MARK BURGESS



Principles of Network and System Administration

WILEY

Principles of Network and System Administration

This page intentionally left blank

Principles of Network and System Administration

Mark Burgess
Oslo College, Norway

JOHN WILEY & SONS, LTD
Chichester · New York · Weinheim · Brisbane · Singapore · Toronto

Copyright © 2000 John Wiley & Sons, Ltd
Baffins Lane. Chichester,
West Sussex PO19 1UD, England

National 01243 779777
International (+44) 1243 779777
e-mail (for orders and customer service enquiries): cs-books@wiley.co.uk
Visit our Home Page on http://www.wiley.co.uk
or http://www.wiley.com

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency, 90 Tottenham Court Road, London, W1P 9HE, UK without the permission in writing of the publisher, with the exception of any material supplied specifically for the purpose of being entered and executed on a computer system for exclusive use by the purchaser of the publication.

Neither the author nor John Wiley & Sons, Ltd accept any responsibility or liability for loss or damage occasioned to any person or property through using the material, instructions, methods or ideas contained herein, or acting or refraining from acting as a result of such use. The author and publisher expressly disclaim all implied warranties, including merchantability or fitness for any particular purpose.

Designations used by companies to distinguish their products are often claimed as trademarks. In all instances where John Wiley & Sons, Ltd is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

Other Wiley Editorial Offices

John Wiley & Sons, Inc., 605 Third Avenue,
New York, NY 10158-0012, USA

Weinheim • Brisbane • Singapore • Toronto

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN 0–471–82303–1

Typeset in ITC Garamond by Kolam Information Services Pvt Ltd, Pondicherry, India
Printed and bound in Great Britain by Biddles Ltd, Guildford and King's Lynn.
This book is printed on acid-free paper responsibly manufactured from sustainable forestry, in which at least two trees are planted for each one used for paper production.

Contents

Preface ix

1 Introduction **1**
1.1 The Scope of System Administration 1
1.2 Is System Administration a Discipline? 2
1.3 A Jigsaw Puzzle 2
1.4 The Goals of System Administration 3
1.5 A Philosophy 3
1.6 The Challenges of System Administration 4
1.7 Common Practice and Good Practice 5
1.8 Bugs 6
1.9 Information Sources for Sysadms 6
Exercises 7

2 The System Components **8**
2.1 What is 'The System'? 8
2.2 Operating Systems 9
2.3 File Systems 16
2.4 Processes and Job control 32
2.5 Logs and Audits 34
2.6 Privileged Accounts 35
2.7 Hardware Awareness 36
2.8 System Uniformity 38

3 Networked Communities **40**
3.1 Communities 40
3.2 User Sociology 41
3.3 Client-Server Cooperation 42
3.4 Host Identities and Name Services 43
3.5 Common Network Sharing Models 46
3.6 Physical Network 49
3.7 TCP/IP Networks 55
3.8 Network Analysis 62
3.9 Planning Network Resources 70

4 Host Management **78**
4.1 Choices 78
4.2 Start-up and Shutdown 80
4.3 Configuring and Personalizing Workstations 81
4.4 Installation of the Operating System 88
4.5 Software Installation 95
4.6 Installing a Unix Disk 104
4.7 Kernel Customization 106

5 User Management **111**
5.1 User Registration 111
5.2 Account Policy 116
5.3 Login Environment 117
5.4 User Support Services 124
5.5 Controlling User Resources 124
5.6 User Well-being 129

6 Models of Network Administration **134**
6.1 Administration Models 134
6.2 Immunity and Convergence 136
6.3 Network Organization 137
6.4 Bootstrapping Infrastructure 139
6.5 Cfengine: Policy Automation 144
6.6 SNMP Network Management 145
6.7 Integrating Multiple OSes 146
6.8 A Model Checklist 149

7 Configuration and Maintenance **151**
7.1 System Policy 151
7.2 Synchronizing Clocks 153
7.3 Executing Jobs at Regular Times 153
7.4 Automation 155
7.5 Preventative Maintenance 161
7.6 Fault Report and Diagnosis 164
7.7 System Performance Tuning 171

8 Services **181**
8.1 High Level Services 181
8.2 Proxies and Agents 182
8.3 Installing a New Service 183
8.4 Summoning Daemons 183
8.5 Setting up the DNS Name Service 187
8.6 Setting up a WWW Server 203
8.7 E-mail Configuration 215
8.8 Mounting NFS Disks 226
8.9 The Printer Service 229

9 Principles of Security 235
9.1 Physical Security 236
9.2 Four Independent Issues 236
9.3 Trust Relationships 237
9.4 Security Policy 237
9.5 Protecting from Loss 239
9.6 System and Network Security 241
9.7 Social Engineering 243
9.8 TCP/IP Security 244
9.9 Attacks 259

10 Security Implementation 265
10.1 The Recovery Plan 265
10.2 Data Integrity 265
10.3 Analysing Network Security 274
10.4 VPNs: Secure Shell and FreeS/WAN 282
10.5 WWW Security 282
10.6 Firewalls 284
10.7 Intrusion Detection and Forensics 290

11 Analytical System Administration 292
11.1 Science vs Technology 292
11.2 Studying Complex Systems 293
11.3 The Purpose of Observation 295
11.4 Evaluation Methods and Problems 295
11.5 Evaluating a Hierarchical System 297
11.6 Faults 298
11.7 Deterministic and Stochastic Behaviour 315
11.8 Observational Errors 326
11.9 Strategic Analyses 334
11.10 Summary 335

12 Summary and Outlook 337
12.1 The Next Generation Internet Protocol (IPv6) 338
12.2 Never-dos in System Administration 338
12.3 Information Management in the Future 339
12.4 Collaboration with Software Engineering 340
12.5 The Future of System Administration 340

A Summary 342
A.1 Summary of Principles 342
A.2 Summary of Suggestions 346

B Some Useful Unix Commands 349

C Programming and Compiling 355
C.1 Make 355
C.2 Perl 359

C.3 WWW and CGI Programming 377
C.4 PHP and the Web 383
C.5 Cfengine 385

D Glossary 393

E Recommended Reading 397

Bibliography 398

Index 410

Preface

This book has grown out of a one semester course in Network and System Administration which has now run successfully for four years at Oslo College, Norway. The course is an introductory course and involves about 30% theory and 70% practical work [29]. It assumes knowledge equivalent to a typical college course on Operating Systems as well as some basic computer skills. The purpose of the book is to provide a mixture of theory and practice for a course in system administration; to extract those principles and ideas of system administration which do not change on a day-to-day basis; and to present them in a defensible manner [159] to a willing audience.

The need for a book providing an overview of the principles of system administration has been clear for some time, and was amply confirmed by the many enthusiastic messages I received from the system administration community after John Wiley & Sons asked me to write this volume. Finding the right approach for a book is never easy, though; in the system administration profession there is a mixture of personalities: there are hard-nosed pragmatists, for whom theoretical discussions are distasteful, and there are the more scientifically inclined who like to delve into the whys and wherefores. My aim in this book has been to write a guide to system administration which has something for both, i.e. a book with a concrete practical value, but which goes beyond a mere collection of recipes to provide a conceptual overview of the field, suitable for a college level course. The extent to which I have succeeded or not remains to be seen.

In assembling this book, I have reviewed the research work of many authors, most of which has centred around the USENIX organization and its many groundbreaking conferences. Since research in the field is growing, I have also included an overview chapter on methods which I hope will provide a useful reference to existing and potential researchers. In spite of a desire for completeness, I have resisted the temptation to include every possible detail and fact which might be useful in the practical world. Several excellent books already exist which cover this need, and I see no reason to compete with them (see the recommended reading list). System administration is sharply divided by a cultural chasm: Windows and Macintosh Vs. Unix and mainframe. It would have been nice to compare and contrast these systems stringently, but alas, there is no room in a teaching volume for such a digression. I have therefore limited myself to examples of each which are either practical or illustrative. If any operating systems have been unfairly brought into focus, I hope it is only the Free operating systems such as GNU/Linux and the BSD's, from which no one other than their users will benefit.

I note that Unix is a registered trademark, referring to a particular type of Unix-like operating system, whose license is owned currently by Novell. The word Unix is often used as a shorthand, however, referring to any Unix-like operating system, e.g. Solaris, AIX, GNU/Linux.

I would like to offer my special thanks to Tina Darmohray for her comments and encouragement, as well as for allowing me to adapt some firewall examples from her excellent notes. Russ Harvey of the University of California, Riverside also made very positive and useful criticisms of the early materials. Special thanks to Per Steinar Iversen for making detailed comments and constructive criticisms on the manuscript from his near-infinite reservoir of technical expertise. Thanks also to David Kuncicky, Sigmund Straumsnes and Kjetil Sahlberg for their careful readings and suggestions for improvement. Any remaining errors must be entirely someone else's fault (but I haven't figured out who I can blame yet). Thanks to Knut Borge of USIT, University of Oslo, for moderating the course on which this book is based, and for teaching me many important things over the years; also to Tore øfsdahl, Harald Hofsæter, our system administrators at Oslo College who constantly help me in often intangible ways. Sigmund generated the graphs which appear in this volume. In addition to them, Runar Jørgensen and Hårek Haugerud commented on the manuscript. Ketil Danielsen has provided me with both tips and encouragement. Thanks to Greg Smith of the NASA Ames Research Center for performance tips, and to Steve Traugott for discussions on infrastructure. A big hug to Cami Edwards of USENIX for making copies of old LISA proceedings available from the archives. I was shocked to discover just how true is the panel debate: why do we keep reinventing the wheel? I should also like to thank all of the students at Oslo College who have attended my lectures and have inspired me to do better than I might otherwise have done. Finally, all credit to the SAGE/USENIX association for their unsurpassed work in spreading state of the art knowledge about computing systems of all sizes and shapes.

Chapter 1

Introduction

1.1 The Scope of System Administration

The task of system administration is a balancing act. It requires patience, understanding, knowledge and experience. It is like working in the casualty ward of a hospital. We need to be the doctor, the psychologist, and – when instruments fail – the mechanic. We need to work with the limited resources we have, be inventive in a crisis, and know a lot of general facts and figures about the way computers work. We need to recognize that the answers are not always written down for us to copy, that machines do not always behave in the way we think they should. We need to remain calm and attentive, and learn a dozen new things a year.

Being a system administrator is as much a state of mind as it is about being knowledgeable. It is the sound of one hand tapping (on the keyboard) while the other is holding the phone, talking to a user and there is a queue of people waiting for help. We must be ready for the unexpected, resigned to the uncertain, and we need to be able to plan for the future. It requires organization and the ability to be systematic. There is no right answer, but there is often a wrong answer. It's about making something robust which works. Stereotypes notwithstanding, today's system administrator is neither haphazard nor messy. Computing systems require the very best of organizational skills and the most professional of attitudes. To start down the road of system administration, we need to know many *facts* and build confidence though experience – but we also need to know our limitations in order to avoid the careless mistakes which are all too easily provoked.

If you have installed Windows, DOS or GNU/Linux on a PC, you might think that you already know a lot about system administration, but in fact you know only the very beginning. Today, no computer system can be examined in isolation from the network. Networking is about cooperation and sharing in an environment with many users. It would be a grave mistake to believe that we know all the answers simply because we know the beginnings of how our own machine works. This book is about much more than that: it is an introduction to the *concepts* of system administration. As system administrators, we have many responsibilities and constraints on our work. Our first responsibility is to the greater network community and then to the users of our system. An administrator's job is to make users' lives bearable and to empower them in the production of real work.

1.2 Is System Administration a Discipline?

System administration practices, worldwide, vary from the haphazard to the state of the art. There is a variety of reasons for this. Most recently, the Internet has grown considerably, operating systems have grown more and more complex, but the number of technically adept system administrators has not grown in proportion. In the past, system administration has been a job which has not been carried out by dedicated professionals, but by interested computer users, as a necessary chore in getting their work done. The focus on making computers easy to use has distracted many vendors from the belief that their computers should also be easy to manage. It is only over the gradual course of time that this has changed, though even today, system administrators are a barely visible race, until something goes wrong. Thanks mainly to Unix user groups and the founding of the independent USENIX organization, technically minded people have come together in order to share and discuss their work in an independent forum: LISA, the Large Installation System Administration conferences. These are unique in being vendor neutral, technical conferences.

The need for a formal discipline in system administration has been recognized for some time, though it has sometimes been met with a certain trepidation by many corners of the Internet community, perhaps because the spirit of free cooperation which is enjoyed by system administrators could easily be shattered by too pompous an academic framework. Nonetheless, there are many closet academics working on system administration, and it is common for system administrators to have scientific background.

The need to formalize the problems of system administration is highlighted by the general ignorance which researchers have of previous research in the community. In one of the earliest conferences on system administration, USENIX/LISA 1990, the question was posed in a panel discussion: why do we keep re-inventing the wheel, and what can we do about it? That question is still pertinent ten years later. The trail of work on system administration exhibits a huge amount of repetition, particularly in tool-building. One of the purposes of this volume is to selectively summarize that body of work, so that the Möbius loop can unravel and we can go forward.

Academic concerns aside, the vast majority of computer systems lie in the private sector, and the Internet is only amplifying this tendency. In order to be good at system administration, a certain amount of dedication is required, with both theoretical and practical skills. For a serious professional, system administration is a career. There is now an appreciable market for consulting services in security and automation of system administrative tasks. Most companies have more money than they have time or expertise; that means that they need to be able to *buy* something in order to satisfy their board members and accountants. For a business, doing something to address a problem means spending money. Companies are looking to pay someone to carry out a service. Even though the best system administration tools are free, companies actively seek to pay consultants to set up and maintain administration tools for them. Not only is system administration a fascinating and varied line of work, it can also be lucrative.

1.3 A Jigsaw Puzzle

Knowledge, in the world of the system administrator, is a disposable commodity – we use it and we throw it away, as it goes out of date. Then we need to find newer, more up-to-date

knowledge to replace it. This is a continual process; the turn-around time is short, the loop endless, the mental agility required considerable. Such a process could easily degenerate into chaos or lapse into apathy. A robust discipline is required to maintain an island of logic, order and stability in a sea of turbulent change.

This book is about the aims and principles involved in maintaining that process. It is supplemented with a reference section of practical recipes and advice. When you master this book you will come to understand why no single book will ever cover every aspect of the problem – you need a dozen others as well[1]. True knowledge begins with understanding, and understanding is a jigsaw puzzle you will be solving for the rest of your life. The first pieces are always the hardest to lay correctly.

1.4 The Goals of System Administration

System administration is about putting together a network of computers (workstations, PCs and supercomputers), getting them running and then *keeping* them running in spite of the activities of *users* who tend to cause the systems to fail. System administration is a service profession, but it is far more than that. System administrators are also mechanics, sociologists and research scientists.

A system administrator works for users, so that they can complete work which is unrelated to the upkeep of the computer system itself. However, a system administrator should not just cater to one or two selfish needs, but also work for the benefit of a whole community. Today, that community is a global community of machines and organizations, which spans every niche of human society and culture, thanks to the Internet. It is often a difficult balancing act to determine the best policy, which accounts for the different needs of everyone with a stake in your local system. One a computer is attached to the Internet, we have to consider the consequences of being directly connected to all the other computers of the world.

1.5 A Philosophy

How to begin and how to behave? How do we come to terms with the idea of riding a lifelong wave of change and complexity? What are bad habits and what are good habits? The value of a solid understanding should not be underplayed. The days when one could fly a system by the seat of one's pants are alas over, for most of us. There is a strong need for practical skills, but nothing can replace real understanding. If we don't understand what we are doing (and why), we will never be able to accomplish two important goals: stability and security. The business of system administration was simpler in the beginning: there were no networks and there were few security issues.

We need to cultivate a way of thinking which embodies some basic principles:

- Independence, or self-sufficiency.
- Systematic and tidy practices.

[1] Later you might want to look at some of the better how-to books such as those in the recommended reading list [187, 95, 96, 177].

- An altruistic view of the system. Users come first: collectively and only then individually[2].
- Balancing a fatalistic view of inevitable errors with a determination to gain firmer control of the system.

Some counter-productive practices could be avoided:

- The belief that there exists a right answer to every problem.
- Getting fraught and upset when things do not work in the way we expect.
- Expecting that every problem has a beginning, a middle and an end (some problems are chronic and cannot be solved without impractical restructurings).

Others are to be encouraged:

- Looking for answers in manuals and newsgroups.
- Using controlled trial and error to locate problems.
- Listening to people who tell us that there is a problem. It might be true, even if we can't see it ourselves.
- Writing down experiences in an A–Z so that we know how to solve the same problem again in the future.
- Taking responsibility for our actions. (Be prepared for accidents. They are going to happen and they will be your fault. You will have to fix them.)
- Remembering the tedious jobs like vacuum cleaning the hardware once a year.
- After learning something new, always pose the question: *how does this apply to me?*

American English is the language of the net. System administrators need it to be able to read documentation, to be able to communicate with others and to ask questions on the Internet. Some sites have even written software tools for training novice administrators. See, for instance, ref. [236].

1.6 The Challenges of System Administration

System administration is not just about installing operating systems. It is about planning and designing an efficient *community* of computers so that real *users* will be able to get their jobs done. That means:

- Designing a network which is logical and efficient.
- Deploying large numbers of machines which can be easily upgraded later.
- Deciding what services are needed.
- Planning and implementing adequate security.
- Providing a comfortable environment for users.
- Developing ways of fixing errors and problems which occur.

[2] 'The needs of the many outweight, the needs of the few (or the one)...'

- Keeping track of, and understanding how to use, the enormous amount of knowledge which increases every year.

Some system administrators are responsible for both the hardware of the network and the computers which it connects, i.e. the cables as well as the computers. Some are only responsible for the computers. Either way, an understanding of how data flow from machine to machine is essential, as is an understanding of how each machine affects every other.

In all countries outside the United States and Canada, there are issues of internationalization, or tailoring the input/output hardware and software to the local language. Internationalization support in computing involves three issue:

- **Choice of keyboard**: e.g. British, German, Norwegian, Thai, etc.
- **Fonts**: Roman, Cryllic, Greek, Persian, etc.
- **Translation of program text messages**.

Inexperienced computer users usually want to be able to use computers in their own language. Experienced computer users, particularly programmers, often prefer the American versions of keyboards and software in order to avoid the awkward placement of commonly used characters on non-US keyboards.

1.7 Common Practice and Good Practice

If this book does nothing else, it should make you think for yourself. You will spend your career as a system administrator hearing advice from many different sources, and not all of it will be good advice. The best generic advice anyone can give in life is: think for yourself, pay attention to experts but don't automatically believe anyone. No authority is self-justified. Every choice needs a reason, even if that reason ends up being an arbitrary choice. That does not undermine the need for a book of this kind: it only cautions us about accepting advice on trust.

This is just the scientific method at work: informed scepticism and constant reappraisal. It is always a good idea to see what others have done in the past. There are three reasons why ideas catch on and 'everyone does it':

- Someone did it once, the idea was copied without thinking and no-one has thought about it since. Now everyone does it because everyone else does it.
- People have thought a lot about it and it really is the best solution.
- An arbitrary choice had to be made, and now it is a matter of convention.

For example, in the British Isles it is a good idea to drive on the left-hand side of the road. That's because someone started doing so and now everyone does it – but it's not just a fad: lives actually depend on this. The choice has its roots in history and in the dominance of right-handed sword-wielding carriage drivers and highwaymen but, for whatever reason, the opposite convention now dominates in other parts of the world and, in Britain, the convention is now mainly preserved by the difficulty of changing. This is not ideal, but it is reasonable.

Some common practices, however, are bizarre but adequate. For instance, in parts of Europe the emergency services Fire, Police and Ambulance have three different numbers

(110, 112 and 113) instead of one simple number like 911 (America) or, even simpler, 999 (UK). The numbers are very difficult to remember, they are not even a sequence! Ours is not to reason why the numbers were chosen, but they are now used because that is what 'is done'. Change, however, would be preferable.

Other practices are simply a result of blind obedience to poorly formulated rules. In public buildings there is a rule that doors should always open outwards from a room. The idea is that in the case of fire, when people panic, doors should 'go with the flow'. This makes eminent sense where large numbers of people are involved. Unfortunately, the building designers of my college have taken this literally, and have done the same thing with every door, even office doors in narrow corridors. When there is a fire (actually all the time), we open our doors into the faces of passers-by (the fleeing masses), injuring them and breaking their noses. The rule could perhaps be reviewed.

In operating systems, many conventions have arisen, e.g. the conventions for naming the 'correct' directory for installing system executables, like daemons, the permissions required for particular files and programs and even the use of particular software. Originally, Unix programs were thrown casually in `usr/bin` or `etc`; nowadays `sbin` or `libexec` are used by different schools of thought, all of which can be discussed.

As a system administrator you often have the power to make your own decisions about your systems. The point is simply this: often rules get made without due thought, by people with insufficient imagination. If you are boss, make logical choices rather than obedient ones.

1.8 Bugs

Operating systems and programs are full of bugs. Learning to tolerate bugs is a matter of survival for system administrators. If one is lucky enough to be using free software from the net, these bugs will usually be solved quickly and one can eliminate them by upgrading. If commercial software is used, it will probably be necessary to wait a lot longer for a patch. Either way, one has to be creative and work around these bugs. Bugs can be caused by many things. They may come from

- Shoddy software.
- Little known problems in the operating system.
- Unfortunate clashes between incompatible software, i.e. one software package destroys the operation of another.
- Totally unexplainable phenomena, cosmic rays and invasions by digital life-forms.

We have to be prepared to live and work around these, no matter what the reason for their existence.

1.9 Information Sources for Sysadms

Information can be found from many sources:

- Printed manuals.

- Unix manual pages (`man` and `apropos` commands).
- The World Wide Web.
- RFCs (Requests for comment), available on the web.
- News groups and discussions.
- Papers from the SAGE/Usenix LISA conferences [16].
- More specialized books.

In order to avoid recursing into a level of detail from which we might never emerge, the number of how-to recipes has been kept to a minimum in the following chapters. Apart from the information which can be found in the many excellent how-to books on specialized topics, a supplement to this book, with a collection of useful recipes and facts, is provided as a resource for system administrators at http://www.iu.hioslo.no/SystemAdmin.

Exercises

Exercise 1.1 Browse through this whole book from start to finish. Browsing information is a skill you will use a lot as a system administrator. Try to get an overall impression of what the book contains and how it is laid out.

Exercise 1.2 List what you think are the important tasks and responsibilities of a system administrator. You will have the opportunity to compare this to your impressions once we reach the end of the book.

Exercise 1.3 Locate other books and information sources which will help you. These might take the form of books (such as the recommended reading list at the end of this book) or newsgroups, or web sites.

Exercise 1.4 Buy an A–Z for noting your solutions to problems.

Exercise 1.5 What is an RFC? Locate a list of RFCs on a WWW or FTP server.

Chapter 2

The System Components

We begin our journey from the top down. Task number one in understanding networked computer systems is to identify the main components which bind them together and make them work. In this chapter, we summarize some of the prerequisites for understanding system administration so as to place ourselves in a frame of mind for the coming chapters.

2.1 What is 'The System'?

In this book, we use the word *system* a lot to refer both to the operating system of a computer and often, collectively the set of computers on a network. The term *operating system* has no rigorous or accepted definition. It can be thought of as the collection of all programs which were bundled with a particular computer.

All contemporary computers are based on the Eckert–Mauchly–von Neumann architecture [195], sketched in Figure 2.1. Each computer has a clock which drives a *central processor unit*

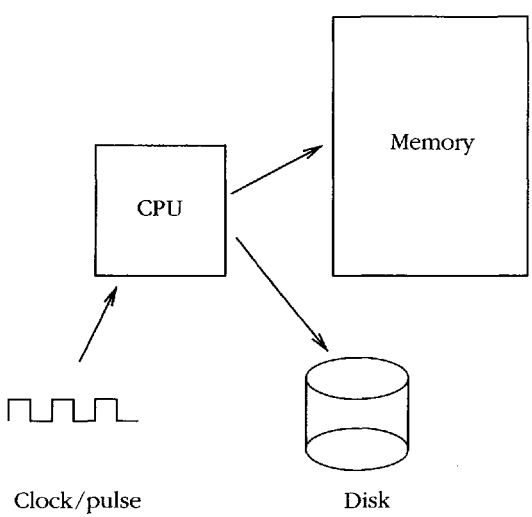


Figure 2.1 The basic elements of the von Neumann architecture

(CPU), *random access memory* (RAM) and an array of other devices, such as disk drives. In order to make these parts work together, the CPU is designed to run programs which can read and write to hardware devices. The most important program is the *operating system* kernel.

2.2 Operating Systems

An operating system is the software which shares and controls the hardware resources of a computer. It is a layer of software which takes care of technical aspects of a computer's operation. It shields the user of the machine from the low-level details of the machine's operation and provides frequently needed facilities. There is no universal definition of what an operating system consists of. We can think of it as being the software which is already installed on a machine, before we add anything of our own. Normally the operating system has a number of key elements: (i) a *technical layer of software* for driving the hardware of the computer, like disk drives, the keyboard and the screen; (ii) a *filesystem* which provides a way of organizing files logically; and (iii) a simple *user interface* which enables users to run their own programs and to manipulate their files in a simple way.

Of central importance to an operating system is a core software system or *kernel* which is responsible for allocating and sharing the resources of the system between several running programs or *processes*. It is supplemented by a number of supporting *services* (paging, RPC, FTP, WWW, etc.) which either assist the kernel or extend its resource sharing to the network domain. The operating system can be responsible for sharing the resources of a single computer, but increasingly we are seeing *distributed operating systems* in which the execution of programs and sharing of resources happens without regard for hardware boundaries; or *network operating systems* in which a central server adds functionality to relatively dumb workstations. Sometimes programs which do not affect the job of sharing resources are called *user programs*.

In short, a computer operating system is composed of many subsystems, some of which are *software systems* and some of which are *hardware systems*. The operating system runs interactive programs for humans, *services* for local and distributed users and support programs which work together to provide the infrastructure that enable machine resources to be shared between many processes. Some operating systems also provide text editors, compilers, debuggers and a variety of other tools. Since the operating system (OS) is in charge of a computer, all requests to use its resources and devices need to go through the OS kernel. An OS therefore provides (iv) *legal entry points* into its code for performing basic operations like writing to devices.

For an operating system to be managed consistently it has to be possible to prevent its destruction by restricting the privileges of its users. Different operating systems vary in their provisions for restricting privilege. In operating systems where any user can change any file, there is little or no possibility of gaining true control over the system. Any accident or whim on the part of a user can make uncontrollable changes.

It is very important to distinguish between a user interface and an operating system. A window system is a *graphical user interface* (GUI), an operating system shares resources and provides functionality. This issue has been confused by the arrival of an operating system called Windows, which includes a single graphical user interface. In principle, an operating

system can sport any number of different windowing interfaces, one for every taste. An operating system can be good or bad independently of whether its windowing system is good or bad.

Operating systems may be classified both by how many tasks they can perform 'simultaneously' and by how many users can be using the system 'simultaneously', i.e. *single-user* or *multi-user* and *single-task* or *multi-tasking*. A multi-user system must clearly be multi-tasking. The table below shows some examples.

OS	Users	Tasks	Processors
MS/PC DOS	S	S	1
Windows 3x	S	QM	1
Macintosh System 7*	S	QM	1
Windows 9x	S	M*	1
AmigaDOS	S	M	1
MTS	M	M	1
Unix-like	M	M	*n*
VMS	M	M	1
NT	S/M	M	*n*

The first of these (MS/PC DOS/Windows 3x) are single user, single-task systems which provide a library of basic functions called the BIOS. Windows also includes a windowing library. These are system calls which write to the screen or to disk, etc. Although all the operating systems can service *interrupts*, and therefore simulate the appearance of multi-tasking in some situations, the DOS environment cannot be thought of as a multi-tasking system in any sense. Only a single user application can be open at any time. Note that Windows 3x is not really a separate operating system from DOS; it is a user interface to DOS.

The Macintosh system 7 can be classified as single-user quasi-multitasking[1]. That means that it is possible to use several user applications simultaneously. A window manager can simulate the appearance of several programs running simultaneously, but this relies on each program obeying specific rules in order to achieve the illusion. Prior to its Mach/Unix incarnation, the Macintosh was not a true multitasking system, in the sense that, if one program crashes, the whole system would crash. Windows 9*x* is purported to be preemptive multitasking but most program crashes also crash the entire system. This might be due to the lack of proper memory protection. Either way the claim is confusing.

MTS (Michigan timesharing system) was the first time-sharing multi-user system[2]. It supports only simple single-screen terminal based input/output, and has no hierarchical file system.

IBM S/370, S/390 and AS/400 mainframe computers are widely used in banks and large concerns for high level processing. These are fully multitasking systems of high caliber.

[1] At the present time, Apple are preparing a new operating system called NextStep or Rhapsody or Mac OS Server X which was based on BSD 4.3 Unix, now BSD 4.4 running on a Mach micro-kernel, and which will run old Macintosh software under emulation.

[2] In Manitoba, Canada, the telephone system is also called MTS. The telephone system is probably more advanced than the original MTS, and certainly faster!

Unix is arguably the most important operating system today, both because of its wide spread use and its historical importance. We shall frequently refer to Unix-like operating systems below. 'Unix', as it is correct to call it now, comes in many forms, developed by different manufacturers. Originally designed at AT&T, Unix split into two camps early on: BSD (Berkeley Software Distribution) and system 5 (AT&T license). The BSD version was developed as a research project at the University of California Berkeley (UCB). Many of the networking and user-friendly features originate from these modifications. With time these two versions have been merged back together, and most systems are now a mixture of both worlds. Historically, BSD Unix has been most prevalent in universities, while system 5 has been dominant in business environments. The trend during the last three years, by Sun Microsystems and Hewlett-Packard amongst others, has been to move towards system 5, keeping only the most important features of the BSD system. A standardization committee for Unix called POSIX, formed by the major vendors and independent user groups, has done much to bring compatibility to the Unix world. Here are some common versions of Unix.

Unix-like OS	Manufacturer	Type
BSD	Univ. California Berkeley	BSD
SunOS (Solaris 1)	Sun Microsystems	BSD/Sys 5
Solaris(2)	Sun Microsystems	Sys 5/BSD
Ultrix	DEC/Compaq	BSD
OSF 1/Digital Unix	DEC/Compaq	BSD/Sys 5
HPUX	Hewlett-Packard	Sys 5
AIX	IBM	Sys 5 / BSD
IRIX	Silicon Graphics	Sys 5
GNU/Linux	GPL Free Software	Posix (Sys V/BSD)
Unixware	Novell	Sys 5

Note that the original BSD source code is now in the public domain, and that the GNU/Linux source code is free software. Unix is generally regarded as the most portable and powerful operating system available today, but NT is improving quickly. Unix runs on everything from laptop computers to CRAY mainframes. It is particularly good at managing large database applications, and can run on systems with hundreds of processors. Most Unix-like operating systems now support symmetric multithreaded processing, and all support simultaneous logins by multiple users.

NT is an operating system from Microsoft, based in part on the old VAX/VMS kernel from Digital Equipment Corporation and the Windows32 API. Initially it reinvented many existing systems, but it is gradually being forced to adopt many open standards from the Unix world. It is fully multitasking, and can support multiple users (but only one at a time – multiple logins by different users is not possible). It has virtual memory and multithreaded support for several processors. NT has a built-in object model and security framework which is amongst the most modern in use. Windows NT has been reincarnated now in the guise of Windows 2000, which is a redesign of Windows NT, adopting many of the successful features of the Novell system, such as consistent directory services.

2.2.1 Multi-User Operating Systems

The purpose of a multi-user operating system is to allow multiple users to share the resources of a single host. In order to do this, it is necessary to protect users from one another by giving them a unique identity or *user name* and a private login area, i.e. by restricting their privilege. In short, we need to simulate a virtual workstation for each individual user, with private files and private processes.

2.2.2 The Legacy of Insecure Operating Systems

The home computer revolution was an important development which spread cheap computing power to a large part of the world. As with all rapid commercial developments, the focus in developing home operating systems was on immediate functionality, not on planning for the future. The home computer revolution preceded the network revolution by a number of years, and home computer operating systems did not address security issues. Operating systems developed during this period include Windows, Macintosh, DOS and Amiga-DOS. All of these systems are completely insecure: they place *no limits* on what a determined user can do.

Fortunately, these systems will slowly be replaced by operating systems which were designed with resource sharing (including networking) in mind. Still, there is a large number of insecure computers in use, and many of them are now connected to the network. This should be a major concern for a system administrator. In an age where one is forced to take security extremely seriously, leaving insecure systems where they can be accessed physically or by network is a potentially dangerous situation. Such machines should not be allowed to hold important data, and they should not be allowed any privileged access to network services. We shall return to this issue in Chapter 10 on security.

2.2.3 Secure Operating Systems

To distinguish them from insecure operating systems, we shall refer to operating systems like Unix and NT as *secure* operating systems. This should not give the impression that Unix and NT are really secure by any stretch of the imagination: complete security is a fairy tale, a pipe dream which will never happen in any operating system. Nevertheless, these operating systems do have the mechanisms which make a basic level of security possible.

The most fundamental tenet of security is the ability to restrict access to certain system resources. The main reason why DOS, Windows 9x and the Macintosh are so susceptible to virus attacks is because any user can change the operating system's files. Properly configured and bug free Unix/NT systems are theoretically immune to such attacks because ordinary users do not have the privileges required to change system files[3]. Unfortunately, the key phrases *properly configured* and *bug-free* highlight the flaw in this dream.

To restrict access to the system we require a notion of *ownership* and *permission*. Ordinary users should not have access to the hardware devices of a secure operating system's files, only their own files, for then they will not be able do anything to compromise the security of the system. System administrators need access to the whole system in order to watch over it,

[3] Not all viruses have to change system files; it is also possible to infect programs directly in memory if process security is weak.

make backups and keep it running. Secure operating systems thus need a privileged account which can be used by the system administrator when he/she is required to make changes to the system.

Secure operating systems are usually multi-user systems, i.e. operating systems where files and processes can be owned by a particular user, and access is restricted on the basis of user identity. Actually, NT has not previously been a true multiuser operating system, because only one user could be logged onto an NT host at any one time. Service Pack 4 introduced the NT Terminal Server to correct this weakness, but did not correct system file permissions which would be inappropriate in such a scenario. What is important for security is that system resources can be shared and protected individually for each user.

2.2.4 Shells or Command Interpreters

Today it is common for operating systems to provide graphical window systems for all kinds of tasks. These are often poorly suited to system administration because they allow us only to choose between pre-programmed operations which the program designers foresaw when they wrote the program. Working with windowing systems is a bit like trying to order a rare steak at a restaurant using semaphore. (There is no flag which means rare steak.)

Most operating systems also provide a command line user interface which has some form of interpreted language, thus allowing the user to express what he or she wants with more freedom. Windows shells are fairly rudimentary; Unix shells are rich in complexity. Many Unix shells are being ported to Windows. Shells can be used to write simple programs called *scripts* or *batch files* which often simplify repetetive administrative tasks.

2.2.5 Comparing Unix-like Operating Systems with NT

The two most popular classes of operating system today are Unix-like operating systems (i.e. those which are either derived from or inspired by System V or BSD) and Microsoft Windows-based operating systems. For reasons which will become clear later, we shall only discuss Windows NT and later derivatives of the Windows family in a network context. Microsoft are now planning to merge all of the Windows operating systems into one release in future. The name of the final product is set to be Windows 2000. For the sake of placing the generalities in this book in a clearer context, it is useful to compare 'Unix' with NT.

Unix-like operating systems are many and varied, but they are basically similar in concept. It is not the purpose of this book to catalogue the complete zoological inventory of the 'Unix' world; our aim is to speak primarily of generalities which rise above such distinctions. Nonetheless, we shall occasionally need to distinguish the special features of these operating systems, and at least distinguish them from NT. This should not detract from the fact that NT has adopted much from the Unix cultural heritage, even though superficial attempts to hide this (e.g. renaming / with \ in filenames, changing the names of some commands, etc.) might obscure the fact.

NT is a multitasking operating system from Microsoft which allows one user at a time to log in to a console or *workstation*. The consoles may be joined together in a network with common resources shared by an NT *domain*. An NT host is either a network domain server or a personal workstation. NT is fairly new, at least in terms of age, and Microsoft has reinvented many well known systems rather than use tried and tested solutions from the

Unix world. This has led to a history of bugs and security issues which has resulted in several developments which bring NT closer to Unix. The basic NT distribution contains only a few tools which can be used for network administration. The NT Resource Kit is an extra package of documentation and unsupported software which, nonetheless, provides many essential tools. Other tools can be obtained free of charge from the network.

NT did not have a remote shell login feature like Unix at the outset, though one may now obtain a Terminal Server which gives NT telnet-like functionality, as of Service Pack 4. The service should be integrated into Windows 2000 server. This correction adds an important possibility: that of remote administration, other than through inheritance from a domain server. The free Perl Win32 package and related tools provides tools for solving a number of problems with NT from a script viewpoint.

Although we are ignoring many important operating systems by comparing just two main players, a comparison of Unix-like operating systems with NT covers most of the important differences. The latest offerings from the Macintosh world, for instance, are based on emulation of BSD 4.4 Unix and MacOS on a Mach kernel, with features designed to compete with NT.

Unix is important, not only for its endurance as the sturdy workhorse of the network, but also for its cultural significance, beyond mere market share. It has influenced so many other operating systems (including NT) that further comparisons would be largely redundant. Let us note, briefly then, for the record, the basic correspondences between Unix-like operating systems and NT. Many basic commands are very similar. Here are some basic commands for file and system control:

Unix-like OS	NT
`chmod`	`CACLS`
`chown`	`CACLS`
`chgrp`	*No direct equivalent.*
`emacs`	*Wordpad* or `emacs` in GNU tools
`kill`	`kill` command in Resource Kit
`ifconfig`	`ipconfig`
`lpq`	`lpq`
`lpr`	`lpr`
`mkfs/newfs`	`format` and `label`
`mount`	`net use`
`netstat`	`netstat`
`nslookup`	`nslookup`
`ps`	`pstat` in Resource Kit
`route`	`route`
`setenv`	`set`
`su`	`su` in resource kit
`tar`	`tar` command in cygnus tools
`traceroute`	`tracer`

The file and directory structures of Unix and NT are rather different, but it is natural that both systems have the same basic elements.

Unix-like OS	NT
/usr	%SystemRoot% usually points to C:\WinNT
/bin or /usr/bin	%SystemRoot%\System32
/dev	%SystemRoot%\System32\Drivers
/etc	%SystemRoot%\System32\Config
/etc/fstab	No equivalent
/etc/group	%SystemRoot%\System32\Config\SAM* (binary)
/etc/passwd	%SystemRoot\%\System32\Config\SAM* (binary)
/etc/resolv.conf	%SystemRoot%\System32\DNS*
/tmp	C:\Temp
/var/spool	%SystemRoot%\System32\Spool

Unix-like OS	NT
Standard libraries	WIN32 API
Unix libraries	Posix compatibility library
Symbolic/hard Links	Hard links (short cuts)
Processes	Processes
Threads	Threads
Long filenames	Long filenames on NTFS
Mount disk on directory	Mount drive A: B: etc
endl is LF	endl is CR LF
UID (User ID)	SID (Subject ID)
groups	groups
ACLs (non standard)	ACLs
Permission bits	(Only in ACLs or with cygwin)
Shared libraries	DLL's
Environment variables	Environment variables

Unix-like OS	NT
Daemons/services/init	Service control manager
DNS/DHCP/bootp (free)	DNS/DHCP (NT server)
X windows	X windows
Various window managers	Windows 95 GUI
System admin GUI (non-standard)	System Admin GUI (Standard)
cfengine	cfengine as of 1.5.0
Any client-server model	Central server model
rsh	limited implementation in server
Free software	Some free software
Perl	Perl + WIN32 module
Scripts	Scripts
Shells	DOS Command window
Primitive security	Primitive security
Dot files for configuration	System registry
Pipes with comm1 | comm2	Combinations comm1 | comm2
Configuration by text/ascii files	Config by binary database

Note: there are differences in nomenclature. What NT refers to as pipes[4] in its internal documentation is not what Unix refers to as pipes in its internal documentation.

At the time of writing, we are in the middle of a war of words between Unix and NT. In 1997, NT gained a lot of ground with newly started companies, eager to get going on the Internet as quickly as possible. Microsoft's efficient marketing and existing dominance of the PC world made NT an interesting option. NT has suffered from setbacks as a result of bugs which affect security and stability, in the face of a vigorous marketing campaign. A major problem is the need for compatibility with DOS, through Windows 9x to NT. Since both DOS and Windows 9x are insecurable systems, this has led to conflicts of interest. Unix vendors have tried to keep step, in spite of the poor public image of Unix (often the result of private dominance wars between different Unix vendors), but the specially designed hardware platforms built by Unix vendors have had a hard time competing with inferior but cheaper technology from the PC world.

In 1998 we saw sales of both Unix and NT increasing, but many users who have tried NT report that they are turning to cheap Unix solutions on Intel platforms (GNU/Linux, FreeBSD, etc.) instead, because of its proven reliability. According to the popular press, Unix seems to work better in a distributed environment and delivers especially good performance as a database server on 64-bit multiprocessor hardware. NT can run on 32-bit multiprocessor systems, but independent benchmarks show that it does not scale as well as many Unix variants do. Also, the issue of remote login is significant for some sites. Using standard Windows user interfaces on NT, one cannot run graphical applications remotely, as with X-windows on Unix. A basic X-windows release 6 is available for NT, however, so one has the option of replacing the Windows interface and choosing a more distributed user interface. Today some systems which are based on the Mach micro-kernel can run native Unix and NT simultaneously. For example, Digital (now Compaq) Unix, GNU/Linux and NT can run on the Alpha processor. It goes without saying that the future will bring many changes.

2.3 File Systems

Files and file systems are the very basis of what system administration is about. Almost every task in host administration or network configuration involves making changes to files. We need to acquire a basic understanding of the principles of file systems: what better way than to examine some of the most important file systems in use today. Specifically, what we are interested in is the user interfaces to common file systems, not the technical details, which are rather fickle. We could, for instance, mention the fact that most file systems (e.g. NT, GNU/Linux) are 32-bit addressable and therefore support a maximum file size of 2GB or 4GB, depending on their implementation details, or that newer file systems like Solaris and Netware 5 are 64-bit addressable, and therefore have essentially no storage limits. We could mention the fact that Unix uses an index node system of block addressing, while DOS uses a tabular lookup system... and so the list goes on. These technical details are of only passing interest since they change at an alarming pace. What is more constant is the user functionality of the file systems: how they allow file access to be restricted to groups of users, and what commands are necessary to manage this.

[4] Ceci n'est pas une pipe!

2.3.1 Unix File Model

Unix has a hierarchical file system which makes use of directories and sub-directories to form a tree. All file systems on Unix-like operating systems are based on a system of *index nodes*, or *inodes*, in which every file has an index entry stored in a special part of the file system. The inodes contain an extensible system of pointers to the actual disk blocks which are associated with the file. The inode contains essential information needed to locate a file on the disk.

The top or start of the Unix file tree is called the root file system or '/'. Although the details of where common files are located differ for different versions of Unix, some basic features are the same.

The File Hierarchy

The main sub-directories of the root directory together with the most important file are shown below. Their contents are as follows:

- `/bin` Executable (binary) programs. On most systems this is a separate directory to /usr/bin. In SunOS, this is a pointer (link) to /usr/bin.
- `/etc` Miscellaneous programs and configuration files. This directory has become very messy over the history of Unix and has become a dumping ground for almost anything. Recent versions of unix have begun to tidy up this directory by creating subdirectories `/etc/mail`, `/etc/services`, etc!
- `/usr` This contains the main meat of Unix. This is where application software lives, together with all of the basic libraries used by the OS.
- `/usr/bin` More executables from the OS.
- `/usr/local` This is where users' custom software is normally added.
- `/sbin` A special area for statically linked system binaries. They are placed here to distinguish commands used solely by the system administrator from user commands, and so that they lie on the system root partition where they are guaranteed to be accessible during booting.
- `/sys` This holds the configuration data which go to build the system kernel. (See below.)
- `/export` Network servers only use this. This contains the disk space set aside for client machines which do not have their own disks. It is like a 'virtual disk' for diskless clients.
- `/dev` and `/devices` A place where all the 'logical devices' are collected. These are called 'device nodes' in Unix and are created by `mknod`. Logical devices are Unix's official entry points for writing to devices. For instance, `/dev/console` is a route to the system console, while `/dev/kmem` is a route for reading kernel memory. Device nodes enable devices to be treated as though they were files.
- `/home` (called `/users` on some systems.) Each user has a separate login directory where files can be kept. These are normally stored under `/home` by some convention decided by the system administrator.
- `/root` On newer Unix-like systems, root has been given a home-directory which is no longer the root of the file system '/'. The name root then loses its logic.

