail, 155

Now. *See* Contact; Up-to-Date

O

ODBC

connections, 231–233

control panel, 131

data source connection configuration, 233

Data Sources window, 232

defined, 214

drivers, 214, 232

Tango and, 235

See also Butler SQL

online MailCenters, 20

Open Database Connectivity. *See* ODBC

Open dialog box (FileMaker Pro), 193, 194

Open File dialog box, 66–67, 135

Open Transport, 296, 372

Options menu (WebSTAR), Verbose Messages com-
mand, 308

Options window (Internet Explorer), 356

P

paper

- sizes, 365
 - special, 366
 - weights, 366
- See also* printers

passwords

- category, 423
- Cumulus, 176–177
- FileMaker Pro security, 190
- Mail*Link gateway, 110
- MailCenter, 24, 48, 52
- PowerTalk signer, 249
- QM-Direct Bridge, 59
- QM-QM Bridge, 62, 64
- QM-Script Gateway, 70
- QuickMail, 35
- realm, 324
- Timbuktu Pro, 313
- Up-to-Date, 420

PDF

- defined, 332
 - documents
 - annotated, 344
 - with bookmarks, 345
 - conversion into, 340–350
 - Internet Explorer configuration, 356–357
 - Mosaic configuration, 353–356
 - Netscape Navigator configuration, 358–359
 - file annotation, 343
 - file creation in background, 334
 - file linking, 345
 - publishing examples, 351
 - read-only replica document in, 342
- See also* Adobe Acrobat

PDFWriter, 341

peer-to-peer e-mail systems, 10–11

- administration, 11
- defined, 10
- illustrated, 10
- network design, 10–11

security, 11

See also e-mail

peer-to-peer hosting, 191

performance, 431–458

- AppleShare, 442
- backup servers, 455
- databases, 452–454
- efficient, 433
- e-mail, 434–439
- file sharing, 440–445
- Personal File Sharing, 441
- printers, 367–368
- printing, 446–447
- results analysis, 432
- test machines, 432
- test parameters, 432
- tests, 431–432
- Web servers, 448–451

Permissions preferences (Cumulus), 173

personal catalog, 142–143

- new, 142
 - opening, 143
- See also* catalogs

Personal File Sharing, 217, 314

- access privileges through, 262
- algorithm, 441
- AppleShare vs., 441–444
- performance, 441
- transmission time, 442
- WAN transmission speeds with, 444

phone book (QuickMail), 66

- entries, 62–63
 - exporting, 66
 - importing, 67
 - using, 66
- See also* QM-QM Bridge

Photograde printing, 363

Pointer (PTR) record, 294–295

- defined, 294
 - Domain Name, 295
 - Pointer Name, 295
- See also* resource records

