



Analyzing Sun Networks

江苏工业学院图书馆
藏书章

Carl Malamud



VAN NOSTRAND REINHOLD
New York

Copyright © 1992 by Carl Malamud

Library of Congress Catalog Card Number 91-462

ISBN 0-442-00366-8

All rights reserved. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems—without written permission of the publisher.

Printed in the United States of America

Van Nostrand Reinhold
115 Fifth Avenue
New York, New York 10003

Chapman and Hall
2-6 Boundary Row
London, SE1 8HN, England

Thomas Nelson Australia
102 Dodds Street
South Melbourne 3205
Victoria, Australia

Nelson Canada
1120 Birchmount Road
Scarborough, Ontario M1K 5G4, Canada

16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1

Cover based on an original perspective drawing by Jan Vredeman de Vries (1527–1604), reprinted in *Perspective*, Dover Pictorial Archive Series (Dover Publications, 1968). Color was added without the participation of the original artist. Interior design by Carl Malamud.

Library of Congress Cataloging-in-Publication Data

Malamud, Carl, 1959—

Analyzing Sun Networks / Carl Malamud.

564 p. cm.

Includes index.

ISBN 0-442-00375-7

1. Local area networks (Computer networks) 2. TCP/IP 3. Computer network architectures. 4. OSI (International standards) I. Title.

TK5105.7.M35 1990

621.39'81—dc20

91-462
CIP



Preface

Sun Microsystems has based their entire business on the concept of open systems. The original Sun workstation was built out of off-the-shelf parts. Using standard components to build state-of-the-art workstations and software is a policy that continues to this day.

This philosophy extends into networks. Sun networks are built on the industry standard TCP/IP protocols, the same protocols used for the international Internet network. This book shows how TCP/IP networks can be used to build heterogeneous, fast, and big networks.

TCP/IP is only the beginning for a network. Also explained are the Network File System, RPC, and all the other protocols that are part of the Open Networking Community (ONC) family. Because of the widespread acceptance of ONC—over one million computers on every major operating system—this book is about much more than just Sun networks; it is a book about any computers that use TCP/IP and NFS.

In addition to the industry standards, also included are network management, security, interoperability, and other topics that are vital to managing networks. We will see how OSI and SNA are built into the network and how this heterogeneous computing environment is managed.

Building viable network architectures is the theme of this book. Throughout the text, case studies from Sun's own network and other leading sites are used. One of the pleasures in writing this book has been working with technology that exists and is being used all over the world.

This is the last volume of a three-volume series on *Analyzing Networks*. For information on OSI-based networks, the reader is directed to *Analyzing DECnet/OSI Phase V*. For information on proprietary networks, the reader is directed to *Analyzing Novell Networks*. For information on heterogeneous networks that are operational today, turn the page.

Carl Malamud
carl@malamud.com

A Note on Trademarks

PostScript is a trademark of Adobe Systems.

Apple, AppleTalk, LocalTalk, and Macintosh are trademarks of Apple Computer.

dbase is a trademark of Ashton-Tate.

Ingres is a trademark of ASK, Inc.

STREAMS, Transport Layer Interface, Unix, and Unix System V Release 4 are trademarks of AT&T.

Auspex is a trademark of Auspex.

CRAY is the registered trademark of Cray Research, Inc.

Cullinet and IDMS are trademarks of Cullinet.

DEC, DECconnect, DECMcc, DECnet, DNA, Message Router, VAXcluster, and VMS are trademarks of Digital Equipment Corporation.

Epoch and InfiniteStorage are trademarks of Epoch.

AS/400, DISOSS, IBM, IBM PC LAN, PC/AT, PC/XT, PROFS, SNA, SNADS, System/370, System/38, and 3270 Display Station are trademarks of International Business Machines, Inc.

MS-DOS and Microsoft are trademarks of Microsoft.

Network General and Sniffer Analyzer are trademarks of Network General Corporation.

HYPERchannel is a trademark of Network Systems Corporation.

Network File System, NFS, Open Network Computing, ONC, SPARC, Sun, SunOS, and TOPS are trademarks of Sun Microsystems, Inc.

Sybase is a trademark of Sybase.

Carl Malamud is a trademark of Carl Malamud.