- `/var` System 5 and mixed systems have a separate directory for spooling. Under old BSD systems, `/usr/spool` contains spool queues and system data. `/var/spool` and `/var/adm`, etc. are used for holding queues and system log files.
- `/vmunix` This is the program code for the unix *kernel* (see below). On HPUX systems with a file it is called `hp-ux`. On Linux it is called `vmlinuz` and on newer systems it is often moved into a subdirectory.
- `/kernel` On newer systems the kernel is built up from a number of modules which are placed in this directory.

Every unix directory contains two 'virtual' directories marked by a single dot and two dots:

```
ls -a
.  ..
```

The single dot represents the directory one is already in (the current directory). The double dots mean the directory one level up the tree from the current location. Thus, if one writes

```
cd /usr/local
cd ..
```

the final directory is `/usr`. The single dot is very useful in C programming if one wishes to read 'the current directory'. Since this is always called '.' there is no need to keep track of what the current directory really is. '`.`' and '`..`' are hard links to the current and parent directories, respectively.

Symbolic Links

A symbolic link is a pointer or an alias to another file. The command

```
ln -s fromfile /other/directory/tolink
```

makes the file `fromfile` appear to exist at `/other/directory/tolink` simultaneously. The file is not copied, it merely appears to be a part of the file tree in two places. Symbolic links can be made to both files and directories.

A symbolic link is just a small invisible file which contains the name of the real file one is interested in. Unlike NT's short-cuts, symbolic links cannot be seen to be files with a text editor; they are handled specially at the level of operating system. Application programs can choose whether they want to treat a symbolic link as a separate file object, or simply as an alias to the file it points to. If we remove the file a symbolic link points to, the link remains – it just points to a non-existent file.

Hard Links

A *hard link* is a duplicate *inode* in the file system which is in every way equivalent to the original file inode. If a file is pointed to by a hard link, it cannot be removed until the link is removed. If a file has n hard links, all of them must be removed before the file can be removed. The number of hard links to a file is stored in the file system *index node* for the file. A hard link is created with the `ln` command, without the `-s` option. Hard links are, in all current Unix-like operating systems, limited to aliasing files on the same disk partition.

Although the POSIX standard specifies the possibility of making hard links across disk partitions, this has presented an insurmountable technical difficulty because it would require inodes to have a global numbering scheme across all disk partitions. This would be an inefficient overhead for an additional functionality of dubious utility, so currently this has been ignored by file system designers.

File Access Control

To restrict privilege to files on the system, and create the illusion of a virtual host for every logged-on user, Unix records information about *who* creates files and also who is allowed to access them later. Unix makes no policy on what names files should have: a file can have any name. A file's contents are classified by *magic numbers* which are codes kept in the file's inode and defined in the magic number file for the system. This is in contrast to systems like NT, where file extensions (e.g. `.EXE`) are used to identify file contents. Under Unix, file extensions (e.g. `.c`) are only discretionary.

Each user has a unique *username* or *loginname*, together with a unique *user id* or *uid*. The user id is a number, whereas the login name is a text string – otherwise the two express the same information. A file belongs to user A if it is *owned* by user A. User A then decides whether or not other users can read, write or execute the file by setting the *protection bits* or the *permission* of the file using the command `chmod`.

In addition to user identities, there are groups of users. The idea of a group is that several named users might want to be able to read and work on a file, without other users being able to access it. Every user is a member of at least one group, called the *login group*, and each group has both a textual name and a number (*group id*). The *uid* and *gid* of each user is recorded in the file `/etc/passwd` (see Chapter 6). Membership of other groups is recorded in the file `/etc/group`, or on some systems `/etc/logingroup`.

The following output is from the command `ls -lag` executed on a SunOS type machine:

```
lrwxrwxrwx  1 root wheel        7 Jun  1 1993  bin -> usr/bin
-r--r--r-- 21 root bin     103512 Jun  1 1993  boot
drwxr-sr-x  2 bin  staff    11264 May 11 17:00 dev
drwxr-sr-x 10 bin  staff     2560 Jul  8 02:06 etc
drwxr-sr-x  8 root wheel      512 Jun  1 1993  export
drwx------  2 root daemon     512 Sep 26 1993  home
-rwxr-xr-x  1 root wheel   249079 Jun  1 1993  kadb
lrwxrwxrwx  1 root wheel        7 Jun  1 1993  lib -> usr/lib
drwxr-xr-x  2 root wheel     8192 Jun  1 1993  lost+found
drwxr-sr-x  2 bin  staff      512 Jul 23 1992  mnt
dr-xr-xr-x  1 root wheel      512 May 11 17:00 net
drwxr-sr-x  2 root wheel      512 Jun  1 1993  pcfs
drwxr-sr-x  2 bin  staff      512 Jun  1 1993  sbin
lrwxrwxrwx  1 root wheel       13 Jun  1 1993  sys->kvm/sys
drwxrwxrwx  6 root wheel      732 Jul  8 19:23 tmp
drwxr-xr-x 27 root wheel     1024 Jun 14 1993  usr
drwxr-sr-x 10 bin  staff      512 Jul 23 1992  var
-rwxr-xr-x  1 root daemon 2182656 Jun  4 1993  vmunix
```

The first column is a textual representation of the protection bits for each file. Column two is the number of hard links to the file (see the exercises below). The third and fourth columns

are the user name and group name, and the remainder show the file size in bytes and the creation date. Notice that the directories `/bin` and `/sys` are symbolic links to other directories.

There are 16 protection bits for a Unix file, but only 12 of them can be changed by users. These 12 are split into four groups of three. Each three-bit number corresponds to one *octal* number.

The leading four invisible bits gives information about the type of file: is the file a *plain file*, a *directory* or a *link*. In the output from `ls` this is represented by a single character: `-`, `d` or `l`.

The next three bits set the so-called *s-bits* and *t-bit* which are explained below.

The remaining three groups of three bits set flags which indicate whether a file can be read `r`, written to `w` or executed `x` by (i) the user who created them, (ii) the other users who are in the group the file is marked with, and (iii) any user at all. For example, the permission

```
Type  Owner  Group  Anyone
 d    rwx    r-x    ---
```

tells us that the file is a directory, which can be read and written to by the owner, can be read by others in its group, but not by anyone else[5].

Here are some examples of the relationship between binary, octal and the textual representation of file modes:

```
Binary  Octal  Text
001       1    --x
010       2    -w-
100       4    r--
110       6    rw-
101       5    r-x
 -      644    rw-r--r--
```

It is well worth becoming familiar with the octal number representation of these permissions.

chmod

The `chmod` command changes the permission or *mode* of a file. Only the owner of the file or the superuser can change the permission. Here are some examples of its use. Try them.

```
# make read/write-able for everyone
chmod a+w myfile

# add the 'execute' flag for directory
chmod u+x mydir/

# open all files for everyone
chmod 755 *

# set the s-bit on my-dir's group
chmod g+s mydir/

# descend recursively into directory opening all files
chmod -R a+r dir
```

[5] Note about directories. It is impossible to `cd` to a directory unless the `x` bit is set. That is, directories must be 'executable' in order to be accessible.

New File Objects: umask

When a new file gets created, the operating system must decide what default protection bits to set on that file. The variable umask decides this. umask is normally set by each user in his or her .cshrc file (see the next chapter). For example,

```
umask 077    # safe
umask 022    # liberal
```

According the Unix documentation, the value of umask is XORed (exclusive OR) with a value of 666 & umask for plain files or 777 & umask for directories in order to find out the standard protection. Actually this is not true: umask only removes bits, it never sets bits which were not already set in 666. For instance,

```
umask     Permission
077       600 (plain)
077       700 (dir)
022       644 (plain)
022       755 (dir)
```

The correct rule for computing permissions is not XOR but NOT umask AND 666/777.

Making Programs Executable

A Unix program is normally executed by typing its pathname. If the x execute bit is not set on the file, this will generate a 'Permission denied' error. This protects the system from interpreting nonsense files as programs. To make a program executable for someone, you must therefore ensure that they can execute the file, using a command like

```
chmod u+x filename
```

This command would set execute permissions for the owner of the file;

```
chmod ug+x filename
```

would set execute permissions for the owner and for any users in the same group as the file. Note that script programs must also be readable in order to be executable, since the shell has the interpret them by reading.

chown and chgrp

These two commands change the ownership and the group ownership of a file. For example:,

```
chown mark ~mark/testfile
chgrp www  ~mark/www/tmp/cgi.out
```

In newer implementations of chown, we can change both owner and group attributes simultaneously, by using a dot notation:

```
chown mark.www ~mark/www/tmp/cgi.out
```

Only the superuser can change the ownership of a file on most systems. This is to prevent users from being able to defeat quota mechanisms. (On some systems, which do not

implement quotas, ordinary users can give a file away to another user but not get it back again.) The same applies to group ownership.

Making a Group

The superuser creates groups by editing the file `/etc/group`. Normally, users other than root cannot define their own groups. This is a weakness in Unix from older times, and one which no one seems to be in a hurry to change. It is possible to 'hack' a solution to this which allows users to create their own groups. The format of the group file is

```
group-name:: group-number: comma-separated-list-of-users
```

The Unix group mechanism is very convenient, but poorly conceived. ACLs go some way to redressing its shortcomings (see below), but at an enormous price, in terms of computer resources. The group mechanism is fast and efficient, but clumsy for users.

s-bit and t-bit (sticky bit)

Apart from the read, write and execute file attributes, Unix has three other flags. The `s` and `t` bits have special uses. They are set as follows:

Name	Octal form	Text form
Setuid bit	`chmod 4000` *file*	`chmod u+s` *file*
Setgid bit	`chmod 2000` *file*	`chmod g+s` *file*
Sticky bit	`chmod 1000` *file*	`chmod +t` *file*

The effect of these bits differs for plain files and directories, and also differs between different versions of Unix. Check particularly the manual page `man sticky` on each system. The following is common behaviour.

For executable files, the setuid bit tells Unix that *regardless of who runs the program* it should be executed with the permissions and rights of owner of the file. This is often used to allow normal users limited access to `root` privileges. A *setuid-root* program is executed as `root` for any user. The setgid bit sets the group execution rights of the program in a similar way.

In BSD Unix, if the setgid bit is set on a directory then any new files created in that directory assume the group ownership of the parent directory, and not the logingroup of the user who created the file. This is standard policy under System 5.

A directory for which the sticky bit is set restricts the deletion of files within it. A file or directory inside a directory with the t-bit set can only be deleted or renamed by its owner or the superuser. This is useful for directories like the mail spool area and `/tmp` which must be writable to everyone, but should not allow a user to delete another user's files.

(Ultrix) If an executable file is marked with a sticky bit, it is held in the memory or system swap area. It does not have to be fetched from disk each time it is executed. This saves time for frequently used programs like `ls`.

(Solaris 1) If a non-executable file is marked with the sticky bit, it will *not* be held in the disk page cache – that is, it is never copied from the disk and held in RAM, but is written to directly. This is used to prevent certain files from using up valuable memory.

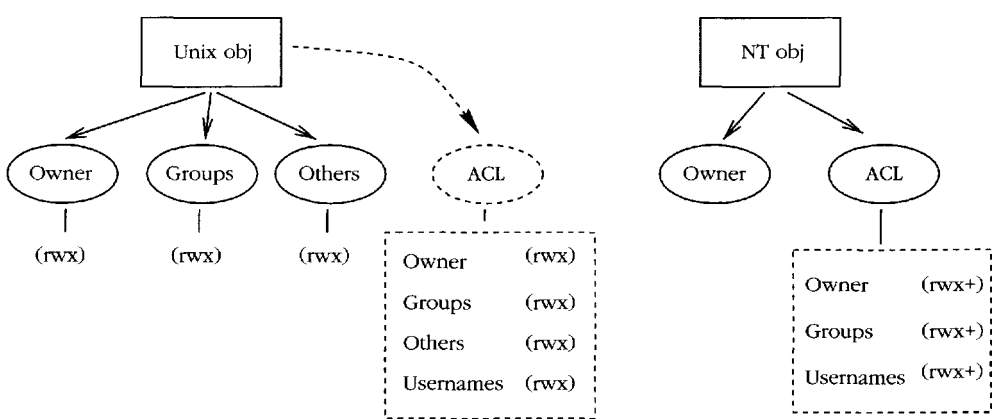
On some systems (e.g. ULTRIX), only the superuser can set the sticky bit. On others (e.g. SunOS) any user can create a sticky directory.

Access Control Lists

ACLs, or access control lists, are a modern replacement for file modes and permissions. With access control lists we can specify precisely the access rights to files for each user individually. Although ACLs are functionally superior to the old Unix group ownership model, experience shows that they are too complicated for most users in practice. Also, the overhead of reading and evaluating ACLs places a large performance burden on a system (see Figure 2.2).

Previously, the only way to grant access to a file to a known list of users was to make a group of those users, and use the group attribute of the file. With ACLs this is no longer necessary. ACLs are both a blessing and a nightmare. They provide a functionality which has long been missing from operating systems, and yet they are often confusing and even hopelessly difficult to understand in some file systems. One reason for this is when file systems attempt to maintain compatibility with older protection models (e.g. Unix/Posix permissions and ACLs, as in Solaris). The complex interactions between creation masks for Unix permissions and inherited properties of ACLs make ACL behaviour non-intuitive. Trying to obtain the desired set of permissions on a file can be like a flirtation with the forces of mysticism. This is partly due to the nature of the library interfaces, and partly due to poor or non-existent documentation.

ACLs were introduced in the DOMAIN OS by Apollo, and were later adapted by Novell, HP and other vendors. A POSIX standard for ACLs has been drafted, but as of today there is no adopted standard for ACLs, and each vendor has a different set of incompatible commands and data-structures. Sun Microsystems' solaris (NFS3) implementation is based on the POSIX draft. We shall follow Solaris ACLs in this section. GNU/Linux and the BSD operating systems do not have ACLs at all. If we grant access to a file which is NFS shared on the network to a



Figrue 2.2 The standard permission model for file objects in Unix and NT

machine which doesn't support ACLs, they are ignored. This limits their usefulness in most cases.

ACLs are literally lists of access rights. Each file has a list of data structures with pairs of names and permissions (see Figures 2.4, 2.5 and 2.6): an ACL is specified by saying what permissions we would like to grant and which user or group of users the permissions should apply to. An ACL can grant access or deny access to a specific user. Because of the amount of time required to check all the permissions in an ACL, ACLs slow down file search operations. Under Solaris, the commands to read and write ACLs have the cumbersome names

- `getfacl` *file* Examine the ACLs for a file
- `setfacl` *file* `-s` *permission* Set ACL entries for a file, replacing the entire list.
- `setfacl` *file* `-m` *permission* Set ACL entries for a file, adding to an existing list.

For example. If we create a new file, it ends up with a default ACL which is based upon the Unix `umask` value and any ACL masks which are set for the parent directory. Suppose `umask` is 077, and no directory ACLs are set, giving minimal rights to others:

```
mercury% touch testfile
mercury% getfacl testfile
# file: testfile
# owner: mark
# group: iu
user::rw-
group::---          #effective:---
mask:---
other:---
```

This tells us that a new file is created with read/write permission for the owner (`mark`) of the file, and no other rights are granted. To open the file for a specific user `ds`, one writes

```
mercury% setfacl -m user:ds:rw- testfile
mercury% getfacl testfile
# file: testfile
# owner: mark
# group: iu
user::rw-
user:ds:rw-         #effective:---
group::---          #effective:---
mask:---
other:---
```

To open a file for reading by a group `iu`, except for one user called `robot`, one would write:

```
mercury% setfacl -m group:iu:r--,user:robot:--- testfile
mercury% getfacl testfile
# file: testfile
# owner: mark
# group: iu
```

```
user::rw-
user:robot:---      #effective:---
user:ds:rw-         #effective:---
group::---          #effective:---
group:iu:r--        #effective:---
mask:---
other:---
```

Notice that this is accomplished by saying that the group has read permission whilst the specific user should have no permissions.

2.3.2 NT File Model

The NT operating system supports a variety of legacy file systems for backward compatibility with DOS and Windows 9x. These older file systems are insecure, in the sense that they have no mechanisms for restricting access to files. The file system NTFS was introduced with NT in order to solve this problem. The file system has gone through a number of revisions, and no doubt will go through many more before it reaches constancy.

NTFS, like the Unix file system, is a hierarchical file system with files and directories. Each file or directory has an owner, but no group membership. Files do not have a set of default permission bits, as does Unix; instead, they all have full-blooded ACLs, which assign a set of permission bits to a specific user. NTFS ACLs are similar to other access control list models, in file systems such as the AFS and DCE/DFS. They have all of the flexibility and all of the confusions which accompany ACLs, such as inheritance of attributes from parent directories and creation masks. The NTFS file system is indexed by a master file table, which serves an analogous function to Unix's inodes, though the details are somewhat different.

File System Layout

Drawing on its DOS legacy, NT treats different disk partitions as independent floppy disks, labelled by a letter of the alphabet:

```
A: B: C: D: ...
```

For historical reasons, drive `A:` is normally the diskette station, while drive `C:` is the primary hard disk partition. Other drive names are assigned at random, but often `H:` is reserved for partitions containing users' home directories. Unlike Unix, different devices are not sewn seamlessly into a unified file tree, though this will probably change in a future release of NT. Originally, DOS chose to deviate from its Unix heritage by changing the sub-directory separator from / to \. Moreover, since each device is treated as a separate entity, there is a root directory on every disk partition:

```
A: B: C: ...
```

and one has a notion of *current working drive*, as well as *current working directory*. These distinctions often cause confusion amongst users who work with both Unix and NT.

The layout of the NT file system has changed through the different versions, in an effort to improve the structure. This description relates to NT 4.0. The system root is usually stored in `C:\WinNT`, and is generally referred to by the system environment variable `%SystemRoot%`:

- `C:\I386` This directory contains binary code and data for the NT operating system. This should normally be left alone.
- `C:\Program Files` This is NT's official location for new software. Program packages which you buy should install themselves in subdirectories of this directory. More often than not, they choose their own locations, often with a distressing lack of discipline.
- `C:\Temp` Temporary scratch space, like Unix's `/tmp`.
- `C:\WinNT` This is the root directory for the NT system. This is mainly for operating system files, so you should not place new files under this directory yourself unless you really know what you are doing. Some software packages might install themselves here.
- `C:\WinNT\config` Configuration information for programs. These are generally binary files, so the contents of NT configuration files is not very interesting.
- `C:\WinNT\system32` This is the so-called system root. This is where most system applications and data files are kept.

File Extensions

Whereas files can go by any name in Unix, Microsoft operating systems have always used the concept of file extensions to identify special file types. For example:

file.EXE	An executable program
file.DOC	Word document
file.JPG	Graphic file format

Links and Shortcuts

Like Unix, NT also has ways of aliasing files in the file system. NT has hard links, or duplicate entries in the master file table, allowing one to associate several names with a given file. This is not a pointer to a file, but an alternative entry point to the same file. Although the file system structure of NTFS is different from the Unix file system, the idea is the same. Hard links are created from the POSIX compatibility subsystem, using the traditional Unix command name `ln`. As with Unix, hard links can only be made to files on the same disk partition.

A *short cut* is a small file which contains the name of another file. It is normally used for aliasing scripts or programs. Unlike Unix's symbolic links, short cuts are not handled transparently by the operating system; they are actual files which can be opened with a text editor. They must be read and dealt with at the application level. Short cuts can be given any name, but they always have the file extension `.LNK`. This suffix is not visible in the graphical user interface. They are created from the graphical user interface by right-clicking on the item one wishes to obtain a pointer to.

Unix compatibility packages like Cygwin32 use short cuts to emulate symbolic links. However, since short cuts work at the application level, what one package does with a short cut is not guaranteed to apply to other software, so the usefulness of short cuts is limited.

Access Control Lists

NT files and directories have the following attributes. Access control lists are composed of Access Control Entries (ACEs), which consist of these:

Permission bit	Files	Directories
R (Read)	See file contents	See directory contents
W (Write)	Modify file contents	Modify directory contents
X (Execute)	Executable program	Can `cd` to directory
D (Delete)	Deletable	Deletable
P (Permission)	Permissions changeable	Permissions changeable
O (Ownership)	Ownership changeable	Ownership changeable

The read, write and execute flags have the same functions as their counterparts in Unix. The execute flag is always set on `.EXE` files. The additional flags allow configurable behaviour, where behaviour is standardized in Unix. The delete flag determines whether or not a particular user has permission to delete an object (note that a user which has write access to the file can destroy its contents independently of this). The permission and ownership flags, likewise determine whether or not a specified user can take ownership or modify the permissions on a file.

Access control lists, or access control entries, are set and checked with either the NT Explorer program (File/Properties/Security/Permissions menu) or the `cacls` command. This command works in more or less the same way as the POSIX `setfacl` command, but with different switches. The switches are

/G	Grant access to user
/E	Edit ACE instead of replacing
/T	Act on all files and subdirectories
/R	Revoke (remove) access rights to a user
/D	Deny access rights to a given user

For example,

```
hybrid> CACLS testfile
C:\home\mark\testfile BUILTIN\Administrators:F
                      Everyone:C
                      MT AUTHORITY\SYSTEM:F

hybrid> CACLS testfile /G ds:F

Are you sure(Y/N)?

hybrid> CACLS testfile
C:\home\mark\testfile HYBRID\ds:F
```

In this example the original ACL consisted of three entries. We then replace it with a single entry for user `ds` on the local machine `HYBRID`, granting full rights. The result is shown in the last line. If, instead of replacing the ACE, we want to supplement it, we write

```
hybrid> CACLS testfile /E /G mark:R
{\var wait for 30 seconds}
Are you sure(Y/N)?

hybrid> CACLS testfile
C:\home\mark\testfile HYBRID\ds:F
                      HYBRID\mark:R
```

New Files: Inheritance

Although the technical details of the NTFS and its masking schemes are not well documented, we can note a few things about the inheritance of permissions. In the absence of any ACL settings on a parent directory, a new file is created, granting all rights to all users. If the parent directory has an ACL, then a new file inherits that ACL at the time of its creation. When a file is moved, it keeps its NTFS permissions, but when a file is copied, the copy behaves like a new file, inheriting the attributes of its new location.

2.3.3 Network File System Models

Unix and NT have two of the most prevalent file system interfaces, apart from DOS itself (which has no interface, since it has no security functionality), but they are both stunted in their development. In recent years, file system designers have returned to an old idea which dates back to a project from Newcastle University, called the Newcastle Connection, an experimental distributed file system which could link together many computers seamlessly into a single file tree [27]. To walk around the disk resources of the entire network, one simply used `cd` to change directory within a global file tree.

This idea of distributed file systems was partially adopted by Sun Microsystems in developing their Network File System (NFS) for Unix-like operating systems. This is a distributed file system, but only to a limited extent within a local area network[6]. Sun's use of open standards and a willingness to allow other vendors to use the technology quickly made NFS a *de facto* standard in the Unix world, overtaking alternatives like RFS. However, owing to vendor disagreement, the Network File System has been limited to the lowest common denominator Unix file system-model. Although many vendor-specific improvements are available, these do not work in a heterogeneous environment, and thus NFS is relatively featureless, by comparison with the functionality available on local disk file systems. In spite of this, there is no denying that NFS has been very effective, as is testified by the huge number of sites which use it unconditionally.

Another major file system, in a similar vein, is the Novell Netware file system. This is an interesting file system which can also create a seamless file tree called the Novell Directory Service (NDS) within an organization. Here files have an owner and an Access Control List, which can grant or restrict access to named users or groups. The NT model was presumably inspired by this. The Netware idea is not unlike NFS in attempting to integrate organizations' disks into a communal file tree, but the user interface is superior, since it is not limited by compatibility issues. However, Netware forces a particular object-oriented interpretation of the network onto disks, whereas NFS does not care about the file tree structure of hosts

[6] Sun have recently been developing NFS over wide area networks.

which incorporate shared file systems. With NFS, hosts do not have to subscribe to a global vision of shared network resources, they simply take what they want and maintain their own private file tree: each host could be kept quite different. Oddly enough, NT has not embraced the model of seamless sharing, choosing instead to mount drives on the old DOS drive letters `A:`, `B:` etc., though it is possible that such seamless integration will come in a future version. Novell too has to deal with this antiquity, since it serves primarily Windows-based machines.

While Solaris' NFS does support its own brand of Access Control Lists, NFS cannot be used to provide inter-platform ACL functionality. Netware does support its own state of the art file system attributes, based on the usual object inheritance model of directories as containers for smaller containers. Each file has an owner and an ACL (see Figure 2.3).

Long before NT's arrival, the developers of distributed operating systems were experimenting with more general distributed operating systems which could span wide area networks. Two file systems which were developed in this context are the Andrew File system (AFS) and the Distributed File System (DFS), which is a part of the Distributed Computing Environment (DCE). These file systems have been driven on by high energy physics laboratories the world over. Physicists' need to share data quickly and efficiently has long made them pioneers of networking technologies. AFS and DFS have been embraced widely in this context, allowing collaborators in Japan, Europe and the United States to be connected simply by changing directory to a new country, organization and site (see section 3.9.2). These file systems also employ Access Control Lists, based on, but not limited by, the Unix permission model (see Figure 2.3).

Flag	Rights acquired by named user in ACL
S	Supervisor grants all rights to a file, directory and all subdirectories
R	Ability to open and read a file or directory contents
W	Ability to open and write to a file or to add files to a directory
C	Ability to create new files and undelete old ones, or create new directories
E	Ability to erase (delete) a file or directory
M	Ability to modify file attributes including rename
F	Ability to see files within a directory when viewing contents
A	Ability to change access rights on file or directory, including granting others access rights. Also change inheritance masks for directories

Figure 2.3 Netware 5 permissions. New file objects inherit the default permissions of their container, minus any flags in the Inherited Rights Filter/Mask (IRF). Permissions can be applied to named users or groups

Flag	Rights acquired by named user, group, other in ACL
r	Ability to open and read to a file or directory contents
w	Ability to open and write to a file or to add files to a directory
x	Ability to execute files as programs or enter directories
d	Ability to erase (delete) a file or directory
c	Ability to modify file attributes including rename
i	Ability to add files to a directory

Figure 2.4 DFS permissions. New files inherit the initial object ACL of their parent directory. These flags can be applied to named lists of users, or groups or others, in the Unix sense

Note that the DCE/DFS file system is not related to NT's DFS file system, though the idea is similar.

As we can see, many of these file systems have drawn on the pioneering ideas of experimental file systems. Today, most file systems work in a similar way, with Unix lagging behind in sophistication, but not in functionality. Ironically, for all the flexibility that ACLs offer, they have proven to be confusing and difficult to understand, and the extra functionality they provide is dwarfed by the feeling of dread which they instill in administrators and users alike (see Figures 2.4, 2.5). On systems with only ACLs, file permissions tend to be set inappropriately more often than on Unix-like systems. Unix's simpler approach, while basically old and simplistic, is a more palatable and manageable alternative for all but the most sophisticated users.

2.3.4 Unix and NT Sharing

File systems can be shared across a network by any of the methods we have discussed above. We can briefly note here the correspondence of commands and methods for achieving

Flag	Rights acquired by named user, group in ACL
r	Ability to open and read to a file or directory contents
l	Lookup within a directory
w	Ability to open and write to a file
i	Ability to insert files in directories
d	Ability to erase (delete) a file or directory
a	Ability to modify file attributes including rename
k	Lock files

Figure 2.5 AFS permissions. These flags can be applied to named lists of users or groups but not 'others'. Four shorthand forms also exist: `write=rlidwk`, `read=rl`, `all=rlidwka`, and `none` removes an entry

network sharing. Unix-like hosts using NFS share file systems by running the daemons `rpc.mountd` and `rpc.nfsd`. File systems are shared by adding them to the file `/etc/exports`, on most systems, or to `/etc/dfs/dfstab` on SVR4-based Unix. The syntax in these files is particular to the flavour of Unix-like operating system one is using. With some operating systems, using `/etc/exports`, it is necessary to run the command `exportfs -a` to make the contents of the export file visible to the daemons which control access. On SVR4 systems, like Solaris, there is a command called `share` for exporting file systems, and the file `/etc/dfs/dfstab` is just a shell script containing a lot of share commands, e.g.

```
allhosts=nomad:vger:nomad.domain.country:vger.domain.
country share -F nfs -o rw=$allhosts /site/server/local
```

Here the command `shareall` is the equivalent for exporting all file systems in this file. It simple runs the shell script. The example above makes the directory tree `/iu/server/local` available to the hosts nomad and vger. Note that, due to different name services implementations and their various behaviours, it is often necessary to use both the unqualified and fully qualified names of hosts when sharing.

On the client or receiving end, we attach a shared file system to a host by 'mounting' it. NFS file systems are mounted in exactly the same way as they mount a local disk, i.e. with the `mount` command. The name of the server-host which shared or exported the file system becomes like the name of the disk, e.g.

```
mkdir -p /site/server/local
mount server:/site/server/local /site/server/local
```

Here we create a directory on which to mount a foreign file system and then mount it on a directory which has the same name as the original on the server. The original name and the new name do not have to be the same, but there is a point to this which we shall return to later. Assuming that the server-host granted us the right to mount the file system on our host, we now have access to the remote file system, as though it were a local disk. The only exception is the superuser root, who is granted no access rights at all to the file system. In fact the user ID of root gets mapped to a special user called `nobody`. The point of this is that the administrator on the client host is not necessarily the administrator on the server host, and has no obvious right to every users' files there. Privileged access can be arranged during NFS sharing, but it is not a recommendable practice.

NT file systems on a server are shared, either using the GUI, or by executing the command

```
net share alias=F:\filetree
```

On the client side, the file tree can then be 'mounted' by executing the command

```
net use X: \\serverhost\alias
```

This attaches the remote file tree, referenced by the alias, to NT drive `X:`. One of the logistical difficulties with the NT drive model is that drive associations are not constant, but might change when new hardware is detected. Drive associations can be made to persist by adding a flag

```
net use X: \\serverhost\alias /persistent:yes
```

to the mount command. This is not a perfect solution, but it works.

2.4 Processes and Job Control

On a multitasking computer, all work on a running program is performed by an abstraction called a *process*. This is a collection of resources such as file handles, allocated memory, program code and CPU registers which is associated with a specific running program. A cursory overview of the way operating systems handle running programs is beneficial to enable us to manage processes for the system and for users. On modern operating systems, processes can contain many concurrent threads which share the processes' resources.

2.4.1 The Unix Process Model

Unix starts new processes by copying old ones, a historical ritual which has to do with the duplication of system services in response to requests. Users start processes from a *shell* command line interface program or by clicking on icons in a window manager. Every Unix process has a Process ID (PID) which can be used to refer to it, suspend it or kill it entirely.

A background process is started from a shell using the special character & at the end of the command line:

```
find / -name '*lib*' -print >& output &
```

The final & on the end of this line means that the job will be run in the background. Note that this is not confused with the redirection operator >&, since it must be the last character on the line. The command above looks for any files in the system containing the string 'lib' and writes the list of files to a file called 'output'.

If we want to see what processes are running, we can use the `ps` command. `ps` without any arguments lists all of your processes, i.e. all processes owned by the user name you logged in with. `ps` takes many options, for instance `ps auxg` will list all processes in detail on BSD-like systems, while `ps -efl` will provide a similar, if not compatible, listing on System V-like systems. Some Unix-like systems support both the BSD and System V flags to the `ps` command.

Processes can be stopped and started, or killed once and for all. The `kill` command does this and more. In fact, it sends generalized signals to running processes, not only the kill signal. There are two versions of the `kill` command: one of them is built into the C-shell and the other is not. If you use the C-shell then you will never care about the difference unless the process table is full. We shall nonetheless mention the special features of the C-shell built-ins below. The kill command takes a number called a *signal* as an argument, and another number called the *process identifier*, or *PID* for short. Kill sends signals to processes. Some of these are fatal and some are for information only. The two commands

```
kill -15 127
kill 127
```

are identical. They both send signal 15 to PID 127. This is the normal *termination* signal, and it is often enough to stop any process from running.

Programs can choose to ignore certain signals by trapping signals with a special handler. One signal they cannot ignore is signal 9:

```
kill -9 127
```

is a sure way of killing PID 127. Even though the process dies, it may not be removed from the kernel's process table if it has a parent (see the next section).

2.4.2 Child Processes and Zombies

When we start a process, the new process becomes a *child* of the original. If one of the children starts a new process then it will be a child of the child (a grandchild). Processes therefore form *hierarchies*. Several children can have a common *parent*. All Unix user-processes are children of the initial process `init`, with process ID 1.

If we kill a parent, then (unless the child has detached itself from the parent) all of its children die too. If a child dies, the parent is not affected. Sometimes when a child is killed, it does not die but becomes *defunct* or a *zombie* process. This means that the child has a parent which is *waiting* for it to finish. If the parent has not yet been informed that the child has died, because it has been suspended itself, for instance, then the dead child is not completely removed from the kernel's process table. When the parent wakes up and receives the message that the child has terminated (and its exit status), the process entry for the dead child can be removed.

Most Unix processes go through a zombie state, but most terminate so quickly that they cannot be seen. A few hang around and use up valuable process slots, which can be a problem. It is not possible to kill a zombie process, since it is already dead. The only way to remove a zombie is to either reactivate the process which is waiting for it, or to kill that process. Persistent zombie processes are usually caused by software bugs.

2.4.3 The Windows/NT Process Model

Like Unix, processes under Windows/NT can live in the foreground or background, though unlike Unix, NT does not fork processes by replicating existing ones. A background process can be started with

```
start /B
```

To kill the process it is necessary to purchase the Resource kit which contains a `kill` command. A background process detaches itself from a login session, and can continue to run even when the user is logged out.

Generally speaking, the reborn PC world of NT and Novell abhors processes. Threads are the preferred method for multitasking. This means that additional functionality is often implemented as modules to existing software, rather than as independent objects.

The shutdown of the whole system must be performed from the Windows menu. Any logged on user can shut down a host. This is not a major problem since only one user can use the system at a time. However, background processes die when this happens, including other users' background processes. A shutdown command also exists for shutting down local or remote systems.

2.4.4 Environment Variables

Environment variables are text-string variables which can be set in any process [253]. Normally they are set by users in shell environments in order to communicate user

preferences or configuration information to software. In the C shell, they are set with the command

```
setenv VARIABLE value
```

and are not to be confused with the C shell's local (non-inherited) variables which are created with `set variable=`*`value`*. In the original Bourne shell they are set by

```
VARIABLE=value
export VARIABLE
```

The `export` command is needed to make the variable global, i.e. to make it inheritable by child processes. In newer Bourne shells like `ksh` and `bash`, one can simply write

```
export VARIABLE=value
```

The values of these variables are later referred to using the dollar symbol:

```
echo VARIABLE
```

When a process spawns a child process, the child inherits the environment variables of its parent. Environment variables are an important way of transmitting preference information between processes.

On NT systems, environment variables are set in the DOS prompt interface by

```
set VARIABLE=value
```

Try not to confuse this with the C shell's `set` command. Environment variables in NT are later dereferenced using the percent prefix and suffix:

```
echo %%VARIABLE%%
```

2.5 Logs and Audits

Operating system kernels share resources and offer services. They can be asked to keep lists of transactions which have taken place so that one can later go back and see exactly what happened at a given time. This is called logging or auditing.

Full system auditing involves logging every single operation which the computer performs. This consumes vast amounts of disk space and CPU time, and is generally inadvisable unless one has a specific reason to audit the system. Part of auditing used to be called system accounting from the days when computer accounts really were accounts for real money. In the mainframe days, users would pay for system time in dollars and thus accounting was important, since it showed who owed what [105], but this practice remains mainly on large supercomputing installations today. Auditing has become an issue again in connection with security. Organizations have becomes afraid of break-ins from system crackers, and want to be able to trace the activities of the system in order to be able to look back and find out the identity of a cracker. The other side of the coin is that system accounting is so resource consuming that the loss of performance might be more important to an organization than the threat of intrusion.

For some organizations auditing is important, however. One use for auditing is so-called *non-repudiation*, or non-denial. If everything on a system is logged, then users cannot back

away and claim that they did not do something: it's all there in the log. Non-repudiation is a security feature which encourages users to be responsible for their actions.

2.6 Privileged Accounts

Operating systems which restrict user privileges need an account which can be used to configure and maintain the system. Such an account must have access to the whole system, without regard for restrictions. It is therefore called a privileged account.

In Unix the privileged account is called *root*, also referred to colloquially as the super-user. In NT, the *Administrator* account is similar to Unix's root, except that the administrator does not have automatic access to everything as does root. Instead he/she must first be granted access to an object. However, the Administrator always has the right to grant him or herself access to a resource, so in practice this feature just adds an extra level of caution. These accounts place virtually no restriction on what the account holder can do. In a sense, they provide the privileged user with a skeleton key, a universal pass to any part of the system.

Administrator and root accounts should never be used for normal work: they wield far too much power. This is one of the hardest things to drill into novices, particularly those who have grown up using insecure operating systems. Such users are used to being able to do whatever they please. To use the privileged account as a normal user account would be to make the systems as insecure as the insecure systems we have mentioned above.

Principle 1 (Privilege) *Restriction of unnecessary privilege protects a system from accidental and malicious damage, and infection by viruses, and prevents users from concealing their actions with false identities. It is desirable to restrict users' privileges for the greater good of everyone on the network.*

The purpose of the root/Administrator accounts is usually misunderstood by newcomers. The romantic notion of being in control of one's own computer system, with many users and a huge network, holds such fascination for many that they are seduced by its lure. Inexperienced users aspire to gain administrator/root privileges as a mark of status. This can generate the myth that the purpose of this account is to gain power over others. In fact the opposite is true: privileged accounts exist precisely because one does *not* want to have too much power, except in exceptional circumstances. The corollary to our first principle is this:

Corollary 2 (Privilege) *No one should use a privileged root/Administrator account as a user account. To do so is to place the system in jeopardy.*

Gigabyte Ethernet	1000
100 Base-TX (Fast ethernet)	100
Etherlink XL 10 BT PCI	133
100 Base-TX	100
100 Base-FX Fibre	100
Token ring	16
Ethernet	10

Figure 2.6 Nominal network protocol speeds in Megabits per second

One of the major threats to Internet security has been the fact that everyone can now be root/Administrator on their own host. Many security mechanisms associated with trusted ports, TCP/IP spoofing, etc., are now broken, since all of the security of these systems lies in the outdated assumption that ordinary users will not have privileged access to network hardware and the kernel. Various schemes for providing limited privilege through special shells, combined with the setuid mechanism in Unix, have been described [124, 45]. See also the amusing discussion by Simmons on use and abuse of the superuser account [245], and an administration scheme where local users have privileges on their own hosts [66].

2.7 Hardware Awareness

To be a system administrator it is not absolutely essential to know much about hardware, but it is very useful to have a basic appreciation of hardware installation procedures and how to nurse-maid hardware later. If you do not feel comfortable handling hardware, then don't do it.

- *Read instructions:* when dealing with hardware, one should always look for and *read* instructions in a manual. It is foolish to make assumptions about expensive purchases. Instructions are there for a reason.
- *Interfaces and connectors:* hardware is often connected to an interface by a cable or connector. Obtaining the correct cable is of vital importance. Many manufacturers use cables which look similar, superficially, but which actually are different. An incorrect cable can result in damage to an interface. Modem cables in particular can damage a computer or modem if they are incorrectly wired, since some computers supply power through these cables which can damage equipment which does not expect to find a power supply coming across the cable.

  Network interfaces are often built-in; they can always be added with expansion cards. Some interfaces are for thin ethernet, some for thick ethernet and others are for twisted pair connectors. Some are for Token Rings. If a computer has the wrong type of interface, it is necessary to buy a transceiver which converts the signal and connection to the right type, or in the worst case, a new interface. The type of network a host uses is determined by the hardware and the protocol which is used by the hardware. This can have a significant effect on the performance of the network (see Figure 2.7).
- *Handling components:* modern day CMOS chips work at low voltages (typically 5 volts or lower). Standing on the floor with insulating shoes, you can pick up a static electric

IDE	2.5
ESDI	3
SCSI-2	5
Fast SCSI-2	10
Fast wide SCSI-2	20
Ultra SCSI	40
Ultra-2 SCSI	80
Ultra-3 SCSI	160

Figure 2.7 Nominal disk drive speeds in Mbytes/sec [70] for various standards

charge of several thousand volts. Such a charge can instantly destroy computer chips. Before touching any computer components, earth yourself by touching the metal casing of the computer. If you are installing equipment inside a computer, wear a conductive wrist strap.

- *Disks:* the most common disk types are IDE (integrated drive electronics) and SCSI (small computer software interface). IDE disks are usually cheaper than SCSI disks, but SCSI disks are more efficient at handling multiple accesses, and are therefore better in multi-tasking systems. They are also much faster (see Figure 2.7). SCSI [174] comes in several varieties, SCSI 1, SCSI 2, wide SCSI, fast-wide, etc. The difference has to do with the width of the data-bus and the number of disks which can be attached to each controller. There are presently three SCSI standards: SCSI-1, SCSI-2 and SCSI-3. The SCSI-2 standard also defines wide, fast and fast/wide SCSI. SCSI-1 defines an 8-bit data bus and limits transfer speeds to 5 Mbytes/sec. Fast SCSI also uses an 8-bit bus, but permits transfer rates of up to 10 Mbytes/sec. Wide SCSI doubles the bus size to 16 bits. When combined with fast SCSI it allows transfer rates of up to 20 Mbytes/sec., etc. The SCSI standard is improving and developing all the time.

  Each SCSI disk has its own address (or number) which must be set by changing a setting on the disk-cabinet, or by changing jumper settings inside the cabinet. Newer disks have programmable identities.

  Disk chains must be terminated with a proper terminating connector. Newer disks often contain automatic termination mechanisms.

- *Memory:* memory chips are sold on small circuit boards called SIMMs. These SIMMs are sold in different sizes, and with different speeds. A computer has a number of slots where SIMMs can be installed. When buying and installing RAM, remember that
  - The physical size of SIMMs is important. Most have 72 pins and some older SIMMs have 30 pins.
  - SIMMs are sold in 1MB, 4MB, 16MB, 64MB sizes, etc. Find out what size you can use in your system. In most cases you are not allowed to mix different sizes.
  - Do not buy slower RAM than that which is recommended for your computer, or it will not work.
  - There are several incompatible kinds of RAM, FP RAM, EDO RAM, SDRAM, such as which work in different ways. ECC/SDRAM RAM (error correcting code, synchronous dynamic RAM) is tolerant to error from external noise sources like cosmic rays, etc. It can be recommended for important servers.
  - On some computers you need to fill up RAM slots in a particular order, otherwise the system will not be able to find them.

Some hosts come with a basic 'monitor' panel which is a ROM-based program that can be used to set the configuration of Non Volatile RAM variables even before the system has booted up a true operating system. On a PC system this corresponds to the BIOS settings. You should familiarize yourself with the kinds of variables which can be set in NVRAM for your system. These could control basic choices about how your machine works, like which network interface is to be used: thick ethernet or twisted pair, etc.

On solaris hosts there is a program called `eeprom` which can be used to set values in NVRAM.

Weather and environment affect computers:

- *Lightning:* strikes can destroy fragile equipment. No fuse will protect hardware from a lightning strike. Transistors and CMOS chips burn out much faster than any fuse. Electronic spike protectors can help here.
- *Power:* failure can cause disk damage and loss of data. A UPS (Uninterruptible Power Supply) can help.
- *Heat:* the blazing summer heat or a poorly placed heater can cause systems to overheat and suddenly black out. One should not let the ambient temperature near a computer rise much above about 25 °C. Heat can cause RAM to operate incorrectly and disks to misread/miswrite. Good ventilation is essential for computers and screens for avoiding electrical faults.
- *Cold:* sudden changes from hot to cold are just as bad. They can cause unpredictable changes in the electrical properties of chips and cause systems to crash. In the long term, these changes could lead to cracks in the circuit boards and irreparable chip damage.

2.8 System Uniformity

Given the chance to choose the hardware at a site, it is wise to spend time picking out reliable, standard hardware and software. The more different kinds of system we have, the more difficult the problem of installing and maintaining them. This is a basic principle.

Principle 3 (Uniformity) *A uniform configuration minimizes the number of differences and exceptions one has to take into account later. This applies to hardware and software alike.*

PC networks are often a melange of random parts from different manufacturers. If possible, one should standardize graphics and network interfaces, disk sizes, mice and any other devices which have to be configured. This means that, not only will it be easier to configure and maintain, but also that it will be easier to buy extra parts or cannibalize systems for parts later. PCs are the biggest problem, since manufacturers of Unix workstations tend to ensure that their parts are either industry standard or that hardware differences are not visible to users.

With software, the same principle applies: a uniform software base is easier to install and maintain than one in which special software needs to be configured in special ways. Few methods are available for handling the *differences* between systems; most administration practices are based on standardization. Cfengine is a tool which can help in maintaining differences between machines in a rational way, but the existence of such a tool should not be seen as an excuse for not simplifying the logistics of maintenance as much as possible.