Portable Document Format. *See* PDF

portable documents, 332

ports (Butler SQL), 222–224

- choosing, 223
- communications toolbox, 222
- program linking, 222
- TCP/IP, 223–224
- See also* Butler SQL
- Postmaster service (ListSTAR), 93–95
 - Account pane, 94
 - defined, 93
 - Error Handling Reply pane, 94
 - See also* ListSTAR
- PostScript, 368
- PowerFax gateway, 144
- PowerShare server, 10, 144
- PowerTalk, 139–154
 - addressing window, 251
 - catalogs, 141–146
 - defined, 141
 - desktop, 142
 - live, 141
 - opening, 147
 - personal, 142–143
 - client/server methodology, 10
 - defined, 141
 - desktop additions, 141
 - digital signature ability, 248
 - electronic addresses, 144–145
 - enclosures, 149
 - included in System 7.5, 140
 - information cards, 145
 - installation folder, 145–146
 - installing, 141
 - Internet mail, 151–154
 - forwarding, 153
 - in In Tray, 153
 - Key Chain and, 151
 - sending, 154
 - Mailbox, 141, 149–150
 - Mailer, 147
 - mail format, 148–149
 - on the desktop, 140–150
 - Personal Internet Gateway, 144, 151
 - planning for, 10
 - with QuickMail, 11
 - sending e-mail, 146–149
 - signature icon, 250
 - signers, 248–249
 - to the world, 151–154
- Pre-Address PowerTalk window, 251
- preferences (Cumulus), 171–173
 - Catalog, 172
 - General, 171
 - Permissions, 173
 - See also* Cumulus
- preferences (Mail*Link), 115–130
 - Remote, 116–126
 - Addressing, 119–120
 - Enclosures, 120–122
 - General, 116–117
 - Naming Hosts, 117–118
 - Remote Receiving, 125–126
 - Remote Systems, 122–125
 - Users, 118–119
 - SMTP, 115–116
 - smtp.daemon, 126–130
 - Destination, 130
 - General, 127
 - Naming Hosts, 127–128
 - SMTP/TCP, 129–130
 - Status Window, 128–129
 - See also* Mail*Link gateways
- Preferences window (Apple Internet Mail Server), 157–159
 - illustrated, 158
 - Outgoing SMTP connections field, 159
 - Password Server pane, 158–159
 - POP3 Server pane, 158–159
 - Server names field, 158
 - SMTP Server pane, 158–159
- Preferences window (FileMaker Pro Server), 197–198
- Preferences window (ListSTAR)
 - About ListSTAR, 88
 - Enclosures, 83–85
 - General, 82–83
 - Naming Hosts, 81–82
 - SMTP/TCP, 86–87
 - Status Window, 85–86
- Preferences window (Mosaic), 353
- pre-indexed CD-ROMs, 280
- Print dialog box (Acrobat), 341–342

- Printer Bridge, 20, 53–58
 - Address button, 58
 - configuring, 54
 - defined, 20
 - Extra Pages button, 56
 - number per server, 53
 - Select Letterhead window side, 54
 - sending out of house and, 58
 - Set margins window side, 54
 - user accounts, 58
 - user accounts address information, 58
 - See also* QuickMail
- printers, 361–370
 - AppleTalk printer drivers and, 369
 - checklist for selecting, 370
 - color, 364–365
 - consumables and, 368–369
 - font technology, 366–367
 - network access of, 362
 - paper sizes, 365
 - PC users and, 367
 - performance, 367–368
 - RAM and, 368
 - resolution of, 362–363
 - saturated pages and, 365–366
 - specialized, 351
 - special papers and, 366
 - support, 368
 - throughput times, 446
 - toner-printing drum, 369
 - user access to, 362
 - VPT, 367
- printing performance, 446–447
 - LAN throughput times, 446
 - WAN throughput times, 447
 - See also* service performance
- Program Linking
 - AppleSearch, 271
 - Butler SQL and, 216–217
 - capability, 169
 - Cumulus, 176
 - enabling, 170, 216–217
 - in guest account, 217–218
 - in user account, 217
 - Informed, 244
 - server selection via, 271
 - WebSTAR, 310–311
- Pub folder, 313–314, 317, 327
- Public Contact Server
 - Configuration screen, 412
 - defined, 409
 - public category installation, 420–425
 - See also* Contact
- Public Event Server
 - Configuration screen, 411
 - configuring, 412–420
 - defined, 409
 - importing file into, 415–416
 - logging into, 416
 - mistakes and, 416
 - name for, 417
 - public category installation, 419–420
 - rooms, resources, 414–415
 - Status window, 412
 - users, 412–414
 - users list, reformatting, 413–414
 - See also* Up-to-Date

Q

- QM Administrator, 20
 - Administrator option, 22
 - Connection option, 22
 - Dead Mail option, 23
 - defined, 20
 - Disk Space option, 23
 - file-based MailCenter from, 50
 - File-Based Server option, 23
 - MailCenter display, 23, 28
 - Memory option, 23
 - NameServer option, 23
 - Next Connection display, 22
 - QM Server option, 23
 - Status display, 24
 - Status option, 22
 - Time option, 23
 - Type display, 23
 - Urgent display, 24
 - using, 22–30
 - Waiting display, 24