Acknowledgments

Geoffrey Baehr, Director of Networking and Data Communications at Sun Microsystems, was instrumental in making the idea for a book on *Analyzing Sun Networks* become a reality. After previous encounters with large computer companies, I was delighted to find in Sun Microsystems a focus on technical excellence and open, nonproprietary systems. Special thanks goes to Terri Davis who provided assistance in ways too numerous to mention.

Reviews of the manuscript and answers to many, many questions were provided by the following engineers at Sun Microsystems: Sally Ahnger, Karl Auerbach, Gigi Babcock, Teresa Barr Beyer, Gabriele Cressman-Hirl, Linda Cwiak, Frank DeMarco, Jonathan Feiber, Dennis Freeman, Burt Fujii, Terence Gibson, Eric Johnson, Tom Kessler, Abhijit Khale, David Kipping, Fred Lowe, Karen Maleski, Milton Mallory, Mike Martinez, Dennis McLain, Chuck McManis, Bill Melohn, Bob Plummer, Ian Pope, Vipin Samar, Warren Smith, Warren Smith, Mark Stein, Richard Thio, Ian Vessey, Keith White, Yinpo Wong, and Dennis Yaro.

Bruce Nelson of Auspex, Joyce K. Reynolds of ISI, William T.C. Kramer and Louise D. Kokinakis of the NASA Ames Research Center, Jay Dombrowski of the San Diego Supercomputer Center, Steve Crocker and Hilarie K. Orman of Trusted Information Systems, Robert Green of Unidata, and Dr. David Mills of the University of Delaware all provided valuable assistance.

Dr. Harry Saal of Network General once again let me use a Sniffer Network Analyzer to help bring protocols to life and provided a valuable review of the manuscript.

Brian Pawlowski of Sun Microsystems, provided an in-depth technical review of the entire manuscript. His attention to detail and his technical expertise contributed immensely to the final product. He asked me to add that he is “a genuinely warm and funny guy.” I would concur.

Dianne Littwin of Van Nostrand Reinhold has been the editor for all three volumes in this series. I'd like to express my sincere thanks to her for all her hard work and incredible patience. She deserves a vacation.



Contents

Preface	xi
A Note on Trademarks	xiii
Acknowledgments	xv
Chapter 1—Introduction	3
Overview	3
Chapter 2—Subnetworks	9
Introduction	9
Ethernet	12
IEEE LANs	15
Bridging Ethernets	21
FDDI and Token Ring	24
Token Ring Support	31
Wide-Area Links	31
SWAN: An Example	36
On the Internet	42
Packet Radio	53
For Further Reading	57
Chapter 3—Networks	61
Networks	61
The Internet Protocol	62
ISO CLNS and TCP's IP	68
Support Protocols: ARP and ICMP	69
ICMP	70

TCP	74
Connections and Options	84
TCP and OSI	88
UDP	90
RIP	90
OSPF and IS-IS	97
Exterior Gateway Protocols	100
Routing Between NSFnet and Milnet	105
TCP on a Sun	106
Starting a Connection	111
For Further Reading	112
 Chapter 4—STREAMS	 117
Protocol Stacks and the Operating System	117
Character I/O Mechanisms	118
The STREAMS Multiplexer	123
Transport Level Interface	123
For Further Reading	125
 Chapter 5—RPC and XDR	 129
Why an RPC?	129
The Portmapper	134
Asynchronous Mechanisms	140
The XDR Library Package	141
For Further Reading	148
 Chapter 6—Network File System	 151
NFS	151
The NFS Protocol	157
The NFS Implementation	163
NFS on the MVS Operating System	169
Hardware NFS Implementations	173
For Further Reading	177
 Chapter 7—RPC Tool	 181
RPC Compilers	181
Netwise RPC Tool	182

Procedure Declarations	185
Process Binding	187
The Client Stub	188
Server Control Procedure	189
Custom Procedures	192
ISO and Sun RPC Mechanisms	193
Other Interfaces to the Network	193
For Further Reading	196
 Chapter 8—Network Time Protocol	 199
Why Bother?	199
NTP	200
Associations	204
NTP Protocol	206
Handling Warped Servers	208
Implementation in the Internet	212
For Further Information	213
 Chapter 9—Names	 217
The New-Name Naming Service: Yellow Pages NIS	217
NIS	218
Typical NIS Operations	226
NIS+	226
Security in NIS+	235
Domain Name System	239
Servers	240
Resource Records	245
DNS Operation	249
Combining Name Services	256
For Further Reading	256
 Chapter 10—Security	 261
Secure Networks?	261
Basic Security	262
Security and RPC	263
Secure RPC	264
Kerberos	270