Exercises

Exercise 2.1 Compare and contrast Windows NT with Unix-like operating systems. If you need a refresher about Unix, consider our online textbook at Oslo College [31].

Exercise 2.2 Under what circumstances is it desirable to use a Graphical User Interface (GUI), and when is it better to use a command language to address a computer? (If you answer *never* to either of these, you are not thinking.)

Exercise 2.3 The purpose of this exercise is to make yourself familiar with a few Unix tools which you will need to use to analyse networks later. Remember that the aim of this course is to make you self-sufficient, not to force-feed you information. This exercise assumes that you have access to a Unix-like operating system.

(a) Use the `rlogin` command to log onto a host in your domain.
(b) Use the command `uname` with all of its options to find out what type of host it is.
(c) Familiarize yourself with the commands `df`, `nslookup`, `mount`, `finger .clients` (GNU finger). What do these commands do and how can you use them?
(d) Start the program `nslookup`. This starts a special shell. Assuming that your local domain is called domain.country, try typing

```
> ls domain.country
```

If you get an error, you should ask your administrator why. The ability to list a domain's contents can be restricted for security reasons. Then try this and explain what you find:

```
> set q=any
> domain.country
```

Exercise 2.4 Review the principal components in a computer. Explain any differences between an electronic calculator and a PC.

Exercise 2.5 Review the concept of *virtual memory*. If you do not have access to a textbook on operating systems, see my on-line textbook [31]. What is swapping and what is paging? Why is paging to a file less efficient than paging to a raw partition?

Exercise 2.6 Explain how a file system solves the problem of storing and retrieving files from a storage medium, such as a disk. Explain how files can be identified as entities on the magnetic surface.

Exercise 2.7 Locate the important log files on your most important operating systems. How do you access them, and what information do they contain? You will need this bird's-eye view of the system error messages when things go wrong. (Hint: there are log files for system messages, services like WWW and FTP, and for mail traffic. Try using `tail -f logfile` on Unix-like hosts to follow the changes in a log file. If you don't know what it does, look it up in the manual pages.)

Exercise 2.8 Describe an access control list. Compare the functionality of the Unix file permission model with that of access control lists. Given that ACLs take up space and have many entries, what problems do you foresee in administrating file security using ACLs?

Chapter 3

Networked Communities

The network is the largest physical appendage to our computer systems, but it is also, the least conspicuous, hiding behind walls and in locked switching rooms. To most users, the network is a piece of magic which they have abruptly learned to take for granted, and yet without it, modern computing practices would be impossible.

We cannot learn anything about a community of networked computer systems without knowing where all the machines are, what their purpose is, and how they interrelate to one another. If we were starting from scratch with a computer network, we would like to look at what work we want to perform, decide what kinds of machines we require to perform it, how their resources need to be shared throughout the network, and finally, put everything together. The alternative (to buy lots of things and see what we can do with them later) lacks a certain vision.

If not starting from scratch, but with an existing network, serviceable or not, the first step is clearly to acquaint ourselves with what material we had to work with. In short, unless we are beginning with a completely blank slate, we need to survey our kingdom, to see how the land lies.

The aim of this chapter is to learn how to navigate network systems using standard tools, and place each piece of the puzzle in context.

3.1 Communities

System administration is not just about machines and individuals, it is about communities. There is the local community of users on multi-user machines; then there is the local area network community of machines at a site; finally, there is the global community of all machines and networks in the world. The fundamental principle of communities is:

Principle 4 (Communities) *What one member of a cooperative community does affects every other member, and vice versa. Each member of the community therefore has a responsibility to consider the well-being of the other members of the community.*

When this principle is ignored, it leads to conflict. We attempt to preserve the well-being of a community by making *rules*, *laws* or *policies*. The difference between these labels only amounts to our opinion of their severity. Rules and laws do not exist because there are fundamental rights and wrongs in the world. They exist because there is a need to summarize the consensus of opinion in a community group. This has two purposes:

- To provide a widely accepted set of conventions which simplify decisions by avoiding the need to think through things from first principles every time.
- To document the will of the community for reference.

Rules can never cover every eventuality. They are a convenient approximation which hopes to cover common situations. In an ideal world, they would never be used as a substitute for thought. However, this is often not the case. We can rewrite the central axiom for the user community of a multi-user host:

Principle 5 (Multi-user communities) *A multi-user computer system does not belong to any one user. All users must share the resources of the system. What one user does affects all other users, and vice versa. Each user has a responsibility to consider the effect of his/her actions on all the other users.*

and also for the worldwide network community:

Principle 6 (Network communities) *A computer which is plugged into the network is no longer just ours. It is part of a society of machines which shares resources and communicates with the whole. What that machine does affects other machines. What other machines do affects that machine.*

The ethical issues associated with connection to the network are not trivial, just as it is not trivial to be a user in a multi-user system, or a member of a civil community. Administrators are responsible for their organization's conduct to the entire rest of the Internet. This great responsibility should be borne wisely.

3.2 User Sociology

Most branches of computer science deal primarily with software systems and algorithms. System administration is made more difficult by the fact that it deals with communities and is therefore strongly affected by what human beings do. In short, a large part of system administration is *sociology*.

A newly installed machine does not usually require attention until it is first used, but as soon as a user starts running programs and storing data, the reliability and efficiency of the system are tested. This is where the challenge of system administration lies.

The load on computers and on networks is a social phenomenon: it peaks in response to patterns of human behaviour. For example, at universities and colleges network traffic usually peaks during lunch breaks, when students rush to the terminal rooms to surf on the web or to read e-mail. In industry the reverse can be true, as workers flee the slavery of their computers for a breath of fresh air (or carbonized air). To understand the behaviour of

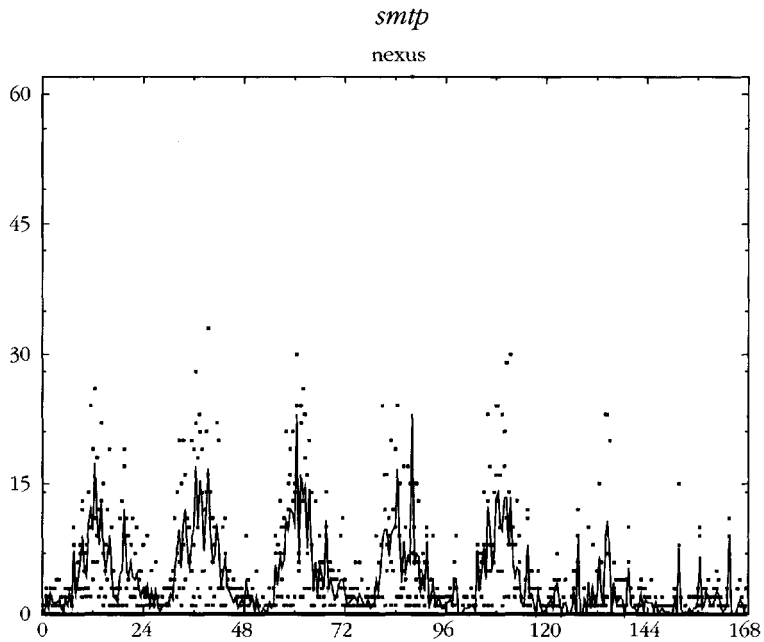


Figure 3.1 E-mail traffic at Oslo College measured over the course of many weeks. The plot shows the weekly average from Monday to Sunday. Over each 24-hour period, there is a daily peak showing users' working hours, and during the week there is a peak around midweek, and little activity during the weekends

the network, the load placed on servers and the availability of resources, we have to take into account the users' patterns of behaviour (see Figure 3.1).

3.3 Client-Server Cooperation

At the heart of all cooperation in a community is a system of *centralization* and *delegation.* No program or entity can do everything alone, nor is it expected to do so. It makes sense for certain groups to specialize in performing certain jobs. That is the function of a society.

Principle7 (Delegation I) *Leave experts to do their jobs. Assigning responsibility for a task to a body which specializes in that task is an efficient use of resources.*

If we need to find out telephone numbers, we invent the directory enquiry service: we give a special body a specific job. They do the phone-number research (once and for everyone) and have the responsibility for dealing out the information on request. If we need a medical service, we train doctors in the specialized knowledge and trust them with the responsibility. That is much more efficient than expecting every individual to have to research phone numbers by themselves, or to study medicine personally. The advantage with a *service* is

that one avoids repeating work unnecessarily and one creates special agents with an aptitude for their task.

In computer systems the same thing applies. Indeed, in recent years the number of client-server systems has grown enormously, because of possibilities offered by networking. Not only can we give a special daemon on one host a special job, but we can say that that daemon will also do the job for every other host on the network. As long as the load placed on the network does not lead to a bottleneck, this is a very efficient centralization of resources. Clearly, the client server model is an extended way of sharing resources. In that sense, it is like a distributed generalization of the kernel[1].

The client-server nomenclature has been confused by history. A server is not a host, but a program or process which runs on a host. A client is any process which requires the services of a server. In Unix-like systems, servers are called daemons. In NT they are just called services. Unfortunately, it is common to refer to the host on which a server process runs as being a server. This causes all sorts of confusion.

The name 'server' was usurped, early on, for a very specific client-server relationship. A server is often regarded as a large machine which performs some difficult and intensive task for the clients (an array of workstations). This prejudice comes from the early days, when many PC-workstations were chained together in a network to a single PC which acted as file server, and printer server, sharing a disk and printer to all of the machines. The reason for this architecture, at the time, was that the operating system of that epoch, MS-DOS, was not capable of multi-tasking, and thus the best solution one could make was to use a new PC for each new task. This legacy of one-machine, one-user, one-purpose, still pervades newer PC operating system philosophy. Meanwhile, Unix and later experimental operating systems have continued a general policy of any machine, any job, as part of the vision of *distributed operating systems.*

In fact, a server-host can be anything from a Cray to a laptop. As long as there is a process which executes a certain service, the host is a server-host.

3.4 Host Identities and Name Services

Whenever more than one computer is coupled together, there is a need for each computer to have a unique identity. As it turns out, this need has been recognized many times, and the result is that today's computer systems can have many different names which identify them in different contexts. The result is a confusion. For Internet-enabled machines, the IP address of the host is usually sufficient for most purposes. A host can have all of the following:

- *Host ID:* circuit board identity number. Often used in software licensing.
- *Install name:* configured at install time. This is often compiled into the kernel, or placed in a file like `/etc/hostname`. Solaris adds to the confusion by also maintaining the install name in `/etc/hostname.le0` or an equivalent file for the appropriate network interface, together with several files in `/etc/net/*/hosts`.

[1] In reality, there are many levels at which the client-server model applies. For example, many system calls can be regarded as client-server interactions, where the client is any program and the server is the kernel.

- *Application level name:* any name used by application software when talking to other hosts.
- *Local file mapping:* originally the Unix `/etc/hosts` file was used to map IP addresses to names, and *vice versa*. Other systems have similar local files, to avoid looking up on network services.
- *Network Information Service:* a local area network database service developed by Sun Microsystems. This was originally called Yellow Pages, and many of its components still bear the 'yp' prefix.
- *Transport level address(es):* each network interface can be configured with an IP address. This numerical converts into a text name through a name service.
- *Network level address(es):* each network interface (Ethernet/FDDI, etc.) has a hardware address burned into it at the factory, also called its MAC address, or Media Access Control address. Some services (e.g. RARP) will turn this into a name or an IP address through a secondary naming service like DNS.
- *DNS name(s):* the name returned by a domain name server (DNS/BIND) based on an IP address key.
- *WINS name(s):* the name returned by a WINS server (Microsoft's name server) based on the IP address.

Different hardware and software systems use these different identities in different ways. The host ID and network level addresses simply exist. They are unique and nothing can be done about them, short of changing the hardware. For the most part they can be ignored by a system administrator. The network level MAC address is used by the network transport system for end-point data delivery, but this is not something which need concern most system administrators. The network hardware takes care of itself.

At boot time, each host needs to obtain a unique identity. In today's networks that means a unique IP address and an associated name for convenience. The only purpose for this name is to uniquely identify the host amongst all of the others on the worldwide network. Although every host has a unique Ethernet address or token ring address, these addresses do not fall into a hardware-independent hierarchical structure. In other words, Ethernet addresses cannot be used to route messages from one side of the planet to the other in a simple way. To make that happen, a system like TCP/IP is required. At boot-time, then, each host needs to obtain an Internet identity. It has two choices:

- Ask for an address to be provided from a list of free addresses (DHCP or BOOTP protocols).
- Always use the same IP address, stored on its system configuration files (requires correct information on the disk).

The first of these possibilities is sometimes useful for terminal rooms containing large numbers of identical machines. In that case, the specific IP address is unimportant as long as it is unique. The second of these is the preferred choice for any host which has special functions, particularly hosts which provide network services. Network services should always be at a well-known, static location.

From the IP address a name can be automatically attached to the host through an Internet *naming service*. There are several services which can perform this conversion. DNS, NIS and

WINS are the three prevalent ones. DNS is the superior service, based on a worldwide database; it can determine hostname to IP address mappings for any host in the world. NIS (Unix) and WINS (Windows) are local network services which are essentially redundant as name services. They continue to exist because of other functions which they can perform.

As far as any host on a TCP/IP network is concerned, a host is its IP address and any names associated with that address. Any names which are used internally by the kernel, or externally, are quite irrelevant. The difficulty with having so many names, quite apart from any confusion which humans experience, is that naming conflicts can cause internal problems. This is an operating system dependent problem, but as a general rule, if we are forced to use more than one naming service, we must be careful to ensure complete consistency between them.

The only worldwide service in common use today is DNS (the Domain Name Service) whose common implementation is called BIND (Berkeley Internet Name Domain). This associates IP addresses with a list of names. Every host in the DNS has a *canonical name*, or official name, and any number of *aliases*. For instance, a host which runs several important services might have the canonical name

```
mother.domain.country
```

and aliases,

```
www.domain.country
ftp.domain.country
```

DNS binds a local network to the worldwide Internet in several important ways. It makes it possible for data to organizations to be spread across the surface of the planet at any location,

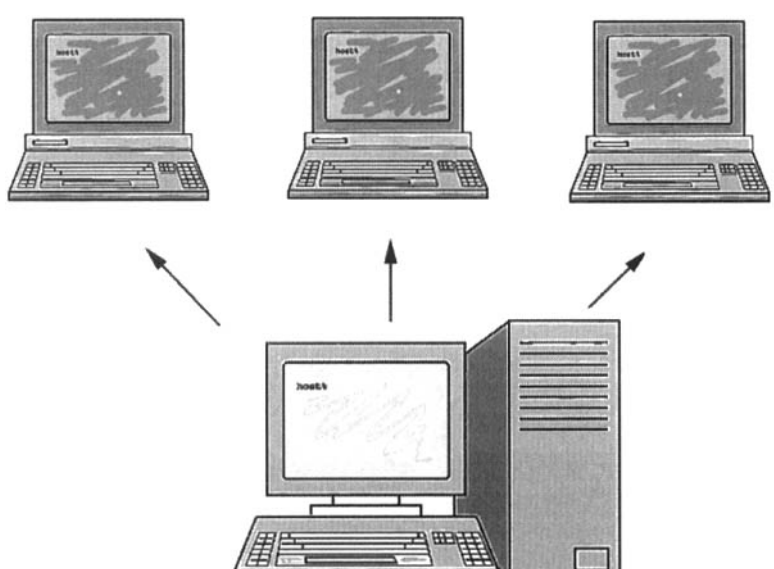


Figure 3.2 Some network infrastructure models single out a special server-host, which is used to consolidate network services and resources. Such a centralization has many administrative advantages, but it concentrates load and can create a bottleneck

and yet still maintain a transparent naming structure. E-mail services use the DNS to route mail.

WINS (Windows Internet Name Service) is a proprietary system built by Microsoft for Windows. Since any local host can register data in this service, it is insecure and is therefore inadvisable in any trusted network. Rumours suggest that WINS will be abandoned as of NT 5.0. NIS also provides a local service which has multiple database functions, not just hostname/IP address conversion. It can also be avoided in favour of DNS for hostname mapping.

Under NT, each system has an alphanumeric name which is chosen during the installation. A domain server will provide an SID (Security ID) for the name which helps prevent spoofing. When NT boots it broadcasts the name across the network to see whether it is already in use. If the name is in use, the user of the workstation is prompted for a new name(!)

The security of a name service is of paramount importance, since so many other services rely on name services to establish identity. If one can subvert a name service, hosts can be tricked into trusting foreign hosts and security crumbles.

3.5 Common Network Sharing Models

During the 1970s it was realized that expensive computer hardware could be used most cost-efficiently (by the maximum number of people) if it was available *remotely*, i.e. if one could communicate with the computer from a distant location. Inter-system communication became possible through the use of modems and UUCP, and later wide area networks.

The large mainframe computers which served sometimes hundreds of users were painfully slow for interactive tasks, although they were efficient at time-sharing. As hardware became cheaper, many institutions moved towards a model of smaller computers coupled to file-servers and printers by a network. This solution was relatively cheap but had problems of its own. At this time the demise of the mainframe was predicted. Today, however, mainframe computers are very much alive for computationally intensive tasks, while the small networked workstation provides access to a world of resources via the Internet.

Dealing with networks is one of the most important aspects of system administration today. The network is our greatest asset and our greatest threat. To be a system administrator it is necessary to understand how and *why* networks are implemented, using a worldwide protocol: the Internet protocol family. Without getting too heavily bogged down in details which do not concern us at an elementary level, we shall explore these themes throughout the remainder of this book.

3.5.1 Constraints on Infrastructure

Different operating systems support different ideas about how networks should be used. We are not always free to use the hardware resources as we would like. Operating system technologies restrict the kind of infrastructures it is possible to build in practice.

- *Unix:* much, if not all, of Unix's success can be attributed to the astonishing freedom which is granted to its users and administrators. Without a doubt, Unix-like operating systems are the most configurable and adaptable ever created. This has kept Unix at the forefront of new technology, but has also created a class of operating systems rather like disorganized piles of treasure in Aladdin's cave.

Unix-like operating systems are not tied to any specific model for utilizing network resources, though vendors sometimes introduce specific technologies for sharing which favour a particular kind of model. (This is almost viewed as a treasonable offense, and is usually quickly rejected in favour of software which offers greater freedom.) Unix lets us decide how we want the network to look. Any Unix system can perform any function, as server, client or both. A Unix network is fully distributed, there is no requirement about centralization of resources, but central models are commonly used. Unix contains troves of tools for making many hosts work together and share resources, but each host can also be configured as a standalone system. Each host either has a fixed IP address, or can be assigned one dynamically at boot time by a service such as BOOTP or DHCP.

- *NT:* NT networks are built around a specific model. There are two types of NT system with separate software licenses: *workstations* and *servers*. NT can work as a standalone system or as a workstation, integrated into a system of network services. NT revolves around a model in which programs are run on a local workstation, but where network services and resources are kept and run on a centralized server. IP addresses may be fixed or may be assigned automatically by a network service such as BOOTP or DHCP. Several NT servers can coexist. Each server serves a logical group of hosts, users and services called a *domain*. Client NT machines subscribe to as many domains as they wish, or have permission to join. NT is not a distributed system in the sense that services are localized on server machines. NT supports two kinds of organizational groups: *workgroups* in which hosts share a simple peer-to-peer network, perhaps with Windows 9x machines, and *domains* which have a central authority through a domain server. Domains provide a common framework including user-id's (SID's in NT language), passwords and user profiles. Domains have a common user-database and a common security policy. Any host which subscribes to a domain inherits the users and the security policy of the domain. NT domains can be simulated by Unix-like hosts [67].
- *Novell Netware:* the Novell Netware software [99] has been through five major versions, each of which has been significantly different. To begin with, Netware was little more than a disk and printer server for small PC networks. It found wide acceptance due to its broad support of different network interface technologies. Today, Netware version 5 is a fully distributed, object-oriented remote procedure service. Novell Netware is not an operating system *per se*. It is a network service for PCs which adds file storage, printing and other network services on top of the basic operating system: Windows, DOS, Macintosh or GNU/Linux. The network protocol for local traffic is IPX, which is lighter than IP and is an inter-networking protocol, but it is not a worldwide protocol, thus Novell-run PCs still need IP configurable interfaces. Each PC can have a fixed or dynamically allocated IP address, with a BOOTP or DHCP broadcast request. In Netware 5, several Novell file servers can coexist to provide a seamless Network Directory Service (NDS), an object-based service model. All services run on these servers, which support a form of modular thread-based multitasking. Novell services are not distributed arbitrarily amongst the PCs which it serves, as with Unix: they require one or more special dedicated machines to work on behalf of users' PCs, more like NT. The client machines must run Netware client software in order to communicate transparently with the servers. Although the nomenclature is different to that of NT domains, all the same functionality and more is available in the Novell software.

- *Windows 2000:* is a reincarnation of Windows NT. It redresses many of the shortcomings of the NT domain model by moving towards Novell-like directory services as its new model for resource organization. It allows remote services such as terminal login, which was only introduced as an afterthought in NT. Upgrading from NT to Windows 2000 can be a non-trivial problem, since it requires a rethink of basic infrastructure.
- *Macintosh:* each Macintosh is an independent system. Simple services like `ftp` can be run in a limited way from a normal machine. Macintosh uses its own network protocol called Appletalk which is incompatible with IP and IPX. Appletalk servers allow networking and disk sharing. IP protocol disk sharing is available, but does not mix well with the Macintosh file system. System administration (actually everything) is by GUI only. Recently, Macintosh have released a new operating system based on emulation of Unix and old-style Macintosh. The Mac OS Server X provides a powerful server based on Mach kernel technology and BSD Unix, to rival Novell's Netware and NT. Its approach seems to be to emulate the policies and practices of its predecessors, MacOS, Windows and Novell, rather than to embrace a new distributed approach. However, Mac OS Server X is based on a more robust and flexible technology than most of the above, so its future is open.

Recently, several companies (e.g. Auspex, Network Appliance) have begun producing solutions for integrating disk storage both for the Unix and for Windows worlds. IBM has traditionally produced software which works both on Unix and Microsoft platforms.

3.5.2 User Preference Storage

Software packages often allow users to store preferences about the way in which software should look and behave. Such data are stored in some kind of information repository. Another issue for networked systems is where software preference data should be stored for users. There are two possibilities here which correspond approximately to the Unix approach and the NT approach:

- *Windows/Mac/Personal:* under Windows (9x/NT) and Macintosh systems, each user is assumed to have his or her own personal workstation which will not normally be used by other users. Configuration data or preferences which the user selects are thus stored locally on the system disk in a location provided by the operating system. This location is common to all users. Only NT distinguishes between different users.
- *Unix/Shared:* under Unix, each user sets up personal preferences in his or her personal *dot files* which are stored in private user space. More general global preferences are stored in a directory of the administrator's choice. Traditionally, this has been the directory `/etc`.

The difficulties associated with the first of these approaches (having a fixed location for the configuration information which lies in the system files) are several. In any single user operating system, one user can overwrite another user's preferences simply by changing them, since the system is not capable of telling the difference between users. This is a fundamental problem which indicates that single user operating systems are basically unsuited to networking. More pertinent to a networked world are the following points:

- When the operating system is reinstalled, configuration information can easily be lost or overwritten if they are stored in an operating system directory.
- In a distributed environment, where users might not sit at the same physical workstation day after day, the user's personal configuration data will not follow him or her from machine to machine.

NT partly solves these problems by maintaining *user profiles* which are stored on the domain server in a `profiles` subdirectory of the system root. These data are copied into the local workstation when a user logs on to a domain server. On a Unix system, it is easy to specify the locations of configuration files, and these can then be kept separate from operating system files, e.g. on a different disk partition so that they are immune to accidental deletion by system re-installation.

The primary difference between Unix and Windows/Macintosh systems is that Unix is a *multiuser* system, i.e. any Unix machine can be used by an arbitrary number of users from an arbitrary location on the network. Microsoft and Apple systems allow only a single user to use a workstation interactively, from the console of the machine itself.

In the future, there might be some semblance of uniformity. RFC 2244 and RFC 2245 describe the Application Configuration Access Protocol, which describes a centralized user application configuration database.

3.6 Physical Network

In recent times, market forces have effected a rough standardization of hardware and software. In the early days of computing, there were as many standards as there were manufacturers. Today, we have settled on a few technologies, in spite of the efforts of some vendors to draw their customers into monopoly traps. Some studies in setting up physical infrastructure have been reported [166, 230]; see also discussion of load [175, 65] in wide area networks [176].

3.6.1 The OSI Model

The International Standards Organization (ISO) has defined a standard model for describing communications across a network, called the OSI model, for *Open Systems Interconnect (reference model)*. This model is a generalized description of how network communication could be implemented. The TCP/IP Internet architecture is somewhat simpler than OSI. Nevertheless, many sources speak of the OSI model and its many layers, so it is useful to remind ourselves of its designations.

The OSI model is seven-layered monster. The layers are described as follows:

7	Application layer	Program which sends data
6	Presentation layer	XDR or user routines
5	Session layer	RPC / sockets
4	Transport layer	TCP or UDP
3	Network layer	IP Internet protocol
2	Data link layer	Ethernet (protocols)
1	Physical layer	Ethernet (electronics)

At the lowest level, the sending of data between two machines takes place by manipulating voltages along wires. This means we need a device driver for the signaller, and something to receive the data at the other end – a way of converting the signals into bytes; then we need a way of structuring the data so that they make sense. Each of these elements is achieved by a different level of abstraction.

1 *Physical layer.* This is the problem of sending a signal along a wire, amplifying it if it gets weak, removing noise, etc. If the type of cable changes (we might want to reflect signals off a satellite or use fibre optics), we need to convert one kind of signal into another. Each type of transmission might have its own accepted ways of sending data (i.e. protocols).
2 *Data link layer.* This is a layer of checking which makes sure that what was sent from one end of a cable to the other actually arrived. This is sometimes called *handshaking.*
3 *Network layer.* This is the layer of software that remembers which machines are talking to other machines. It establishes connections and handles the delivery of data by manipulating the physical layer. The network layer needs to know something about addresses, i.e. where the data are going, since data might flow along many cables and connections to arrive at their destination.
4 *Transport layer.* We shall concentrate on this layer for much of what follows. The transport layer builds 'packets' or 'datagrams' so that the network layer knows what is data and how to get the data to their destination. Because many machines could be talking on the same network all at the same time, data are broken up into short 'bursts'. Only one machine can talk over a cable at a time, so we must have *sharing.* It is easy to share if the signals are sent in short bursts. This is analogous to the sharing of CPU time by the use of time-slices.
5 *Session layer.* This is the part of a host's operating system which helps a user program set up a connection. This is typically done with *sockets* or the RPC, CORBA or DCOM.
6 *Presentation layer.* How are the data to be sent by the sender and interpreted by the receiver, so that there is no doubt about their contents? This is the role played by the external data representation (XDR) in the RPC system.
7 *Application layer.* The program which wants to send data.

As always, the advantage of using a layered structure is that we can change the details of the lower layers without having to change the higher layers. Layers 1 to 4 are those which involve the transport of data across a network. We could change all of these without doing serious damage to the upper layers – thus as new technology arrives, we can improve network communication without having to rewrite software.

Most of these layers are quite static – only the physical layer is changing appreciably.

3.6.2 Cables and Interface Technologies

A network is a line of communications between two or more hosts. Since it is impractical to have a private cable between every pair of hosts on a network (this would require a cat's cradle of N network interfaces and N cables per host, and would be quite unmanageable, not to say expensive), it is usually some kind of shared cable which is attached to several hosts simultaneously by means of a single *network interface.*

Different vendors have invested in different networking technologies, with different Media Access Control (MAC) specifications. Most Unix systems use some form of Ethernet interface. IBM systems have employed Token Ring networking technology very successfully for their mainframes and AS/400 systems; they now also support Ethernet on their RS/6000 systems. Most manufacturers now provide solutions for both technologies, though Ethernet is undoubtedly popular for local area networks.

- *Bus/Ethernet approach:* Ethernet technology was developed by Xerox, Intel and DEC in 1976, at the Palo Alto Research Center (PARC) [76]. In the Ethernet bus approach, every host is connected to a common cable or bus. Only one host can be using a given network cable at a given instant. It is like a conference telephone call: what goes out onto a network reaches all hosts on that network (more or less) simultaneously, so everyone has to share the line by waiting for a suitable moment to say something. We should not think of data transmission over the network as being a little stream of bytes rollin' down the track one behind the other, from origin to destination, like Thomas the Tank Engine. Every bit, every 1 or 0, is a signal (a voltage or light pulse) on a cable which fills the entire cable at a good fraction of the speed of light. It's like sending Morse code with a lighthouse. Everyone sees the signal, but only the recipient bothers to read it. Ethernet is defined in the IEEE 802.3 standard documents. An Ethernet network is available to any host at any time, provided the line isn't busy. This is called CSMA/CD, or Carrier Sense Multiple Access/Collision Detect. A collision occurs when two hosts attempt to send signals simultaneously. CSMA/CD means that if a card has something to send, it will listen until no other card is transmitting, then start transmitting and listen if no other card starts transmitting at the very same time. If another card began transmitting it will stop, wait for a random interval and try again. The original Ethernet, with a capacity of 10 Mbits per second, could carry packets of 1518 bytes.

  Today, Ethernet is progressing in leaps and bounds. Switched Ethernet running on twisted pair cables can deliver up to 100 Mbits/sec (100BaseT, fast Ethernet). The main limitation of Ethernet networks is the presence of collisions. When many hosts are talking, performance degrades quickly due to time wasted by hosts waiting to get a word in. To avoid collisions, packet sizes are limited. With a large number of small packets, it is easier to share the time between more hosts. Ethernet interfaces are assigned a unique MAC address when they are built. The initial numbers of the address identify each manufacturer uniquely. Full-duplex connections at 100 Mbits are possible with fast Ethernet on dedicated cables. This disables the CSMA/CD protocol.

- *Token ring/FDDI approach:* in the token ring approach [214], hosts are coupled to hubs or nodes, each of which has two network interfaces, and the hosts are connected in a unidirectional ring. The token ting is described in IEEE 802.5. The token ring is a deterministic protocol; if Ethernet embraces chaos, then the token ring demands order. No matter when a host wishes to transmit, it must wait for a passing token, in a specified time-slot. If a signal (token) arrives, a host can append something to the signal. If nothing is appended, the token is passed on to the next host, which has the opportunity to do the same. Similarly, if the signal arriving at one of the interfaces is for the host itself, then it is read. If it is not intended for the host itself, the signal is forwarded to the next host where the same applies. A common token ring based interface in use today is the optical FDDI (Fiber Distributed Data Interface). Token rings can pass 16 Mbits/sec, with packet

sizes of 18 kilobytes. The larger packet sizes are possible, since there is no risk of collisions.

Like Ethernet interfaces, token ring interfaces are manufactured with a uniquely assigned address, though these can be overridden.

- *ATM*: Asynchronous Transfer Mode technology [18] is a high capacity transmission technology developed by telephone companies to exploit existing copper telephone networks. ATM is a radical departure from previous approaches. It takes a switched approach to time-sharing data. Current ATM implementations use 48-byte cells with 5 bytes of routing information. The cells are of a fixed length, allowing predictable switching times, and each cell has enough information to route it to its destination. ATM is believed to be able to reach transfer rates as high as 10 Gbits/sec. Its expense, combined with the increasing performance of fast Ethernet, has made ATM most attractive for high speed Internet backbones, though some local area networks have been implemented as proof of principle.

All of these network approaches are susceptible to spying devices. If a host on the network wants to overhear a conversation between two others, they have only to listen. In some cases, extra hardware *switches* can be used to isolate private connections (see below). In what follows, an Ethernet type of network will be assumed, since this is the commonest form in user-workstation environments.

Even with the bus approach, any host can be connected to several independent network segments. It must have a network interface for each network it is attached to. Each network interface then has a separate network address; thus a host which is connected to several networks will have a different address on each network. A device which is coupled to several networks and which forwards data from one network to another is called a *router*.

Network signals are carried by a variety of means. These days, copper cables are being replaced by fibre-optic glass transmission for long distance communication, and even radio links. In local area networks, it is still copper cables which carry the signals. These cables usually carry Ethernet protocols. Thick yellow cables (about 1 cm thick) carry thick Ethernet, thin black cables (0.5 cm) with BNC connectors carry thin Ethernet. Fibre optic cables (FDDI) have varying appearances. Twisted pair lines are sometimes referred to as 10baseT, 100baseT, etc., or for short T1, T10, T100. The numbers indicate the capacity of the line, 'base' indicates that the cable is used in a *baseband* system, and the 'T' stands for twisted-pair. Twisted pair cables are very thin, so they are not tapped as are thin and thick Ethernet cables. Instead, each host has a single cable connecting it to a *multi-way repeater* or *hub*.

3.6.3 Connectivity

The cables of the network are joined together in segments by hardware which makes sure that messages are transmitted from cable segment to cable segment in the right direction to reach their destinations. A host which is coupled to several network segments and which forwards data, from one network to another is called a *router*. Routers not only forward data, but they prevent the spread of network messages which other network segments do not need to know about. This limits the number of hosts which are sharing any given cable segment, and thus limits the traffic which any given host sees. Routers can also filter unwanted traffic for security purposes [55]. A router knows which destination addresses lie

on which of the networks it is connected to, and it does not let message traffic spread onto irrelevant cables.

A *bridge* is a hardware device which acts like a filter on busy networks. A bridge works like a 'mini-router', and separates two segments of the same cable. A bridge knows which parts of the cable do *not* contain a destination address, and prevents traffic from spreading to this part of a cable. A bridge is used to isolate traffic on busy sections of a network, or conversely to splice networks together.

A *repeater* is an amplifier which strengthens the network signal over long stretches of cable. A *multi-port repeater*, also called a *hub*, does the same thing, and also splits one cable into N sub-cables for convenience. Hubs are common in twisted-pair networks, where it is necessary to fan a cable out into a star pattern from the hub to send one cable to each host. A *switch* is a hub which can direct a message from one host cable directly to the intended host by routing the signal directly. The advantage with this is that other machines do not have to see the traffic between two hosts. Each pair of hosts has a virtual private cable. Switched networks are immune to spies, net-sniffing or network listening devices. A switch performs many of the tasks of a router, and *vice versa*. The difference is that a switch works at layer 2 of the OSI model (i.e. with MAC addresses), whereas a router works at layer 3 (IP addresses). A switch cannot route data on a worldwide basis.

When learning about a new network, one should obtain a plan of the physical setup. If we have done our homework, then we will know where all of these boxes are on the network.

3.6.4 Protocols and Encapsulation

Information transactions take place by agreed standards or *protocols*. Protocols exist to make sure that transmitted data are understood by the receiver in the way that the sender intended. On a network, protocols are required to make sure that data are understood, not only by the receiver, but by all the network hardware which carry them between source and destination. The data are wrapped up in envelope information which contains the address of the destination. Each transmission layer in the protocol stack (protocol hierarchy) is prefixed with a some header information which contains the destination address and other data which identify it. The Ethernet protocol also has a trailer, see Figure 3.3.

Wrapping data inside envelope information is called *encapsulation*, and it is important to understand the basics of this mechanism. Ten years ago network administrators did not need to concern themselves with protocols and their like; today, however, network attacks make clever use of the features and flaws in these protocols, and system administrators need to understand them in order to protect systems from the attacks.

The Internet family of protocols has been the basis of Unix networking for many years, since they were implemented as part of the Berkeley Software Distribution (BSD) Unix. The hierarchy is shown in Figure 3.4.

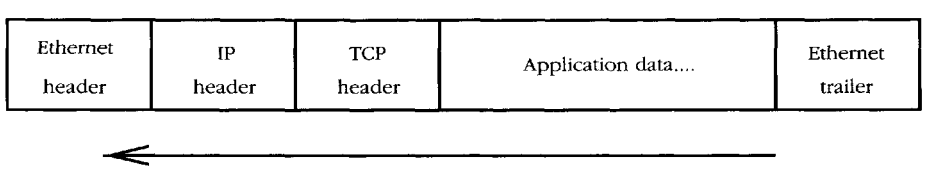


Figure 3.3 Protocol encapsulation

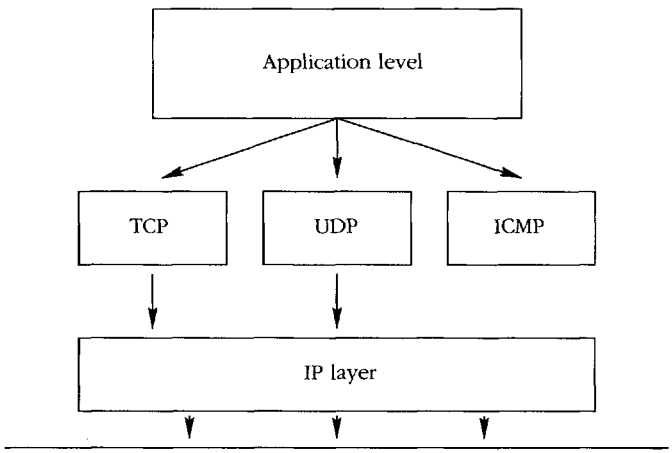


Figure 3.4 The Internet protocol hierarchy

The Transmission Control Protocol (TCP) is for reliable, connection oriented transfer. The User Datagram Protocol (UDP) is a rather cheaper connectionless service, and the Internet Control Message Protocol (ICMP) is used to transmit error messages and routing information for TCP/IP. These protocols have an address structure which is hierarchical and *routable*, which means that IP addresses can find their way from any host in the world to any other so long as they are connected. The Ethernet protocol does not know much more about the world than the cable it is attached to.

NT supports three network protocols, running on top of Ethernet:

- *NETBEUI:* NETBIOS Extended User Interface, Microsoft's own network protocol. This was designed for small networks and is not routable. It has a maximum limit of 20 simultaneous users, and is thus hardly usable.
- *NWLink/IPX:* Novell/Xerox's IPX/SPX protocol suite. Routable. Maximum limit of 400 simultaneous users.
- *TCP/IP:* Standard Internet protocols. The default for NT 4 and Unix-like systems. Novell Netware and Apple Macintosh systems also support TCP/IP. There is no in-built limit to the number of simultaneous users.

Novell's Netware PC server software is based mainly on the IPX suite running on Ethernet hardware; Macintosh networks use their own proprietary Appletalk, which will run on Ethernet or token ring hardware. All platforms are converging on the use of TCP/IP for its open standard and its generality.

3.6.5 Data Formats

There are many problems which arise in networking when hardware and software from different manufacturers have to exist and work together. Some of the largest computer

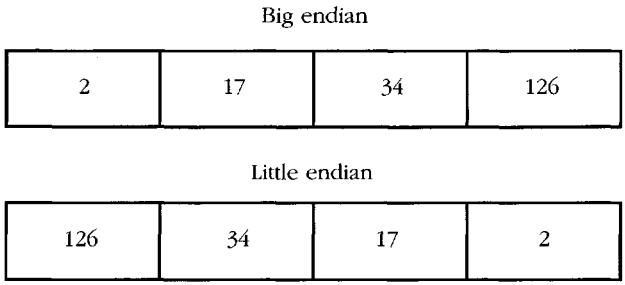


Figure 3.5 Byte ordering sometimes has to be specified when compiling software. The representation of the number 34,677,374 has either of these forms

companies have tried to use this to their advantage on many occasions in order to make customers buy only their products. An obvious example is the choice of network protocols used for communication. Both Apple and Microsoft have tried to introduce their own proprietary networking protocols. TCP/IP has won the contest because it was an *inter*-network protocol (i.e. capable of working on and joining together any hardware type), and also because it is a freely open standard. Neither the Appletalk nor NETBIOS protocols have either of these features.

This illustrates how networking demands standards. That is not to say that some problems do not still remain. No matter how insistently one attempts to fuse operating systems in a network melting pot, there are basic differences in hardware and software which cannot be avoided. One example, which is occasionally visible to system administrators when compiling software, is the way in which different operating systems represent numerical data. Operating systems (actually the hardware they run on) fall into two categories known as *big endian* and *little endian*. The names refer to the *byte-order* of numerical representations. The names indicate how large integers (which require, say, 32 bits or more) are stored in memory. Little endian systems store the least significant byte first, while big endian systems store the most significant byte first. For example, the representation of the number 34,677,374 has either of the forms shown in Figure 3.5. Obviously, if one is transferring data from one host to another, both hosts have to agree on the data representation, otherwise there would be disastrous consequences. This means that there has to be a common standard of *network byte ordering*. For example, Solaris (SPARC hardware) uses network byte ordering (big endian), while NT or Unix-like operating systems on Intel hardware use the opposite (little endian). This means that Intel systems have to convert their data format every time something is transmitted over the network.

3.7 TCP/IP Networks

TCP/IP networking is so important to networked hosts that we shall return to it several times during the course of this book. Its significance is cultural, historical and practical, but the first item in our agenda is to understand its logistic structure.

3.7.1 IP Addresses

Every network interface on the Internet needs to have a unique number which is called its address. At present, the Internet Protocol is at version 4, and this address consists of four bytes, or 32 bits. In the future this will be extended, in a new version of the Internet Protocol, IPv6, to allow more IP addresses, since we are rapidly using up the available addresses. The addresses will also be structured differently. The form of an IP address in IPv4 is

```
aaa.bbb.ccc.mmm
```

Some IP addresses represent networks, whereas others represent hosts, or actually (if we are being 100% correct) they represent network interfaces. The usual situation is that an IP address represents a host attached to a network. Such an address has a host part and a network part, but the format of the address is not necessarily the same for all hosts and networks. There is some freedom to define local conventions.

In every IPv4 address there are 32 bits. We can choose to use these bits in different ways: we could use all 32 bits for host addresses and keep every host on the same enormous cable, without any routers (this would be physically impossible in practice), or we could use all 32 bits for network addresses and have only one host per network (i.e. a router for every host). Both these extremes are silly. After all, we are trying to save resources by sharing a cable between convenient groups of hosts, but shielding other hosts from irrelevant traffic. What we want instead is to group hosts into clusters so as to restrict traffic to localized areas.

Networks fall historically into three classes, called class A, class B and class C networks. Sometimes class D and E networks are also defined. The rigid distinction between different types of network has proven to be a costly mistake for the IPv4 protocol. Amongst other things, it means that only somewhere of the order of 2% of the actual number of IP addresses can actually be used in practice.

The difference between class A, B and C networks concerns which bits in the IP addresses on that network refer to the network itself, and which bits refer to actual hosts.

Class A Networks

IP addresses from `1.0.0.0` to `127.0.0.0` are *class A* networks. The first byte is a network part and the last three bytes are the host address. This allows 126 possible networks (since network 127 is reserved for the loopback service). The number of hosts per class A network is 256^3 minus reserved host addresses on the network. Since this is a ludicrously large number, none of the owners of class A networks are able to use all of their host addresses. Class A networks are no longer issued; they are all assigned, and all the free addresses are wasted. Class A networks were intended for very large organizations (the U.S. government, Hewlett Packard, IBM), and are only practical with the use of a netmask which divides up the large network into manageable subnets. The default subnet mask is `255.0.0.0`.

Class B Networks

IP addresses from `128.0.0.0` to `191.255.0.0` are *class B* networks. There are 16,384 such networks. The first two bytes are the network part and the last two bytes are the host part. This gives a maximum of 256^2 minus reserved host addresses, or 65,534 hosts per

network. Class B networks are typically given to large institutions such as universities and Internet providers, or to institutions such as Sun Microsystems, Microsoft and Novell. All the class B addresses have now been allocated to their parent organizations, but many of these lease out these addresses to third parties. The default subnet mask is `255.255.0.0`.

Class C Networks

IP addresses from `192.0.0.0` to `223.255.255.0` are *class C* networks. There are 2,097,152 such networks. Here the first three bytes are network addresses and the last byte is the host part. This gives a maximum of 254 hosts per network. The default subnet mask is `255.255.255.0`. Class C networks are the most numerous, and there are still a few left to be allocated, though they are disappearing with alarming rapidity.

Class D (Multicast) Addresses

Multicast networks form what is called the MBONE, or multicast backbone. These include addresses from `224.0.0.0` to `239.255.255.0`. These addresses are not used for sending data to specific hosts, but rather for routing data to multiple destinations. Multicast is like a restricted broadcast. Hosts can 'tune in' to multicast channels by subscribing to MBONE services.

Class E (Experimental) Addresses

Addresses `240.0.0.0` to 255.255.255.255 are unused and are considered experimental.

Other Addresses

Some IP addresses are reserved for a special purpose. They do not necessarily refer to hosts or networks.

```
0.0.0.0       Default route
0.*.*.*       Not used
127.0.0.1     Loopback address
127.*.*.*     Loopback network
*.*.*.0       Network addresses (or old broadcast)
*.*.*.255     Broadcast addresses
*.*.*.1       Router or gateway (conventionally)
224.*.*.*     Multicast addresses
```

Note that older networks used the network address itself for broadcasting. This practice has largely been abandoned, however. The default route is a default destination for outgoing packets on a subnet, and is usually made equal to the router address.