- window illustration, 22
 - See also* QuickMail
- QM Administrator Configuration window (QuickMail), 48–49
- QM-AOL Gateway, 21
- QM-Direct Bridge, 20, 58–61
 - agreements before using, 59
 - configuring, 60
 - Custodian, 68–69
 - defined, 20, 58
 - maximum number per server, 59
 - password, 59
 - security, 61
 - setting up, 59–60
 - star configuration, 60
 - user accounts, 60
 - See also* Bridges (QuickMail); MailCenters
- QM-Link Gateway, 21
- QM-QM Bridge, 21, 61–69
 - Create button, 62
 - defined, 21, 61
 - modem information, 65
 - outgoing mail password, 64
 - password, 62
 - phone book, 62–63, 66
 - Preferences window, 64
 - setting up, 61–62
 - Setup window, 65–66
 - user list, 66
 - window, 62
 - See also* bridges (QuickMail); QuickMail
- QM-QM menu (QuickMail)
 - Export Phone Book command, 66
 - Import Phone Book command, 67
 - Preferences command, 64
 - Setup command, 65
- QM Remote, 74–76
 - defined, 74
 - mail status checking, 76
 - Main window, 75
 - receiving mail with, 75
 - using, 75–76
- QM-Script Gateway, 13, 21, 69–74
 - configuring, 70
 - Custodian, 73
 - defined, 21, 69
 - Display connection checkbox, 71
 - Include QuickMail headers option, 71
 - maximum MailCenters per server, 70
 - Multiple Enclosure Uploading option, 70
 - newly created account, 72
 - password, 70
 - Script option, 71
 - security, 74
 - setting up, 70
 - Strip 8th (high) bit option, 71–72
 - user account names in, 73
 - user list, 72
 - window, 70
 - See also* gateways (QuickMail); MailCenters; QuickMail
- QM Server control panel, 18
 - illustrated, 18
 - in message path, 45, 46–47, 50
 - setting up, 439
 - See also* QuickMail
- Quarterdeck Mail, 109
- query builder documents, 237
- Queued Messages window (Mail*Link), 136
 - Delete button, 137
 - Finished button, 137
 - Get Info button, 136
 - illustrated, 136
 - Return button, 137
- QuickDNS Pro, 296–300
 - activity display, 300
 - Admin application, 296
 - A resource record creation, 298
 - CNAME resource record creation, 299
 - completed resource records, 299
 - defined, 285
 - folder, 297
 - MX resource record creation, 298
 - NS resource record creation, 297
 - requirements for using, 296
 - resource record creation, 297
 - reverse domain mappings, 300
 - Server, 300
 - SOA values in, 296
 - See also* DNS

- QuickDraw, 368, 369
 - QuickMail, 1–76
 - administrators, 25–26
 - costs, 8
 - custodians, 25, 26
 - efficiency, 436–437
 - over LANs, 437
 - over WANs, 438
 - forms, 37
 - functioning of, 43–48
 - MailCenters, 20–21
 - Memo form, 38, 57
 - messages. *See* messages (QuickMail)
 - NameServer, 19
 - network bandwidth, 436
 - network design, 8
 - Notifier, 51
 - over ARA, 74
 - passwords, 25
 - Personal Gateway, 11
 - PowerTalk with, 11
 - Printer Bridge, 53–58
 - QM Administrator, 20
 - QM-AOL Gateway MailCenter, 21
 - QM-Direct Bridge MailCenter, 20, 58–61
 - QM-Link Gateway MailCenter, 14, 21
 - QM-Printer Bridge MailCenter, 20
 - QM-QM Bridge MailCenter, 21, 61–69
 - QM Remote, 74–76
 - QM-Script Gateway MailCenter, 13, 21, 69–74
 - QM Server, 18, 45–47, 50, 439
 - remote connections, 74–76
 - sending mail with, 37–42
 - server
 - advanced setup, 48–76
 - dial-in, 76
 - setup, 18–22
 - signing into, 434–435
 - testing, 437
 - WAN performance of, 445
 - See also* LAN client/server e-mail systems
 - QuickMail Administrator Manual*, 52, 125
 - QuickMail client, 31–42
 - DOS, 32–33
 - Macintosh, 33–42
 - Windows, 31–32
 - QuickMail for DOS, 49
 - QuickMail for OS/2, 49
 - QuickMail for Windows, 49–51
- ## R
- RAIC (Redundant Array of Independent Computers), 329
 - RAID5 storage, 401, 402
 - RAIDbank, 402
 - RAM
 - backup servers and, 401
 - Cumulus*Client and, 181
 - Cumulus and, 181
 - database setup and, 191
 - FileMaker Pro allocation, 195–196
 - printer, 368
 - Web server performance and, 451
 - readers, 332
 - realms, 323–325
 - defined, 323
 - grouping by name, 324
 - passwords, 324
 - user names, 324
 - See also* WebSTAR
 - Realms window (WebSTAR), 324
 - Receiving Configuration window (Mail*Link), 126
 - Record menu (Cumulus), Catalog command, 174, 178
 - recursion, 288
 - recursive name service, 289
 - recursive queries, 288
 - Registered Users window (Mail*Link), 131–132
 - Delete button, 132
 - Edit button, 132
 - Export button, 132
 - illustrated, 131
 - Import button, 131
 - New button, 131
 - Purge button, 131
 - Sort by pop-up menu, 132
 - relational databases, 186–187
 - advantages of, 187