Administering Security	273
For Further Reading	277
Chapter 11—The Automounter	281
The Problem with /etc/fstab	281
Automounter Overview	282
Emulating an NFS Server	284
Maps	285
Multiple Mounts and Multiple Locations	287
Built-In Maps	291
When to Use the Automounter	292
For Further Reading	292
Chapter 12—Distributed Applications	295
Overview	295
License Service	296
License Service Operation	301
Network Lock Manager Protocol	303
REX	309
The R* Commands	313
The PC	315
For Further Reading	322
Chapter 13—Mail	325
Overview	325
SMTP	325
RFC 822	331
Using DNS to Forward Messages	338
UUCP	339
Munging Messages	341
Case Study: Sun's Mail System	342
X.400	347
IPMS	350
Message Transfer Agents	353
Distribution Lists	360
Message Stores	363
Gateways	364

Munging to X.400	365
For Further Reading	367
 Chapter 14—Files, Printers, and Terminals	 371
Overview	371
Telnet	371
The File Transfer Protocol	377
FTAM	386
Resource Location	393
Trivial FTP	396
Miscellaneous TCP/IP Services	398
The Line Printer Daemon	400
For Further Reading	404
 Chapter 15—Management	 407
Managing a Network	407
Structure of Management Information	409
Names	409
SNMP	412
SunNet Manager	415
The Management Database	418
The Discover Tool	421
SNMP Support	423
Operation of SunNet Manager	426
Basic Management Tools	429
For Further Reading	431
 Chapter 16—Interoperability	 435
Overview	435
The Special Case: SNA	436
SNA Case Studies	442
The Rest of the World	452
 Glossary	 461
 Index	 563

Introduction

Introduction

Overview

For some, the network is the equivalent to a freeway. It's there, you have to use it, so you learn one or two exits and shuttle back and forth to work every day. For more inquisitive types, the freeway is just the backbone with a wide variety of different exits, each providing different services.

For this inquisitive reader, this book is intended to provide a road map. We assume the reader has some interest in knowing how a network based on the Transmission Control Protocol and the Internet Protocol (TCP/IP) operates.

TCP/IP is much broader than just Sun workstations, so the title of this book is a bit of a misnomer. The first portion of the book applies to any computer which supports TCP/IP, which is any computer still operating that doesn't use vacuum tubes.

We will see how TCP/IP can be used to integrate a variety of networks together into a coherent whole. Each network may be composed of a variety of different links—FDDI, Ethernet, X.25, ISDN, for example. These different data links are examined in more detail in Chapter 2. (The reader may also consult the glossary for quick explanations of acronyms.) Chapter 3 then shows how TCP and IP (and several other protocols) are built on top of the subnetworks, integrating them into internetworks.

Built on top of TCP/IP are a variety of services that were originally developed by Sun Microsystems; they are grouped together under the official name of Open Network Computing (ONC) or the unofficial moniker of the Network File System (NFS), a protocol that allows a file system on a remote computer to appear as a local drive. NFS opens the door for a wide variety of distributed computing, including the diskless workstation.

ONC is much more than NFS, however. It includes a remote procedure call protocol on which several other services are built. The Network Information Service (NIS) is a name service; there is a lock manager and many

other support services. Other applications, such as SunNet Manager or 3270 color screen emulation also use this ONC infrastructure.

ONC is the subject of much of this book. Chapters 5 and 7 discuss the RPC and XDR protocols, Chapter 6 describes NFS. Other ONC applications are also discussed: the Automounter in Chapter 11; the License Server and lock manager in Chapter 12.

Describing ONC and the TCP/IP protocols still doesn't really qualify this as a book on "Sun" networks. Although Sun developed NFS, a liberal licensing policy has seen it migrate to every major operating system.

This book is thus really about building heterogeneous networks with a slant on Sun workstations. Most of the technology described, however, has applicability to a wide selection of operating systems and hardware platforms.

The emphasis is on understanding how this heterogeneous network, of which the Sun is one type of server or workstation, works. A wide variety of topics are examined, organized in a way that attempts to make some sense out of the wide range of options that are available. The choice of topics is, of course, arbitrary. Not every issue is covered, and we have not attempted to deal with any issues in a definitive fashion. Points to sources for further reading are provided for people who want to implement software, tune systems, or do any of the other hands-on tasks of managing a network.