The *loopback address* is an address which every host uses to refer to itself internally. It points straight back to the host. It is a kind of internal pseudo-address which allows programs to use network protocols to address local services without anything being transmitted on an actual network.

The zeroth address of any network is reserved to mean the network itself, and the 255th (or on older networks sometimes the zeroth) is used for the broadcast address. Some Internet

addresses are reserved for a special purpose. These include *network addresses* (usually xxx.yyy.zzz.0), *broadcast addresses* (usually xxx.yyy.zzz.255, but in older networks it was xxx.yyy.zzz.0) and *multicast addresses* (usually 224.xxx.yyy.zzz).

Since obtaining unassigned IP addresses is rapidly becoming more difficult, several solutions have been designed to 'fudge' matters for local networks. It is now possible to purchase a device called a *network address translator*, which allows private networks to use any IP address they wish, behind their own closed doors [284]. As soon as they wish to send something over the Internet, however, they must go through a NAT gateway, which translates these private addresses into legally assigned addresses. This allows a local domain to have many IP addresses for internal use, even though they have only be allocated a small number for actual Internet access. It is important, however, that all hosts use the NAT. The outside world (i.e. the true Internet) should not be able to see the private addresses, only the legitimate addresses of the organization using a NAT. Using an already allocated IP address is not just bad manners, it could quickly spoil routing protocols and preclude us from being able to send to the real owners of those addresses. NATs are often used in conjunction with a firewall.

3.7.2 Subnets and Broadcasts

In practice, what we refer to as a network might consist of very many separate cable systems, coupled together by routers, switches and bridges. One problem with very large networks, like class B or class A networks, is that *broadcast messages* (i.e. messages which are sent to every host) create traffic which can slow a busy network. In most cases, broadcast messages only need to be sent to a subset of hosts which have some logical or administrative relationship, but unless something is done a broadcast message will by definition be transmitted to all hosts on the network. What is needed, then, is a method of assigning groups of IP addresses to specific cables and limiting broadcasts to hosts belonging to the group, i.e. breaking up the larger community into more manageable units. The purpose of subnets is to divide up networks into regions which naturally belong together and to isolate regions which are independent. This reduces the propagation of useless traffic, and it allows us to *delegate* and distribute responsibility for local concerns.

This logical partitioning can be achieved by dividing hosts up, through routers, into subnets. Each network can be divided into *subnets* by using a *netmask*. Each address consists of two parts: a *network address* and a *host address*. A system variable called the *netmask* decides how IP addresses are interpreted locally. The netmask decides the boundary between how many bits of the IP address will be kept for hosts and how many will be kept for the network location name. There is thus a trade-off between the number of allowed domains and the number of hosts which can be coupled to each subnet. Subnets are usually separated by routers, so the question is how many machines do we want on one side of a router?

The netmask only has a meaning as a binary number. When looking at the netmask, we have to ask which bits are ones and which are zeroes? The bits which are *ones* decide which bits can be used to specify the domain and the subnets within the domain. The bits which are zeroes decide which are hostnames on each subnet. The local network administrator decides how the netmask is to be used.

The host part of an IP address can be divided up into two parts by moving the boundary between network and host part. The netmask is a variable which contains zeroes and ones.

Every one represents a network bit and every zero represents a host bit. By changing the value of the netmask, we can trade many hosts per network for many subnets with fewer hosts. A subnet mask can be used to separate hosts which also lie on the same physical network, thereby forcing them to communicate through the router. This might be useful for security or administrative purposes.

3.7.3 Netmask Examples

The most common subnet mask is 255.255.255.0. This forces a separation where three bytes represent a network address and one byte is reserved for hosts. For example, consider the class B network 128.39.0.0. With a netmask of 255.255.255.0 everywhere on this network, we divide it up into 255 separate subnets, each of which has room for 254 hosts (256 minus the network address, minus the broadcast address):

```
128.39.0.0
128.39.1.0
128.39.2.0
128.39.3.0
128.39.4.0
...
```

We might find, however, that 254 hosts per subnet is too few. For instance, if a large number of client hosts contact a single server, then there is no reason to route traffic from some clients simply because the subnet was too small. We can therefore double the number of hosts by moving the bit pattern of the netmask one place to the left (see Figure 3.6). Then we have a netmask of 255.255.254.0. This has the effect of pairing the addresses in the previous example. If this netmask were now used throughout the class B network, we would have single subnets formed as follows:

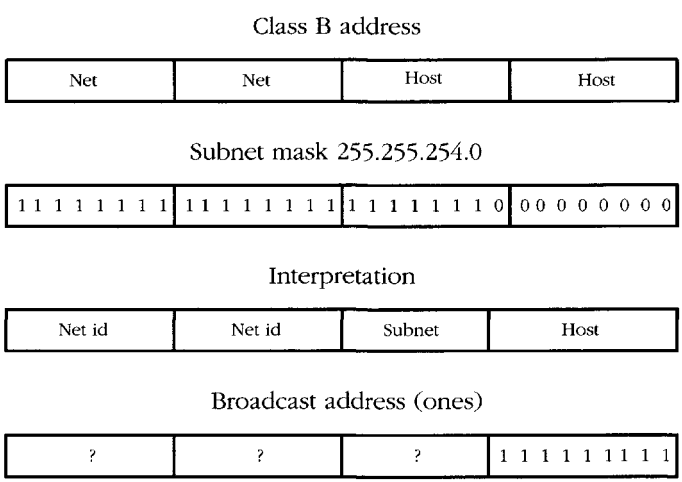


Figure 3.6 Example of how the subnet mask can be used to double up the number of hosts per subnet by pairing host parts. The boundary between host and subnet parts of the address is moved one bit to the left, doubling the number of hosts on the subnets which have this mask

```
128.39.0.0
128.39.1.0

128.39.2.0
128.39.3.0

128.39.4.0
128.39.5.0
...
```

Each of these subnets now contains 510 hosts ($256 \times 2 - 2$), with two addresses reserved: one for the network and one for broadcasts. Similarly, if we moved the netmask again one place to the left, we would multiply by two again, and group the addresses in fours, i.e. netmask 255.255.252.0:

```
128.39.0.0
128.39.1.0
128.39.2.0
128.39.3.0

128.39.4.0
128.39.5.0
128.39.6.0
128.39.7.0
...
```

It is not usually necessary for every host on the entire class B network to share the same subnet mask, though certain types of hardware could place restrictions upon the freedom allowed (e.g. multi-homed hosts). It is only necessary that all hosts within a self-contained group share the same mask. For instance, the first four groups could have netmask 255.255.252.0, the two following could have mask 255.255.254.0, the next two could have separately 255.255.255.0 and 255.255.255.0, and then the next four could have 255.255.252.0 again. This would make a pattern like this:

```
128.39.0.0  (255.255.252.0)
128.39.1.0
128.39.2.0
128.39.3.0

128.39.4.0 (255.255.254.0)
128.39.5.0

128.39.6.0 (255.255.255.0)

128.39.7.0 (255.255.255.0)

128.39.8.0 (255.255.252.0)
128.39.9.0
128.39.10.0
128.39.11.0
...
```

3.7.4 Interface Settings

The IP address of a host is set in the network interface. The Unix command `ifconfig` (interface-configuration) or the NT command `ipconfig` are used to set this. Normally, the

address is set at boot time by a shell script executed as part of the `rc` start-up files. These files are often constructed automatically during the system installation procedure. The `ifconfig` command is also used to set the broadcast address and netmask for the subnet. Each system interface has a name. Here are the network interface names commonly used by different Unix types:

```
Sun          le0 / hme0
DEC ultrix   ln0
DEC OSF/1    ln0
HPUX         lan0
AIX          en0
GNU/Linux    eth0
IRIX         ec0
FreeBSD      ep0
Solarisx86   dnet0
```

Look at the manual entry for the system for the `ifconfig` command, which sets the Internet address, netmask and broadcast address. Here is an example on a Sun system with a Lance-Ethernet interface:

```
ifconfig le0 192.0.2.10 up netmask 255.255.255.0 broadcast
192.0.2.255
```

Normally we do not need to use this command directly, since it should be in the startup-files for the system, from the time the system was installed. However, we might be working in single-user mode or trying to solve some special problem. A system might have been incorrectly configured.

3.7.5 Routing

Unless a host operates as a router in some capacity, it only requires a minimal routing configuration. Each host must define a *default route* which is a destination to which outgoing packets will be sent for processing when they do not belong to the subnet. This is the address of the router or gateway on the same network segment. It is set by a command like this:

```
route add default my-gateway-address 1
```

The syntax varies slightly between systems. On GNU/Linux systems one writes:

```
/sbin/route add default gw my-gateway-address metric 1
```

The default route can be checked using the `netstat -r` command. The result should just be a few lines like this:

```
Kernel routing table
Destination Gateway Genmask        Flags Metric Ref Use Iface
localnet    *       255.255.255.0  U     0      0   932 eth0
loopback    *       255.0.0.0      U     0      0    38 lo
default     my-gw   0.0.0.0        UG    1      0  1534 eth0
```

where `my-gw` is the address of the local gateway (usually subnet address 1).

If this default route is not set, a host will not know where to send packets and will therefore attempt to build a table of routes, using a different entry for every outgoing address.

This consumes memory rapidly and leads to great inefficiency. In the worst case, the host might not have contact with anywhere outside its subnet at all.

3.7.6 ARP/RARP

ARP is the (IP) address resolution protocol. ARP takes an IP address and turns it into an Ethernet address (hardware, MAC address). The ARP service is mirrored by a Reverse ARP (RARP) service. RARP takes a hardware address and turns it into an IP address.

Ethernet (or generally hardware) addresses are required when routing traffic from one device to another. While it is the IP addresses which contain the structure of the Internet and permit routing, it is the hardware address to which one must deliver packets in the final instance; this is the address which is burned into the network interface.

The hardware addresses are cached by each host on the network so that repeated calls to the service ARP translation service are not required. Addresses are checked later, however, so that if an address from a host claiming to have a certain IP address originates from an incorrect hardware address (i.e. the packet does not agree with the information in the cache) then this is detected and a warning can be issued to the effect that two devices are trying to use the same IP address. ARP sends out packets on a local network, asking the question 'Who has IP address xxx.yyy.zzz.mmm?' The host concerned replies with its hardware address.

For hosts which know their own IP address at boot-time, these services only serve as confirmations of identity. Diskless clients (which have no place to store their IP address) do not have this information when they are first switched on and need to ask for it. All they know originally is the unique hardware (Ethernet) address which is burned into their network interface. To bring up and configure an Internet interface, they must first use RARP to find out their IP addresses from a RARP server. Services like BOOTP or DHCP are used for this. Also, the Unix file `/etc/ethers` and `rarpd` can be used.

3.8 Network Analysis

A top down approach is useful for understanding network inter-relationships. We therefore begin at the network level, i.e. at the level of the collective society of machines.

In most daily situations, we start with a network already in place, i.e. we do not have to build one from scratch. It is important to know what hardware one has to work with and where everything is to be found; how it is organized (or not) and so on. Here is a checklist:

- How does the network fit together? (What is its topology?)
- How many different subnets does the network have?
- What are their network addresses?
- Find the router addresses (the default routes) on each segment.
- What is the netmask?
- What hardware is there in the network? (Hosts, printers, etc.)
- Which function does each host/machine have on the network?

- Where are the key network services located?

Hardware can be efficiently identified using SNMP technology. Most newer network hardware supports some kind of querying using SNMP protocols (see section 6.6). This is a form of network communication which talks directly to the device and extracts its hardware profile. Without SNMP, identifying hardware automatically is problematical. One author has proposed using the Unix log service `syslogd` to track hardware configurations [211]. An overview of network services can sometimes be obtained using port-scanning software, such as `nmap`, though this should be agreed in advance to avoid misunderstandings. Many network intrusion attempts begin with port scans; these can make security-conscious administrators nervous.

Of course, when automated methods fail, we can always resort to a visual inspection. In any event, an organization needs some kind of inventory list for insurance purposes, if not merely for good housekeeping. A rough overview of all this information needs to be assembled in the system administrator's mind, in order to understand the challenge ahead.

Having thought about the network in its entirety, we can drop down a level and begin to think about individual host machines. We need to know hosts both from the viewpoint of hardware and software:

- What kind of machines are on the network? What are their names and addresses and where are they? Do they have disks. How big? How much memory do they have? If they are PCs, which screen cards do they have?
- How many CPUs do the hosts have?
- What operating systems are running on the network? MS-DOS, Novell, NT or Unix (if so, which Unix? GNU/Linux, Solaris, HPUX?)
- What kind of network cables are used? Is it thin/thick Ethernet? Is it a star net (hubs/twisted pair), or fibre optic FDDI net?
- Where are hubs/repeaters/the router or other network control boxes located? Who is responsible for maintaining them?
- What is the hierarchy of responsibility?

There is information about the local environment:

- What is the local time zone?
- What broadcast address convention is used? 255 or the older 0?
- Find the key servers on these networks:
  - Where are the NFS network disks located? Which machine are they attached to?
  - Which name service is in use (DNS, NIS or NIS plus)?
  - Where is the inevitable WWW/http service? Who is running pirate servers?

Finding and recording this information is an important learning process, and the information gathered will prove invaluable for the task ahead. Of course, the information will change as time goes by. Networks are not static; they grow and evolve with time, so we must remain vigilant in pursuit of the moving target.

3.8.1 Network Orientation

Familiarizing oneself with an organization's network involves analysing the network's hosts and all of their inter-relationships. It is especially important to know who is responsible for maintaining different parts of the network. It might be us or it might be someone else. We need to know who to contact when something is going wrong, over which we have no control ourselves. The most obvious way to view an organization is by its logical structure. This is usually reflected in the names of different machines and domains. Who do we call if the Internet connection is broken? What service contracts exist on hardware, what upgrade possibilities are there on software? What system is in use for making backups? How does one obtain a backup should the need arise? In short, it is essential to know where to begin in solving any problem which might arise, and who to call[2] if the responsibility for a problem lies with someone else.

The Internet is permeated by a *naming scheme* which, naturally, is used to describe organizational groupings of Internet addresses. We can learn a lot by inspecting the name data for an organization. Indeed, many organizations now see this as a potential security hazard and conceal their naming strategies from outsiders. The Domain Name Service (DNS) is the Internet's primary naming service. It not only allows us to name hosts, but also whole organizations, placing many different IP addresses under a common umbrella. The DNS is thus a hierarchical organization of machine names and addresses. Organizations are represented by *domains* and a domain is maintained either by or on behalf of each organization. Global domains are divided into countries, or groupings like `.com` and `.org`, and *subdomains* are set up within larger domains, so that a useful name can be associated with the function of the organization. To analyse our own network, we begin by asking: who runs the DNS domain above ours?

For our organizational enquiry, we need an overview of the hosts which make up our organization. A host list can be obtained from the DNS using `nslookup` (unless that privilege has been revoked by the DNS administrator, see section 8.5.4). If there are Unix systems on the network, one can learn a lot without physical effort by logging onto each machine and using the `uname` command to find out what OS is being used:

```
nexus% uname -a
SunOS nexus 5.5 Generic sun4m

borg% uname -a
Linux borg 1.3.62 #2 Mon Feb 12 11:06:19 MET 1996 i586
```

This tells us that host `nexus` is a SunOS kernel version 5.5 (colloquially known as the Solaris 2.5) system with a sun4m series processor, and that host `borg` is a GNU/Linux system kernel version 1.3.62. If the `uname` command doesn't exist, then the operating system is an old dinosaur from BSD 4.3 days, and we have to find out what it is by different means. Try the commands `arch` and `mach`.

Knowing the operating system of a host is not sufficient. We also need to know what kind of resources the host has to offer the network, so that we can later plan the distribution of services. Thus we need to dig deeper:

[2] Hostbusters!

- How much memory does a host have? (Most systems print this when they boot. Sometimes the information can be coaxed out of the system in other ways.) What disks and other devices are in use?
- Use `locate` and `find` and `which` and `whereis` to find important directories and software. How is the software laid out?
- What software directories exist? `/usr/local/bin`, `/local/bin`?
- Do the Unix systems have a C compiler installed? This is often needed for installing software. Finding out information about other operating systems, such as Windows, which we cannot log onto is a tedious process. It must be performed by manual inspection, but the results are important nonetheless.

3.8.2 Using `nslookup`

`nslookup` is a program for querying the Domain Name Service (DNS). The name service provides a mapping or relationship between Internet numbers and Internet names, and contains useful information about domains: both our own and others. The first thing we need to know is the domain name. This is the suffix part of the internet names for the network. For instance, suppose our domain is called `example.org`. Hosts in this domain have names like `hostname.example.org`.

If you don't know your DNS domain name, it can probably be found by looking at the file `/etc/resolv.conf` on Unix hosts. For instance:

```
borg% more /etc/resolv.conf
domain example.org
nameserver 192.0.2.10
nameserver 192.0.2.17
nameserver 192.0.2.244
```

Also, most UNIX systems have a command called `domainname`. This prints the name of the local Network Information Service (NIS) domain, which is not the same thing as the DNS domain name (though, in practice, many sites would use the same name for both). Do not confuse the output of this command with the DNS domain name.

Once we know the domain name, we can find out the hosts which are registered in your domain by running the name service lookup program `nslookup, or dig.`

```
borg% nslookup
Default Server:  mother.example.org
Address:  192.0.2.10
```

`nslookup` always prints the name and address of the server from which it obtains its information. Then we get a new prompt > for typing commands. Typing help provides a list of the commands which `nslookup` understands.

hostname/IP lookup

Type the name of a host or Internet (IP) address and `nslookup` returns the equivalent translation. For example:

```
dax% nslookup
Default Server:  mother.example.org
Address:  192.0.2.10

> www.gnu.org
Server:  mother.example.org
Address:  192.0.2.10

Name:    www.gnu.org
Address:  206.126.32.23

> 192.0.2.238
Server:  mother.example.org
Address:  192.0.2.10

Name:    dax.example.org
Address:  192.0.2.238
```

In this example we look up the Internet address of the host called `www.gnu.org` and the name of the host which has Internet address `192.0.2.238`. In both cases, the default server is the name server `mother.example.org` which has Internet address `192.0.2.10`.

Note that the default server is the first server listed in the file `/etc/resolv.conf` which answers for queries on starting `nslookup`.

Special Information

The domain name service identifies certain special hosts which perform services like the name service itself and mail-handlers (called mail exchangers). These servers are identified by special records so that people outside of a given domain can find out about them. After all, the mail service in one domain needs to know how to send mail to a neighbouring domain. It also needs to know how to find out the names and addresses of hosts for which it does not keep information personally.

We can use `nslookup` to extract this information by setting the 'query type' of a request. For instance, to find out about the mail exchangers in a domain we write

```
> set q=mx
>  domain name
```

For example

```
 > set q=mx
> otherdomain.org
Server:  mother.example.org
Address:    192.0.2.10

Non-authoritative answer:
otherdomain.org    preference = 0, mail exchanger =
  mercury.otherdomain.org

Authoritative answers can be found from:
otherdomain.org    nameserver = mercury.otherdomain.org
otherdomain.org    nameserver = delilah.otherdomain.org
mercury.otherdomain.org       internet address = 158.36.85.10
delilah.otherdomain.org  internet address = 129.241.1.99
```

Here we see that the only mail server for `otherdomain.org` is `mercury.otherdomain.org`.

Another example, is to obtain information about the name servers in a domain. This will allow us to find out information about hosts which is not contained in our local database (see section 3.8.2). To get this, we set the query-type to `ns`:

```
> set q=ns
> otherdomain.org
Server:  mother.example.org
Address:   192.0.2.10

Non-authoritative answer:
otherdomain.org      nameserver = delilah.otherdomain.org
otherdomain.org      nameserver = mercury.otherdomain.org

Authoritative answers can be found from:
delilah.otherdomain.org  internet address = 192.0.2.78
mercury.otherdomain.org   internet address = 192.0.2.80
>
```

Here we see that there are two authoritative name servers for this domain, called `delilah.otherdomain.org` and `mercury.otherdomain.org`.

Finally, if we set the query type to 'any', we get a summary of all this information.

Listing Hosts Belonging to a Domain

To list every registered Internet address and hostname for a given domain one can use the `ls` command inside `nslookup`. For instance

```
> ls example.org
[mother.example.org]
 example.org.            server = mother.example.org
 example.org.            server = mercury.otherdomain.org
 pc61                    192.0.2.61
 pc59                    192.0.2.59
 pc59                    192.0.2.59
 pc196                   192.0.2.196
etc...
```

Newer name servers can restrict access to prevent others from obtaining this list all in one go, since it is now considered a potential security hazard. First, the name servers are listed and then the host names and corresponding IP addresses are listed.

If we try to look up hosts in a domain for which the default name server has no information, we get an error message. For example, suppose we try to list the names of the hosts in the domain over ours:

```
> ls otherdomain.org
[mother.example.org]
*** Can't list domain otherdomain.org: Query refused
>
```

This does not mean that it is not possible to get information about other domains, only that we cannot find out information about other domains from the local server (see section 3.8.2).

Changing to a Different Server

If we know the name of a server which contains authoritative information for a domain, we can tell `nslookup` to use that server instead. In that way it might be possible to list the hosts in a remote domain and find out detailed information about it. At the very least, it is possible to find out about key records, like name servers and mail exchangers (MX). To change the server we simply type

```
> server new-server
```

or

```
> lserver new-server
```

There is a subtle difference between these two commands. If one uses the first command to change the server to another host which is not running a `named` daemon (the DNS daemon), one finds oneself in a situation where it is no longer possible to look up host names or IP addresses. The second form always uses the default (the first) server to look up the names we use. Assuming that no security barriers have been erected, we can now use this to list all of the data for a remote domain. First we change server; once this is done we use `ls` to list the names:

```
> server ns.college.edu
Default Server:  ns.college.edu
Address:  192.0.2.10

> ls college.edu

  (listing..)
```

Another advantage to using the server which is directly responsible for the DNS data, is that we obtain extra information about the domain, namely a contact address for the person responsible for administrating the domain. For example:

```
> server ns.college.edu
Default Server:  ns.college.edu
Address:  192.0.2.10

> college.edu
Server:  ns.college.edu
Address:  192.0.2.10

college.edu     preference = 0, mail exchanger = ns.college.edu
college.edu     nameserver = ns.college.edu
college.edu
    origin = ns.college.edu
    mail addr = postmaster.ns.college.edu
    serial = 1996120503
    refresh = 3600 (1 hour)
    retry = 900 (15 mins)
    expire = 604800 (7 days)
    minimum ttl = 86400 (1 day)
college.edu       nameserver = ns.college.edu
ns.college.edu    internet address = 192.0.2.10
```

This is probably more information than we are interested in, but it does tell us that we can address queries and problems concerning this domain to `postmaster@ns.college.edu`. (Note that DNS does not use the @ symbol for 'at' in these data.)

3.8.3 Contacting Other Domains

Sometimes we need to contact other domains, perhaps because we believe there is a problem with their system, or perhaps because an unpleasant user from another domain is being a nuisance and we want to ask the administrators there to put that person through a long and painful death. We now know how to obtain one contact address using `nslookup`. Another good bet is to mail the one address which every domain must have: postmaster@-domain. Any domain which does not define this mail address deserves to have its wires cut[3]. Various unofficial standards also encourage sites to have the following mail addresses which one can try:

```
webmaster
www
ftp
abuse
info
security
hostmaster
```

Apart from these sources, there is little one can do to determine who is responsible for a domain. A number of domains are registered with another network database service called the *whois* service. In some cases it is possible to obtain information this way. For example:

```
host% whois moneywurld.com
   Financial Connections, Inc (MONEYWURLD-COM)
     2508 5th Ave, #104
     Mars,

     Domain Name: MONEYWURLD.COM

     Administrative Contact, Technical Contact, Zone Contact:
       Willumz, Bob  (BW747)  willy@MONEYWURLD.COM
       206 269 0846

     Record last updated on 13-Oct-96.
     Record created on 26-Oct-95.

     Domain servers in listed order:

     NSH.WORLDHELP.NET 129.0.0.1
     NSS.MONEYWURLD.COM 129.0.0.2

The InterNIC Registration Services Host contains ONLY Internet
Info (Networks, ASN's, Domains, and POC's).
Please use the whois server at nic.ddn.mil for MILNET
Information.
```

[3] Some obnoxious domains which send out unsolicited mail do not define this address because they are afraid that annoyed users will actually mail them back.

3.9 Planning Network Resources

A network community is an organism, working through the orchestrated cooperation of many parts. We need to understand its operation carefully in order to make it work well. The choices we make about the system can make it easy to understand, or difficult to understand, efficient or inefficient. This is the challenge of community planning.

3.9.1 Mapping Out Services

Existing network services have to be analysed, so that we know where we are starting from, and new networks need to be planned or extended. If the obligatory school frog dissections never appealed, then one can at least take comfort in the fact that dissecting the network organism is, if nothing else, a cleaner operation. Starting with the knowledge we have already gained about host types and operating systems, we must now identify all of the services which are running on all of the hosts.

Location can be performed by a visual inspection of the process tables, or from configuration files. There are tools for port scanning networks in order to locate services, e.g. the `nmap` program. We should be careful about using these, however, since port scans normally signify a network intrusion attempt, so others might misconstrue. If a network is well run, local administrators will know what services are running on which hosts. The information we gather is then open to scrutiny. Our aim is to arrange for the machines in the network to work together optimally, so we begin by thinking:

- How to choose the right hardware for the right job.
- Which hosts should be servers and for which services.
- How to make disks available to the network.
- How to share tasks between machines.
- How clock/time synchronization will work.

What roles do the hosts play now? How might this be improved in the future? Is everything already working satisfactorily or do we need to rewire our frog? In the ideal universe, we would always have unlimited resources for every task, but when reality bites, some kind of compromise is usually needed.

The efficiency of a network can be improved greatly by planning carefully how key networks services are organized: particularly file servers and name services, which form the basic infrastructure of a network. Here is a partial checklist:

- Which hosts keep the physical disks for NFS disk servers? It makes sense to keep all file services which use those disks on that same host. If the source data are on host A, then we run all file services for those data on host A, otherwise data will first have to be copied from A to B, over the network, in order to be served back over the network to host C, i.e. there will be an unnecessary doubling of traffic.
- Normally we shall want to use a powerful system for the servers which provide key disk and WWW services, since these are at the heart of network infrastructure. Other hosts depend upon these. However, if resources are limited we might need to reserve the

fastest host for running some especially heavy software. This has to be a site dependent calculation.

- File servers always benefit from a large amount of RAM. This is a cheap form of optimization which allows caching. Fast network interfaces and hard disks are also amongst the most effective optimizations one can invest in. If we are going to buy RAM and fast disks, don't give it all away for users' selfish workstations; treat the server-host to the biggest share.
- If we can, it helps to separate users' home directories over several disks and keep problem disk-users on a partition for themselves, away from honest users.
- Shall we consolidate many services on one host, or distribute them across many? The first possibility is easier to administrate, but the second might be more efficient and less prone to host crashes.
- Any binary or software servers we set up to share software are individual to each operating system type we maintain. A Sun machine cannot run software compiled on a GNU/Linux host, etc.
- Dependency can be a source of many insidious problems. Try not to create deadlocks whereby host A needs host B and host B needs host A. This is a particularly common mistake with NFS file system mounts. It can cause a hanging loop.
- If high availability is an issue, will one server per service be enough? Do we need a backup server? Backup name service servers (DNS, NIS, WINS) could be considered a must. Without a name-service, a network is paralyzed.

There is no textbook solution to these issues. There are only recipes and recommendations based on trial-and-error experience. If we want our frog to win the high-jump, we need to strike a balance between concentrating muscle in key areas, and spreading the load evenly. We are unlikely to get everything just right, first time around, so it is important to construct a solid system, at the same time as anticipating future change.

Principle 8 (Adaptability) *Optimal structure and performance are usually found only with experience of changing local needs. The need for system revision will always come. Make network solutions which are adaptable.*

3.9.2 Uniform Resource Locators (URLs)

Each operating system has a model for laying out its files in a standard pattern, but user files and local additions are usually left unspecified. Choosing a sound layout for data can make the difference between an incomprehensible chaos and a neat orderly structure. An orderly structure is useful not only for the users of the system, but also when making backups. Some of the issues are:

- Disk partitions are associated with drives or directory trees when connected to operating systems. These need names.
- Naming schemes for files and disks are operating system dependent.
- The name of a partition should reflect its function or contents.

- In a network the name of a partition ought to be a URL, i.e. contain the name of the host.
- It is good practice to consolidate file storage into a few special locations rather than spreading it out all over the network. Moreover, a basic principle in cataloging resources is:

Principle 9 (One name for one object I) *Each unique resource should have a unique name which labels it and describes its function.*

with the corollary:

Corollary 10 (Aliases) *Sometimes it is advantageous to use aliases or pointers to unique objects so that a generic name can point to a specific resource.*

Data kept on many machines can be difficult to manage, compared to data collected on a few dedicated file servers. Also, insecure operating systems offer files on a local disk no protection.

The URL model of file naming has several advantages. It means that one always knows the host-provider and function of a network resource. Also, 'it falls nicely into a hierarchical directory pattern. A simple but effective scheme is to use a three level mount-point for adding disks: each user disk is mapped onto a directory with a name of the form

```
/site/host/content
```

(see Figure 3.7). This scheme is adequate even for large organizations, and can be extended in obvious ways.

In DOS-derived operating systems one does not have the freedom to 'mount' network file systems into the structure of the local disk; network disks always become a special 'drive', like H: or I, etc. It is difficult to make a consistent view of the disk resources with this system, although it is rumoured that NT 5.0 will have this possibility and one can already use file systems like the DFS on NT which do support this model.

Within an organization a URL structure provides a global naming scheme, like those used in true network file systems like AFS and DFS. These use the name of the host on which a resource is physically located to provide a point of reference. This is also an excellent way of labelling backups of partitions, since it is then immediately clear where the data belong. A few rules of thumb allow this naming scheme to live painlessly along-side traditional Unix naming schemes:

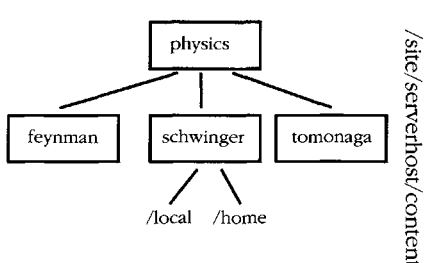


Figure 3.7 A universal naming scheme (URL) for network resources makes distributed data comprehensible

- When mounting a remote file system on a host, the client and server directories should always have exactly the same name. Anything else only causes confusion and problems later [186].
- The name of every file system mount-point should be unique and tell us something meaningful about where it is located and what its function is.
- To preserve tradition, we can invoke the corollary and use an alias to provide a generic reference point for a specific resource. For instance, names like `/usr/local` can be used to point to more accurate designations like `/site/host/local`. On different clients, the alias `/usr/local` might point to a file system on a single server, or to file systems on many servers. The purpose of an alias is to hide this detail, while the purpose of the file system designation is to identify it.
- It doesn't matter whether software compiles in the path names of special directories into software as long as we follow the points above.

For example, the following scheme was introduced at Oslo University and later copied at the College. The first link in the mount point is the department of the organization or, in our case, the university faculty to which the host belongs; the second link is the name of the host to which the disk is physically connected, and the third and final link is a name which reflects the contents of the partition. Some examples:

```
/site/hostname/content

/research/grumpy/local
/research/happy/home1
/research/happy/home2

/sales/slimy/home1

/physics/einstein/data
/biology/pauling/genome-db
```

The point of introducing this scheme was twofold:

- To instantly be able to identify the NFS server on which the disk resource physically resided.
- To instantly be able to identify the correct locations of files on backup tapes, without any special labelling of the tapes (see section 10.2.3).

The problem of drive names in NT and Windows is an awkward one if one is looking for Unix/NT inter-operability. In this case, Sun's PCNFS might be an answer. In practice, many networks based on NT and Windows will use Microsoft's model throughout, and while it might not gleam with elegance it does the job. The problem of backups is confined to the domain servers, so the fact that Windows is not a fully distributed operating system restricts the problem to manageable proportions.

3.9.3 Choosing Server-Hosts

Choosing the best host for a service is an issue with several themes. The main principles have to do with efficiency and security, and can be summarized by the following questions:

- Does traffic have to cross subnet boundaries?
- Do we avoid unnecessary network traffic?
- Have we placed insecure services on unimportant hosts?

Service requests made to servers on different subnets have to be routed. This takes time and uses up switching opportunities which might be important on a heavily loaded network. Some services (like DNS) can be mirrored on each subnet, while others cannot be mirrored in any simple fashion. Unnecessary network traffic can be reduced by eliminating unnecessary dependencies of one service on another. For example, suppose we are setting up a file server (WWW or FTP). The data which these servers will serve to clients lie on a disk which is physically attached to some host. If we place the file-server on a host which does not have direct physical access to the disks, then we must first use another network service (e.g. NFS) as a proxy in order to get the data from the host with the disk attached. Had we placed the file server directly on the host with the disk, the intermediate step would have been unnecessary and we could halve the amount of network traffic.

Principle 11 (Inter-dependency) *Avoid making one service reliant on another. The more independent a service is, the more efficient it will be, and the fewer possibilities there will be for its failure.*

Some services are already reliant on others by virtue of their design. For example, most services are reliant on the DNS.

Suggestion 1 (File servers with common data) *Place all file servers which serve the same data on a common host, e.g. WWW, FTP and NFS serving user files. Place them on the host which physically has the disks attached. This will save an unnecessary doubling of network traffic and will speed up services. A fast host with a lot of memory and perhaps several CPUs should be used for this.*

3.9.4 Distributed File Systems and Mirroring

The purpose of a network is to share resources amongst many hosts. Making files available to all hosts from a common source is one of the most important issues in setting up a network community. There are three types of data which we have to consider separately:

- Users' home directories.
- Software or binary data (architecture specific).
- Other common data (architecture unspecific).

Since users normally have network accounts which permit them to log onto any host in the network, user data clearly have to be made available to all hosts. The same is not true of software, however. Software only needs to be shared between hosts running comparable operating systems. A windows program will not run under GNU/Linux (even though they share a common processor and machine code), nor will a SCO Unix program run under Free BSD. It does not make sense to share binary file systems between hosts, unless they share a common architecture. Finally, sharable data, such as manual information or architecture-

independent databases, can be shared between any hosts which specifically require access to them.

How are network data shared? There are two strategies:

- Use of a shared file system (e.g. NFS or Novell Netware).
- Disk mirroring.

Using a network file system is always possible, and it is a relatively cheap solution, since it means that we can minimize the amount of disk space required to store data, by concentrating the data on just a few servers. The main disadvantage with the use of a network file system is that network access rates are usually much slower than disk access rates, because the network is slow compared to disks, and a server has to talk to many clients concurrently, introducing *contention* or competition for resources. Even with the aggressive caching schemes of NFS, there is a noticeable difference in loading files from the network and loading files locally.

Bearing in mind the principles of the previous section, we would if possible like to minimize load to the network. A certain amount of network traffic can be avoided by mirroring software rather than sharing with a network file system. Mirroring means copying every file from a source disk to the local disk of another host. This can be done during the night when traffic is low and, since software does not change often, it does not generate much traffic for upgrades after the initial copy. Mirroring is cheap on network traffic, even during the night. During the daytime, when users are accessing the files, they collect them from the mirrors. This is both faster and requires no network bandwidth at all.

Mirroring cannot apply to users' files since they change too often, while users are logged onto the system, but it applies very well to software. If we have disk space to spare, then mirroring software partitions can relieve the load of sharing. There are various options for disk mirroring. On Unix hosts we have `rdist`, `rsync` and `cfengine`; variations on these have also been discussed [224, 267, 89, 72]. The use of `rdist` can no longer be recommended (see section 6.4.6) for security reasons. Cfengine can also be used on NT. Network file systems can be used for mirroring, employing only standard local copy commands; file systems are first mounted and then regular copy commands are used to transfer the data as if they were local files.

The benefits of mirroring can be considerable, but it is seldom practical to give every workstation a mirror of software. A reasonable compromise is to have a group of file servers, synchronized by mirroring from a central source. One would expect to have at least one file server per subnet, to avoid router traffic, money permitting.

Exercises

Exercise 3.1 What are the advantages and disadvantages of making access to network disks transparent to users?

Exercise 3.2 What is the Domain Name Service? How do hosts depend upon this service? Many security mechanisms make use of host names to identify hosts. Suppose that the data in the DNS could be corrupted. Would this be a security risk?

Exercise 3.3 In what way is using a name service better than using static host tables? In what way is it worse?

Exercise 3.4 Draw a diagram of the physical topology of your local network, showing routers, switches, cables and other hardware.

Exercise 3.5 Determine all of the subnets which comprise your local network. (If there are many, consider just the closest ones in your department.) What is the netmask on these subnets? (You only need to determine the subnet mask on a representative host from each subnet, since all hosts must agree on this choice. Hint, try: `ifconfig -a`.)

Exercise 3.6 If the network xxx.yyy.74.mmm has subnet mask 255.255.254.0, what can you say about the subnet mask for the addresses on xxx.yyy.75.mmm. (Hint: how many hosts are allowed on the subnet?) Which IP addresses does the subnet consist of?

Exercise 3.7 If the network xxx.yyy.74.mmm has subnet mask 255.255.255.0, what can you say about the subnet mask for the addresses on xxx.yyy.75.mmm?

Exercise 3.8 Using `nslookup`, determine the answers to the following questions:

(a) Find the IP address of the host `www.gnu.org`.

(b) Find the name of the name servers for the domain `gnu.org`.

(c) Are `axis.iu.hioslo.no` and `thistledown.iu.hioslo.no` two different hosts?

(c) Are `ftp.iu.hioslo.no` and `www.iu.hioslo.no` two different hosts?

(d) Find the name of the mail exchanger for the domain `iu.hioslo.no`.

Exercise 3.9 The purpose of this next problem is to make you think about the consequences of cloning all hosts in a network, so that they are all alike. The principles apply equally well to other societies. Try not to get embroiled in politics; concentrate on practicalities rather than ideologies.

(a) Discuss the pros and cons of uniformity. In a society, when is it advantageous for everyone in a group to have equal access to resources? In what sense are they equal? What special characteristics will always be different, i.e. why are two persons never completely equal (e.g. their names are different)?

(b) When is it advantageous for some members of a community to have more resources and more power than others? You might like to consider what real power is. For instance, would you say that garbage disposal workers and water engineers have power in a society?

(c) What is meant by delegation? How is delegation important to cooperation?

(d) What is meant by dependency? How does delegation lead to dependency? Can you foresee any problems with this, for network efficiency?

(e) What is meant by a network service?

(f) Discuss each of the above points in connection with computers in a network.

Exercise 3.10 Design a universal naming scheme for directories, for your site. Think about what types of operating system you have and how the resources will be shared; this will affect your choices. How will you decide drive names on windows hosts?

Exercise 3.11 What are ARP/RARP? Why can't we use Ethernet addresses instead of IP addresses? Why are IP addresses needed at all?

Exercise 3.12 At some sites, it was common practice to use disk mirroring to synchronize the system disks of hosts, where compiled software had been mixed in with the operating system's own files. This solves the problem of making manual changes to one host, and keeping other hosts the same as the source machine. Discuss whether this practice is advisable, with respect to upgrades of the operating system.

Chapter 4

Host Management

We have invested some effort in understanding the basics of how hosts function within a network community. Only now are we sufficiently prepared to turn our attention to hosts within such a network. It should be clear from the previous chapter that it would be a mistake to think of the host as being the fundamental object in the network. If we initially focus on too small a part of the entire system, time and effort can be wasted configuring hosts in a way which does not take into account the cooperative aspects of the network. That would be a recipe for failure and only a prelude to later re-installation.

4.1 Choices

We can make life easy or difficult for ourselves by the decisions we make at the outset of host installation. The first step in setting up hosts is to make some basic choices. Should we:

- Follow the OS designer's recommended setup? (Often not good enough.)
- Create our own setup?
- Make all machines alike?
- Make all machines different?

Most vendors will only provide immediate support for individual hosts or, in the best case, clusters of hosts manufactured by them. They will almost never address the issue of total network solutions, without additional cost, so their recommendations often fall notably short of the recommendable in a real network. We have to be aware of the big picture when installing and configuring hosts.

4.2 Start-up and Shutdown

The two most fundamental operations which one can perform on a host are to start it up and to shut it down. With any kind of mechanical device with moving parts, there has to be a procedure for shutting it down. One does not shut down any machine in the middle of a crucial operation, whether it be a washing machine in the middle of a program, an aircraft in mid-flight, or a computer writing to its disk.

With a multi-tasking operating system, the problem is that it is never possible to predict when the system will be performing a crucial operation in the background. For this simple reason, every multi-tasking operating system provides a procedure for shutting down safely. A safe shutdown avoids damage to disks by mechanical interruption, but it also synchronizes hardware and memory caches, making sure that no operation is left incomplete.

4.2.1 Booting Unix

Normally it is sufficient to switch on the power to boot a Unix-like host. Sometimes you might have to type 'boot' or 'b' to get it going. Unix systems can boot in several different modes or *run levels*. The most common modes are called multi-user mode and single-user mode. On different kinds of Unix, these might translate into run-levels with different numbers, but there is no consensus. In single-user mode, no external logins are permitted. The purpose of single-user mode, is to allow the system administrator access to the system without fear of interference from other users. It is used for installing disks or when repairing file systems, where the presence of other users on the system would cause problems.

In the olden days, the Unix boot procedure was controlled entirely by a file called `/etc/rc`, meaning 'run commands', and subsidiary files like `rc.local`. These files were no more than shell scripts. Newer versions of Unix have made the boot process insufferably complex by introducing a program called `init`. `init` reads a configuration file called `/etc/inittab`, and a directory called `/etc/rc.d`. `/etc/inittab` defines a number of run-levels and starts scripts depending on what run-level you choose. The idea behind `inittab` is to make Unix installable in packages, where each package can be started or configured by a separate script. Which packages get started depends upon the run-level you choose.

The default form for booting is to boot in multi-user mode. We have to find out how to boot in single-user mode on our system, in case we need to repair a disk at some point. Here are some examples. Under SunOS and Solaris, one interrupts the normal booting process by typing stop a, where stop represents the 'stop key' on the left-hand side of the keyboard. If you do this, you should alway give the `sync` command to synchronize disk caches and minimize file system damage:

```
Stop a

ok? sync
ok? boot -s
```

If the system does not boot right away, you might see the line

```
type b) boot, c) continue or n) new command
```

In this case, you should type

```
b -s
```

in order to boot in single-user mode. Under the GNU/Linux operating system, using the LILO boot system, we interrupt the normal boot sequence by pressing the 'SHIFT key when the LILO prompt appears. This should cause the system to stop at the prompt:

```
Boot:
```

To boot, we must normally specify the name of a kernel file, normally `linux`. To boot in single-user mode, we then type

```
Boot: linux single
```

Or at the LILO prompt, it is possible type ? in order to see a list of kernels. There appears to be a bug in some versions of GNU/Linux so that this does not have the desired effect. In some cases, one is prompted for a run-level. The correct run-level should be determined from the file `/etc/inittab`. It is normally called `S` or `1` or even `1S`.

Once in single-user mode, we can always return to multi-user mode just by exiting the single-user login.

4.2.2 Shutting Down Unix

Anyone can start a Unix-like system, but we have to be the superuser to shut one down correctly. Of course, one could just pull the plug, but this can ruin the disk file system. Even when no users are touching a keyboard anywhere, a Unix system can be writing something to the disk – if we pull the plug, we might interrupt a crucial write-operation which destroys the disk contents. The correct way to shut down a Unix system is to run one of the following programs:

- `halt`: stops the system immediately and without warning. All processes are killed with the TERM-inate signal 15 and disks are synchronized.
- `reboot`: as `halt`, but the system reboots in the default manner immediately.
- `shutdown`: this program is the recommended way of shutting down the system. It is just a friendly user-interface to the other programs, but it warns the users of the system about the impending shutdown and allows them to finish what they are doing before the system goes down.

Here are some examples of the `shutdown` command. The first is from BSD Unix:

```
shutdown -h +3 "System halting in three minutes"
shutdown -r +4 "System rebooting in four minutes"
```

The option `-h` implies that the system will halt and not reboot automatically. The option `-r` implies that the system will reboot automatically. The times are specified in minutes.

System V unix R4 (e.g. Solaris) has a different syntax which is based on its system of run-levels. The shutdown command allows one to switch run-levels in a very general way. One of the run-levels is the 'not running' or 'halt' run-level. To halt the system, we have to call this:

```
shutdown -i 5 -g 120 "Powering down os...."
```

The `-i 5` option tells SVR4 to go to run-level 5, which is the power-off state. Run level 0 would also suffice here. The `-g 120` option tells `shutdown` to wait for a grace-period of 120 seconds before shutting down. Note that Solaris also provides a BSD version of shutdown in `/usr/ucb`.

Never assume that the run levels on one system are the same as those on another.