- data files and, 189
 - defined, 186
 - development time, 186
 - illustrated, 187
 - many-to-one relationship, 186
 - planning, 189–190
 - See also* databases
 - Remark, 335
 - remote connections (QuickMail), 74–76
 - over ARA, 74
 - receiving mail with, 75
 - See also* QM Remote
 - remote information sources, 267–270
 - Remote Receiving preferences (Mail*Link Remote), 125–126
 - illustrated, 125
 - New button, 125
 - when to use, 125
 - See also* Mail*Link gateways
 - Remote Systems preferences (Mail*Link Remote), 122–125
 - illustrated, 123
 - New button, 123
 - Open button, 123
 - See also* Mail*Link gateways
 - reporters
 - defined, 259
 - server and client, 260
 - See also* AppleSearch
 - reserved addresses, 98–99
 - resolution, 362–363
 - explained, 363
 - helpers, 363
 - RAM and, 363
 - See also* printers
 - resource records, 290–295
 - A record, 291–292
 - CNAME record, 293–294
 - completed, 299
 - creating, 297
 - HINFO record, 293
 - MX record, 292–293
 - NS record, 292
 - PTR record, 294–295
 - SOA record, 290–291
 - TXT record, 294
 - See also* DNS
 - Retrospect, 397
 - archive scripts, 405–406
 - backup planner, creating, 399–401
 - backup scripts, 402–404
 - Backup Server, setting up, 401–404
 - duplicate scripts, 404
 - See also* backup management
 - reverse domain mappings
 - defined, 294
 - illustrated, 300
 - Round Robin CGI, 329
 - rules
 - defined, 91
 - See also* ListSTAR
- ## S
- saturation printing, 365–366
 - Schema menu (Butler SQL)
 - New Column command, 226
 - New Table command, 226
 - Script Editor, 102
 - Scripts menu (Informed Manager), 253
 - SCSI Manager 4.3, 200
 - search requests
 - entering, 273
 - reading, 274
 - returned, 273
 - saving, 274
 - scheduling, 275
 - See also* AppleSearch
 - security
 - CGIs and, 211
 - electronic bulletin board system, 6
 - LAN client/server e-mail system, 9
 - LAN/WAN e-mail system, 15
 - peer-to-peer e-mail system, 11
 - QM-Direct Bridge, 61
 - QM-Script Gateway, 74
 - signature, 250
 - WebSTAR, 309, 325–327
 - Security Toolkit, 308
 - Select MailCenter window (QuickMail), 434