A fairly detailed overview does help, however, even if you will have to go chasing after some manual when a problem occurs. If you understand what pieces are involved, you have a valuable sense of perspective on the details listed in a reference manual or protocol specification. Throughout the book, illustrations bring concepts to life: network configurations, network analyzer screen dumps, lists of system calls, or reference tables. One reader termed this information "raw data," which this author has chosen to interpret in the good sense of the term: concrete examples that illustrate theory.

This raw data is presented to the reader because it gives an understanding of the capabilities of a given protocol or module. Looking at dumps from a network analyzer is useful for debugging error situations, but it is also useful as a learning tool to see how a packet is constructed. Knowing what system calls are available, or what management information is kept by a module, is equally instructive. Of course, all this raw data is supplemented by a healthy dose of commentary. A book of all pictures would be nice, but the picture book of network analysis will have to wait for another time.

The Sun Wide-Area Network (SWAN) is used throughout this book as a case study. In addition to the case study material from Sun, material from the San Diego Supercomputer Center, NASA-Ames, and the Internet is in-

PID	TT	STAT	TIME	COMMAND
0	?	D	0:00	swapper
1	?	IW	0:00	/sbin/init -s
2	?	D	0:00	pagedaemon
55	?	S	1:02	portmap
58	?	IW	0:00	ypbind
60	?	IW	0:00	keyserv
68	?	S	18:51	in.routed
71	?	I	0:00	(biode)
72	?	I	0:00	(biode)
73	?	I	0:00	(biode)
74	?	I	0:00	(biode)
84	?	IW	0:01	syslogd
92	?	IW	0:02	/usr/lib/sendmail -bd -qlh
96	?	IW	0:00	rpc.statd
98	?	IW	0:00	rpc.lockd
105	?	S	0:00	automount
112	?	S	10:46	update
115	?	IW	0:00	cron
120	?	IW	0:16	inetd
123	?	IW	0:00	/usr/lib/lpd
253	?	IW	0:00	rpc.rquotad
383	co	S	0:04	-csh (csh)
519	co	R	0:00	ps -ax

1-1 Processes on a Typical Workstation

cluded. Supplementing the case studies are a series of screen dumps taken from a Network General Sniffer Analyzer. This device is a modified PC that can be attached to token ring, Ethernet, ARCnet, and other networks to filter and capture data.

This captured data can then be analyzed. The Sniffer Analyzer is used because the information that is captured is displayed in a format that, given the technical material, is remarkably close to clear English.

For some readers, this book will provide all they need, a conceptual overview of how networking works in an environment with TCP/IP, NFS, or other protocols. Most readers, however, will probably have some contact with a Sun workstation or a TCP/IP network, or they wouldn't be reading this book. For those readers, this book is a supplement to the raw data of manuals, functional specifications, hardware diagrams, and all the other data that accompanies the computer. One nice thing about Sun is that much of this data is no longer shipped as large walls of documentation but is available electronically.

This book provides the overview. If the reader is a user of the network, the overview might help understand what is happening when a remote file system is mounted with NFS or why the license server is not letting you proceed with your application. Others will be interested in writing applications for use on the network. It is important to understand what tools are available so that an appropriate choice can be made. Using existing services, such as using NIS to register the location of a database server, means that the programmer doesn't have to reinvent the wheel.

Users and software developers are two classes that may find this information useful, but there is also one more class: the people who have to buy, install, and manage networks. For this class of readers, some material may be rudimentary compared to their current level of expertise. An engineer who designs FDDI chips is not going to learn anything about FDDI (or Ethernet) from reading this book but might learn something about message systems.

This is a bootstrap attempt: significantly more than a marketing white paper or the trade press but certainly less than all the primary resource documents. Hoping that human nature holds, the author is counting on the fact that a 1-inch thick book will be a more attractive target than the dozen or so linear feet that make up a decent set of source documents.

In Figure 1-1, the reader will see the first of many screen dumps. This dump is from a Sun workstation and is a listing of currently active processes on the computer. Eleven different processes in that listing are discussed in this book. An interesting test for the reader (and consequently of the author) is to see if, upon completion of the text, he or she can identify which processes are involved in a real network and what the processes do.