4.2.3 Booting and Shutting Down NT

Booting and shutting down NT is a trivial matter. To boot the system, it is simply a matter of switching on the power. To shut it down, one chooses shutdown from the Start Menu.

There is no direct equivalent of single-user mode for NT. To switch off network access on an NT server so that disk maintenance can be performed, one must normally choose a reboot and connect new hardware while the host is down. File System checks are performed automatically if errors are detected. The plug'n'play style automation of NT removes the need for manual work on file systems, but it also limits flexibility.

The NT boot procedure on a PC begins with the BIOS, or PC hardware. This performs a memory check and looks for a bootable disk. A bootable disk is one which contains a Master Boot Record (MBR). Normally, the BIOS is configured to check the floppy drive `A:` first and then the hard-disk `C:` for a boot block. The boot block is located in the first sector of the bootable drive. It identifies which partition is to be used to continue with the boot procedure. On each primary partition of a bootable disk, there is a boot program which 'knows' how to load the operating system it finds there. NT has a menu-driven boot manager program which makes it possible for several OSes to coexist on different partitions.

Once the disk partition containing NT has been located, the program `NTLDR` is called to load the kernel. The file `BOOT.INI` configures the defaults for the boot manager. After the initial boot, a program is run which attempts to automatically detect new hardware and verify old hardware. Finally, the kernel is loaded and NT starts proper.

4.3 Configuring and Personalizing Workstations

Permanent, read-write storage changed PCs from expensive ping-pong games into tools for work as well as pleasure. Today, disk is so cheap that it is not uncommon for even personal workstations to have several gigabytes of local storage.

Flaunting wealth is the sport of the modern computer owner: more disk, more memory, better graphics. Why? Because it's there. This is the game of free enterprise, encouraged by the availability of home computers and personal workstations. Not so many years before such things existed, however, computers only existed as large multi-user systems, where hundreds of users shared a few kilobytes of memory and a processor no more powerful than a now arthritic PC. Rational resource sharing was not just desirable, it was the only way to bring computing to ordinary users. In a network, we have these two conflicting interests in the balance.

4.3.1 Personal Workstations or Networkstations?

Today we are spoiled, often with more resources than we know what to do with. Disk space is a valuable resource which can be used for many purposes. It would be an ugly waste to allow huge areas of disk to go unused simply because small disks are no longer manufactured; but at the same time, we should not simply allow anyone to use disk space as they please, just because it is there.

Operating systems which have grown out of home computers (Windows and Macintosh) take the view that whatever is left over of disk resources is for the local owner to do with as

he or she pleases. This is symptomatic of the idea that one computer belongs to one user. In the world of the network, this is an inflexible model. Users move around organizations; they ought not to be forced to take their hardware with them as they move. Allowing users to personalize workstations is thus a questionable idea in a network environment.

Network sharing allows us to make disk space available to all hosts on a network, e.g. with NFSs, Netware or DFS. This allows us to make disk space available to all hosts. There are positives and negatives with sharing, however. If sharing was a universal panacea, we would not have local disks: everything would be shared by a network. This approach has been tried: diskless workstations, network computers and X-terminals have all flirted with the idea of keeping all disk resources in one place and using the network for sharing. Such systems have been a failure: they perform badly and they simply waste a different resource: network bandwidth. Some files are better placed on a local disk: namely the files which we need often, such as the operating system and temporary scratch files, such as those created in the processing of large amounts of data.

In organizing disk space, we can make the best use of resources, and separate:

- Space for the operating system.
- Space which can be shared and made available for all hosts.
- Space which can be used to optimize local work, e.g. temporary scratch space, space which can be used to optimize local performance (avoid slow networking).
- Space can be used to make distributed backups, for multiple redundancy.

These independent areas of use need to be separated from one another, by partitioning disks.

4.3.2 Partitioning

Disks can be divided up into partitions. Partitions physically divide the disk surface into separate areas which do not overlap. The disk controller makes sure that partitions behave as independent, logical disks. The main difference between two partitions on one disk and two separate disks is that partitions can only be accessed one at at time, whereas multiple disks can be accessed in parallel.

Disks are partitioned so that files with separate purposes cannot be allowed to spill over into one anothers' space. Partitioning a disk allows us to reserve a fixed amount of space for a particular purpose, safe in the knowledge that nothing else will encroach on that space. For example, it makes sense to place the operating system on a separate partition, and user data on another partition. If these two independent areas shared common space, the activities of users could quickly choke the operating system by using up all of its workspace.

In partitioning a system, we have in mind the issues described in the previous section, and try to size partitions appropriately for the tasks they will fulfil. Here are some practical points to consider when partitioning disks:

- Size partitions appropriately for the jobs they will perform. Bear in mind that operating system upgrades are almost always bigger than previous versions, and that there is a general tendency for everything to grow.
- Bear in mind that RISC (e.g. Sun Sparc) compiled code is much larger than CISC compiled code (e.g. GNU/Linux), so software will take up more space on a RISC system.

- If we are sharing partitions, there might be limits on partition sizes. Some older implementations of NFS could handle file systems larger than 2GB, owing to the 32-bit limitation. This is not normally a problem now, though individual files cannot exceed 4GB on a 32-bit OS.
- Consider how backups will be made of the partitions. It might save many complications if disk partitions are small enough to be backed up in one go with a single tape, or other backup device.

Choosing partitions optimally requires both experience and forethought. Thumb-rules for sizing partitions change constantly, in response to changing RAM requirements and operating system sizes, disk prices, etc. In the early 1990s, many sites adopted diskless or partially diskless solutions [9], thus centralizing disk resources. In today's climate of ever cheaper disk space, there are few limitations left. Disk partitioning is performed with a special program. On PC hardware this is called `fdisk` or `cfdisk`. On Solaris systems the program is called, confusingly, `format`. To repartition a disk, we first edit the partition tables, then we have to write the changes to the disk itself. This is called *labelling* the disk. Both of these tasks are performed from the partitioning programs. It is important to make sure manually that partitions do not overlap. The partitioning programs do not normally help us here. If partitions overlap, data will be destroyed and the system will sooner or later get into deep trouble, as it assumes that the overlapping area can be used legitimately for two separate purposes.

Partitions are labelled with logical device names in Unix. As one comes to expect, these are different in every flavour of Unix. The general pattern is that of a separate device node for each partition, in the `/dev` directory, e.g. `/etc/sd1a`, `/etc/sd1b`, `/dev/dsk/c0t0d0s0`, etc. The meaning of these names is described in section 4.6.

The introduction of meta-devices and logical volumes in many operating systems allows one to ignore disk partitions to a certain extent. Logical volumes provide seamless integration of disks and partitions into a large virtual disk which can be organized without worrying about partition boundaries. This is not always desirable, however. Sometimes partitions exist for protection, rather than merely for necessity.

4.3.3 Formatting and Building File Systems

Disk formatting is a way of organizing and finding a way around the surface of a disk. It is a little bit like painting parking spaces in a car park. We could make a car park in a field of grass, but everything would get rapidly disorganized. If we paint fixed spaces and number them, then it is much easier to organize and reuse space, since people park in an orderly fashion and leave spaces of a standard, reusable size. On a disk surface, it makes sense to divide up the available space into sectors or blocks. The way in which different operating systems choose to do this differs, and thus one kind of formatting is incompatible with another.

The nomenclature of formatting is confused by differing cultures and technologies. Modern hard disks have intelligent controllers which can map out the disk surface independently of the operating system which is controlling them. This means that there is a kind of factory formatting which is inherent to the type of disk. For instance, a SCSI disk surface is divided up into *sectors*. An operating system using a SCSI disk then groups these sectors into new units

called *blocks* which are a more convenient size to work with, for the operating system. With the analogy above, it is a little like making a car park for trucks by grouping parking spaces for cars. It also involves a new set of labels. This regrouping and labelling procedure is called *formatting* in PC culture, and is called *making a file system* in Unix culture[1]. Making a file system also involves setting up an infrastructure for creating and naming files and directories. A file system is not just a labelling scheme, it also provides functionality.

If a file system becomes damaged, it is possible to lose data. Usually, file system checking programs called disk doctors, e.g. the Unix program `fsck` (file system check), can be used to repair the operating system's map of a disk. In Unix file systems, data which lose their labelling get placed for human inspection in a special directory which is found on every partition, called `lost+found`.

The file system creation programs for different operating systems go by various names. For instance, on a Sun host running SunOS/Solaris, we would create a file system on the zeroth partition of disk 0, controller zero with a command like this to the raw device:

```
newfs -m 0 /dev/rdsk/c0t0d0s0
```

Use `newfs` command is a friendly front-end to the `mkfs` program. The option `-m 0`, used here, tells the file system creation program to reserve zero bytes of special space on the partition. The default behaviour is to reserve 10% of the total partition size which ordinary users cannot write to. This is an old mechanisms for preventing file systems from becoming too full. On today's disks, 10% of a partition size can be many files indeed, and if we partition our cheap, modern disks correctly, there is no reason not to allow users to fill them up completely. This partition is then made available to the system by mounting it. This can either be performed manually:

```
mount /dev/dsk/c0t0d0s0 /mountpoint/directory
```

or by placing it is the file system table `/etc/vfstab`.

GNU/Linux systems have a simpler file system, with a straightforward `mkfs` command, e.g.

```
mkfs /dev/hda1
```

The file systems are registered in the file `/etc/fstab`. Other Unix variants register disks in equivalent files with different names, e.g. HPUX in `/etc/checklist` (prior to 10.x) and AIX in `/etc/file systems`.

On NT systems, disks are detected automatically and partitions are assigned to different logical drive names. Drive letters `C:` to `Z:` are used for non-floppy disk devices. NT assigns drive letters based on what hardware it finds at boot-time. Primary partitions are named first; then each secondary partition is assigned a drive letter. The `format` program is used to generate a file system on a drive. The command

```
format /fs:ntfs /v:spare F:
```

[1] Sometimes Unix administrators speak about reformatting a SCSI disk. This is misleading. There is no reformatting at the SCSI level; the process referred to here amounts to an error correcting scan, in which the intelligent disk controller reevaluates what parts of the disk surface are undamaged and can be written to. All disks contain unusable areas which have to be avoided.

would create an NTFS file system on drive `F:` and give it a volume label 'spare'. The older, insecure file system FAT can also be chosen; however, this cannot be recommended. The GUI can also be used to partition and format inactive disks.

4.3.4 Swap Space

In Windows operating systems, virtual memory uses file system space for saving data to disk. In Unix-like operating systems, a preferred method is to use a whole, unformatted partition for virtual memory storage. This is more efficient than using space within a file system.

A virtual memory partition is traditionally called the swap partition, though few modern Unix-like systems 'swap' out whole processes, in the traditional sense. The swap partition is now used for *paging*. It is virtual memory scratch space, and uses direct disk access to address the partition. No file system is needed, because no functionality, in terms of files and directories, is needed for the paging system. As a rule of thumb, it makes sense to allocate a partition twice the size of the total amount of RAM for paging. On heavily used login servers, this might not be enough. More swap partitions can be added later, however. Swap partitions are listed in the file system table.

4.3.5 File System Layout

We have no choice about the layout of the software and support files which are installed on a host as part of the 'operating system'. This is decided by the system designers and cannot easily be changed. Software installation, user registration and network integration all make changes to this initial state, however. Such additions to the system are under the control of the system administrator, and it is important to structure these changes according to logical and practical principles which we shall consider below.

A working computer system has several facets:

- The operating system software distribution.
- Third party software.
- Users' files.
- Information databases.
- Temporary scratch space.

These are logically separate because

- They have different functions.
- They are maintained by different sources.
- They change at different rates.
- A different policy of backup is required for each.

Most operating systems have hierarchical file systems with directories and subdirectories. This is a powerful way of organizing data. Disks can also be divided up into partitions. In many operating systems the largest supported partition size is 2GB or 4GB, since that is the maximum size which can be represented in a 32-bit register. This might be a limitation which you have to live with in designing your file structure. Another issue in sizing partitions is how

you plan to make a backup of those partitions. To make a backup you need to copy all the data to some other location, traditionally tape. The capacities of different kinds of tape varies quite a bit, as does the software for performing backups.

The point of directories and partitions is to separate files so as not to mix together things which are logically separate. There are many things which we might wish to keep separate: for example,

- User home directories.
- Development work.
- Commercial software.
- Free software.
- Local scripts and databases.

One of the challenges of system design is in finding an appropriate directory structure for all data which are not a part of the operating system, i.e. all those files which are locally maintained.

Principle 12 (Separation I) *Data which are separate from the operating system should be kept in a separate directory tree, preferably on a separate disk partition. If they are mixed with the operating system file-tree it makes re-installation or upgrade of the operating system unnecessarily difficult.*

The essence of this is that it makes no sense to mix logically separate file trees. For instance, users' home directories should never be on a common partition with the operating system. Indeed, file systems which grow with a life of their own should never be allowed to consume so much space as to throttle the normal operation of the machine.

These days there are few reasons for dividing the files of the operating system distribution into several partitions (e.g. `/`, `/usr` in Unix). Disks are large enough to install the whole operating system distribution on a single independent disk or partition. If you have done a good job of separating your own modifications from the system distribution, then there is no sense in making a backup of the operating system distribution itself, since it is trivial to reinstall from source (CD-ROM or ftp file base). Some adminstrators like to keep `/var` on a separate partition, since it contains files which vary with time, and should therefore be backed up.

Operating systems often have a special place for installed software. Regrettably they often break the above rule and mix software with the operating system's file tree. Under Unix-like operating systems, the place for installed third-party software is traditionally `/usr/local`, or simply `/local`. Fortunately under Unix, separate disk partitions can be woven anywhere into the file tree on a directory boundary, so this is not a practical problem as long as everything lies under a common directory. In NT software is often installed in the same directory as the operating system itself; also, NT does not support partition mixing in the same way as Unix, so the re-installation of NT means re-installation of all the software as well.

Data which are installed or created locally are not subject to any constraints, though they may be installed anywhere. One can therefore find a naming scheme which gives the system logical clarity. This benefits users and management issues. Again, we may use directories for this purpose. Operating systems which descended from DOS also have the concept of drive numbers like `A:`, `B:`, `C:`, etc. These are assigned to different disk partitions. Some Unix

operating systems have virtual file systems which allow one to add disks transparently without ever reaching a practical limit. Users never see partition boundaries. This has both advantages, and disadvantages, since small partitions are a cheap way to contain groups of misbehaving users, without resorting to disk quotas.

4.3.6 Object Orientation: Separation of Independent Issues

The computing community is currently riding a wave of affection for object orientation as a paradigm in computer languages and programming methods. Object orientation in programming languages is usually presented as a fusion of two independent ideas: classification of data types and access control based on scope. The principle from which this model has emerged is simpler than this, however: it is simply the observation that information can be understood and organized most efficiently if *logically independent* items are kept separate[2]. This simple idea is a powerful discipline, but like most disciplines it requires a strong will on the part of a system administrator in order to avoid a decline into chaos. We can now restate the earlier principle about operating system separation more generally:

Principle 13 (Separation II) *Data which are logically separate belong in separate directory trees, perhaps on separate disk partitions.*

The basic file system objects, in order of global to increasingly local, are *disk partition*, *directory* and *file*. As system administrators, we are not usually responsible for the contents of files, but we do have some power to decide on their organization by placing them in carefully labelled directories, within partitions. Partitions are useful because they can be dumped (backed-up to tape, for instance) as independent units. Directories are good because they hide and group related files into units.

Many institutions make backups of the whole operating system partition because they do not have a system for separating the files which they have modified, or configured specially. The number of such files is usually small. For example,

- the password and group databases,
- kernel configuration,
- files in `/etc` like services, default configurations files,
- special start-up scripts.

It is easy to make a copy of these few files in a location which is independent of the locations where the files actually need to reside, according to the rules of the operating system.

A good solution to this issue is to make *master copies* of files like `/etc/group`, `/etc/services`, `/etc/sendmail.cf`, etc., in a special directory which is separate from the OS distribution. For example, you might choose to collect all of these in a directory such as `/local/custom` and to use a script, or cfengine, to make copies of these master files in the actual locations required by the operating system. The advantages to this approach are

[2] It is sometimes claimed that object orientation mimics the way humans think. This, of course, has no foundation in the cognitive sciences. A more careful formulation would be that object orientation mimics the way in which humans organize and administrate. That has nothing to do with the mechanisms by which thoughts emerge in the brain.

- RCS version control of changes is easy to implement.
- Automatic backup and separation.
- Ease of distribution to other hosts.

The exception to this rule must be the password database `/etc/passwd`, which is actually altered by an operating system program `/bin/passwd` rather than the system administrator. In that case the script would copy from the system partition to the custom directory.

Keeping a separate disk partition for software, which you install from third parties, makes clear sense. It means that you will not have to reinstall that software later when you upgrade your operating system. The question then arises as to how such software should be organized within a separate partition.

Traditionally, third party software has been installed in a directory under `/usr/local`, or simply `/local`. Software packages are then dissected into libraries, binaries and supporting files which are installed under `/local/lib`, `/local/bin` and `/local/etc`, to mention just a few examples. This keeps third party software separate from operating system software, but there is no separation of the third party software. Another solution would be to install one software package per directory under `/local`.

4.4 Installation of the Operating System

The installation process is one of the most brutal and destructive things we can do to a computer. It doesn't matter whether it is a first-time installation of a new host, or an operating system upgrade from scratch: it involves deleting everything and building it up again from scratch. Obviously this is also constructive, but *everything* on the disk will disappear during the installation. You need to accept this and learn to live with the idea. It has a positive side: if you do something wrong, you cannot do any more damage to the system than you have already done!

Before deleting everything on a machine you should step back and sweat a bit – feel the responsibility. You should be thinking: have I kept a copy of important information? Which information will I be deleting? How could I, if necessary, get it back?

Installing a new machine is a simple affair these days. The operating system comes on some removable medium (like a CD). You put it in the player and run a program, usually called install. Alternatively, you boot a machine directly from the CD and the program is run automatically. Then you simply answer a few questions and the installation is done for you.

Operating systems are getting big so they are split up into packages. You are expected to choose whether you want to install everything or whether you just want to install certain packages. Most operating systems provide a fancy package installation program which helps you with this. In most cases these programs are quite stupid, they don't tell you that something will break if you don't install certain packages. For that reason it is strongly recommended that you always install the complete operating system: every single package. Whether you know it or not, you almost certainly need the whole thing – and the stuff you don't need probably doesn't take up much space anyway. Disk is cheap, but time spent trying to find out what went wrong with an installation is expensive. The exception here is GNU/Linux which bundles all available software with the OS.

To answer the questions about installing a new host, you need to collect some information and make some choices:

- You must decide a name for your machine.
- You will need an unused Internet address.
- You need to decide how much swap space to allocate. A good rule of thumb is at least twice the amount of RAM you have installed.
- You need to know the local netmask.
- You need to know the local timezone.
- You need to know the name of your local domain.
- You need to know whether you are using the Network Information Service (NIS) or other directory service and, if so, what the name of the server is.

When we have this information, we are ready to begin.

4.4.1 Solaris

Solaris can be installed in a number of ways. The simplest is from CD-ROM. At the boot prompt, we simply type

```
? boot cdrom
```

This starts a graphical user interface which leads one through the steps of the installation from disk partitioning to operating system installation. The procedure is well described in the accompanying documentation; indeed it is quite intuitive, so we needn't belabour the point here. The installation procedure proceeds through the standard list of questions, in this order:

- Preferred language and keyboard type.
- Name of host.
- Net interfaces and IP addresses.
- Subscribe to NIS or NIS plus domain, or not.
- Subnet mask.
- Timezone.
- Choose upgrade or install from scratch.

Solaris installation addresses an important issue, namely that of customization and integration. As part of the installation procedure, Solaris provides a service called Jumpstart, which allows hosts to execute specialized scripts that customize the installation. In principle, the automation of hosts can be completely automated using Jumpstart. Customization is extremely important for integrating hosts into a local network. As we have seen, vendor standard models are almost never adequate in real networks. By making it possible to adapt the installation procedure to local requirements, Solaris makes a great contribution to automatic network configuration.

Installation from CD-ROM assumes that every host has a CD-ROM from which to install the operating system. This is seldom the case, so Solaris also enables hosts with CD-ROM players to act as network servers for their CD-ROMS, thus allowing the operating system to be installed directly from the network. Again, these hosts have access to Jumpstart procedures. The only disadvantage of the network installation is Sun's persistence in relying on NIS and NIS plus.

4.4.2 GNU/Linux

GNU/Linux is not one, but a family of operating systems. There are many distributions, maintained by different organizations and they are installed in different ways. Usually, one balances ease of installation with flexibility of choice. It is common to opt for a package system, whereby each software item is bundled as a package which can either be included or excluded from the installation.

What makes GNU/Linux installation unique amongst operating system installations is the sheer size of the program base. Since every piece of Free Software is bundled, there are literally hundreds of packages to choose from. This presents GNU/Linux distributors with a dilemma. To make installation as simple as possible, package maintainers make software self-installing with some kind of default configuration. This applies to user programs and to operating system services. Here lies the problem: installing network services which we don't intend to use presents a security risk to a host. A service which is installed is a way into the system. A service which we are not even aware of could be a huge risk. If we install everything, then, we are faced with a uncertainty in knowing what the operating system actually consists of, i.e. what we are getting ourselves into.

As with most operating systems, GNU/Linux installations assume that you are setting up a standalone PC which is yours to own and do with as you please. Although GNU/Linux is a multi-user system, it is treated as a single user system. Little thought is given to the effect of installing services like news servers and web servers. The scripts which are bundled for adding user accounts also treat the host as a little microcosm, placing users in `/home` and software in `/usr/local`. To make a network workstation out of GNU/Linux, we need to override many of its idiosyncrasies.

4.4.3 NT 4

The installation of NT is similar to both of the above. One starts with three boot diskettes and an already partitioned hard-drive (one is not asked about repartitioning during the installation procedure). On rebooting with the first of the boot diskettes, we are asked whether we wish to install NT anew, or repair an existing installation. This is rather like the GNU/Linux rescue disk. Next we choose the file system type for NT to be installed on, either DOS or NTFS. There is clearly only one choice: installing on a DOS partition would be irresponsible with regard to security. Choose NTFS.

NT reboots several times during the installation procedure. The first time around, it converts its default DOS partition into NTFS and reboots again. Then the remainder of the installation proceeds with a graphical user interface. There are several installation models for NT workstation, including regular, laptop, minimum and custom. Having chosen one of these, one is asked to enter a license key for the operating system. The installation procedure asks us whether we wish to use DHCP to configure the host with an IP address dynamically, or whether a static IP address will be set. After various other questions, the host reboots and we can log in as Administrator.

NT service packs are patch releases which contain important upgrades. These are refreshingly trivial to install on an already-running NT system. One simply inserts them into the CD-ROM drive and up pops the Explorer program with instructions and descriptions of contents. Clicking on the install link starts the upgrade. After a service pack upgrade, NT reboots

predictably and then we are done. Changes in NT configuration require one to reinstall service packs, however.

4.4.4 Configuring the Name Service

The name service must be configured for a system to be able to look up host names and Internet addresses. The most important file in this connection is `/etc/resolv.conf`. Ancient IRIX systems seem to have placed this file in `/usr/etc/resolv.conf`. This old location is obsolete. Without the resolver configuration file, a host will probably stop dead whilst trying hopelessly to look up Internet addresses. Hosts which use NIS or NIS plus might be able to look up local names. The most important features of this file are the definition of the domain-name and a list of name servers which can perform the address translation service. These name servers must be listed as IP numerical addresses (since DNS can't look up any names until it knows the name of a server to look them up on, and that's what we're trying to do). The format of the file is as shown below:

```
domain domain.country
nameserver 192.0.2.10
nameserver 158.36.85.10
nameserver 129.241.1.99
```

Some prefer to use the search directive in place of the domain directive, since it is more general allows several domains to be searched in special circumstances:

```
search domain.country
nameserver 192.0.2.10
nameserver 192.0.2.85
nameserver 192.0.2.99
```

On the host which is itself a name server, the first name server should be listed as the loopback address, so as to avoid sending traffic out onto the network when none is required:

```
search domain.country
nameserver 127.0.0.1
nameserver 192.0.2.10
nameserver 192.0.2.99
```

DNS has several competing services. A mapping of host names to IP addresses is also performed by the `/etc/hosts` database, and this file can be shared using NIS or NIS plus.

Windows NT has the WINS service. Modern Unix-like systems allow us to choose the order in which these competing services are given priority when looking up host name data. Unfortunately, there is no standard way of configuring this. GNU/Linux and public domain resolver packages for old SunOS (`resolv+`), use a file called `/etc/hosts.conf`. The format of this file is

```
order hosts,bind,nis
multi on
```

This example tells the lookup routines to look in the `/etc/hosts` file first, then to query DNS/BIND, and then finally, look at NIS. The resolver routines quit after the first match they find, they do not query all three databases every time. Solaris, and now also some GNU/Linux

distributions, use a file called `/etc/nsswitch.conf` which is a general configuration for all database services, not just the hostname service.

```
#           files,nis,nisplus,dns

passwd:     files
group:      files
hosts:      files dns
networks:   files
protocols:  files
rpc:        files
ethers:     files
netmasks:   files
bootparams: files
```

4.4.5 Marrying Unix and NT

If we are installing a personal PC workstation which does not have a special role in the network, it could be of interest to be able to choose between operating systems. We might want DOS to play games, Unix for development work, NT or Windows for its well-known applications, and perhaps OS/2, out of respect. Each of these operating systems has to live on a different partition of the host's disk(s). The question is: how do we marry these operating systems in such a way that they do not try to kill each other?[3]

A word of warning: because of the wide range of system cracking tools available to users, it can be risky to install dual-boot operating systems on any important host. NT is particularly vulnerable to password editing on dual boot systems.

The installation programs for NT and GNU/Linux do not always respect each other's independence. Experience says: install NT first, then Unix-like OSes, being careful *not* to choose to install a 'master boot record' as we move through the installation menus.

If you do not have access to the OS2 boot manager, you can use NT's own boot manager if you trick it by copying the boot blocks from the unix file system to a DOS file system. The command

```
dd if=/dev/hdaLINUX of=/dev/hdaDOS/bootsect.linux bs=512
  count=1
```

copies the bootsector to `C:\bootsect.linux`. It can now be used by the boot manager, by editing `C:\boot.ini`

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition({\var NT partition
  number})\WINNT
[operating systems]
c:\bootsect.linux="Linux"
multi(0)disk(0)rdisk(0)partition({\var NT Partition})
  \Winnt="Windows Nt"
```

An alternative method is to copy the program `loadlin` and kernel to the DOS drive so that UNIX can be started from a windows command (which you might like to put in a batch file)

[3] Of course, marrying them does not imply that they will be able to understand or talk to one another!

```
loadlin vmlinuz root="/dev/hdaLinux-partition"
```

Note–it might not be a good idea to use this method from a multi-tasking Windows NT machine, since it does not give the system a chance to shut down properly and damage to the NT partition might result. This method was meant for DOS PCs.

Another problem is that disk boot information might not get written to the GNU/Linux partition. To fix this problem, boot with the GNU/Linux boot diskette and edit the file `/etc/lilo.conf`. Make sure that the correct partition is bootable (use `fdisk`). Once done, we execute the command `lilo`. If the lilo boot information is not correctly stored, the boot manager might not recognize the partitions and claim that the GNU/Linux partition is not formatted.

4.4.6 Diskless Clients

Diskless workstations are, as per the name, workstations which have no disk at all. Diskless workstations know absolutely nothing other than the MAC address of their network interface (Ethernet address). In earlier times, when disks were expensive, diskless workstations were seen as a cheap option. Today, diskless workstations are no longer a good idea. Disks are cheap, and diskless workstations are expensive in terms of network bandwidth. They are also awkward to maintain.

Diskless clients require disk space on a server-host in order to function, i.e. some other host which does have a disk needs to be a disk server for the diskless clients. Most vendors supply a script for creating diskless workstations. This script is run on the server-host.

A diskless workstation needs its own root file system and its own swap space, but it can share system files under `/usr`. The script creates disk areas for a root partition, `/export/swap/clientname` and `/export/root/clientname`. These areas need to be exported to the clients with root privileges granted. The file `/etc/ethers` on the server must contain the Ethernet addresses of the clients.

When a diskless system is switched on for the first time, it has no files and knows nothing about itself except the Ethernet address on its network card. It proceeds by sending a RARP (Reverse Address Resolution Protocol), BOOTP or DHCP request out onto the local subnet in the hope that a server (`in.rarpd`) will respond by telling it its Internet address. The server hosts must be running two services: `rpc.bootparamd` and `tftpd`, the trivial file transfer program. This is another reason for arguing against diskless clients: these services are rather insecure and could be a security risk for the server host. A call to the `rpc.bootparamd` daemon transfers data about where the diskless station can find a server, and what its swap-area and root directory are called in the file tree of this server. The root directory and swap file are mounted using the NFS. The diskless client loads its kernel from its root directory, and thereafter everything proceeds as normal. Diskless workstations swap to files rather than partitions. The command `mkfile` is used to create a fixed-size file for swapping. This is also less than maximally efficient.

4.4.7 Dual Homed Host

A host with two network interfaces, both of which is coupled to a network, is called a dual-homed host. Dual homed hosts are important in building *firewalls* for network security. Most

vendor operating systems will configure dual network interfaces automatically. Free Unix-like operating systems will not, however. Briefly, here is a GNU/Linux setup for two network interfaces:

1 Compile a new kernel with inbuilt support for both types of interface.

2 Change the lilo configuration to detect both interfaces, by adding

```
append="ether=0,0,eth0 ether=0,0,eth1"
```

to `/etc/lilo.conf`.

3 The new interface can be assigned an IP address in the file `/etc/init.d/network`.

One must then decide how the IP addresses are to be registered in the DNS service? Will the host have the same name on both interfaces, or will it have a different name? This is essentially a cosmetic issue. Packet routing on dual-homed hosts has been discussed elsewhere [231].

4.4.8 Cloning Systems

We are almost never interested in installing every machine separately. A system administrator usually has to install ten, twenty or even a hundred machines at a time. He or she would also like them to be as far as possible the same, so that users will always know what to expect. This might sound like a straightforward problem, but it is not. There are several approaches:

- A few Unix-like operating systems provide a solution to this using package templates so that the installation procedure becomes standardized.
- The hard disks of one machine can be physically copied, and then the host name and IP address can be edited afterwards.
- All software can be placed on one host and shared using NFS, or another shared file system.

Each of these approaches has its attractions. The NFS/shared file system approach is without doubt the least amount of work, since it involves installing the software only once, but it is also the slowest in operation for users.

As an example of the first, here is how Debian GNU/Linux tackles this problem using the Debian package system:

```
Install one system
 dpkg --get-selections > file
On the remaining machines type
 dpkg --set-selections < file
Run install packages program.
```

Alternatively, one can install a single package with:

```
dpkg -i package.deb
```

In RedHat Linux, a similar mechanism looks like this:

```
rpm -ivh package.rpm
```

Disks can be mirrored directly, using some kind of cloning program. For instance, the Unix tape archive program can be used to copy the entire disk of one host. To make this work, we first have to perform a basic installation of the OS, with zero packages and then copy over all remaining files which constitute the packages we require. In the case of the Debian system above, there is no advantage to doing this, since the package installation mechanism can do the same job more cleanly. For example, with a GNU/Linux distribution:

```
tar --exclude /proc --exclude /lib/libc.so.5.4.23 \
    --exclude /etc/hostname --exclude /etc/hosts -c -v \
    -f host-imprint.tar /
```

Note that several files must be excluded from the dump. The file `/lib/libc.so.5.4.23` is the C library; if we try to write this file back from backup, the destination computer will crash immediately. `/etc/hostname` and `/etc/hosts` contain definitions of the host name of the destination computer, and must be left unchanged. Once a minimal installation has been performed on the destination host, we can access the tar file and unpack it to install the image:

```
(cd / ; tar xfp /mnt/dump/my-machine.tar; lilo)
```

Afterwards, we have to install the boot sector, with the `lilo` command. The cloning of Unix systems has been discussed in refs. [256, 291].

Note that NT systems cannot be cloned without special software (e.g. Norton Ghost or PowerQuest Drive Image). There are fundamental technical reasons for this. One is the fact that many host parameters are configured in the impenetrable *system registry*. Unless all of the hardware and software details of every host are the same, this will fail with an inconsistency. Another reason is that users are registered in a binary database with security IDs which can have different numerical values on each host. Finally, domain registration cannot be cloned. A host must register manually with its domain server.

4.5 Software Installation

With the notable exception of the GNU/Linux operating system, most standard operating system installations will not leave us in possession of an immediately usable system. We also need to install third-party software in order to get useful work out of the host. Software installation is a very similar problem to that of operating system installation – after all, the operating system is software. However, third party software originates from a different source than the operating system. It is often bound by license agreements, and it needs to be distributed around the network. Some software has to be compiled from source. We therefore need a thoughtful strategy for dealing with software. Specialized schemes for software installation were discussed in refs. [62, 168], and a POSIX draft was discussed in ref. [12], though little seems to have come of it.

4.5.1 Free and Proprietary Software

Unlike other popular operating systems, Unix grew up around people who write their own software rather than relying on off-the-shelf products. The Internet contains gigabytes of

software for Unix systems which costs nothing. Large companies like the oil industry and newspapers can afford off-the-shelf software for Unix, but most people can't.

There are therefore two kinds of software installation: the installation of software from binaries, and the installation of software from source. Commercial software is usually installed from a CD by running an installation program and following the instructions carefully; the only decision we need to make is where we want to install the software. Free software and Open Source software usually comes in an source form and must therefore be compiled. Unix programmers have gone to great lengths to make this process as simple as possible for system administrators.

4.5.2 Structuring Software

The first step in installing software is to decide where we want to keep it. We could, naturally, locate software anywhere we like, but consider the following:

- Software should be separated from the operating system's installed files, so that the OS can be reinstalled or upgraded without ruining a software installation.
- Unix-like operating systems have a naming convention. Compiled software can be collected in a special area, with a `bin` directory and a `lib` directory so that binaries and libraries conform to the usual Unix conventions. This makes the system consistent and easy to understand. It also keeps the program search `PATH` variable simple.
- Home-grown files and programs which are special to our own particular site can be kept separate from files which could be used anywhere. In that way, we define clearly the validity of the files and we see who is responsible for maintaining them.

The directory traditionally chosen for installed software is called `/usr/local`. One then makes sub-directories `/usr/local/bin` and `/usr/local/lib` and so on [119]. Unix has a *de facto* naming standard for directories which we should try to stick to as far as reason permits, so that others will understand how our system is built up:

- `bin` Binaries or executables for normal user programs.
- `sbin` Binaries or executables for programs which only system administrators require. Those files in `/sbin` are often statically linked to avoid problems with libraries which lie on unmounted disks during system booting.
- `lib` Libraries and support files for special software.
- `etc` Configuration files.
- `share` Files which might be shared by several programs or hosts. For instance, databases or help information; other common resources.

One suggestion for structuring installed software is shown in Figure 4.1. Another is shown in Figure 4.2. Here we divide these into three categories: regular installed software, GNU software (i.e. free software) and site-software. The division is fairly arbitrary. The reason for this is as follows:

- `/usr/local` is the traditional place for software which does not belong to the OS. We could keep everything here, but we will end up installing a lot of software after a while, so it is useful to create two other sub-categories.

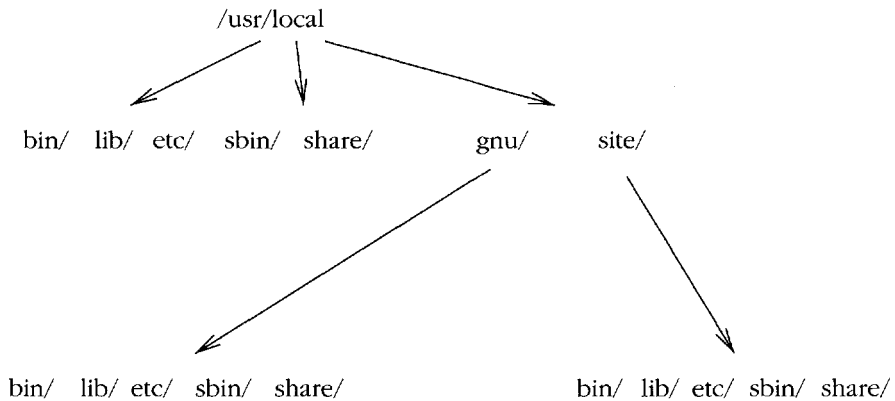


Figure 4.1 One way of structuring local software. There are plenty of things to criticize here. For instance, is it necessary to place this under the traditional `/usr/local` tree? Should GNU software be underneath `/usr/local`? Is it even necessary or desirable to formally distinguish GNU software from other software?

- GNU software, written by and for the Free Software Foundation, forms a self-contained set of tools which replace many of the older UNIX equivalents, like `ls` and `cp`. GNU software has its own system of installation and set of standards. GNU will also eventually become an operating system in its own right. Since these files are maintained by one source, it makes sense to keep them separate. This also allows us place GNU utils ahead of others in a user's command PATH.
- Site specific software includes programs and data which we build locally to replace the software or data which follows with the operating system. It also includes special data like the database of `aliases` for e-mail and the DNS tables for our site. Since it is

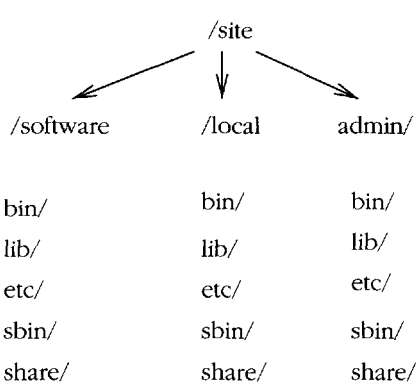


Figure 4.2 Another, more rational way of structuring local software. Here we drop the affectation of placing local modifications under the operating system's `/usr` tree and separate it completely. Symbolic links can be used to alias `/usr/local` to one of these directories for historical consistency

special to our site, created and maintained by our site, we should keep it separate so that it can be backed up often and separately.

A similar scheme to this was described in refs. [170, 49, 282, 220], in a system called *Depot*. In the Depot system software is installed under a file node called `/depot` which replaces `/usr/local`. In the depot scheme, separate directories are maintained for different machine architectures under a single file tree. This has the advantage of allowing every host to mount the same file system, but the disadvantage of making the single file system very large. Software is installed in a package-like format under the depot tree, and is linked in to local hosts with symbolic links. A variation on this idea from the University of Edinburgh was described in ref. [8], and another from the University of Waterloo uses a file tree `/software` to similar ends in ref. [232]. In the Soft environment [82], software installation and user environment configuration are dealt with in a combined abstraction.

4.5.3 GNU Software Example

Let us now illustrate the GNU method of installing software which has become widely accepted. This applies to any type of Unix, and to NT if one has a Unix compatibility kit, such as Cygwin or UWIN. To begin compiling software, one should always start by looking for a file called `README` or `INSTALL`. This tells us what we have to do to compile and install the software. In most cases, it is only necessary to type a couple of commands, like in the following example. When installing GNU software, we are expected to give the name of a *prefix* for installing the package. The prefix in the above cases is `/usr/local` for ordinary software, `/usr/local/gnu` for GNU software and `/usr/local/site` for site-specific software. Most software installation scripts place files under `bin` and `lib` automatically. The steps are as follows:

1 Make sure we are working as a regular, unprivileged user. The software installation procedure might do something which we do not agree with. It is best to work with as few privileges as possible until we are sure.

2 Collect the software package by `ftp` from a site like `ftp.uu.net` or `ftp.funet.fi`, etc. Use a program like `ncftp` for painless anonymous login.

3 Unpack the file using `tar zxf software.tar.gz`, if using GNU tar, or `gunzip software.tar.gz; tar xf software.tar` if not.

4 Enter the directory which is unpacked, `cd software`.

5 Type: `configure --prefix=/usr/local/gnu`. This checks the state of our local operating system and other installed software, and configures the software to work correctly there.

6 Type: `make`

7 If all goes well, type `make -n install`. This indicates what the make program will install and where. If we have any doubts, this allows us to make changes or abort the procedure without causing any damage.

8 Finally, switch to privileged root/Administrator mode with the `su` command and type `make install`. This should be enough to install the software. Note, however, that this

step is a security vulnerability. If one blindly executes commands with privilege, one can be tricked into installing back-doors and Trojan horses (see Chapter 9).

9 Some installation scripts leave files with the wrong permissions so that ordinary users cannot access the files. We might have to check that the files have a mode like 555 so that normal users can access them. This is in spite of the fact that installation programs attempt to set the correct permissions [246].

Today this procedure should be more or less the same for just about any software we pick up. Older software packages sometimes provide only Makefiles which you must customize yourself. Some X11-based windowing software requires you to use the `xmkmf` X-make-makefiles command instead of configure. Always look at the README file.

4.5.4 Proprietary Software Example

If we are installing proprietary software, we will have received a copy of the program on a CD-ROM, together with licensing information, i.e. a code which activates the program.

The steps a somewhat different:

1 To install from CD-ROM we must start work with root/Administrator privileges, so the authenticity of the CD-ROM should be certain..

2 Insert the CD-ROM into the drive. Depending on the operating system, the CD-ROM might be mounted automatically or not. Check this using the `mount` command with no arguments, on a Unix-like system. If the CD-ROM has not been mounted, then for standard CD-ROM formats, the following will normally suffice:

```
mkdir /cdrom                    if necessary
mount /dev/cdrom /cdrom
```

For some manufacturers, or on older operating systems, we might have to specify the type of file system on the CD-ROM. Check the installation instructions.

3 On an NT system a click-able icon appears to start the installation program. On a Unix-like system we need to look for an installation script

```
cd /cdrom/cd-name
less README
./install-script
```

4 Follow instructions.

Some proprietary software requires the use of a license server, such as `lmgrd`. This is installed automatically, and we are required only to edit a configuration file, with a license key which is provided, in order to complete the installation. Note, however, that if we are running multiple licensed products on a host, it is not uncommon that these require different and partly incompatible license servers which interfere with one another. If possible, one should keep to only one license server per subnet.

4.5.5 Recommended Free Software

A few software packages can be highly recommended for any Unix-like system:

- *Perl:* a system independent script language which is invaluable for system automation tasks, and for WWW programming.
- *GNU fileutils:* This is GNU's set of standard shell commands for file handling, including a replacement `ls` and the superior `tar` command, alluded to above.
- *GNU binutils:* This is GNU's set of compilation tools. If we include the GNU C compiler `gcc` then we have an extensive set of development tools, and everything we need to compile third-party software.

These packages make diverse operating systems look and behave similarly. They solve many problems of standardization and absent functionality in the outdated shell-commands distributed with many types of Unix.

Suggestion 2 (GNU fileutils on Unix) *The GNU fileutils programs for Unix-like operating systems are superior in functionality than their corresponding vendor versions. Moreover, they work on every platform, bringing a pleasant dose of uniformity to a heterogeneous network. They can be placed in the users'* `PATH` *variable so as to override the vendor commands. In some instances, vendor programs have specially adapted features. One example is the* `ls` *command. Some Unix-like systems have ACLs (Access Control Lists) which give extended file permissions. These are invisible with the GNU version of* `ls`*, but are marked with an additional '+' to the left of the access bits, when using the vendor* `ls` *command. In the case of* `ls`*, it is probably worth removing or renaming the GNU* `ls` *to, say,* `gls`.

4.5.6 Installing Shared Libraries

Systems which use shared libraries or shared objects sometimes need to be reconfigured when new libraries are added to the system. This is because the names of the libraries are cached to fast access. The system will not look for a library if it is not in the cache file:

- SunOS (prior to Solaris 2): after adding a new library, one must run the command ldconfig *`lib-directory`*. The file `/etc/ld.so.cache` is updated.
- GNU/Linux: new library directories are added to the file `/etc/ld.so.conf`. Then one runs the command `ldconfig`. The file `/etc/ld.so.cache` is updated.

4.5.7 Configuration Security

In the preceding sections we have looked at some examples and suggestions for dealing with software installation. Let us now take a step back from the details to analyse the principles underlying these.

The first is a principle which we shall return to umpteen times in this book. It is one of the key principles in computer science, and we shall be repeating it with slightly different words again and again.

Principle 14 (Separation III) *Independent systems should not interfere with one another, or be confused with one another. Keep them in separate storage areas.*

The reason is clear: if we mix up files which do not belong together, we lose track of them. They become obscured by a lack of structure. They vanish into anonymity. The reason why all modern computer systems have directories for grouping files, is precisely so that we do not have to mix up all files in one place. This was discussed in section 4.3.5. The application to software installation is clear: we should never consider installing software in `/usr/bin` or `/bin` or `/lib` or `/etc` or any directory which is controlled by the system designers. To do so is like lying down in the middle of a freeway and waiting for a new operating system or upgrade to roll over us. If we mix local modifications with operating system files, we lose track of the difference the system, others will not be able to see what we have done. All our hard work will be for nothing when a new system is installed.

Suggestion 3 (Vigilance) *Be on the lookout for software which is configured, by default, to install itself on top of the operating system. Always check the destination using* `make -n` *install before actually committing to an installation. Programs which are replacements for standard operating system components often break the principle of separation*[4].

The second important point above is that we should never work with root privileges unless we have to. Even when we are compiling software from source, we should not start the compilation with superuser privileges. The reason is clear: why should we trust the source of the program? What if someone has placed a command in the build instructions to destroy the system, plant a virus or open a back-door to intrusion? As long as we work with low privilege then we are immune from such problems.

Principle 15 (Limited privilege) *No process or file should be given more privileges than it needs to do its job. To do so is a security hazard.*

Another use for this principle arises when we come to configure certain types of software. When a user executes a software package, it normally gets executed with the user privileges of that user. There are two exceptions to this:

- *Services which are run by the system:* daemons which carry out essential services for users or for the system itself, run with a user ID which is independent of who is logged on to the system. Often, such daemons are started as root or the Administrator when the system boots. In many cases, the daemons do not need these privileges and will function quite happily with ordinary user privileges after changing the permissions of a few files. This is a much safer strategy than allowing them to run with full access. For example, the `httpd` daemon for the WWW service uses this approach. In recent years, bugs in many programs which run with root privileges have been exploited to give intruders access to the system. If software is run with a non-privileged user ID, this is not possible.
- *Unix setuid programs:* Unix has a mechanism by which special privilege can be given to a user for a short time, while a program is being executed. Software which is installed with the Unix setuid bit set, and which is owned by root, runs with root's special privileges. Some software producers install software with this bit set with no respect

[4] Software originating in BSD Unix is often an offender, since it is designed to be a part of BSD Unix, rather than an add-on, e.g. sendmail and BIND.

for the privilege it affords. Most programs which are setuid root do not need to be. A good example of this is the Common Desktop Environment (a multi-vendor desktop environment used on Unix systems). In a recent release, almost every program was installed with setuid root. Within only a short time, a list of reports about users exploiting bugs to gain control of these systems appeared. In the next release, none of the programs had setuid root.

All software servers which are started by the system at boot time are started with root/Administrator privileges, but daemons which do not require these privileges can relinquish them by giving up their special privileges and running as a special user. This approach is used by the Apache WWW server and by MySQL, for instance. These are examples of software which encourage us to create special user IDs for server processes. To do this, we create a special user in the password database, with no login rights (this just reserves a UID). In the above cases, these are usually called `www` and `mysql`. The software allows us to specify these user IDs so that the process owner is switched right after starting the program. If the software itself does not permit this, we can always force a daemon to be started with lower privilege using:

```
su -c 'command' user
```

The management tool cfengine can also be used to do this. Note, however, that server processes which run on reserved (privileged) ports 1–256 have to be started with root privileges in order to bind to their sockets.

On the topic of root privilege, a related security issue has to do with programs which write to temporary files.

Principle 16 (Temporary files) *Temporary files or sockets which are opened by any program should not be placed in any publicly writable directory like* `/tmp`. *This opens for the possibility of race conditions and symbolic link attacks. If possible, configure them to write to a private directory.*

Users are always more devious than software writers. A common mistake in programming is to write to a file which ordinary users can create, using a privileged process. If a user is allowed to create a file object with the same name, then he or she can direct a privileged program to write to a different file instead, simply by creating a symbolic or hard link to the other file. This could be used to overwrite the password file or the kernel, or the files of another user. Software writers can avoid this problem by simply unlinking the file they wish to write to first, but that still leaves a window of opportunity after unlinking the file and before opening the new file for writing, during which a malicious user could replace the link (remember that the system time-shares). The lesson is to avoid making privileged programs write to directories which are not private, if possible.

Before closing this section, a comment is in order. Throughout this chapter, and others, we have been advocating a policy of building the best possible, most logical system by tailoring software to our own environment. Altering absurd software defaults, customizing names and locations of files and changing user identities is no problem as long as everyone who uses and maintains the system is aware of this. If a new administrator started work and, unwittingly, reverted to those software defaults, then the system breaks.

Principle 17 (Flagging customization) *Customizations and deviations from standards should be made conspicuous to users and administrators. This makes the system easier to understand both for ourselves and our successors.*

4.5.8 When Compilation Fails

Today, software producers who distribute their source code are able to configure it automatically to work with most operating systems. Compilation usually proceeds without incident. Occasionally, though, an error will occur which causes the compilation to halt. There is a few things we can try to remedy this:

- A previous configuration might have been left lying around; try

  ```
  make clean
  make distclean
  ```

  and start again, from the beginning.
- Make sure that the software does not depend upon the presence of another package, or library. Install any dependencies, missing libraries and try again.
- Errors at the linking stage about missing functions are usually due to missing or unlocatable libraries. Check that the

  ```
  LD_LIBRARY_PATH
  ```

  variable includes all relevant library locations. Are any other environment variables required to configure the software?
- Sometimes an extra library needs to be added to the Makefile. To find out whether a library contains a function, we can use the following C-shell trick:

  ```
  host% cd /lib
  host% foreach lib ( lib* )
  > echo Checking lib ----------------------
  > nm lib | grep function
  >end
  ```
- Carefully try to patch the source code to make the code compile.
- Check in news groups whether others have experienced the same problem.
- Contact the author of the program.

4.5.9 Upgrading Software

Some software (especially free software) gets updated very often. We could easily spend an entire life just chasing the latest versions of favourite software packages. Avoid this.:

- It is a waste of time.
- Sometimes new versions contain more bugs than the old one, and an even-newer version is just around the corner.
- Users will not thank us for changing things all the time. Stability is a virtue. Everyone likes time to get used to the system before change strikes.

4.6 Installing a Unix Disk

Adding a new disk or device to a Unix-like host involves some planning. The first concern is what type of hard disk. There are several types of disk interface used for communicating with hard disks. Some are cheap and cheerful (IDE), while others are more expensive, fast and reliable (SCSI).:

- *IDE:* a maximum of two disks can be used and the size of the disks may be limited to less than a gigabyte, but these disks are cheap and cheerful.
- *EIDE:* an extended version of the IDE interface. Allows four disks.
- *SCSI:* most small Unix systems use SCSI disks. This interface can be used for devices other than disks too. It is better than IDE at multitasking. The original SCSI interface was limited to seven devices in total per interface. Wide SCSI can deal with 14 disks. See also the notes in Chapter 2.

To connect a new disk to a Unix host, we have to power down the system. Here is a typical checklist for adding a SCSI disk to a Unix system:

- Power down the computer.
- Connect disk and terminate SCSI chain with proper terminator.
- Set the SCSI ID of the disk so that it does not coincide with any other disks. On Solaris hosts, SCSI ID 6 of controller zero is reserved for a CD-ROM player.
- On Sun machines one can use the ROM command `probe-scsi` (or `probe-scsi-all`), if there are several disk interfaces) to probe the system for disks. This shows which disks are found on the bus. It can be useful for troubleshooting bad connections, or accidentally overlapping disk IDs, etc.
- Partition and label the disk. Update the defect list.
- Edit the `/etc/fstab` file system table, or equivalent to mount the disk. See also the next section.

4.6.1 `mount` and `umount`

To make a disk partition appear as part of the file tree it has to be *mounted.* We say that a particular file system is *mounted on* a directory or *mountpoint.* The command `mount` mounts file systems and disks defined in the file system table file. This is a file which holds data for mount to read.

The file system table has different names on different implementations of Unix.

Solaris 1 (SunOS)	/etc/fstab
Solaris 2	/etc/vfstab
HPUX	/etc/checklist or /etc/fstab
AIX	/etc/file systems
IRIX	/etc/fstab
ULTRIX	/etc/fstab
OSF1	/etc/fstab
GNU/Linux	/etc/fstab

These files also have different syntax on different machines. The eventual standard which most systems comply with (SunOS, HPUX, OSF1) is

```
#
# SunOS 4* / Solaris 1
# /dev/sd0a / 4.2 rw 1 1
/dev/sd0g /usr 4.2 rw 1 2
# NFS
gluino:/site/gluino/pc  /site/gluino/pc  nfs rw,bg,hard,intr  0 0
proton:/var/spool/mail  /var/spool/mail  nfs rw,bg,hard,intr  0 0
proton:/site/proton/u1  /site/proton/u1  nfs rw,bg,hard,intr  0 0
proton:/site/proton/u2  /site/proton/u2  nfs rw,bg,hard,intr  0 0
```

In HPUX:

```
#
# HPUX
#
/dev/dsk/c201d6s0        /                    hfs defaults  0 1
/dev/dsk/c201d5s0        /site/hope/disk      hfs defaults  0 2
proton:/site/proton/fys /site/proton/fys     nfs rw,nosuid  0 0
proton:/usr/spool/mail   /usr/mail            nfs rw,suid  0 0
proton:/site/proton/u1   /site/proton/u1      nfs rw,suid  0 0
proton:/site/proton/u2   /site/proton/u2      nfs rw,suid  0 0
polaron:/site/polaron/u1/site/polaron/u1      nfs rw,suid  0 0
```

The syntax of the command is

```
mount file system directory type (options)
```

There are two main types of file system–a disk file system (called ufs, hfs, etc.) (which means a physical disk) and the *NFS* network file system. If we mount a 4.2 file system it means that it is, by definition, a local disk on our system and is described by some logical device name like `/dev/something`. If we mount an NFS file system, we must specify the name of the file system and the name of the host to which the physical disk is attached.

Here are some examples, using the SunOS file system list above:

```
mount -a                 # mount all in fstab
mount -at nfs            # mount all in fstab which are type nfs
mount -at 4.2            # mount all in fstab which are type 4.2
mount /var/spool/mail    # mount only this fs with options given
                         # in fstab
```

(The `-t` option does not work on all Unix implementations.) Of course, we can type the commands manually too, if there is no entry in the file system table. For example, to mount an nfs file system on machine 'wigner' called `/site/wigner/local` so that it appears in our file system at `/mounted/wigner`, we would write

```
mount wigner:/site/wigner/local /mounted/wigner
```

The directory `/mounted/wigner` must exist for this to work. If it contains files, then these files will no longer be visible when the file system is mounted on top of it, but they are not destroyed. Indeed, if we then unmount using

```
umount /mounted/wigner
```

(the spelling `umount` is correct) then the files will reappear again. Some implementations of NFS allow file systems to be merged at the same mount point, so that the user sees a mixture of all the file systems mounted at the same point.

4.6.2 Disk Partition Device Names

The convention for naming disk devices in BSD and system 5 Unix differs. Let us take SCSI disks as an example. Under BSD, the SCSI disks have names according to the following scheme:

`/dev/sd0a`	First partition of disk 0 of the standard disk controller. This is normally the root file system `/`.
`/dev/sd0b`	Second partition of disk 0 on the standard disk controller. This is normally used for the swap area.
`/dev/sd1c`	Third partition of disk 1 on the standard disk controller. This partition is usually reserved to span the entire disk, as a reminder of how large the disk is.

System 5 Unix employs a more complex, but also more general, naming scheme. Here is an example from Solaris 2:

`/dev/dsk/c0t3d0s0`	Disk controller 0, target (disk) 3, device 0, segment (partition) 0
`/dev/dsk/c1t1d0s4`	Disk controller 1, target (disk) 1, device 0, segment (partition) 4

Not all systems distinguish between target and device. On many systems you will find only `t` or `d`, but not both.

4.7 Kernel Customization

The operating system kernel is that most important part of the system, it drives the hardware of the machine and shares it between multiple processes. If the kernel does not work well, the system as a whole will not work well. The main reason for making changes to the kernel is to fix bugs and to upgrade system software; performance gains can also be achieved, however, if one is patient. We shall return to the issue of performance again in section 7.7. Kernel configuration varies widely between operating systems. Some systems require kernel modification for every miniscule change, while others live quite happily with the same kernel unless major changes are made to the hardware of the host.

Until quite recently, most operating system kernels were statically compiled programs which were specially built for each host. Even today, many of these operating systems still exist and flourish, but static programs are inflexible and the current trend is to replace them with software configurable systems which can be manipulated without the need to recompile the kernel. System V Unix has blazed the trail of adaptable, configurable kernels, in its quest to build an operating system which will scale from laptops to mainframes. It introduces

kernel modules which can be loaded on demand. By loading parts of the kernel only when required, one reduces the size of the resident kernel memory image, which can save memory. This policy also makes upgrades of the different modules independent of the main kernel software, which makes patching and reconfiguration simpler. SVR4 Unix and its derivatives, like Solaris and Unixware, are testimony to the flexibility of SVR4.

NT has also taken a modular view to kernel design. Configuration of the NT kernel also does not require a re-compilation, only the choice of a number of parameters, accessed through the system editor in the Performance Monitor, followed by a reboot. GNU/Linux switched from a static, monolithic kernel to a modular design quite quickly. The Linux kernel strikes a balance between static compilation and modular loading. This balances the convenience of modules with the increased speed of having statically compiled code forever in memory. Typically, heavily used kernel modules are compiled in statically, while infrequently used modules are accessed on demand.

Solaris

Neither Solaris nor NT require or permit kernel re-compilation in order to make changes. In Solaris, for instance, one edits configuration files and reboots for an auto-reconfiguration. First we edit the file `/etc/system` to change kernel parameters, then reboot with the command

```
reboot -- -r
```

which reconfigures the system automatically. There is also a large number of system parameters which can be configured on the fly (at run time) using the `ndd` command.

GNU/Linux

The Linux kernel is subject to more frequent revision than many other systems, owing to the pace of its development. It must be recompiled when new changes are to be included, or when an optimized kernel is required. Many GNU/Linux distributions are distributed with older kernels, while newer kernels offer significant performance gains, particularly in kernel intensive applications like NFS, so there is a practical reason to upgrade the kernel.

The compilation of a new kernel is a straightforward, if not time consuming, process. The standard published procedure for installing and configuring a new kernel is this. New kernel distributions are obtained from any mirror of the Linux kernel site [148]. First we back up the old kernel, unpack the kernel sources into the operating system's files (see the note below) and alias the kernel revision to `/usr/src/linux`. Note that the `bash` shell is required for kernel compilation:

```
$ cp /boot/vmlinuz /boot/vmlinux.old
$ cd /usr/src
$ tar zxf /local/site/src/linux-2.2.9.tar.gz
$ ln -s linux-2.2.9 linux
```

There are often patches to be collected and applied to the sources. For each patch file:

```
$ zcat /local/site/src/patchX.gz | patch -p0
```

Then we make sure that we are building for the correct architecture (Linux now runs on several types of processor):

```
$ cd /usr/include
$ rm -rf asm linux scsi
$ ln -s /usr/src/linux/include/asm-i386 asm
$ ln -s /usr/src/linux/include/linux linux
$ In -s /usr/src/linux/include/scsi scsi
```

Next we prepare the configuration:

```
$ cd /usr/src/linux
$ make mrproper
```

The command `make config` can now be used to set kernel parameters. A more user friendly windows-based program `make xconfig` is also available, though this does require one to run X11 applications as root, which is a potential security *faux pas*. The customization procedure has defaults which one can fall back on. The choices are `Y` to include an option statically in the kernel, `N` to not include and `M` to include as module support. The capitalized option indicates the default. Although there are defaults, it is important to think carefully about the kind of hardware we are using. For instance, is SCSI support required? One of the questions prompts us to specify the type of processor, for optimization:

```
Processor type (386, 486, Pentium, PPro) [386]
```

The default, in square brackets, is for generic 386, but Pentium machines will benefit from optimizations if we choose correctly. If we are compiling on hosts without CD-ROMs and tape drives, there is no need to include support for these, unless we plan to copy this compiled kernel to other hosts which *do* have these.

After completing the long configuration sequence, we build the kernel:

```
$ make dep
$ makeclean
$ make zImage
```

and move it into place:

```
$ mv arch/i386/boot/zImage /boot/vmlinuz-2.2.9
$ ln -s /boot/vmlinuz-2.2.9 /boot/vmlinuz
```

The last step allows us to keep track of which version is running, while still having the standard kernel name:

```
lilo
```

After copying a kernel loader into place, we have to update the boot blocks on the system disk so that a boot program can be located before there is an operating kernel which can interpret the file system. This applies to any operating system, e.g. SunOS has the `installboot` program. After installing a new kernel in GNU/Linux, we update the boot records on the system disk by running the `lilo` program. The new loader program is called by simply typing `lilo`. This reads a default configuration file `/etc/lilo.conf` and writes loader data to the Master Boot Record (MBR). One can also write to the primary Linux partition, in case something should go wrong:

```
lilo -b /dev/hda1
```

so that we can still boot, even if another operating system should destroy the boot block.

Logistics of Kernel Customization

The standard procedure for installing a new kernel breaks a basic principle: don't mess with the operating system distribution, as this will just be overwritten by later upgrades. It also potentially breaks the principle of reproducibility: the choices and parameters which we choose for one host do not necessarily apply for others. It seems as though kernel configuration is doomed to lead us down the slippery path of making irreproducible, manual changes to every host.

We should always bear in mind that what we do for one host must usually be repeated for many others. If if were necessary to recompile and configure a new kernel on every host individually, it would simply never happen. It would be a project for eternity.

The situation with a kernel is not as bad as it seems, however. Although, in the case of GNU/Linux, we collect kernel upgrades from the net as though it were third-party software, it is rightfully a part of the operating system. The kernel is maintained by the same source as the kernel in the distribution, i.e. we are not in danger of losing anything more serious than a configuration file if we upgrade later. However, reproducibility across hosts is a more serious concern. We do not want to repeat the job of kernel compilation on every single host. Ideally, we would like to compile once and then distribute to similar hosts. Kernels can be compiled, cloned and distributed to different hosts provided they have a common hardware base (this comes back to the principle of uniformity). Life is made easier if we can standardize kernels; to do this we must first have standardized hardware. The modular design of newer kernels means that we also need to upgrade the modules in `/lib/modules` to the receiving hosts. This is a logistic problem which requires some experimentation in order to find a viable solution for a local site.

These days it is not usually *necessary* to build custom kernels. The default kernels supplied with most OSes are good enough for most purposes. Performance enhancements are obtainable, however, particularly on busy servers. See section 7.7 for more hints.

Exercises

Exercise 4.1 If you have a PC to spare, install a GNU/Linux distribution, e.g. Debian, or a commercial distribution. Consider carefully how you will partition the disk. Can you imagine repeating this procedure for 100 hosts.

Exercise 4.2 Install NT. You will probably want to repeat the procedure several times to learn the pitfalls. Consider carefully how you will partition the disk. Can you imagine repeating this procedure for 100 hosts.

Exercise 4.3 If space permits, install GNU/Linux and NT together on the same host. Think carefully, once again, about partitioning.

Exercise 4.4 For both of the above installations, design a directory layout for local files. Discuss how you will separate operating system files from locally installed files. What will be

the effect of upgrading or reinstalling the operating system at a later time? How does paritioning of the disk help here?

Exercise 4.5 Imagine the situation in which you install every independent software package in a directory of its own. Write a script which builds and updates the PATH variable for users automatically, so that the software will be accessible from a command shell.

Exercise 4.6 Describe what is meant by a URL or universal naming scheme for files. Consider the location of software within a directory tree: some software packages compile the names of important files into software binaries. Explain why the use of a universal naming scheme guarantees that that the software will always be able to find the files even when mounted on a different host, and conversely, why cross mounting a directory under a different name on a different host is doomed to break the software.

Exercise 4.7 Upgrade the kernel on your GNU/Linux installation. Collect the kernel from ref. [148].

Exercise 4.8 Determine your Unix/NT current patch level. Search the web for mor recent patches. Which do you need? Is it always right to patch a system?

Chapter 5

User Management

Computer systems exist for their users. The system manager's predilection for the machinery of computing systems, although a natural side-effect of a deep fascination with complex machinery, is secondary to the needs of users. It draws the technically inclined and the scientifically curious into the field, and provides us with the motivation to keep computing systems alive and well, but when it comes down to it, the computer is a tool for its users. Besides, without users, there would be few challenges in system administration. Users are both the reason that computers exist and their greatest threat.

The role of the computer as a tool has changed extensively throughout its history. From John Von Neumann's vision of the computer as a device for predicting the weather, to a calculator for atomic weapons, to a desktop typewriter, to a means of global communication, computers have changed the world and have reinvented themselves in the process. System administrators need to cater to all needs, and ensure the stability and security of the system as as whole.

5.1 User Registration

One of the first issues on a new host is to issue accounts for users. Surprisingly, this is an area where operating system designers provide virtually no help. The tools provided by operating systems for this task are, at best, primitive and are rarely suitable for the task without considerable modification. For small organizations, user registration is a relatively simple matter. Users can be registered at a centralized location by the system manager, and made available to all of the hosts in the network by some sharing mechanism, such as a login server, distributed authentication service or by direct copying of the data. There are various mechanisms for doing this, and we shall return to them below.

For larger organizations, with many departments, user registration is much more complicated. The need for centralization is often in conflict with the need for delegation of responsibility. It is convenient for autonomous departments to be able to register their own users, but it is also important for all users to be registered under the umbrella of the organization, to ensure unique identities for the users and flexibility of access to different parts of the organization. What is needed is a solution which allows local system managers to be able to register new users in a global user database. User account administration has been discussed many times [2, 258, 56, 161, 136, 165, 51, 182]. The special problems of each institution and work environment are reflected in these works.

PC server systems like NT and Netware have an apparent advantage in this respect. By forcing a particular administration model onto the hosts in a network, they can provide straightforward delegation of user registration to anyone with domain credentials. Registration of single users under NT can be performed remotely from a workstation, using the

```
net user username password /ADD /domain
```

command. While most Unix-like systems do not provide such a ready-made tool, many solutions have been created by third parties. The price one pays for such convenience is an implicit *trust relationship* between the hosts. Assigning new user accounts is a security issue, thus to grant the right of a remote user to add new accounts requires us to trust the user with access to that facility.

It is rather sad that no acceptable, standardized user registration methods have been widely adopted. This must be regarded as one of the unsolved problems of system administration. Part of the problem is that the requirements of each organization are rather different. Many Unix-like systems provide shell scripts or user interfaces for installing new users, but most of these scripts are useless because they follow a model of system layout which is woefully inadequate for a network environment. Also, they have different ideas about how the system should work than the organization with specialized requirements.

5.1.1 Local and Network Accounts

Most organizations need a system for centralizing passwords, so that each user will have the same password on each host on the network. In fixed model computing environments such as NT or Novell Netware, where a login or domain server is used, this is a simple matter. In larger organizations with many departments or sub-domains it is more difficult [59, 264, 273].

Both Unix and NT support the creation of accounts locally on a single host, or 'globally' within a network domain. With a local account, a user has permission to use only the local host. With a network account, the user can use any host which belongs to a network *domain*. Local accounts are configured on the local host itself. Unix registers local users by added them to the files `/etc/passwd` and `/etc/shadow`. In NT the Security Accounts Manager (SAM) is used to add local accounts to a given workstation.

For network accounts, Unix-like systems have widely adopted Sun Microsystems' Network Information Service (NIS), formerly called Yellow Pages, or simply YP. The NIS-plus service was later introduced to address a number of weaknesses in NIS, but this has not been widely adopted. NIS is reasonably effective at sharing passwords, but it has security implications: encrypted passwords are distributed in the old password format, clearly visible, making a mockery of shadow password files. NIS users have to be registered locally as users on the master NIS server; there is no provision for remote registration, or for delegation of responsibility. Variations on the NIS theme have been discussed in refs. [53, 213, 118]. NT uses its model of domain servers, rather like a NIS, but including a registration mechanism. A user in the SAM of a primary domain, controller is registered within that domain, and has an account on any host which subscribes to that domain. An approach to user accounts based on SQL databases was discussed in ref. [14].

An NT domain server involves not only shared databases, but also shared administrative policies and shared security models. A host can subscribe to one or more domains, and one domain can be associated with one another by a trust relationship. When one NT domain

'trusts' another, then accounts and groups defined in the *trusted* domain can be used in the *trusting* domain. NIS is indiscriminating in this respect. It is purely an authentication mechanism, implying no side-effects by the login procedure.

Another model of network computing is the Open Software Foundation's Distributed Computing Environment (DCE). This is a distributed user environment which can be used to provide a seamless worldwide distributed network domain. The DCE has been ported to both Unix and NT, and requires a special login authentication after normal login to Unix/NT.

To summarize, rationalized user registration is a virtually unsupported problem in most operating systems. The needs of different organizations are varied, and no successful solution to the problem has been devised and subsequently adopted as a standard. Networks are so common now that we have to think of the network first. Whether it happens today or tomorrow, at any given site, users will be moving around from host to host. They will need access to system resources wherever they are. It follows that they need distributed accounts. In creating a local solution, we have to bear in mind some basic constraints.

Principle 18 (Distributed accounts) *Users move around from host to host, share data and collaborate. They need easy access to data and workstations all over an organization.*

Standardizing user names across all platforms simplifies both the logistics of user management and opens for cross-platform compatibility. User names longer than eight characters can cause problems with Unix-like systems and FTP services. Users normally expect to be able to use the same password to log onto any host and have access to the same data, except for hosts with special purposes.

Suggestion 4 (Passwords) *Give users a common username on all hosts, of no more than eight characters. Give them a common password on all hosts, unless there is a special reason not to do so. Some users never change their passwords unless forced to, and some users never even log in, so it is important to assign good passwords initially. Never assign a simple password and assume that it will be changed.*

Perl scripts are excellent ways of making user installation scripts which are tailored to local needs. See ref. [177] for an excellent discussion of this on NT. Interactive programs are almost useless since users are seldom installed one by one. At universities, hundreds of students are registered at the same time. No system administrator would type in all the names by hand. More likely they would be input from some administrative list generated by the admissions department. The format of that list is not a universal standard, so no off-the-shelf software package is going to help here.

Sites which run special environments, such as the Andrew File System (AFS), the Distributed Computing Environment (DCE), Athena or Kerberos, often require extra authentication servers and registration procedures [60, 273].

5.1.2 Unix Accounts

To add a new user to a Unix-like host we have to

- Find a unique user name, user-id (uid) number and password for the new user.

- Update the system database of user accounts, e.g. add a line to the file `/etc/passwd` for Unix (or on the centralized password server of a network) for the new user.
- Create a login directory (home directory) for the user.
- Choose a shell for the user (if appropriate).
- Copy some configuration files like `.cshrc` or `.profile` into the new user's directory, or update the system registry.

Because every site is different, user registration requires different tools and techniques in almost every case. For example: where should users' home directories be located? GNU/Linux has an `adduser` script which assumes that the user will be installed on the local machine under `/home/user`, but many users belong to a network and their disk space lies physically on a different host which is mounted by NFS.

Just to make life difficult, Unix developers have created three different password file formats which increase the awkwardness of distributing passwords. The traditional password file format has the following format:

```
mark:Ax7Wc1Kd8ujo2:123:456:Mark Burgess:/home/mark:/bin/tcsh
```

The first field is a unique user name (up to eight characters) for the user. The second is an encrypted form of the user's password; then comes the user-id (a unique number which represents the user and is equivalent to the user name) and the default group-id (a unique number which represents the default group of users to which this user belongs). The fifth column is the so-called GECOS field, which is usually just the full name of the user. On some systems, comma separated entries may be given for full name, office, extension and home phone number. The sixth is the home directory for the user (the root directory for the user's private virtual machine). Finally, the seventh field is the user's default shell. This is the command interpreter which is started when the user logs in.

Newer Unix-like systems make use of *shadow password files*, which conceal the encrypted form of the password for ordinary users. The format of the password file is then the same as above, except that the second password field contains only an 'x', e.g.

```
mark:x:123:456:Mark Burgess:/home/mark:/bin/tcsh
```

There is then a corresponding line in `/etc/shadow` with the form

```
mark:Ax7Wc1Kd8ujo2:6445::::::
```

The shadow file is not readable by ordinary users. It contains many blank field which are reserved for the special purpose of password aging and other expiry mechanisms. See the manual page for 'shadow' for a description of the fields. The only number present by default is the time at which the password was last changed, measured in the number of days since Jan. 1. 1970.

The third form of password file is used by the BSD 4.4 derived operating systems:

```
mark:Ax7Wc1Kd8ujo2:3232:25::0:0:Mark Burgess:/home/mark:/bin/tcsh
```

It has extra fields which are not normally used. These systems also have an optimization: in addition to the master password file base, they have a compiled binary database for rapid lookup. Administrators edit the file `/etc/master.password` and then run the

pwd_mkdb command to compile the database which is actually used for lookups. This generates text and binary versions of the password database.

5.1.3 NT Accounts

Single NT accounts are added with the command

```
net user username password /ADD /domain
```

or using the GUI. NT does not provide any assistance for mass registration of users. The additional Resource Kit package contains tools which allow lists of users to also be registered from a standard file format, with `addusers.exe`, but only at additional cost.

NT users begin in the root directory by default. It is customary to create a `\users` directory for home directories. Network users usually have their home directory on the domain server mapped to the drive `H:`. Needless to say, there is only a single choice of shell (command interpreter) for NT, so this is not specified in the user registration procedure.

5.1.4 Groups of Users

Both Unix and NT allow users to belong to multiple groups. A group is an association of user names which can be referred to collectively by a single name. File and process permissions can be granted to a group of users. Groups are defined statically by the system administrator.

On Unix-like systems they are defined in the `/etc/group` file, like this:

```
users::100:user1,mark,user2,user3
```

The name of the group, is this case, is `users`, with group-id 100 and members user1, mark, user2 and user3. The second, empty field provides space for a password, but this facility is seldom used. A number of default groups are defined by the system, for instance

```
root::0:root
other::1:
bin::2:root,bin,daemon
```

The names and numbers of system groups vary with different flavours of Unix. The root group has superuser privileges.

Unix groups can be created for users or for software which runs under a special user-id. In addition to the names listed in the group file, a group also accrues users from the default group membership in field four of `/etc/passwd`. Thus, if the group file had the groups:

```
users::100:
msql::36:
ftp::99:
www::500:www
www-data::501:www,toreo,mark,geirs,sigmunds,mysql
privwww::502:
```

and every user in `/etc/passwd` had the default group 100, then the `users` group would still contain every registered user on the system. By way of contrast, the group www contains no members at all, and is to be used only by a process which the system assigns that group

identity, whereas `www-data` contains a specific named list and no others, as long as all users have the default group 100.

NT also allows the creation of groups. Groups are created by command, rather than by file editing, using

```
net group groupname /ADD
```

Users may then be added with the syntax

```
net group groupname username1 username2... /ADD
```

They can also be edited with the GUI on a local host. NT distinguishes global groups (consisting only of domain registered users) from local groups, which may also contain locally registered users. Some standard groups are defined by the system, e.g.

```
Administrators
Users
Guest
```

The Administrators group has privileged access to the system.

5.2 Account Policy

Most organizations need a strict policy for assigning accounts and opening the system for users. Users *are* the foremost danger to a computing system, so the responsibility of owning an account should not be dealt out lightly. There are many ways in which accounts can be abused. Users can misuse accounts for villainous purposes, and they can abuse the terms on which the account was issued, wasting resources on personal endeavours. For example, in Norway, where education is essentially free, students have been known to undergo semester registration simply to have an account, giving them essentially free access to the Internet and a place to host their web sites.

Any account policy should contain a clause about weak passwords. If weak passwords are discovered, it must be understood by users that their account can be closed immediately. Users need to understand that this is a necessary security initiative. Closing Unix accounts can be achieved simply by changing their default shell in `/etc/passwd` with a script such as

```
#!/bin/sh

echo "/local/bin/blocked.passwd was run" | mail sysadm
/usr/bin/last -10 | mail sysadm

message='
You account has been closed because your password was found to
be vulnerable to attack. To reopen your account, visit the
admin office, carrying some form of personal identification.
'
echo "$message"

sleep 10
exit 0
```

Although this does not prevent them from doing simple things on a X-windows console, it does prevent them from logging in remotely, and it gets their attention. A more secure method is to simply replace their encrypted password with NP or *, which prevents them from being authenticated.

It is occasionally tempting to create guest accounts for visitors and transient users. NT has a ready-made guest account, which is not disabled by default on some versions of NT. Guest accounts are a bad idea, because they can be used long after a visitor has gone, they usually have weak or non-existent passwords and therefore are an open invitation to attack the system. Shared accounts are also a bad idea, since they are inherently more fragile from a security perspective, though the use of shared Unix accounts, in which users could not log in as a shared user, are described in ref. [26]. This is similar to the ability in Unix to set a password on a group.

5.3 Login Environment

When a new user logs in for the first time, he or she expects the new account to work straight away. Printing should work, programs should work and there should be no strange error messages about files not being found or programs not existing. Most users want to start up a window environment. If users will be able to log on to many different kinds of operating system, we have to balance the desire to make systems look alike, with the need to distinguish between different environments. Users need to understand the nature of their work environment at all times in order to avoid hapless errors. The creation of default login environments has been discussed in refs. [247, 255, 280], though this is now somewhat out of date.

5.3.1 Unix Environment

Unix and its descendents have always been about the ability to customize. Everything in Unix is configurable, and advanced users like to play around; many create their own setups, but many users simply want basics. The use of multitudinous 'dot' files for setting defaults in Unix has led to its being criticized for a lack of user friendliness. Various attempts have been made to provide interfaces which simplify the task of editing these configuration files [97, 73], though the real problem is not so much the fact one has to edit files, as the fact that every file has its own syntax. A system administrator has to ensure that everything works properly with acceptable defaults, right from the start. Here is a simple checklist for configuring a user environment. Gradually, the appearance of newer and better user interfaces like KDE and GNOME is removing the need for users to edit their own window configuration files.

- `.cshrc` If the default shell for users is a C shell or derivative, then we need to supply a default 'read commands' file for this shell. This should set a path which searches for commands, a terminal type and any environment variables which a local system requires.
- `.profile` If the default shell is a Bourne-again shell like `bash` or `ksh`, then we need to supply this file to set a PATH variable, terminal type and environment variables which the system requires.

- `.xsession` This file is read by the Unix `xdm` login service. It specifies what windows and what window manager will be used when the X-windows system is started. The file is a shell script which should begin by setting up applications in the background (with a & symbol after them) and end up `exec`-ing a window manager in the foreground. If the window manager is called as a background process, the script will be able to exit immediately and users will be logged out immediately. Some systems use the file called `.xinitrc`, though this file is officially obsolete. The official way to start the X11 window system is through the `xdm` program, which provides a login prompt window. GNU/Linux seems to have revived the use of the obsolete command `startx` which starts the X windows system from a tty-shell. The older `startx` system used the `.xinitrc` file, whereas `xdm` uses `.xsession`. Most GNU/Linuxes hack this so that one only needs a `.xsession` file.
- `.mwmrc` This file configures the default menus, etc., for the `mwm` window manager.
- `.fvmwrc` This file customizes the behaviour of the `fvwm` window manager.
- `.fvwm2rc` This file customizes the behaviour of the `fvwm2` window manager.
- `.fvwm95rc` This file customizes the behaviour of the `fvwm95` window manager. This is a mock windows-95 interface.

A shell setup should define a terminal type, a default prompt and appropriate environment variables, especially a command path.

Principle 19 (Environment) *It should always be clear to users which host they are using and what operating system they are working with. Default environments should be kept simple both in appearance (prompts, etc.,) and in functionality (specially programmed keys, etc.). Simple environments are easy to understand.*

We need to aim a default environment at an average user and ensure that basic operating system functions work unambiguously. The visual clarity of a work environment is particularly important. In a windowing environment this is usually not a problem. Command shells require some extra thought, however. A command shell can, in principle, be opened on any Unix-like host. A user with many windows open, each with a shell running on a different host, could easily become confused. Suppose we wish to copy a newer version of a file on one host to a repository on another host. If we mix the hosts up, we could risk wiping the new version with an old version, instead of the other way around.

Suggestion 5 (Clear prompts) *Try to give users a command prompt which includes the name of the host they are working on. This is important, since different hosts might have different operating systems, or different files. Including the current directory in the prompt, like DOS, is not always a good idea. It uses up half the width of the terminal and can seem confusing. If users want the name of the current directory in the prompt, let them choose that. Don't assign it as a default.*

Some systems offer global shell configuration files which are read for every user. These files are usually located in `/etc` or `/etc/default`. The idea of a global default file has attractive features in principle, but it is problematic in practice. The problem has to do with the separation of local modifications from the operating system, and also the standardi-

zation of defaults across all hosts. These files could be distributed from a central source to every host, but a better approach is to simply place an equivalent defaults file on the same distributed file systems which contain users' home directories. This is easily achieved by simply ignoring the global defaults, and giving every user a default shell configuration file which reads a site-dependent file instead.

Suggestion 6 (Unix shell defaults) *Avoid the host-wide files for shell setup in* `/etc`. *They are mixed up in the operating system distribution, and changes here will be lost at upgrade time. Use an over-ridable include strategy in the user's own shell setup to read in global defaults. Do not link a file on a different file system to these in case this causes problems during system boot-up.*

Here is an example configuration file for the C shell, which would be installed for all users in their home directories:

```
#
# cshrc file (for tcsh)
#

source ../.setupfiles/cshrc-global

#
# Place own definitions below
#

alias f finger
alias ed emacs
```

Note that we use the `source` directive to read in a file of global C-shell definitions which we have copied into place from a central repository for all important system master files. Notice also that, by copying this onto the same file system as the home directory itself (the directory over the user's home directory, see Figure 5.1), we make sure that the file is always NFS exported to all hosts together with the home directory. This allows us to change the global setup for everyone at the same time, or separately for different classes of user on different partitions. For each separate home partition, we could have a different set of defaults. This is probably not recommended, however, unless users are distinguished in some important way.

One of the functions of a local shell configuration is to set up a command path and a library path for software. Since the command path is searched in order, we can override

```
/home1

.
..
lost+found
.setupfiles - cshrc-global
user1
user2
mark - .cshrc
user4
```

Figure 5.1 File system structure on a home directory file system

operating system commands with local solutions simply by placing site-dependent binaries at the start of the path. GNU file utilities and binary utilities can also be placed ahead of operating system standard utilities. They are often more standard and more functional.

5.3.2 Example Shell Configuration

Here is an example shell configuration for the `tcsh`.

```
#!/bin/csh -f
##############################################################
#
# C Shell startup file
# System Wide Version.
#
##############################################################

umask 077 # default privacy on new files
setenv HOSTTYPE 'uname'

##############################################################

switch (\$HOSTTYPE)

  case SunOS:
      set path = (                        \
                 /local/site/bin          \
                 /local/kde/bin           \
                 /local/gnu/bin           \
                 /usr/ccs/bin             \
                 /local/jdk1.1.6/bin      \
                 /local/bin               \
                 /local/qt/bin            \
                 /usr/ucb                 \
                 /bin                     \
                 /usr/bin                 \
                 /usr/openwin/bin         \
                 .                        \
                 )
               breaksw

  case Linux:
    set path =   (                        \
                 /local/site/bin          \
                 /local/bin               \
                 /local/jdk1.1.6/bin      \
                 /local/bin/X11           \
                 /local/qt/bin            \
                 /local/kde/bin           \
                 /local/gnu/bin           \
                 /local/bin/X11           \
                 /usr/bin/X11             \
                 /usr/bin                 \
                 /bin                     \
                 .                        \
```

```
                )
              breaksw
endsw

#############################################################
#
# set TERM for "at" batches in non-interactive shells
# tcsh wants to write something to stdout, but I
# can't see what => term has to be set even though its
# irrelevant )
#
if (! $?TERM) setenv TERM vt100;
if (! $?term) set term = vt100;
if (! $?prompt) exit 0;

#
# End for non-interactive shells (batch etc.)
#

setenv TERM vt100 # Many shell types do not work
set term = $TERM # This is a safe default, omit it if you dare

#############################################################
# set
#############################################################

set history=100 savehist=100
set prompt = "'hostname'% "
set prompt2 = "%m %h> "
set fignore = (.o \~ .BAK .out \%)

#############################################################
# Common Environment
#############################################################

setenv EDITOR       emacs
setenv ESHELL       tcsh
setenv NNTPSERVER   nntp-server.domain.country
setenv QTDIR        /usr/local/qt
setenv CLASSPATH    /usr/local/jdk1.1.6/lib/classes.zip:.
setenv JAVA_HOME    /usr/local/jdk1.1.6
setenv MYSQL        /usr/local/bin/mysql

#############################################################
# platform specific environment (overrides common)
#############################################################

switch ($HOSTTYPE)

    ###############

    case SunOS*:
    case solaris:
      setenv LD_LIBRARY_PATH /usr/openwin/lib:/local/lib/X11:\
```

```
/local/gnu/lib:/usr/dt/lib:/local/qt/lib:/local/lib:
      setenv LPATH /usr/lib:/local/lib:

      if ( $?DISPLAY || $TERM == "sun" ) then
        setenv MOTIFHOME /usr/dt
        setenv X11HOME /usr/openwin
        setenv FONTPATH /usr/openwin/lib/X11/fonts:\
/usr/openwin/lib/locale/iso_8859_5/X11/fonts:\
/usr/openwin/share/src/fonts:/usr/openwin/lib/X11/fonts:\
/local/sdt/sdt/fonts/SDT3/X11
        setenv OPENWINHOME /usr/openwin
        setenv XKEYSYMDB /local/site/X11/XKeysymDB
        setenv XAPPLRESDIR /usr/openwin/lib/X11/app-
          defaults
        setenv GS_FONTPATH /local/share/ghostscript/fonts
        setenv GS_LIB_PATH /local/share/ghostscript/4.03
        endif
        setenv MANPATH /local/gnu/man:/usr/man:/local/man:\
/usr/openwin/share/man
        limit coredumpsize 0
      breaksw
  ###############
  case Linux:
  case i486:
      setenv MANPATH /local/man:/local/site/man:/local/
        man:\
/usr/man:/usr/man:/usr/man/preformat:/usr/X11/man
      setenv XAPPLRESDIR /local/site/X11/app-defaults:\
/var/X11R6/lib/app-defaults
      stty erase '^?' intr '^C' kill '^U' susp '^Z'
      setenv LD_LIBRARY_PATH /usr/X11R6/lib:/local/lib:\
/local/qt/lib:/local/kde/lib
      setenv XNLSPATH /usr/X11R6/lib/X11/nls
    breaksw
endsw

############################################################
# aliases
############################################################
alias del       'rm -i '
alias dir       'ls -lg \!* | less -E'
alias .         'echo $cwd'
alias f finger
alias h history
alias go a.out
alias cd.. cd ..
```

```
alias grant setfacl
alias cacls getfacl
alias rlogin ssh
alias rsh ssh

##############################################################
#
# Check message of the day
#
##############################################################

# Not always necessary

if ( -f /etc/motd ) then
  /bin/cat /etc/motd
endif

##############################################################
#
# Check whether user has a vacation file
#
##############################################################

if ( -f ~/.forward ) then

  if ( "`grep vacation ~/.forward`" != "" ) then

    echo '**********************************************'
    echo '          YOU ARE RUNNING THE vacation SERVICE      '
    echo '            RUN vacation AGAIN TO CANCEL IT !        '
    echo '**********************************************'

  endif

endif
```

5.3.3 The Superuser's Environment

What kind of user environment should the superuser have? As we know, a privileged account has potentially dangerous consequences for the system. From this account, we have the power to destroy the system, or sabotage it. In short, the superuser's account should be configured to avoid as many casual mistakes as possible.

There is no harm in giving Unix's root account an intelligent shell like `tcsh` or `bash` provided that shell is physically stored on the root partition. When a Unix system boots, only the root partition is mounted. If we reference a shell which is not available, we can render the host unbootable.

The superuser's `PATH` variable should not include '`.`', i.e. the current directory. To give root easy access to commands which are left lying around in directories is to open the system to attack. For instance, suppose an ordinary user left a file called `ls` in the `/tmp` directory, and suppose the root account had the path