- Sending Setup window (Apple Internet Mail Server), 163
- Server Manager, 199, 310
- Server menu (Apple Internet Mail Server)
 - Account Information command, 159
 - Connection Statistics command, 164
 - Preferences command, 157
 - Sending Setup command, 163
- Server menu (AppleSearch)
 - Information Sources command, 269
 - Preferences command, 265, 270
 - Sources command, 263
- Server Preferences window (AppleSearch), 266, 270
- servers. *See* backup servers; building servers; file servers; list servers; Web servers
- Server Tools, 393
- service performance, xii, 431–458
 - backup servers, 455
 - databases, 452–454
 - efficient, 433
 - e-mail, 434–439
 - Ethernet and, 456
 - file sharing, 440–445
 - printing, 446–447
 - results analysis, 432
 - test machines, 432
 - test parameters, 432
 - tests, 431–432
 - Web servers, 448–451
- services (ListSTAR), 89–95
 - account establishment, 97–98
 - Administration, 93
 - default, 89–93
 - demo, 96–97
 - Mailer, 89
 - Postmaster, 93–95
 - Timer, 89
 - See also* ListSTAR
- Setup window (FTPd), 31
- shared server folders, 261
- shared volumes, 372
- Sharing Setup control panel, 216–217, 271
- signatures, 248–250
 - electronic, 248
 - icon variations, 250
 - multiple, 249
 - reasons for, 248
 - security and, 250
 - signers, 248–249
 - valid, 252
 - verification of, 248, 250, 252
 - See also* forms; PowerTalk
- sign-in, 434–435
 - comparison, 436
 - performance and, 434–435
- Skyline Satellite, 432
- smtp.daemon preferences, 126–130
 - Destination, 130
 - General, 127
 - Naming Hosts, 127–128
 - SMTP/TCP, 129–130
 - Status Window, 128–129
 - See also* Mail*Link gateways
- SMTP Status window (Mail*Link), 128
- SMTP/TCP preferences (ListSTAR), 86–87
 - # of Incoming Connections option, 86
 - illustrated, 86
 - NameResolver pop-up menu, 87
 - Ping before opening connection option, 80, 87
 - TCP Command Timeout pop-up menu, 86
 - See also* Preferences window (ListSTAR)
- SMTP/TCP preferences (smtp.daemon), 129–130
 - # of Incoming Connections pop-up menu, 129
 - illustrated, 129
 - NameResolver pop-up menu, 130
 - Ping before opening connection checkbox, 130
 - TCP Command Timeout pop-up menu, 129
 - See also* Mail*Link gateways
- Source, 4
- SQL
 - defined, 213
 - sample program, 234
 - See also* Butler SQL
- StarNine Resolver, 130
- Start of Authority (SOA) record, 290–291
 - Domain Name, 290
 - Expire, 291
 - Minimum Time to Live, 291
 - MName, 290
 - QuickDNS Pro values, 296

- Refresh, 291
 - Retry, 291
 - RName, 290
 - Serial, 290–291
 - See also* resource records
 - Status Window preferences (ListSTAR), 85–86
 - Display ADMIN Status Window checkbox, 86
 - illustrated, 85
 - Limit mail count to a max of option, 86
 - See also* Preferences window (ListSTAR)
 - Status Window preferences (smtp.daemon), 128–129
 - Display SMTP Status Window, 128
 - illustrated, 128
 - Limit mail count to max of field, 129
 - See also* Mail*Link gateways
 - storage evaluation, 394–395
 - Store & Forward MailCenter, 21, 46, 47
 - D file, 47
 - In Mail folder, 47
 - Structured Query Language. *See* SQL
 - StuffIt Expander, 84
 - suffix mappings, 320
 - Suffix Mapping window (WebSTAR), 322
 - Summary window (FTPd), 317
 - Sysop, 5
 - System-to-Call Configuration window (Mail*Link), 123–125
 - Call this system every... pane, 124
 - Connection button, 125
 - Enable outgoing calls checkbox, 124
 - illustrated, 123
 - Login field, 124
 - Login script button, 124
 - System Name field, 123–124
- T**
- Tango, 214, 216, 235–237
 - ACGI, 236
 - configuring, 235
 - defined, 235
 - Editor, 236, 237
 - folder, 236
 - ODBC support, 235
 - query builder documents, 237
 - Test Drive, 237
 - using, 235–237
 - See also* Butler SQL
 - TCP/IP control panel, 113
 - Technical Information Library*, 269
 - technical support numbers (ListSTAR), 88
 - tests, 431–432
 - database performance, 453
 - list servers, 108
 - machines, 432
 - parameters, 432
 - self, 457–458
 - See also* service performance
 - Text (TXT) record, 294
 - TEXT files, 135
 - threads, 5
 - concurrent, 129
 - main, 129
 - Timbuktu Pro, 310
 - AppleTalk and, 312
 - Control function, 312
 - exchanging files with, 311
 - passwords, 312
 - TCP/IP and, 312, 313
 - Timer services (ListSTAR), 89
 - tracking, 251, 254, 255
 - TurboGopher, 317
- U**
- Universal Resource Locators. *See* URLs
 - UNIX mail systems, 119
 - “Unnecessary but Mandatory Recriminations Cycle,” 254
 - UPS (Uninterruptible Power Supply), 385
 - Up-to-Date, 407–430
 - calendar distribution files, 429
 - Calendar window, 408
 - client installation, 426
 - client setup, 426–430
 - customizing, 426
 - defined, 408
 - distributing, 426
 - distribution files, 429

- file creation, 426–430
 - group scheduling abilities, 420
 - initial configuration, 411–412
 - installing, 409
 - launching, 417
 - password protection, 420
 - personal category, 419
 - preferences file, 429–430
 - public category, 419–420
 - Public Event Server, 409, 411–420
 - recurring maintenance reminder, 429
 - requirements for using, 409
 - server location, 409
 - setting up, 409–425
 - using, 429
 - See also* Contact
 - URLs, 303, 326
 - user account planning form, 386
 - group information, 388
 - illustrated, 386
 - name, action, authorization, 387
 - network information, 387–388
 - removal, 387
 - training, 387
 - user accounts (Mail*Link), 131–135
 - automatic setup, 132–134
 - MailCenter users list, 134–135
 - manual setup, 131–132
 - See also* Mail*Link gateways
 - User dialog box (QuickMail), 27
 - user list, 28
 - adding accounts to, 49
 - deleting, 48
 - Mail*Link, 134–135
 - QM-QM Bridge, 66
 - QM-Script Gateway, 72
 - See also* MailCenters
 - User menu (QuickMail)
 - Create command, 27, 58, 65, 72, 134
 - Import command, 135
 - Users & Groups, 217
 - access privileges, 391–392
 - basic information, 389–390
 - folder hierarchy, 391–392
 - planning, 389–392
 - system information, 390–391
 - See also* Apple Workgroup Server
 - users
 - account review, 393–394
 - Apple Workgroup Server access planning, 386–388
 - availability of, 418
 - defined, 372
 - deleting, 413
 - edited list of, 413
 - editing information, 418
 - imported, 417
 - list reformatting, 413–414
 - locating (QuickMail), 40–41
 - scheduling, 407–430
 - text list generated AppleShare, 413
 - Users preferences (Mail*Link Remote), 118–119
 - Administrator's Address, 118
 - Database button, 119, 131
 - illustrated, 118
 - Registered Users Database, 119
 - Send "Host Down" Messages, 119
 - See also* Mail*Link gateways
 - Users window (Butler SQL Server), 218
 - Utilities menu (QuickMail), Configure File-based command, 52
 - UUencode format, 84, 85, 121, 122
- ## V
- viewers, 332
 - View menu (Informed Manager), Tracking Status command, 253
 - ViewPoint, 393–394
 - defined, 393
 - privileges, 393
 - user privileges, 394
 - Views control panel, 458
 - Virtual Printer Technology (VPT) printers, 367
- ## W
- WAIS, 267–270
 - AppleSearch gateway, 268
 - choosing sources, 269