```
setenv PATH .:/bin:/usr/bin
```

If the superuser does the following

```
host# cd /tmp
host# ls
```

then because the path search looks in the current directory first, it would find and execute the program which had been left by the user. That program gets executed with root privileges and could be used to give the user concerned permanent privileged access to the system, for instance by installing a special account for the user which has root privileges. It should be clear that this is a security hazard.

A common mistake which is frequently perpetrated by inexperienced administrators, and which is actually encouraged by some operating systems, is to run X11 applications with root privileges. Root should never run X11 or any other complex applications. There are just too many uncertainties involved. There are so many applications for X11 which come from different sources. There could be a Trojan horse in any one of them. If possible, root should only use a few trusted application programs.

5.4 User Support Services

All users require help at some time or another. The fact that normal users are not *privileged users* means that they occasionally need to ask a superuser to clean up a mess, or fix a problem which is beyond their control. If we are to distinguish between privileged and non-privileged users, we cannot deny users this service.

The amount of support which one offers users is a matter of policy. System administrator's time is usually in short supply, though increased automation is steadily freeing us to concentrate on higher level problems, like support. The ability to support a system depends upon its size in relation to the available resource personell. Supporting hardware and software involves fixing errors, upgrading and perhaps providing tuition or telephone help-desks. E-mail help desks such as Rust, Gnats, Nearnet, Netlog, PTS, QueueMH can assist in the organization of support services, but they are mainly task tracking tools. Sometimes hosts and software packages are labelled *unsupported* in order to emphasize to users that they are on their own if they insist on using those facilities.

One of the challenges system administrators sometimes have to endure on coming to a new site, where chaos reigns, is the transition from anarchy to a smaller set of supported platforms and software. See, for instance, refs. [188, 153]. This can be a tough problem, since users always prefer freedom to restriction. Support services need to be carefully considered and tailored to each local environment.

A recent development in user assistance is the Virtual Network Computing model from AT&T [19]. This is a way to allow a remote user duplicate access to a graphical user interface. Thus, an administrator can log onto an existing user session and have dual controls, allowing users to be nurse-maided through difficulties on-line.

5.5 Controlling User Resources

Every system has a mixture of passive and active users. *Passive users* use the system, quietly accepting the choices which have been made for them. They are often minimal users, who place no great demands on the system. They do not follow the progress of the system with

any zeal and they are often not even aware of what files they have. They seldom make demands other than when things go wrong. *Active users*, on the other hand, follow every detail of what happens. They find every error in the system and contact system administrators frequently demanding upgrades of their favourite programs.

System administrators have a responsibility to find a balance which addresses both users' needs but which keeps the system stable and functional. If we upgrade software too often, users will be annoyed. New versions of software function differently, and this can hinder people in their work. If we do not upgrade often enough, we can also hinder work by restricting possibilities.

5.5.1 Disk Space

Disks fill up at an alarming rate. Users almost never throw away files unless they have to. If one is lucky enough to have only very experienced and extremely friendly users on the system, then one can try asking them nicely to tidy up their files. Most administrators do not have this luxury, however. Most users never think about the trouble they might cause others by keeping lots of junk around. After all, multi-user systems and network servers are designed to give every user the impression that they have their own private machine. Of course, some users are problematical by nature.

Suggestion 7 (Problem users) *Keep a separate partition for problem users' home directories, so that they only cause trouble for one another, not for more considerate users.*

No matter what we do to fight the fire, users still keep feeding the flames. To keep hosts working it is necessary to remove files, not just add them. Quotas limit the amount of disk space users can have access to, but this does not solve the real problem. The real problem is that in the course of using a computer, many files are created as temporary data but are never deleted afterwards. The solution is to delete them.

- Some files are temporary by definition. For example, the byproducts of compilation, `*.o` files, files which can easily be regenerated from source like TeX `*.dvi` files, cache files in `.netscape/` loaded in by Netscape's browser program, etc.
- Some files can be defined as temporary as a matter of policy. Files which users collect for personal pleasure like `*.mp3`, video formats and pornography.

When a Unix program crashes, the kernel dumps its image to disk in a file called `core`. These files crop up all over the place and have no useful purpose. To most users they are just fluff on the upholstery and should be removed. A lot of free disk space can be claimed by deleting these files. Many users will not delete them themselves, however, because they do not even understand why they are there.

Disk quotas mean that users have a hard limit to the number of bytes they are allowed to use on the disk. They are an example of a more general concept known as system accounting, whereby you can control the resources used by any user, whether they be the number of printed pages sent to the printer or the number of bytes written to the disk. Disk quotas have advantages and disadvantages:

- The advantage is that users really cannot exceed their limits. There is no way around this.

- Disk quotas are very restrictive, and when a user exceeds their limit they do not often understand what has happened. Usually, users do not even get a message unless they are logging in. Quotas also prevent users from creating large temporary files which can be a problem when compiling programs. They carry with them a system overhead, which makes everything run a little slower.

In some environments the idea of deleting a user's files is too horrifying to contemplate. In a company or research laboratory one might want to be extremely careful in such a practice. In other cases, like schools and universities, this is pure necessity. Deciding whether to delete files automatically must be a policy decision. It might be deemed totalitarian to delete files without asking. On the other hand, this is often the only way to ever clear anything up. Many users will be happy if they do not have to think about the problem themselves. A tidy policy, rather than a quota policy, gives users a greater illusion of freedom, which is good for system morale. We must naturally be careful never to delete files which cannot be regenerated or reacquired if necessary. File tidying was first suggested by Zwicky in ref. [288], within a framework of quotas. See also refs [106, 37].

A useful strategy is to delete files one is not sure about only if they have not be *accessed* for a certain period of time, say a week. This allows users to use files freely as long as they need to, but prevents them from keeping the files around for ever. Cfengine can be used to perform this task.

For example, a simple cfengine program would look like:

```
control:

  actionsequence = ( tidy )

  mountpattern = ( /site/host )
  homepattern = ( home? )

  #
  # 3 days minimum, remember weekends
  #

tidy:

     home                 pattern=core   recurse=inf age=0
     home                 pattern=a.out  recurse=inf age=3
     home                 pattern=*%     recurse=inf age=3
     home                 pattern=*      recurse=inf age=3
     home                 pattern=*.o    recurse=inf age=1
     home                 pattern=*.aux  recurse=inf age=3
     home                 pattern=*.mp3  recurse=inf age=14

     home/Desktop/Trash   pattern=*      recurse=inf age=14
     home/.netscape/cache pattern=*      recurse=inf age=0
```

This script iterates automatically over all users' home directories, and recurses into them, deleting files if the time since they were last accessed exceeds the time limits specified.

Care should be always taken in searching for and deleting patterns containing 'core'. Some operating systems keep directories called core, while others have files called `core.h`. As long as the files are plain files with an exact name match, one if usually safe.

5.5.2 Quotas and Limits

Although we should never forget that computers exist for their users, it is also important to understand that users are the greatest threat to the stability of their computers. Two or three generations have now grown up with computers in their homes, but these computers were private machines which were not (until recently) attached to a network. They were not part of an organization with many users – they were used by perhaps one or two family members. In short, users have grown up thinking that what they do with their computers is nobody's business but their own. That is not a good attitude in a network community.

In a shared environment, all users share the same machine resources. If one user is selfish, that affects all of the other users. Given the opportunity, users will consume all of the disk space and all of the memory and CPU cycles somehow, whether through greed or simply through inexperience. Thus it is in the interests of the *user community* to limit the ability of users to spoil things for other users.

One way of protecting operating systems from users and from faulty software is to place quotas on the amount of system resources which they are allowed:

- *Disk quotas:* place fixed limits on the amount of disk space which can be used per user. The advantage of this is that the user cannot use more storage than this limit; the disadvantage is that many software systems need to generate/cache large temporary files (e.g. compilers or web browsers), and a fixed limit means that these systems will fail to work as a user approaches his/her quota.
- *CPU time limit:* some faulty software packages leave processes running which consume valuable CPU cycles to no use. Users of multi-user computer systems occasionally steal CPU time by running huge programs which make the system unusable for others. The C-shell `limit cputime` function can be globally configured to help prevent accidents.
- *Policy decisions:* users collect garbage. To limit the amount of it, one can specify a system policy which includes items of the form: 'Users may not have mp3, wav, mpeg, etc., files on the system for more than one day'. To enforce such a policy, see section 5.5.

Quotas have an unpleasant effect on system morale, since they restrict personal freedom. They should probably only be used as a last resort. There are other ways of controlling the build up of garbage, see section 5.5.

Principle 20 (Freedom) *Quotas, limits and restrictions tend to antagonize users. Users place a high value on personal freedom. Restrictions should be minimized. Workaround solutions which avoid rigid limits are preferable, if possible.*

5.5.3 Killing Old Processes

Processes sometimes do not get terminated when they should. There are several reasons for this. Sometimes users forget to log out, sometimes poorly written terminal software does not properly kill its processes when a user logs out. Sometimes background programs simply crash or go into loops from which they never return. One way to clean up processes in a work environment is to look for user processes which have run for more than a day. (Note that the assumption here is that everyone is supposed to log out each day and then log in again the next day – that is not always the case.) Cfengine can also be used to clean up old

processes. Cfengine's processes commands are used to match processes in the process table (which can be seen by running ps ax on Unix). Here is an example:

```
control:

  actionsequence = ( processes )

processes:

  "Jan|Feb|Mar|Apr|May|Jun|Jul|Aug|Sep|Oct|Nov|Dec"

      signal=kill

      include=tcsh
      include=xterm
      include=netscape
      include=ftp
      include=tkrat
      include=pine
      include=irc
      include=kfm
      include=java
```

This rule works by noticing that, when processes are more than a day old, the date (i.e. the name of the month) appears in the process listing. Thus to find processes which are more than a day old, we only need to search for any line containing the name of a month and then pick out a subset of those which contain the strings 'tcsh', 'xterm', etc. This is an extremely useful way of cleaning up old processes which poor terminal software leave behind, or which forgetful users leave behind when they forget to log out. In the latter case, user security is also protected by this clean up operation.

5.5.4 Moving Users

When disk partitions become full, it is necessary to move users from old partitions to new ones[1]. Moving users is a straightforward operation, but it should be done with some caution. A user who is being moved should not be logged in while the move is taking place, or files could be copied incorrectly. We begin by looking for an appropriate user, perhaps one who has used a particularly large amount of disk space. On Unix-like systems we have all the tools we require:

```
cd /site/host/home-old
du -s *
```

Having chosen a user, with username *user*, we copy the directory to its new location,

```
tar cf - user | (cd /site/host/home-new; tar xpvf - )
```

edit the new location in the password file,

```
emacs /etc/passwd
```

[1] Some systems might be equipped with virtual volume managers which provide the illusion of infinitely large partitions, but not everyone can afford this luxury.

and finally remove the old data:

```
rm -r user
```

Users need to be informed about the move: we have to remember that they might hard-code the names of their home directories in scripts and programs, e.g. CGI-scripts.

5.5.5 Deleting Old Users

Users who leave an organization eventually need to be deleted from the system. For the sake of certainty, it is often advisable to keep old accounts for a time in case the user actually returns, or wishes to transfer data to a new location. Whether or not this is acceptable must be a question of policy. Clearly it would be unacceptable for company secrets to be transferred to a new location. Before deleting a user completely, a backup of the data can be made for safe-keeping. Then we have to remove the following:

- Account entry from the password database.
- Personal files.
- E-mail.
- Removal from groups and lists.

5.6 User Well-being

Because computer systems are communities, populated by real people, there are issues in system administration which are directly connected with users' well-being. Contented users work well and treat the system well; disgruntled users cause trouble for the system and for their neighbours. This is not to say that system administrators are (or should be) responsible for the psychological well-being of all the system's users, but there are some simple precautions which the system staff can observe in order to promote the smooth running of the community. In some countries, an organization might be sued by a user who believed he or she had not been sufficiently looked after.

5.6.1 Health

Frequent computer users are not usually aware of how they can be damaging their own health. Unlike cigarettes, computers do not have a government health warning. Whether or not this is an issue for system administrators is open for discussion, but often the system administrator is the only person who thinks about the users and the hardware they use. Certainly every administrator needs to look after his/her own health and, along the way, it is natural to think of the health of others. Fortunately, it is not difficult to avoid the worst problems.

Eyes should be protected, We only have one pair and they must last our entire lives. Ironically, users who wear glasses (not contact lenses) suffer less from computer usage, because their eyes are partially protected from the radiation from the screen.

A computer screen works by shooting charged electrons at a phosphorescent surface. If one touches the screen one notices that it is charged with static electricity. The effect of this is

to charge dust particles and throw them out into users' faces. This can cause irritation to the eyes over long periods. Solution: wear glasses or obtain an anti-static screen with an Earth-wire which counteracts this problem.

Another major cause of eye strain is through reflection. If there is a light source behind a user, it will reflect in the screen and the eyes will be distracted by the reflection. The image on the screen lies on the screen surface, any reflected images lie behind the screen (as far behind the screen as the source is in front of the screen). This confuses the eyes into focusing back and forth between the reflection and the image. The result is eye-strain. The solution is to (i) eliminate all sharp light sources which can cause reflections, and (ii) obtain an anti-reflective screen cover. This can be combined with an anti-static screen, and it is probably the best investment a user can make.

Prolonged eye strain can lead to problems reading and focusing. It can lead to headaches and neck ache from squinting:

- *Back.* The back (spine) is one of the most complex and important parts of the body. It supports the upper body and head, and is attached to the brain (where applicable). The upper body is held up by muscles in the stomach and lower back. If these muscles are relaxed by slouching for long periods, unnecessary strain in placed on muscles and bones which were not meant to bear the weight of the body.

  To avoid back problems, users should (i) sit in a good chair, and (ii) sit upright, using those all important flat-tummy muscles and lower back muscles to support your upper body. Don't sit in a draft. Cold air blowing across the back and neck causes stiffness and tension.
- *Mouse strain.* Mouse strain is a strain in the tendons of the finger and forearm, which spreads to the shoulder and back and can be quite painful. It comes from using the mouse too much. The symptoms can be lessened by making sure that users do not sit too far away from the desk where the mouse lies, and by having a support for the mouse forearm. The ultimate solution is simple: don't use the mouse. Use of the keyboard is far less hazardous. Learning keyboard shortcuts is good for prolonged work.
- *Pregnancy and cancer.* Some studies recommend that pregnant women wear protective aprons when sitting in front of computer screens. It is unclear whether this has any real purpose, since any radiation from the screen would be easily stopped by normal clothing.
- *Generally.* Users should not sit for long periods without taking a break. Looking away from the screen (to a far away object) at regular intervals relaxes the eyes. Walking around exercises the back and relaxes the shoulders. Use of anti-static, anti-reflective screens is recommended.

5.6.2 Dealing with Users: Etiquette

Although even the most stoical administrator's convictions might occasionally be called into question, system administration is a social service and it is important to remain calm and reasonable. Users frequently believe that the system administrator has nothing better to do than to answer every question and execute every whim and fancy. Dealing with users is no small task. In ref. [145], user friendly administrators are likened to user-friendly software!

5.6.3 Ethics and Responsibility

A system administrator wields great power. He or she has the means to read everyone's mail, change anyone's files, to start and kill anyone's processes. This power can easily be abused, and that temptation could be great.

Another danger is the temptation for an administrator to think that the system exists primarily for him or her, and that the users are simply a nuisance to the smooth running of things; if network service is interrupted, or if a silly mistake is made which leads to damage in the course of an administrator's work, that is okay: the users should accept these mistakes because they were made trying to improve the system. When wielding such power there is always the chance that such arrogance will build up. Some simple rules of thumb are useful.

The ethical integrity of a system administrator is clearly an important issue. Administrators for top secret government organizations and administrators for small businesses have the same responsibilities towards their users and their organizations. One has only to look at the governing institutions around the world to see that power corrupts. Few individuals, however good their intentions, are immune to the temptations of such power at one time or other. As with governments, it is perhaps a case of: those who wish the power are least suited to deal with it.

Administrators 'watch over' backups, e-mail, private communications and they have access to everyone's files. While it is almost never necessary to look at a user's private files, it is possible at any time, and users do not usually consider the fact that their files are available to other individuals in this way. Users need to be able to trust the system and its administrator.

- What kind of rules can you fairly impose on users?
- What responsibilities do you have to the rest of the network community, i.e. the rest of the world?
- Censoring of information or views.
- Restriction of personal freedom.
- Taking sides in personal disputes.
- Extreme views (some institutions have policies about this).
- Unlawful behaviour.

Objectivity of the administrator means avoiding taking sides in ethical, moral, religious or political debates, in the role of system administrator. Personal views should be kept separate from professional views. However, the extent to which this is possible depends strongly upon the individual, and organizations have to be aware of this. Some organizations dictate policy for their employees. This is also an issue to be cautious with: if a policy is too loose it can lead to laziness and unprofessional behaviour; if it is too paranoid or restrictive it can lead to bad feelings in the organization. Historically, unhappy employees have been responsible for the largest computer crimes. For other references see [77, 78].

5.6.4 Propaganda and Misinformation

Computers can lie with flawless equanimity, sufficient to convince any inexperienced user that they always tell the truth. A computer has a perceived authority which makes it a very dangerous tool for abuse. An ill-thought out remark in a login message, or a deliberate

attempt to steer users with propaganda, can have equally insidious results. One might argue that this is no worse than the general reliance of a large proportion of the population on television and media, and indeed this is true. Information warfare plays on our vulnerabilities to authority symbols, and it is on the rise.

In the Paramount film the Wrath of Khan, a questioning lieutenant Saavik queries Spock about his use of a verbal code to mislead the enemy: "You lied?" she says. Spock replies: "I exaggerated." Although the scene is amusing, it highlights another way in which computers can convince us of incorrect information. A sufficient exaggeration might also be enough to convince us of a lie. Information can always be presented misleadingly. Where do we draw the line? Software which is incorrectly configured and delivers incorrect information is perhaps the worst example. For example, an early version of Mathematica (a tool for mathematical manipulation) gave an incorrect answer for the derivative of a well-known function. It would have been easy to have simply used this answer, knowing that Mathematica performs many complex manipulations flawlessly. Fortunately, the main users of Mathematica, at the time, were scientists, who are a naturally sceptical breed, and so the error was discovered. In a CD-ROM encyclopaedia, produced by Microsoft, a Norwegian right-wing political party was listed as a neo-Nazi organization. This was clearly an unfair exaggeration of the truth, with potentially damaging consequences abroad had this party ever been elected to government. The fact that the information was on a CD-ROM, containing a body of essentially correct information, would tend to convince readers of its general truth.

The book you are reading, by virtue of being in print, also has an authority and the power to mislead. If I were to write, 'the correct way to do X is Y', it might appear that that was the only correct solution to the problem. That might be true, but it might also only be my flawed opinion. That is one of the reasons why the emphasis of this book is on becoming independent and thinking for ourselves. To summarize: most users look up to computers in awe; for that reason, the computer is an authority symbol with a great potential for abuse. System administrators need to be on the look out for problems like this, which can damage credibility and manipulate users.

Principle 21 (Mind control) *Computers have a perceived authority. We need to be on the look out for abuses of that authority, whether by accident or by design.*

5.6.5 User Age Groups

Today, network communities consist of all age groups. It is a basic fact of life that different age groups have different attitudes and concerns, and that they behave differently towards one another and amongst themselves. In the anonymous world of electronic communication, age is not usually apparent except through behaviour. While as pre-teenagers we tend to be careful and polite, as teenagers we are often rude and arrogant. As we grow older, the sharp edges get worn down and we are once again nice rounded shapes which fit into our nice rounded holes.

The art of communication between age groups is a difficult one. While teenagers into their early twenties can be abrasive in their manner, they also have many important ideas, since they have not yet been worn down by peer pressure into subservience, or had their imaginations erased by media conservatism.

The way in which age groups use computers reflects their interests and attitudes, and we have to consider this relation to the rules and policies for use of a computer system. We need to separate recreational use from professional use and consider to what extent recreational use could damage an organization professionally. It is not uncommon to see employees sign their e-mail with a phrase of the form

> *The opinions expressed here are purely my own, and should not be identified in any way with my employer.*

This is one way of clarifying the point, but it might not be sufficient. If a user expresses radical or discomforting opinions about something publically, this could colour others' views of the organization which the individual works for. It might not be fair, but it is unavoidable. System policy has to take into account the human differences between age groups. Whatever seems to be acceptable behaviour for one group in a community can be unacceptable for another.

Exercises

Exercise 5.1 (This problem is most easily solved on a Unix-like host.) Imagine that it is the start of the university semester and a hundred new students require an account. Write an `adduser` script which uses the file system layout which you have planned for your host to install home-directories for the users and which registers them in the password database. The script should be able to install the accounts from a list of users provided by the university registration service.

Start either by modifying an existing script (e.g. GNU/Linux has an adduser package) or from scratch. Remember that installing a new user implies the installation of enough configuration to make the account work satisfactorily at once, e.g. Unix dot files.

Exercise 5.2 One of the central problems in account management is the distribution of passwords. If we are unable (or unwilling) to use a password distribution system like NIS, passwords have to be copied from host to host. Assume that user home-directories are shared amongst all hosts. Write a script which takes the password file on one host and converts it into all of the different file formats used by different Unix-like OSes, ready for distribution.

Exercise 5.3 Write a script to monitor the amount of disk space used by each user.

Exercise 5.4 Consider the terminal room at your organization. Review its layout critically. Does the lighting cause reflection in the screens, leading to eye-strain. How is the seating? Is the room too warm or too cold? How could the room be redesigned to make work conditions better for its users?

Exercise 5.5 Describe the available support services for users at your site. Could these be improved? What would it cost to improve support services (can you estimate the number of man-hours, for instance) to achieve the level of support which you would like?

Exercise 5.6 Analyse and comment on the example shell configuration in section 5.3.2.

Chapter 6

Models of Network Administration

Until recently, computer systems were organized either by inspired local ingenuity or through an inflexible prescription, dictated by a vendor. With the success of Unix, as the backbone of the Internet, and indeed our new era of worldwide communication, Unix system managers have had more or less complete freedom to find optimal solutions to the problem of Unix administration. Some of the fruits of this have since filtered back into PC administration models, and while this has not increased their flexibility, it has revolutionized PC server technology. Meanwhile, Unix remains the anarchist amongst rigid regiments of PCs, where new ideas are forged.

6.1 Administration Models

Models of network administration have evolved by distilling locally acquired experience from many sites. In latter years, attempts have been made to build software systems which apply certain principles to the problem of management. Network management has, to some extent, been likened to the process of software development [158] in the System Administration Maturity Model, by Kubicki. This work was an important step in formalizing system administration. Later, a formalization was introduced by describing system administration in terms of automatable primitives. Several approaches to the management of computers in a network have emerged:

- *Reboot:* with the rapid expansion of networks the number of local networks has outgrown the number of experienced technicians. The result is that there are many administrators who are not skilled in the systems they are forced to manage. A disturbing but common belief, which originated in the 1980s microcomputer era, is that problems with a computer can be fixed by simply rebooting the operating system. Since home computer systems tend to crash with alarming regularity, this is a habit which has been acquired from painful experience. Unfortunately, it is somewhat analogous to ancient peoples sacrificing their young to the gods in order to appease them and cure their current ills (a rather unreliable mysticism). The reboot habit should, of course, be stifled. More mature preemptive systems like Unix and perhaps NT should *almost never* need to be

rebooted[1], indeed it can be damaging to do so. Rebooting a multi-user system is dangerous since users might be logged in from remote locations and could lose data. Moreover, the number of problems which can be fixed by rebooting a system is small compared to the number of problems which eventually arise. For this reason, we classify this as a regrettable but real development in network management.

- *Manual:* the default approach to system management is to allow qualified humans do everything by hand. This approach suffers from a lack of scalability. It suffers from human flaws and a lack of intrinsic documentation. Humans are not well-disciplined at documenting their work, or their intended configurations. There are also issues concerned with communciation and work in a team which can interfere with the smooth running of systems. When two manual administrators have a difference of opinion, there can be contention. The relevance of interpersonal skills in system administration teamwork was considered in ref. [140], and a cooperative shell environment for helping to discipline work habits was considered in ref. [3].
- *Control:* another approach to system administration is the use of control systems. Tivoli, HP OpenView and Sun Solstice are examples of these. In the control approach, the system administrator follows the state of the network by defining error conditions to look for. A process on each host reports errors as they occur to the administrator. In this way the administrator has an overview of every problem on the network from his/her single location, and can either fix the problems by hand as they occur (if the system supports remote login), or distribute scripts and antidotes which provide a partial automation of the process. The disadvantage with this system is that a human administrator usually has to start the repair procedures by hand, and this creates a bottleneck: all the alarms go to one place to be dealt with serially. With this approach, the amount of work required to run this system increases roughly linearly with the number of hosts on the network.
- *Immunology (self-maintenance):* a relatively new approach to system management which is growing in popularity is the idea of equipping networked operating systems with a simple immune system. By analogy with the human body, an immune system is an automatic system that every host possesses which attempts to deal with emergencies. An immune system is the Fire, Police and Paramedic services as well as the garbage collection agencies. In an immune system, every host is responsible for automatically repairing its own problems, without crying warnings about what is going on to a human. This avoids a serial bottleneck created by a human administrator. The time spent on implementing and running this model is independent of the number of hosts on the network.

  Unix administrators have run background scripts to perform system checks and maintenance for many years. Such scripts (often called *sanity checking* scripts) run daily or hourly, and make sure that each system is properly configured, perform garbage cleaning and report any serious problems to an administrator. In an immunological model, the aim is to minimize the involvement of a human being as far as possible. The tool cfengine introduced a technology which promotes this approach. Note that there is nothing to prevent this model from working along side a control system.

[1] NT still seems to have some stablility problems, but in principle it ought to be as stable as Unix.

These latter two approaches are easily implemented on Unix systems because of the large number of scripting languages and tools available. On Microsoft systems the amount of freely available languages ported from Unix has now increased to the point where there are few problems there either. For Macintosh systems there seems to have been little development, and one is locked into the commercial products available there. This will change with Apple's newest server product, Mac OS Server X, which is based on emulation technology, including BSD 4.3 running on a Mach kernel.

The main problem in implementing monitoring software on the insecure operating systems is that they are too often unstable (this has nothing to do with their security). If a network administrator were informed every time a Windows 9x machine crashed or was rebooted on a network of hundreds, there would be a huge amount of irrelevant traffic. Since there is nothing an administrator can do to change this fact, it has to be accepted as a characteristic of these operating systems and tolerated.

NT can be both easier and harder to administrate than Unix. It can be easier because the centralized model of having a domain server running all the network services means that all configuration information can be left in one place (on the server), and that each workstation can be made (at least to a limited degree) to configure itself from the server's files. It is harder to administrate because the tools provided for system administration tasks work mainly by the Graphical User Interface (GUI) and this is not a suitable tool for addressing the issues of hundreds of hosts. The Resource Kit and free tools go some way to helping with this problem, but the Resource kit adds a substantial sum to the cost of running NT. It provides tools analogous to `cron` and script languages for automating tasks.

6.2 Immunity and Convergence

The immunity model is about self-sufficient maintenance and is of central importance to all scalable approaches to network management, since it is the only model which scales trivially with the number of networked hosts. The idea behind immunity is to automate host maintenance in such a way as to give each host responsibility for its own configuration. A level of automation is introduced to every host, in such as way as to bring each host into an *ideal state*. What we mean by an ideal state is not fixed: it depends upon local system policy, but the central idea of the immunity model is to keep hosts as close to their ideal state as possible.

The immunity model has its origins in the work of John von Neumann, the architect of modern computer systems. He was the first person to recognize the analogy between living organisms and computers [271, 272], and clearly understood the conceptual implications of computing machines which could repair and maintain themselves, as early as 1948.

All collective systems (including all biological phenomena) are moderated and stabilized by a cooperative principle of *feedback regulation*. This regulating principle is sometimes called the Prey-Predator scenario, because it is about *competition* between different parts of a system. When one part of the system starts to grow out of control, it tends to favour the production of an antidote which keeps that part in check. Similarly, the antidote cannot exist without the original system, so it cannot go so far as to destroy the original system, since it destroys itself in the process. A balance is therefore found between the original part of the system and its antidote. The classical example of a Prey-Predator model is that of populations

of foxes and rabbits. If the number of rabbits increases suddenly, then foxes feed well and grow in numbers, eating more rabbits, thus stabilizing the numbers. If rabbits grow scarce, then foxes die and thus an equilibrium is maintained. Another example of this type of behaviour is to be found in the body's own repair and maintenance systems. The name 'immunity' is borrowed from the idea that systems of biological complexity are able to repair themselves, in such a way as to maintain an equilibrium called *health*. The relative immunity of, for instance, the human body to damage and disease is due to a continual equilibrium between death, cleanup and renewal. Immunity from disease is usually attributed to an *immune system*, which is comprised of cells which fight invading organisms, though it has become clear over the years that the phenomenon of immunity is a function of many cooperating systems throughout the entire human organism, and that disease does not distinguish between *self* and *non-self* (body and invader), as was previously thought. In the immunity model, we apply this principle to the problem of system maintenance.

Automatic systems maintenance has been an exercise in tool-building for many years. The practice of automating basic maintenance procedures has been commonplace in the Unix world (see section 7.4.1). Following von Neumann's insights, the first theoretical work on this topic, addressing the need for convergence, appears to be by Burgess [32, 37]. The biological analogy between computers and human immune systems has been used to inspire models for the detection of viruses, principally in insecure operating systems. This was first discussed in 1994 by Kephart of IBM [147], and later expanded upon by Forrest *et al.* [91, 249, 94, 92, 129, 128, 274, 127, 69, 93, 198, 68]. The analogy between system administration and immunology was discussed independently by Burgess [34, 35] in the wider context of general system maintenance. References [35, 33] also discuss how computer systems can be thought of as statistical mechanical systems, drawing on a wide body of knowledge from theoretical physics. Interestingly, refs. [35] and [249], which appeared slightly earlier, point out many of the same ideas independently, both speculating freely on the lessons learned from human immunology, though the latter authors do not seem to appreciate the wider validity of their work to system maintenance.

The idea of immunity requires a notion of convergence. *Convergence* means that maintenance work (the counter force or antidote) tends to bring a host to a state of *equilibrium*, i.e. a stable state, which is the state we would actually like the system to be in. The more maintenance which is performed, the closer we approach the ideal state of the system. When the idea state is reached, maintenance work stops, or at least has no further effect. The reason for calling this the immunity model is that this is precisely the way that biological maintenance works. As long as there is damage or the system is threatened, a counter force is mobilized, followed by a garbage collection and a repair team. There is a direct analogy between medicine and computer maintenance. Computer maintenance is just somewhat simpler.

6.3 Network Organization

As we have already mentioned in section 3.1, a network is a community of cooperating and competing players. A system administrator has to choose the players and assign them their roles on the basis of the job which is intended for the computer system. There are two

aspects of this to consider: the machine aspect and the human aspect. The machine aspect relates to the use of computing machinery to achieve a functional infrastructure; the human aspect is about the way people are deployed to build and maintain that infrastructure.

Identifying the purpose of a system is the first step to building a successful computer system. Choosing hardware and software is the next. If we are only interested in word processing, we do not buy a mainframe running Unicos. On the other hand, if we are interested in high volume distributed database access, we do not buy a laptop running Windows. There is always a balance to be achieved, a right place to spend money and a right place to save money. For instance, since the CPU of most computers is idle some 90% of the time, simply waiting for input, money spent on fast processors is often wasted; conversely, the greatest speed gains are usually to be made in extra RAM memory, so money spent on RAM is always well spent. Of course, it is not always possible to choose the hardware we have to work with. Sometimes we inherit a less than ideal situation and have to make the best of it. This also requires ingenuity and careful planning.

Assuming that we can choose hardware, it is particularly prudent to weigh the convenience of keeping to a single type of hardware and operating system (e.g. just PCs with NT), with the possible advantages of choosing the absolutely best hardware for the job. Of course, vendors always want to sell us a solution based on their own products, so they cannot be trusted to evaluate an organization's needs objectively. The great advantage of uniformity is ease of administration and automatic hardware redundancy. For many issues, keeping to one type of computer is more important than what the type of computer is.

Principle 22 (Homogeneity/Uniformity I) *System homogeneity or uniformity means that all hosts appear to be essentially the same. This makes hosts predictable for users and manageable for administrators. It allows for reuse of hardware in an emergency.*

If we have a dozen machines of the same type, we can establish a standard routine for running them and for using them. If one fails, we can replace it with another. The downside of uniformity is that there are sometimes large performance gains to be made by choosing special machinery for a particular application. For instance, a high availability server requires multiple, fast processors, lots of memory and high bandwidth interfaces for disk and network. In short, it has to be a top quality machine. A word processor does not. Purchasing such a machine might complicate host management slightly. Tools such as cfengine can help integrate hosts with special functions painlessly.

Having chosen the necessary hardware and software, we have to address the function of each host within the community, i.e. the *delegation* of specialized tasks called *services* to particular hosts, and also the competition between users and hosts for resources, both local and distributed. For all of this to work with some measure of equilibrium, it has to be carefully planned and orchestrated.

In the deployment of machinery, there are two opposing philosophies: one machine, one job, and the consolidated approach. In the first case, we buy a new host for each new task on the network. For instance, there is a mail server and a printer server and a disk server, and so on. This approach was originally used in PC networks running DOS, because each host was only capable of running one program at a time. That does not mean that it is redundant today: the distributed approach still has the advantage of spreading the load of service across several hosts. This is useful if the hosts are also workstations which are used interactively by

users, as they might be in small groups with few resources. Making the transition from mainframe to a distributed solution was discussed in a case study in ref. [266].

On the whole, however, modern computer systems have more than enough resources to run several services simultaneously. Indeed, a lot of unnecessary network traffic can be avoided by placing all file services (disk, web and FTP) on the same host (see chapter 8). It does not make sense to keep data on one host and serve them from another, since the data first have to be sent from the disk to the server and then from the server to the client, resulting in twice the amount of network traffic.

The consolidated approach to services is to place them all on just a few server-hosts. This can lead to better security, since it means that we can exclude users from the server itself. Today most PC network solutions make this simple by placing all of the burden of services on specialized machines. PC server-hosts are not meant to be used by users themselves: they stand apart from workstations. With Unix-based networks, we have complete freedom to run services wherever we like. There is no principal difference between a workstation and a server-host. This allows for a rational distribution of load.

Of course, it is not just machine duties which need to be balanced throughout the network; there is also the issue of human tasks, such as user registration, operating system upgrades, hardware repairs, and so on. This is all made simpler if there is a team of humans. Based on the principle of delegation, we can make the following suggestion:

Suggestion 8 (Delegation II) *For large numbers of hosts, distributed over several locations, consider a policy of delegating responsibility to a local administrator with closer knowledge of the hosts' patterns of usage. Zones of responsibility allow local experts to do their jobs.*

It is important to understand the function of a host in a network. For small groups in large organizations, there is nothing more annoying than to have central administrators mess around with a host which they do not understand. They will make inappropriate changes and decisions.

Zones of responsibility have as much to do with human limitations as with network structure. Human psychologists have shown that each of us has the ability to relate to no more than around 150 people. There is no reason to suppose that this limitation does not also apply to other objects which we assemble into our work environment. If we have 4000 hosts which are identical, then that need not be a psychological burden to a single administrator, but if those 4000 consist of 200 different groups of hosts, where each group has its own special properties, then this would be an unmanageable burden for a single person to follow. Even with special software, a system administrator needs to understand how a local milieu uses its computers, in order to avoid making decisions which work against that milieu.

6.4 Bootstrapping Infrastructure

Until recently, little attention was given to analysing methodologies for the construction of efficient and stable networks from the ground up, although some case studies for large scale installations were made earlier [142, 84, 248, 43, 181, 151, 101, 234, 121, 178, 137, 80]. One interesting exception is a discussion of human roles and delegation in network

management in refs. [173, 107]. With the explosion in numbers of hosts combined in networks, several authors have begun to address the problem of defining an infrastructure model which is stable, reproducible and robust to accidents and upgrades [32, 81, 263, 35].

The term 'Bootstrapping an infrastructure' was coined by Traugott and Huddleston [263], and nicely summarizes the basic intent. Both Evard [81] and Traugott and Huddleston have analysed practical case studies of system infrastructures both for large networks (4000 hosts) and for small networks (as few as three hosts). Interestingly, Evard's conclusions, although researched independently of Burgess [40, 32, 37, 33, 34], seem to vindicate the theoretical model used in constructing the tool cfengine.

6.4.1 Principles of Stable Infrastructure

The principles on which we would like to build an infrastructure are straightforward. These summarize both common sense and the experiences of the authors cited above.

Principle 23 (Scalability) *Any model of system infrastructure must be able to scale efficiently to large numbers of hosts (and perhaps subnets, depending on the local netmask).*

A model which does not scale efficiently with numbers of hosts is likely to fail quickly, as networks tend to expand rapidly beyond expectations.

Principle 24 (Reliability) *Any model of system infrastructure must have reliability as one of its chief goals. Down-time can often be measured in real money.*

Reliability is not just about the initial quality of hardware and software, but also about the need for preventative maintenance. The issue of convergence is central here.

Corollary 25 (Redundancy) *Reliability is safeguarded by redundancy, or backup services running in parallel, ready to take over at a moment's notice [244].*

Although redundancy does not prevent problems, it aids swift recovery. Barber has discussed improved server availability through redundancy [20]. High availabilty clusters and mainframes are often used for this problem. Gomberg *et al.* have compared scalable software installation methods on Unix and NT [104]. A refinement of the principle of homogeneity can be stated here, in its rightful place:

Principle 26 (Homogeneity/Uniformity II) *A model in which all hosts are basically similar is (i) easier to understand conceptually both for users and administrators, (ii) cheaper to implement and maintain, and (iii) easier to repair and adapt in the event of failure.*

and finally:

Corollary 27 (Reproducibility) *Avoid improvising system modifications, on the fly, which are not reproducible. It is easy to forget what was done, and this will make the functioning of the system difficult to understand and predict, for you and for others.*

6.4.2 Virtual Machine Model

As Traugott and Huddleston have eloquently put it, we need to think of a network solution not so much as a loose association of hosts, but rather as a large virtual machine, composed of associated organs. It is a small step from viewing a multi-tasking operating system as a collaboration between many specialized processes, to viewing the entire network as a distributed collaboration between specialized processes on different hosts. There is little or no difference in principle between an internal communication bus and an external communication bus. Many sites adopt specific policies and guidelines in order to create this seamless virtual environment [42], by limiting the magnitude of the task. Institutions with a history of managing large numbers of hosts have a tradition of either adapting imperfect software to their requirements or creating their own. Tools such as `make`, which have been used to jury-rig configuration schemes [263], can now be replaced by more specific tools like `cfengine` [32, 37]. As with all things, getting started is the hard part.

6.4.3 Creating Uniformity Through Automation

Simple, robust infrastructure is created by planning a system which is easy to understand and maintain. If we want hosts to have the same software and facilities, creating a general uniformity, we need to employ automation of keep track of changes [126, 32, 37]. To begin, we must formulate the needs and potential threats to system availability. That means planning resources, as in the foregoing sections, and planning the actually motions required to implement and maintain and system. If we can formalize those needs by writing them in the form of a policy, program or script, then half the battle is already won, and we have automatic reproducibility.

Principle 28 (Abstraction generalizes) *Expressing tasks in an operating-system independent language reduces time spent debugging, promotes homogeneity and avoids unnecessary repetition.*

A script implies reproducibility, since it can be rerun on any host. The only obstacle to this is that not all script languages work on all systems.

Suggestion 9 (Platform independent languages) *Use languages and tools which are independent of operating system peculiarities, e.g.* `cfengine` `perl`, `python`. *More importantly, use the right tool for the right job.*

Perl is particularly useful, since it runs on most platforms and is about as operating system independent as it is possible to be. The disadvantage of Perl is that it is a low level programming language, which requires us to code with a level of detail which can obscure the purpose of the code. Cfengine was invented to address this problem. The cfengine is a very high level interface to system administration. It is also platform independent, and runs on most systems. Its advantage is that it hides the low level details of programming, allowing us to focus on the structural decisions. We shall discuss this further below.

6.4.4 Revision Control

One approach to the configuration of hosts is to have a standard set of files in a file-base which can be simply copied into place. Several administration tools have been built on

this principle, e.g. Host Factory [83]. The Revision Control System (RCS), designed by Tichy [261] was created as a repository for files, where changes could be traced through a number of evolving versions. RCS was introduced as a tool for programmers, to track bug fixes and improvements through a string of versions. The CVS system is an extended front-end to this system. System configuration is a similar problem, since it involves modifying the contents of many key files. Many administrators have made use of the revision control systems to keep track of configuration file changes, though little has been written about it. PC management with RCS has been discussed by Rudorfer [221]. Revision control is a useful way of keeping track of text file changes, but it does not help us with other aspects of system maintenance, such as file permissions, process management or garbage collection.

6.4.5 Software Synchronization

In section 3.9.4 we discussed the distribution of data amongst a network community. This technique can be used to maintain a level of uniformity in the software used around the network. Software synchronization has been discussed in refs. [21, 119, 240]. Distribution by package mechanisms were pioneered by Hewlett Packard [215], with the `ninstall` program. For some software packages, Hewlett Packard use cfengine as a software installation tool [37]. Distribution by placement on network file systems like the AFS has been discussed in ref. [154].

6.4.6 Push Models and Pull Models

Revision control does not address the issue of uniformity unless the contents of the file-base can be distributed to many different hosts. There are two types of distribution mechanism, which are generally referred to as *push* and *pull* models of distribution:

- *Push:* the model is epitomized by the `rdist` program. Pushing files from a central location to a number of hosts is a way of forcing a file to be written to a group of hosts. The central repository decides when changes are to be sent, and the hosts which receive the files have no choice about receiving them [172]. In other words, control over all of the hosts is forced by the central repository. The advantage of this approach is that it can be made efficient. A push model is more easily optimized than a pull approach. The disadvantage of a push model is that hosts have no freedom to decide their own fate. A push model forces all hosts to open themselves to a central will. This could be a security hazard. In particular, `rdist` requires a host to grant not just file access, but full complete privilege to the distributing host. Another problem with push models is the need to maintaina list of all the hosts to which data will be pushed. For large numbers of hosts, this can become unwieldy.
- *Pull:* the pull model is represented by `cfengine` and `rsync`. With a pull model, each host decides to collect files from a central repository, of its own volition. The advantage of this approach is that there is no need to open a host to control from outside, other than the trust implied by accepting configuration files from the distributing host. This has significant security advantages. It was recommended as a model of centralized system administration in refs. [226, 37, 263]. The main disadvantage to this method is that

optimization is harder. `rsync` addresses this problem by using an ingenious algorithm for transmitting only file changes, and thus achieves a significant compression of data, while `cfengine` uses multi-threading to increase server availability.

6.4.7 Reliability

One of the aims of building a sturdy infrastructure is to cope with the results of failure. Failure can encompass hardware and software. It includes downtime due to physical error (power, net cables and CPUs) and also downtime due to software crashes. The net result of any failure is loss of service. Our only defence against actual failure is parallelism, or redundancy. When one component fails, another can be ready to take over. Often it is possible to *prevent failure* with pro-active maintenance (see the next chapter for more on this issue). For instance, it is possible to vacuum clean hosts, to prevent electrical short-circuits. It is also possible to perform garbage collection which can prevent software error. System monitors (e.g. cfengine) can ensure that crashed processes get restarted, thus minimizing loss. Reliability is clearly a multi-faceted topic. We shall return to discuss reliability more quantitatively in section 11.6.9.

Component failure can be avoided by parallelism or redundancy. One way to think about this is to think of a computer system as providing a service which is characterized by a flow of information. If we consider Figure 6.1, it is clear that a flow of service can continue, when servers work in parallel, even if one or more of them fails. In Figure 6.2 it is clear that systems which are dependent on other series are coupled in series and a failure prevents the flow of service. Of course, servers do not really work in parallel. The normal citation is to employ a *fail-over* capability. This means that we provide a backup service. If the main service fails, we replace it with a backup server. The backup server is not normally used, however. Only in a few cases can one find examples of load-sharing by switching between (de-multiplexing) services.

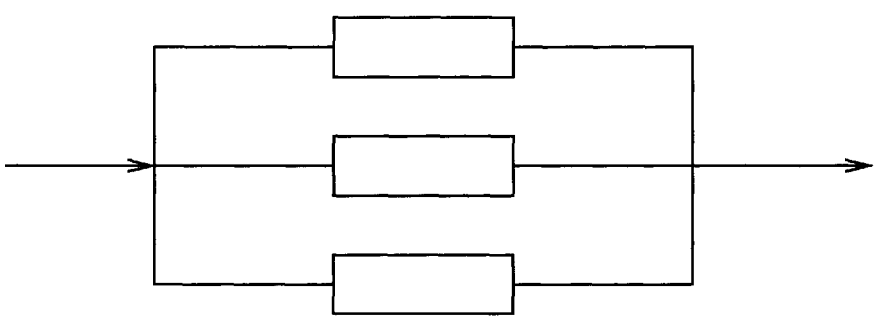


Figure 6.1 System components in parallel, implies redundancy

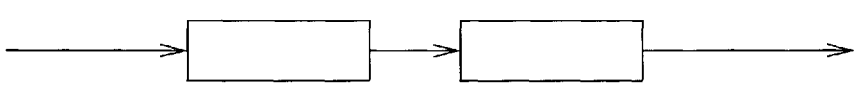


Figure 6.2 System components in series, implies dependency

6.5 Cfengine: Policy Automation

The idea of being able to automate the configuration from a high level policy was the idea behind cfengine. Prior to cfengine, several authors had explored the possibilities for automation and abstraction without combining all the elements into an integrated framework [110, 86, 126, 143]; most of these were too specific or too low level to be generally useful.

Cfengine is a system administration tool consisting of two elements: a language and a configuration engine. Together these are used to instruct and enable all hosts on a network about how to configure and maintain themselves. Rather than being a cloning mechanism, cfengine takes a broader view of system configuration, enabling host configurations to be built from scratch on *classes* of host.

Cfengine is about defining the way we want all the hosts on our network to be configured, and having them do the work themselves. It is a tool for automation and for definition. Because it includes a language for describing system configuration at a high level, it can also be used to express system policy in formal terms. The correct way to use cfengine is therefore to specify and automate system policy in terms of concrete actions. See section C.5.

What makes cfengine different from scripting languages is the high level at which is operates. Rather than allowing complete programming generality, cfengine provides a set of *intelligent primitives* for configuring and maintaining systems. An important feature of cfengine primitives is that they satisfy, as far as possible, the principle of convergence (see section 6.2). This means that a policy expressed by a cfengine program can easily be made to embody a convergent behaviour. As a system inevitably drifts from its ideal state, a cfengine policy brings it back to that ideal state. When it reaches that state, cfengine becomes impotent and does no more.

Cfengine works from a central configuration, maintained from one location. That central configuration describes the entire network by referring to classes and types of host. Many abstraction mechanisms are provided for mapping out the networks. The work of configuration and maintenance is performed by each host separately. Each host is thus given responsibility for its own state and the work of configuration is completely distributed. This means that a cfengine configuration scales trivially with the number of hosts, or put another way, the addition of extra hosts does not affect the ability of other hosts to maintain themselves. Traffic on servers increases at most linearly with the number of hosts, and the network is relied upon as little as possible. This is not true of network based control models, for instance, where network resource consumption increases at least in proportion to the total number of hosts, and is completely reliant on network integrity.

Cfengine programs make it easy to specify general rules for large groups of host and special rules for exceptional hosts. Here is a summary of cfengine's capabilities:

- Check and configure the network interface on network hosts.
- Edit text files for the system or for all users.
- Make and maintain symbolic links, including multiple links from a single command.
- Check and set the permissions and ownership of files.
- Tidy (delete) junk files which clutter the system.
- Systematic, automated (static) mounting of NFS file systems.
- Checking for the presence or absence of important files and file systems.

- Controlled execution of user scripts and shell commands.
- Process management.

The full details of cfengine are described in ref. [30].

Because cfengine runs locally as a host process, and can be started any number of times without harming a host, it can be combined with other host control mechanisms. Complex commercial management products like Tivoli and OpenView, to name two examples, rely on scripts at the end host for most of their operations. Such scripts could easily be exchanged with cfengine programs to simplify and improve convergence properties of hosts.

The scalability of the cfengine model means that it can be deployed on a single host or on networks with thousands of hosts.

6.6 SNMP Network Management

The ability to read information about the performance of network hardware via the network itself is an attractive idea. Suppose we could look at a router on the second floor of a building half a mile away and immediately see the load statistics, or number of rejected packets it has seen; or perhaps the status of all printers on a subnet. That would be useful diagnostic information. Similar information could be obtained about software systems on any host.

SNMP (Simple Network Management Protocol) is a protocol designed to do just this. SNMP was spawned in 1987 as a Simple Gateway Monitoring Protocol, but was quickly extended and became a standard for network monitoring. SNMP was designed to be small and simple enough to be able to run on even minor pieces of network technology like bridges and printers. It now exists in two versions: version 1 and version 2.

SNMP supports two operations: *get* and *put*. It can therefore read and modify the data stored on a device. SNMP access is mediated by a server process on each hardware node, which normally communicates by UDP/IP on ports 161 and 162. Modern operating systems often run SNMP daemons or services which advertise their status to an SNMP capable manager. The services are protected by a rather weak password which is called the *community string*.

SNMP information is stored in data structures called MIBs (Management Information Bases). The MIBs are catalogued in RFC 1213 and give hardware and software profiles. They are dynamically updated. An SNMP request specifies the information it wants to read/write by giving the name of a MIB. In SNMP v2 there are standard MIBs for address translation tables, TCP/IP statistics and so on. There are 27 default parameters that can be altered by SNMP: system name, location and human contact; interface state (up/down), hardware and IP address, IP state (forwarding gateway/not) IP TTL, IP next HOP address, IP route age and mask, TCP state, neighbour state, SNMP trap enabling.

Operating systems define their own MIBs for system performance data. Some commercial network management systems like Hewlett Packard's OpenView work by reading and writing MIBs using SNMP client-server technology. Most Unix variants now support SNMP. NT also supports for SNMP; its MIBs can be used to collect information from NT systems, such as the names of users who are logged on. This can be considered a security risk. Novell Netware 5 has SNMP support for network monitoring.

SNMP seems to be increasing in popularity, but like any public information database, it can be abused by network attackers. SNMP has very weak security; it is a prime target for abuse.

Some sites choose to disable SNMP services altogether on hosts, using it only for monitoring network transport hardware. All sites should filter SNMP packets to and from external networks to avoid illegal access of these services from intruders.

6.7 Integrating Multiple OSes

Combining radically different operating systems in a network environment is a challenge both to users and administrators. Each operating system services a specific function well, and if we are to allow users to move from operating system to operating system, with access to their personal data, we need to balance the convenience of availability with the caution of differentiation. It ought to be clear to users where they are, and what system they are using, to avoid unfortunate mistakes. Combining different Unix-like systems is challenge enough, but adding Windows hosts or Macintosh technology to a primarily Unix based network, or vice versa, requires careful planning [28]. Integrating radically different network technologies is not worth the effort unless there is some particular need. It is always possible to move data between two hosts using the universally supported FTP protocol. But do we need to have open file sharing or software compatibility?

6.7.1 File System Sharing

Sharing of file systems between different operating systems can be useful in a variety of circumstances. File-servers, which host and share users' files, need to be fast, stable and capable machines. Workstations for end-users, on the other hand, are chosen for quite different reasons. They might be chosen to run some particular software, or on economic grounds, or perhaps for user-friendliness. The Macintosh has always been a favourite workstation for multimedia applications. It is often the preferred platform for music and graphical applications. Windows operating systems are cheap and have a wide and successful software base.

There are other reasons for wanting to keep an inhomogeneous network. An organization might need a mainframe or vector processor for intensive computation, whose disks need to be available to workstations for collecting data. There might be legacy systems waiting to be replaced with new machinery, which we have to accommodate in order to run old software, or development groups supporting software across multiple platforms. There is a dozen reasons for integration.

What about solutions? Most solutions to the file sharing problem are software based. Client and server software is available for implementing network sharing protocols across platform boundaries. For example, client software for the Unix NFS file system has been implemented for both Windows (PCNFS) and Macintosh System 7/8. This enables Windows and Macintosh workstations to use Unix-like hosts as file and printer servers, in much the same way as NT servers or Novell Netware servers provide those services. These services are adequate for insecure operating systems, since there is no need to map file permissions across foreign file systems. NT is more of a problem, however. NT ACLs cannot be represented in a simple fashion on a Unix file system.

The converse, that of making Unix files available to PCs, has the reverse problem. While NT is capable of representing Unix file permissions, Windows 9x and the Macintosh are not.

Insecure operating systems are always a risk in network sharing. The Samba software is a free software package which implements Unix file semantics in terms of the Windows SMB (Server Message Block) protocols.

Specialized hardware can be used to implement heterogeneous sharing. Network Appliance or Auspex intelligent servers can imitate Unix and NT file systems from a common server.

Netware provides an NT client called NDS (Network Directory Services) for NT which allows NT domain servers to understand the Novell object directory model. Clearly, there is already file system compatibility between PC servers. Conversely, NT provides Netware clients, and other server products can be purchased to provide access to AS/400 mainframes. Both Novell and NT provide Macintosh clients, and Macintosh products can also talk to NT and Unix servers. GNU/Linux has made a valiant attempt to link up with most existing sharing protocols on Unix, PCs and Apple hosts.

Mechanisms clearly exist to implement cross-platform sharing. The main question is, how easy are these systems to implement and maintain? Are they worth the cost in time and money?

6.7.2 User IDs and Passwords

If we intend to implement sharing across such different operating systems as Unix and NT, we need to have common user names on both systems. Cross-platform user authentication is usually based on the understanding that user name text can be mapped across operating systems. Clearly, numerical user IDs and security IDs cannot map meaningfully between systems without some glue to match them: that glue is the user name. To achieve sharing, then, we must standardize user names. Unix-like systems often require user names to be no more than eight characters, so this is a good limit to keep to if Unix-like operating systems are involved or might become involved.

Principle 29 (One name for one object II) *Each user should have the same unique name on every host. Multiple names lead to confusion and mistaken identity. A unique user name makes it clear which user is responsible for which actions.*

Common passwords across multiple platforms is much harder than disk sharing, and it is a much more questionable practice (see below).

6.7.3 Authentication

Making passwords work across different operating systems is often a pernicious problem in a scheme for complete integration. The password mechanisms for Unix and Windows are completely different and basically incompatible. The new Mac OS Server X is based on BSD4.3 emulation, so its integration with other Unix-like operation systems should be relatively painless. Windows, however, remains the odd one out. Whether or not it is correct to merge the password files of two separate operating systems is a matter for policy. The user bases of one operating system are often different from the user bases of another. From a security perspective, making access easy is not always the right thing to do. Owing to the cultural backgrounds of their user bases, Windows accounts are not always held in the same regard as Unix accounts. Windows provides the illusion of privacy, our own inviolable personal computer, whereas Unix feels more open and vulnerable.

Passwords are incompatible between Windows and Unix for two reasons: NT passwords can be longer than Unix passwords, and the form of encryption used to store them is different. The encryption mechanisms which are used to store passwords are one way transformations, so it is not possible to convert one into the other. There is no escaping the fact that these systems are basically incompatible.

A fairly recent development on the Unix side is Sun Microsystems' invention of Pluggable Authentication Modules (PAM). GNU/Linux has adopted several features from SunOS recently, and also supports PAM. The PAM mechanism is a little-documented method of allowing the standard Unix password mechanism to be exchanged or supplemented by any number of other password mechanisms, simply by adding modules to a configuration file `/etc/pam.conf`. Instead of being prompted for a Unix password on login, users are connected to one or more password modules. Each module prompts for a password and grants security credentials if the password is correctly received. Thus, for instance, users could be immediately prompted for a Unix password, a Kerberos password and a DCE password on login, thus removing the necessity for a manual login to these extra systems later. PAM also supports the idea of mapped passwords, so that a single strong password can be used to trigger the automatic login to several stacked modules, each with its own private password stored in a PAM database. This is a very exciting possibility, mitigated only by a conspicuous lack of documentation about how to write modules for PAM. PAM could clearly help in the integration of Unix with Windows if a module for Windows style authentication could be written for Unix. At the present time, I am not aware of anyone who has accomplished such an integration, so we must wait in anticipation for the details of PAM to become sufficiently lucid as to make it a useful integration tool.

6.7.4 Samba

Samba is a free software solution to the problem of making Unix file systems available to Windows operating systems. Windows NT uses a system of network file sharing based on their own SMB (Server Message Block) protocol. `Samba` is a Unix daemon-based service which makes Unix disks visible to Windows NT. Samba maps user names, so to use Samba we need an account with the same name on the NT server and on the Unix server. It maps user name textually, without much security. Samba configuration is in Unix style, by editing the text-file `/etc/smb.conf`. Here is an example file. Note carefully the 'hosts allow' line which restricts access to disks to specific IP addresses, like TCP wrappers.