- custom database file, 283
 - database representation, 270
 - defined, 267
 - directory, 268
 - queries, 268
 - servers, 268
 - site scanning times, 270
 - sources, 268
 - unpublished sites, 282–283
- WANs
 - AppleShare transmission speeds, 445
 - file sharing over, 444–445
 - logs and, 436
 - performance, 438
 - Personal File Sharing transmission speeds, 444
 - printer throughput times, 447
 - QuickMail efficiency over, 438
 - QuickMail performance, 445
 - Web server throughput, 450
 - See also* LANs
- WebLink, 359
- WebLink Preferences window (Acrobat), 359
- WebMap, 321, 322
- Webmaster Macintosh*, 330
- Web server performance, 448–451
 - CGIs and, 450
 - data transfer slowdown problems, 450–451
 - LAN throughput, 448
 - RAM and, 451
 - WAN throughput, 450
 - See also* service performance
- Web servers, 303
- WebSTAR, 77, 203, 303–330
 - add-on products, 308
 - Admin application, 309, 329
 - main window, 318
 - running remotely, 311
 - Timeout counter, 319
 - aliases, 326
 - allow and deny, 325–326
 - basics, 308–317
 - browser interaction with, 236
 - Butler SQL integration, 235
 - defined, 303
 - file access with, 309
 - firewalls and, 323
 - folders, 304, 309
 - alias, 309
 - setting up, 309
 - structure, 309
 - free memory, 319
 - index file, 309
 - installing, 308
 - internal functions, 318
 - multiple copies of, 327–328
 - names and, 323
 - port numbers, 328
 - Program Linking, 310–311
 - realms, 323–325
 - remote management, 310–313
 - security, 309, 325–327
 - Security Toolkit, 308
 - serial numbers, 328
 - server logon, 308
 - Status window, 319, 329
 - suffix mappings and, 320
 - throughput rules and, 328
- Wide Area Information Servers. *See* WAIS
- wildcard character (*), 163
- Windows client (QuickMail), 31–32
 - CE.INI, 31
 - CONFIG.EXE, 31
 - connection screen, 32
 - QMAILWIN.EXE, 31
 - QMNOTE.EXE, 31
 - See also* QuickMail
- Windows menu (ListSTAR), General Preference command, 81
- World Wide Web (WWW)
 - Adobe Acrobat on, 351–359
 - AppleSearch via, 276–279
 - CGIs, 188
 - RAIC, 329
 - See also* Web server performance; Web servers



AppleTalk Network Services

by Dorian J. Cougias, Tom Dell,
& E.L. Heiberger



Developers of



Apple Certified
Server Engineer
Program

This is the fourth book in the Network Frontiers Field Manual Series. This series, written by the developers of the Apple Certified Server Engineer Program, is intended for the Macintosh networking professional and each book is designed to help you *do* the business of networking. No hype. No product marketing. Just the plain and simple way to build and manage real-world networks.

AppleTalk Network Services covers the most commonly requested services running on AppleTalk and TCP/IP networks. AppleShare, AppleSearch, WebSTAR, ListSTAR, PowerTalk, MailLink, QuickMail, Acrobat, Informed, and QuickDNS are just a few of the services covered. In conclusion, the effects that each service has on your network utilization are analyzed in order to maximize your network efficiency.

About the Authors:

Dorian J. Cougias is the founder of Network Frontiers, Inc., a leader in Macintosh connectivity. They develop course material and implement training programs for all network support staff. They also offer the AppleTalk Network Design, Management, and Troubleshooting Series to dealers, network administrators, VARs, and trainers at various Apple Market Centers across the country, as well as on-site at their larger clients.

Tom Dell is a journalist and consultant who has been involved in the use of print production networks for more than 10 years. He was Senior Editor for the Macintosh programming and training firm of Bear River Associates, and at Network Frontiers has applied his experience to the needs of such clients as Apple Computer, Inc., Regis McKenna Public Relations, and Ketchum Advertising.

E. L. Heiberger, a programmer for both the Macintosh and the PC, was project manager at Bear River Associates for the network technology division, where she laid the groundwork for the national training program. Today she continues to combine her knowledge of marketing, business management, and technology at Network Frontiers as Managing Partner. She coordinates all facets of Network Frontiers' training program and oversees the AppleTalk Certified Server Engineer Program. As a member of BICSI, she teaches design of cabling schemes for CAT5 networks and consults extensively for clients in need of cable design specifications.

System Requirements:

Any Macintosh or compatible with CD-ROM drive, running Macintosh System 7 or later.

Skill Level:

Intermediate. Basic Macintosh and networking experience is assumed.



clockwise from top left:
Lynn Heiberger,
Tom Dell, Cass Kovel,
and Dorian Cougias



Includes free
network tools



ISBN 0-12-192570-6

> \$34.95

315598

299