```
[global]
  printing = bsd
  printcap name = /etc/printcap
  load printers = yes
  guest account = nobody
  invalid users = root
  workgroup = UNIX
  hosts allow = 128.39.

[homes]
  comment = Home Directories
  browseable = no
  read only = no create mode = 0644
```

```
[printers]
  comment = All Printers
  browseable = no
  path = /tmp
  printable = yes
  public = no
  writable = no
  create mode = 0644
```

Once the Samba server is active, the disks are available for use with the `net use` command, e.g.

```
C:\> net use F: \\host\directory
```

This example maps the named directory on the named host to NT drive letter `F:`. The reverse problem of mounting NT file systems on a Unix host works only for GNU/Linux hosts at present:

```
linux% smbmount //nthost/dir /mountpoint -U admin
```

6.8 A Model Checklist

Having decided on some model for network cooperation, it is only proper to take a moment to evaluate its implications. Does it pass the following tests?

- Will our installation survive a re-installation or upgrade of the OS?
- What is more important: user freedom or system well-being?
- Will the network survive the loss of any host?
- Do any choices compromise security or open any back-doors?
- Do users understand their responsibilities with regard to the network? (Do they need to be educated?)
- Have we observed all system responsibilities with respect to the network?
- Is the system easy to understand?
- Is the solution general for all operating systems?

If it fails one of these tests, one could easily find oneself starting again in a year or so.

Exercises

Exercise 6.1 Explain what is meant by Traugott and Huddleston's virtual machine view of the network. Compare this view of a computer system to that of a living organism, formed from many cooperating organs.

Exercise 6.2 Explain what is meant by *convergence*. What are the advantages of convergence?

Exercise 6.3 In an administrative environment, it is often important to have the ability to *undo* changes which have been made. What tools can help maintain versions of system configuration?

Exercise 6.4 Explain the difference between a *push* model and a *pull* model of system administration. What are the security implications of these, and how well do they allow for delegation of responsibility in the network?

Exercise 6.5 Discuss what problems are to be solved in a heterogeneous network, i.e. one composed of many different operating system types.

Exercise 6.6 Evaluate the cfengine primitives: are these sufficient for maintaining any operating system?

Exercise 6.7 What is the advantage of a central point of control and configuration in network management?

Exercise 6.8 Suppose you have at your disposal four Unix workstations, all of the same type. One of them has twice the amount of memory. What would you use for DNS? Which would be a web server? Which would be an NFS server?

Chapter 7

Configuration and Maintenance

We are now faced with two overlapping issues: how to make a computer system operate in the way we have intended, and how to keep it in that state over a period of time.

Configuration and maintenance are clearly related issues. Maintenance is simply configuration in the face of creeping decay. All systems tend to decay into chaos with time. There are many reasons for this decline, from deep theoretical reasons about thermodynamics, to the more intuitive notions about wear and tear. To put it briefly, it is clear that the number of ways in which a system can be in order is far fewer than the number of ways in which a system can be in a state of disorder, thus statistically, any random change in the system will, statistically, move it into disorder, rather than the other way around. We can even elevate this to a principle to emphasize its inevitability:

Principle 30 (Disorder) *All systems will eventually tend to a state of disorder unless a rigid and automated policy is maintained.*

Whether by creeping laziness or through undisciplined cooperation in a team [204, 229, 292, 79], poor communication or whatever, the system *will* degenerate as small errors and changes drive it forward. That degeneration can be counteracted by repair work which either removes or mitigates the errors.

Principle 31 (Equilibrium) *Deviation from a system's ideal state can be smoothed out by a counteractive response. If these two effects are in balance, the system will stay in equilibrium.*

The time scales over which errors occur and which repairs are made are clearly important. If we correct the system too slowly, it will run away from us. There is thus an inherent instability in computer networks.

7.1 System Policy

So far our analysis of networks has been about mapping out which machines performed which function on the network. Another side of network setup is the policies, practices and

procedures which are used to make changes to or to maintain the system as a whole, i.e. what humans do as part of the system administration process.

System administration is often a collaborative effort between several administrators. It is therefore important to have agreed policies for working so that everyone knows how to respond to 'situations' which can arise, without working against one another. A system policy also has the role of summarizing the attitudes of an organization to its members and its surroundings, and often embodies security issues. As Howell cites from Pogo [134], *We have met the enemy, and he is us!* A system policy should contain the issues we have been discussing in the foregoing chapters. There are issues to be addressed at each level: network level, host level, user level.

Principle 32 (Policy) *A clear expression of goals and responses prepares a site for future trouble and documents intent and procedure.*

It is crucial that everyone agrees on policy matters. Although a policy can easily be an example of blind rule-making, it is also a form of communication. A policy documents acceptable behaviour, but it should also document what response is appropriate in a crisis. Only then are we assured of an orchestrated response to a problem, free of conflicts and disagreements. What is important is that the document does not simply become an exercise in beauracracy, but is a living guide to the *practice* of network community administration. A system policy can include some or all of the following:

- *Organization:* what responsibility will the organization take for its users' actions? What responsibility will the organization take for the users' safety. Who is responsible for what? Has the organization upheld its responsibilities to the wider network community? Measures to prevent damage to others and from others.
- *Users:* allowing and forbidding certain types of software. Rigid control over space (quotas) or allow freedom, but police the system with controls. Choice of default configuration. A response to software piracy. A response to anti-social behaviour and harassment of others (spamming, obnoxious news postings, etc.) Are users allowed to play games, if so when? Are users allowed to chat online? Are users allowed to download files such as MP3 or pornography. Policy on sharing of accounts (i.e. preferably not). Policy on use of IRC robots and other automatic processes which collect large amounts of data off-line. Policy on garbage collection when disks become full: what files can legitimately be deleted?
- *Network:* will the network be segmented, with different access policies on different subnets? Will a firewall be used? What ports will be open on which subnets, and which will be blocked at the router. What services will be run?
- *Mail:* limit the size of incoming and outgoing mail. Spam filtering. Virus controls.
- *WWW:* allowing or forbidding user CGI scripts. Guidelines for allowed content of web pages. Policy regarding advertising on web pages. Load restrictions: what to do if certain pages generate too much traffic. Policy on plagiarism and illegal use of imagery.
- *Printing:* how many pages can be printed. Is printing of personal documents allowed. Should there be a limit to the number of pages which can be printed at one time (large documents hold up the print queue)?
- *Security:* physical security of hosts. Backup schedule. Who is allowed to be master of their own hosts? Can arbitrary users mount other users' home directories or mailboxes

with NFS on their private PCs (this means that they have automatic access to everyone's personal files)? What access controls should be used on files? Password policy (aging, how often should passwords change) and policy on closing accounts which have been compromised. 'Redundancy is our last avenue of survival' [242].

- *Privacy:* is encryption allowed? What tools will be provided for private communication?

See also the discussion of policy as a system administration tool in refs. [239, 112].

7.2 Synchronizing Clocks

One of the most fundamental tasks in a network is to keep the clocks on all hosts synchronized. Many security and maintenance issues depend upon clocks being synchronized correctly. Clock reliability varies enormously. The clocks on cheap PC hardware tend to drift very quickly, whereas clocks on more expensive workstations are rather better at keeping time. This is therefore a particular problem for cheap PC networks.

One option for most Unix-like systems is the `rdate` command, which sets the local clock according to the clock of another host. The `rdate` was absent from earlier GNU/Linux distributions. It can be simulated by a script:

```
#!/bin/sh
#
# Fake rdate script for linux - requires rsh access on server
#

echo Trying time server

DATE='/bin/su -c '/usr/bin/rsh time-server date' remote-user'

echo Setting date string...

/bin/date --set="DATE"
```

A more reliable way of keeping clocks synchronized, which works both for Unix and for NT, is the use the NTP protocol, or Network Time Protocol. A time-server is used for this purpose. The network time protocol daemon `xntpd` is used to synchronize clocks from a reliable time server. Two configuration files are needed to set up this service on a Unix-like host: `/etc/ntp.conf` and `/etc/ntp.drift`. `/etc/ntp.conf` looks something like this, where the IP address is that of the master time server, whose clock we trust:

```
driftfile /etc/ntp.drift
authdelay 0.000047
server 128.39.89.10
```

The `/etc/ntp.drift` file must exist, but its contents are undetermined. Commercial and shareware NTP clients are available for virtually all operating systems [47].

7.3 Executing Jobs at Regular Times

The ability of a host to execute jobs at predetermined times lies at the heart of keeping control over a changing, dynamical system.

7.3.1 The Unix cron Service

Unix has a time daemon called cron: it's chronometer. Cron reads a configuration file called a crontab file which contains a list of shell-commands to execute at regular time intervals. On modern Unix-like systems, every user may create and edit a crontab file using the command

```
crontab -e
```

This command starts a text editor allowing the file to be edited. The contents of a user's crontab file may be listed at any time with the command crontab -l. The format of a crontab file is a number of lines of the form

```
minutes 0-59 hours 0-23 day 1-31 month 1-12 weekday
  Mon-Sun Shellcommand
```

An asterisk or star * may be used as a wildcard, indicating 'any'. For example:

```
# Run script every weekday morning Mon-Fri at 3:15 am:

15 3 * * Mon-Fri /usr/local/bin/script
```

A typical root crontab file looks like this:

```
#
# The root crontab
#
0 2 * * 0,4 /etc/cron.d/logchecker
5 4 * * 6 /usr/lib/newsyslog
0 0 * * * /usr/local/bin/cfwrap /usr/local/bin/cfdaily
30 * * * * /usr/local/bin/cfwrap /usr/local/bin/cfhourly
```

The first line is executed at 2:00 a.m. on Sundays and Wednesdays, the second at 4:05 on Saturdays; the third is executed every night at 00:00 hours and the final line is executed one per hour on each half-hour.

In old BSD 4.3 Unix, it was only possible for the system administrator to edit the crontab file. In fact, there was only a single crontab file for all users, called /usr/lib/crontab or /etc/crontab. This contained an extra field, namely the user name under which the command was to be executed. This type of crontab file is largely obsolete now, but may still be found on some older BSD 4.3 derivatives such as DEC's ULTRIX.

```
0,15,30,45 * * * *root /usr/lib/atrun
00 4 * * *         root /bin/sh /usr/local/sbin/daily    2> 1 | mail root
30 3 * * 6         root /bin/sh /usr/local/sbin/weekly   2> 1 | mail root
30 5 1 * *         root /bin/sh /usr/local/sbin/monthly  2> 1 | mail root
```

A related service under Unix is the at command. This executes specific batch processes once only at a specific time. The at command does not use a configuration file, but a command interface. On some Unix-like systems, at is merely a front-end which is handled by a cron-scheduled program called atrun.

Suggestion 10 (Cron management) *Maintaining cron files on every host individually is awkward. We can use cfengine as a front-end to cron, to give us a global view of the task list (see section 7.4.4).*

7.3.2 NT Schedule Service

The NT scheduling service is similar to `cron`, except that the concept running multiple user processes simultaneously is a stumbling block for NT, since it is not a true multi-user operating system. By default, only the Administrator has access to the scheduling service. All jobs started by the scheduling service are executed with the same user rights as the user who runs the service itself (normally the Administrator). Some commercial replacements for the schedule service exist, but these naurally add extra cost to the system.

When the scheduling service is running, the `at` command provides a user interface for managing the queue of tasks to execute. The scheduling service is coordinated for all hosts in a domain by the domain server, so the host name on which a batch job is to run can be an argument to the scheduling command.

To schedule a new job to be executed either once or many times in the future, we write:

```
at host time command
```

The host argument is optional and refers to the host's NT name, not its DNS name (hopefully the two cooincide, up to a domain name). The time argument is written in the form `3:00pm` or `15:00`, etc. It may be followed by a qualifier which specifies the day, date and/or how many times the job is to be scheduled. The qualifier is a comma separated list of days or dates. `/next` means execute once for each date or day in the following list. `/every` means execute every time one of the items in the following list matches. For example:

```
at 3:00pm /next Friday,13 C:\crsite\host \local \myscript
```

would execute `myscript` at 15:00 hours on the first Friday following the date on which the command was typed in, and then again on the first 13th of the month following the date at which the command was typed in. It does not mean execute on the first coming Friday 13th. The items in the list are not combined logically with AND as one might be tempted to believe. The `at` command without any arguments lists the active jobs, like its Unix counterpart. Here one finds that each job has its own identity number. This can be used to cancel the job with a command of the form:

```
at ID /delete
```

7.4 Automation

The need for automation has become progressively clearer, as sites grow and the complexity of administration increases. Some advocates have gone in for a distributed object model [130, 257, 61]. Others have criticized a reliance on network services [35].

7.4.1 Tools for Automation

Most system administration tools developed and sold today (insofar as they exist) are based either on the idea of *control interfaces* (interaction between administrator and machine to make manual changes) or on the cloning of existing reference systems (mirroring) [143]. One sees user graphical user interfaces of increasing complexity, but seldom any serious attention to autonomous behaviour.

Many ideas for automating system administration have been reported [110, 86, 152, 163, 15, 162, 10, 88, 218, 85, 61, 217, 210, 54, 191, 183, 67, 117, 219, 146]. Most of these have been ways of generating or distributing simple shell or Perl scripts. Some provide ways of cloning machines by distributing files and binaries from a central repository. In spite of the creative effort spent developing the above systems, few if any of them can survive in their present form in the future. As indicated by Evard [81], analysing many case studies, what is needed is a greater level of abstraction. Although developed independently, cfengine [30, 32, 37] satisfies Evard's requirements quite well.

Vendors have also built many system administration products. Their main focus in commerical system administration solutions has been the development of man-machine interfaces for system management. A selection of these projects are described below. They are mainly control-based systems which give responsibility to humans, but some can be used to implement partial immunity type schemes by instructing hosts to execute automatic scripts. However, they are not comparable to cfengine in their treatment of automation; they are essentially management frameworks which can be used to activate scripts.

Tivoli [257] is probably the most advanced and wide-ranging product available. It is a Local Area Network (LAN) management tool based on CORBA and X/Open standards; it is a commercial product, advertised as a complete management system to aid in both the logistics of network management and an array of configuration issues. As with most commercial system administration tools, it addresses the problems of system administration from the viewpoint of the business community, rather than the engineering or scientific community. Tivoli admits bidirectional communication between the various elements of a management system. In other words, feedback methods could be developed using this system. The apparent drawback of the system is its focus on application level software rather than core system integrity. Also, it lacks abstraction methods for coping with with real-world variation in system setup.

Tivoli's strength is in its comprehensive approach to management. It relies on encrypted communications and client-server inter-relationships to provide functionality including software distribution and script execution. Tivoli can activate scripts, but the scripts themselves are a weak link. No special tools are provided here; the programs are essentially shell scripts with all of the usual problems. Client-server reliance could also be a problem: what happens if network communications are prevented?

Tivoli provides a variety of ways for activating scripts, rather like cfengine:

- Execute by hand when required.
- Schedule tasks with a cron-like feature.
- Execute an action (run a task on a set of hosts, copy a package out) in response to an event.

Tivoli's Enterprise Console includes a language Prolog for attaching actions to events. Tivoli is clearly impressive but also complex. This might also be a weakness. It requires a considerable infrastructure in order to operate, an infrastructure which is vulnerable to attack.

HP OpenView [192] is a commercial product based on SNMP network control protocols. OpenView aims to provide a common configuration management system for printers, network devices, Windows and HPUX systems. From a central location, configuration data may be sent over the local area network using the SNMP protocol The advantage of OpenView is

a consistent approach to the management of network services; its principal disadvantage, in the opinion of the author, is that the use of network communication opens the system to possible attack from hacker activity. Moreover, the communication is only used to alert a central administrator about perceived problems. No automatic repair is performed, and thus the human administrator is simply overworked by the system.

Sun's Solstice [180] system is a series of shell scripts with a graphical user interface which assists the administrator of a centralized LAN, consisting of Solaris machines, to initially configure the sharing of printers, disks and other network resources. The system is basically old in concept, but it is moving towards the ideas in HP OpenView.

Host Factory [83] is a third party software system, using a database combined with a revision control system [261] which keeps master versions of files for the purpose of distribution across a LAN. Host Factory attempts to keep track of changes in individual systems using a method of revision control. A typical Unix system might consist of thousands of files comprising software and data. All of the files (except for user data) are registered in a database and given a version number. If a host deviates from its registered version, then replacement files can be copied from the database. This behaviour hints at the idea of an immune system, but the heavy-handed replacement of files with preconditioned images lacks the subtlety required to be flexible and effective in real networks. The blanket copying of files from a master source can often be a dangerous procedure. Host Factory could conceivably be combined with cfengine in order to simplify a number of the practical tasks associated with system configuration and introduce more subtlety into the way changes are made. Currently, Host Factory uses shell and Perl scripts to customize master files where they cannot be used as direct images. Although this limited amount of customization is possible, Host Factory remains essentially an elaborate cloning system. Similar ideas for tracking network heterogeneity from a database model were discussed in refs. [260, 255, 190].

In recent years, the GNU/Linux community has been engaged in an effort to make GNU/Linux (indeed Unix) more user-friendly by developing any number of graphical user interfaces for the system administrator and user alike. These tools offer no particular innovation other than the novelty of a more attractive work environment. Most of the tools are aimed at configuring a single stand-alone host, perhaps attached to a network. Recently, two projects have been initiated to tackle clusters of Linux workstations [41, 209]. A GUI for heterogeneous management was described in ref. [200].

7.4.2 Monitoring Tools

Monitoring tools have been in proliferation for several years [116, 238, 150, 114, 122, 193, 223, 113]. They usually work by having a daemon collect some basic auditing information, setting a limit on a given parameter and raising an alarm if the value exceeds acceptable parameters. Alarms might be sent by mail, they might be routed to a GUI display or they may even be routed to a system admin's pager [113].

Network monitoring advocates have done a substantial amount of work in perfecting techniques for the capture and decoding of network protocols. Programs such as `etherfind`, `snoop`, `tcpdump` and bro [196], as well as commercial solutions such as Network Flight Recorder [75], place computers in 'promiscuous mode' allowing them to follow the passing data-stream closely. The thrust of the effort here has been in collecting data [7], rather than analysing them in any depth. The monitoring school advocates storing the huge

amounts of data on removable media such as CD to be examined by humans at a later date if attacks should be uncovered. The analysis of data is not a task for humans, however. The level of detail is more than any human can digest, and the rate of its production and the attention-span and continuity required are inhuman. Rather, we should be looking at ways in which machine analysis and pattern detection could be employed to perform this analysis – and not merely after the fact. In the future, adaptive neural nets and semantic detection will be used to analyse these logs in real time, avoiding the need to even store the data.

Unfortunately, there is currently no way of capturing the details of every action performed by the local host, analogous to promiscuous network monitoring without drowning the host in excessive auditing. The best one can do currently is to watch system logs for conspicuous error messages. Programs like SWATCH [113] perform this task. Another approach which we have been experimenting with at Oslo college is the analysis of system logs at a statistical level. Rather than looking for individual occurrences of log message, one looks for patterns of logging behaviour. The idea is that logging behaviour reflects (albeit imperfectly) the state of the host [74].

Visualization is now being recognized as an important tool in understanding the behaviour of network systems [57, 135, 100]. This reinforces the importance of investing in a documentable understanding of host behaviour, rather than merely relating experiences and beliefs [36]. Network traffic analysis has been considered in refs. [11, 278, 189].

7.4.3 Formalizing System Policy: cfengine

Experience indicates that simple rules are always preferable, though this is so far unsubstantiated by any specific studies. A rule which says "If the user concerned has consumed more than *X* megabytes of diskspace and it is not Friday (when there is little activity) and his/her boss has not said anythings in advance, and..." may seem smart, but most users will immediately write it off as being stupid. It is ill-advised because it is opaque.

Principle 33 (Simplest is best) *Simple rules make system behaviour easy to understand. Users tolerate rules if they understand them.*

7.4.4 Using `cfengine` as a Front-end to `cron`

One of cfengine's strengths is its use of classes to identify systems from a single file or set of files. Distributed resource administration would be much easier if the `cron` daemon also worked in this way. One way of setting this up is to use cfengine's time classes to work like a user interface for `cron`. This allows us to have a single, central file which contains all the `cron` jobs for the whole network without losing any of the fine control which cron affords us. All of the usual advantages apply:

- It is easier to keep track of what cron jobs are running on the system when they are all registered in one place.
- Groups and user-defined classes can be used to identify which host should run which programs.

The central idea behind this scheme is to set up a regular cron job on every system which executes cfengine at frequent intervals. Each time cfengine is started, it evaluates time classes

and executes the shell commands defined in its configuration file. In this way, we use cfengine as a wrapper for the cron scripts, so that we can use cfengine's classes to control jobs for multiple hosts. Cfengine's time classes are at least as powerful as `cron`'s time specification possibilities, so this does not restrict us in any way. The only price is the overhead of parsing the cfengine configuration file.

To be more concrete, imagine installing the following `crontab` file onto every host on the network:

```
#
# Global Cron file
#
0,15,30,45 * * * * /usr/local/cfengine/inputs/run-cfengine
```

This file contains just a single cron job, namely a script which calls cfengine. Here we are assuming that it will not be necessary to execute any cron script more often than every fifteen minutes. If this is too restrictive, the above can be changed. We refer to the time interval between runs of the script `run-cfengine` as the 'scheduling interval', and discuss its implications in more detail below.

The script `run-cfengine` can as simple as this:

```
#!/bin/sh
#
# Script run-cfengine

export CFINPUTS=/usr/local/cfengine/inputs

/usr/local/gnu/bin/cfengine

#
# Should we pipe mail to a special user?
#
```

or it could be more fancy. We could also use the `cfwrap` script to pipe mail to the mail address described in the cfengine file.

```
#
# Global Cron file
#
0,15,30,45 * * * * path/cfwrap  path/run-cfengine
```

Time Classes

Each time cfengine is run, it reads the system clock and defines the following classes based on the time and date:

- Yr*xxx*: the current year, e.g. `Yr1997`, `Yr2001`. This class is probably not useful very often, but it might help us to turn on the new-year lights, or shine up your systems for the new millennium (1st Jan 2001)!
- *Month*:: the current month can be used for defining very long term variations in the system configuration, e.g. `January`, `February`. These classes could be used to determine when students have their summer vacation, for instance, to perform extra tidying, or to specially maintain some administrative policy for the duration of a conference.

- *Day*:: the day of the week may be used as a class, e.g. `Monday`, `Sunday`.
- Day*xx*:: a day in the month (date) may be used to single out by date, e.g. the first day of each month defines `Day1`, the 21st `Day21`, etc.
- Hr*xx*:: an hour of the day, in 24-hour clock notation: `Hr00`...`Hr23`.
- Min*xx*:: the precise minute a which cfengine was started: `Min00` ... `Min59`. This is probably not useful alone, but these values may be combined to define arbitrary intervals of time.
- Min*xx_xx*:: the five-minute interval in the hour at which cfengine was executed, in the form `Min0_5`, `Min5_10` .. `Min55_00`.

Time classes based on the precise minute at which cfengine started are unlikely to be useful, since it is improbable that we will want to ask cron to run cfengine every single minute of every day: there would be no time for anything to complete before it was started again. Moreover, many things could conspire to delay the precise time at which cfengine were started. The real purpose in being able to detect the precise start time is to define composite classes which refer to arbitrary intervals of time. To do this, we use the `group` or `classes` action to create an alias for a group of time values. Here are some creative examples:

```
classes: # synonym groups:
   LunchAndTeaBreaks = ( Hr12 Hr10 Hr15 )
   NightShift        = ( Hr22 Hr23 Hr00 Hr01 Hr02 Hr03 Hr04 Hr05 Hr06 )
   ConferenceDays    = ( Day26 Day27 Day29 Day30 )
   QuarterHours      = ( Min00 Min15 Min30 Min45 )
   TimeSlices        = ( Min01 Min02 Min03 Min33 Min34 Min35)
```

In these examples, the left-hand sides of the assignments are effectively the OR-ed result of the right-hand side. This if any classes in the parentheses are defined, the left-hand side class will become defined. This provides an excellent and readable way of pinpointing intervals of time within a program, without having to use `|` and `.` operators everywhere.

Choosing a Scheduling Interval

How often should we call a global cron script? There are several things to think about:

- How much fine control do we need? Running cron jobs once each hour is usually enough for most tasks, but we might need to exercise finer control for a few special tasks.
- Are we going to run the entire cfengine configuration file or a special lightweight file?
- System latency. How long will it take to load, parse and run the cfengine script?

Cfengine has an intelligent locking and timeout policy which should be sufficient to handle hanging shell commands from previous crons so that no overlap can take place.

7.4.5 A Generalized Scripting Language: Perl

Customization of the system requires us to write programs to perform special tasks. Perl was the first of a group of scripting languages including python, tcl and scheme, to gain

acceptance in the Unix world. It has since been ported to Windows operating systems. Perl programming has to some extent replaced shell programming as the Free Software *lingua franca* of system administration.

The Perl language (see Appendix) is a curious hybrid of C, Bourne shell and C-shell, together with a number of extra features which make it ideal for dealing with text files and databases. Since most system administration tasks deal with these issues; this places Perl squarely in the role of system programming. Perl is semi-compiled at runtime, rather than interpreted line-by-line like the shell, so it gains some of the advantages of compiled languages, such as syntax check before execution, and so on. This makes it a safer and more robust language. It is also portable (something which shell scripts are not [13]). Although introduced as a scripting language, like all languages, Perl has been used for all manner of things for which it was never intended. Scripting languages have arrived on the computing scene with an alacrity which makes them a favourable choice to anyone wanting to get code running quickly. This is naturally a mixed blessing. What makes Perl a winner over many other special languages is that it is simply too convenient to ignore, for a wide range of frequently required tasks. By adopting the programming idioms of well-known languages, as well as all the basic functions in the C library, Perl ingratiates itself to system administrators and becomes an essential tool. At the time of writing, only the GNU/Linux operating system bundles Perl as standard software. However, with a huge program base written in Perl, it is already indispensible.

7.5 Preventative Maintenance

In some countries doctors do not get paid if their patients get sick. This motivates them to practice preventative medicine, thus keeping the population healthy and functional at all times. A computer system which is healthy and functional is always equipped to perform the task it was intended for. A sick computer system is an expensive loss, in downtime and in human resources spent fixing the problem. It is surprising how effective a few simple measures can be toward stabilizing a system.

The key principle which we have to remember is that system behaviour is a social phenomenon, an interaction between users' habits and resource availabilty. In any social or biological system, survival is usually tied to the ability of the system to respond to threats. In biology we have immunity and repair systems; in society we have emergency services like fire, police, paramedics and the garbage collection service, combined with routines and policy ('the law'). We scarely notice these services until something goes wrong, but without them our society would quickly decline into chaos.

7.5.1 Policy Decisions

A policy of prevention requires system managers to make several important decisions. Let's return for a moment to the idea that users are the greatest danger to the stability of the system; we need to strike a balance between restricting their activities and allowing them freedom. Too many rules and restrictions leads to unrest and bad feelings, while too much freedom leads to anarchy. Finding a balance requires a policy decision to be made. The policy must be digested, understood and, not least, obeyed by users and system staff alike:

- *Determine the system policy.* This is the prerequisite for all system maintenance. Know what is right and wrong and know how to respond to a crisis. Again, as we have reiterated throughout, no policy can cover every eventuality, not should it be a substitute for thinking. A sensible policy will allow for sufficient flexibility (fault tolerance). A rigid policy is more likely to fail.
- *Sysadm team agreement*: the team of system administrators need to work together, not against one another. That means that everyone must agree on the policy and enforce it.
- *Expect the worst*: be prepared for system failure and for rules to be broken. Some kind of police service is required to keep an eye on the system. We can use a script, or an integrated approach like cfengine for this.
- *Educate users in good and bad practice*: ignorance is our worst enemy. If we educate users in good practice, we reduce the problem of policy trangressions to a few 'criminal' users, looking to try their luck. Most users are not evil, just uninformed.
- *Special users*: do some users require special attention, extra resources, or special assistance? An initial investment catering to their requirements can save time and effort in the long run.

7.5.2 General Provisions

Damage and loss can come in many forms: by hardware failure, resource exhaustion (full disks, excessive load), by security breaches and by accidental error. General provisions for prevention mean planning ahead in order to prevent loss, but also minimizing the effects of inevitable loss:

- Do not rely exclusively on service or support contracts with vendors. They can be unreliable and unhelpful, particularly in an organization with little economic weight. Vendor support helpdesks usually cannot diagnose problems over the phone, and a visit can take longer than is convenient, particularly if a larger customer also has a problem at the same time. Invest in local expertise.
- Educate users by posting information in a clear and friendly way.
- Make rules and structure as simple as possible, but no simpler.
- Keep valuable information about configuration securely, but readily available.
- Document all changes and make sure that co-workers know about them, so that the system will survive, even if the one who made the change is not available.
- Do not make changes just before going away on holiday: there are almost always consequences which need to be smoothed out.
- Be aware of system limitations, hardware and software, capacity. Do not rely on something to do a job it was not designed for.
- Work defensively and follow the pulse of the system. If something looks unusual, investigate and understand what is happening.
- Avoid gratuitous changes to things which already work adequately. 'If it ain't broke, don't fix it'.

- Redundancy is our last avenue of survival [242]. Duplication of service and data gives us a fallback which can be brought to bear in a crisis.

Vendors often like to pressure sites into signing expensive service contracts. Modern computer hardware is pretty reliable these days, the only exception being cheap PC kits. For the cost of a service contract it might be possible to buy several new machines each year, so one can ask the question: should we write off seldom hardware failure as acceptable loss, or pay the one-off repair bill? If one chooses this option, it is important to have another host which can step in and take over the role of the old one, while a replacement is being procured. Again, this is the principle of redundancy. The economics of service contracts need to be considered carefully.

7.5.3 Garbage Collection

Computer systems have no natural waste disposal system. If computers were biological life, they would have perished long ago, poisoned by their own waste. No system can continue to function without waste disposal. It is a thermodynamical impossibility to go on using resources forever, without releasing some of them again. That process must come to an end.

Garbage collection, in a computer system, refers to two things: disk files and processes. Users seldom clear garbage of their own accord, either because they are not really aware of it, or because they have an instinctive fear of throwing things away. Administrators have to enforce and usually automate garbage collection as a matter of policy. Cfengine can be used to automate this kind of garbage collection.

- *Disk tidying*: many users are not even aware that they are building up junk files. Junk files are often the by-product of running a particular program. Ordinary users will often not even understand all of the files which they accumulate, and will therefore be afraid to remove them. Moreover, few users are educated to think of their responsibilities as individuals to system community of all users, when it comes to computer systems. It does not occur to them that they are doing anything wrong by filling the disk with every bit of scrap they take a shine to.
- *Process management*: processes or running programs do not always complete in a timely fashion. Some buggy processes go amok and consume CPU cycles by executing infinite loops, others simply hang and fail to disappear. On multi-user systems, terminals sometimes fail to terminate their login processes properly, and will leave whole hierarchies of idle processes which do not go away by themselves. This leads to a gradual filling of the process table. In the end, the accumulation of such processes will prevent new programs from being started. Processes are killed with the `kill` command on Unix-like systems, or with the NT Resource Kit's `kill` command, or the graphical user interface.

7.5.4 Productivity or Throughput

Throughput is the how much real work actually gets done by a computer system. How efficiently is the system fulfilling its purpose or doing its job? The policy decisions we make can have an important bearing on this. For instance, we might think that the use of disk

quotas would be beneficial to the good of the system community, because then no user would be able to consume more than his or her fair share of disk space. However, this policy can be misguided. There are many instances (during compilation, for instance) where users have to create large temporary files which can later be removed. Rigid disk quotas can prevent a user from performing legitimate work; they can get in the way of the system throughput. Limiting users resources can have exactly the opposite effect of that which was intended.

Another example is in process management. Some jobs require large amounts of CPU time and take a long time to run: intensive calculations are an example of this. Conventional wisdom is to reduce the process priority of such jobs so that they do not interfere with other users' interactive activities. On Unix-like systems this means using the `nice` command to lower the priority of the process. However, this procedure can also be misguided. Lowering the priority of a process can lead to process *starvation*. Lowering the priority means that the heavy job will take even longer, and might never complete at all. An alternative stategy is to do the reverse: increasing the priority of a heavy task will get rid of it more quickly. The work will be finished and the system will be cleared of a demanding job, at the cost of some inconvenience for other users over a shorter period of time. We can summarize this in a principle:

Principle 34 (Resource restriction) *Restriction of resources can lead to poor performance and low productivity. Free access to resources prevents bottlenecks.*

But there is an obvious warning to be added:

Corollary 35 (Resource restriction) *With free access to resources, resource usage needs to be monitored to avoid the problem of runaway consumption, or the exploitation of those resources by malicious users.*

7.6 Fault Report and Diagnosis

While it is probably true that faults and problems are inevitable, a solid effort towards prevention can reduce them to a scarcity[1]. When problems arise, we need to develop a systematic approach to diagnosing the errors and getting the system on its feet again. As in the field of medicine, there is only a limited number of symptoms which a body or computer system can express (sore throat, headache, fever, system runs sluggishly, hangs, etc.). What makes diagnosis difficult is that virtually all ailments therefore lead to the same symptoms. Without further tests, it is thus virtually impossible to determine the cause of symptoms.

As mentioned in section , a distressing habit acquired from the home computer revolution is the tendency to give up before even attempting a diagnosis, and simply reboot the computer. This solves nothing and we learn nothing about why the problem arose. It is like killing a patient and replacing him with another. It is a habit which has to be unlearned in a multi-user community. The act of rebooting a computer can have unforseen effects on what

[1] This presumes, naturally, that the basic operating system software is sound. If it is not, then system crashes will be the rule rather than the exception.

other users are doing, disrupting their work and perhaps placing the security of data in jeopardy. Rather, we need to carefully examine the evidence on a process by process, and file by file basis.

7.6.1 Error Reporting

Reporting a health problem is the first step to recognizing its importance and solving it. Users tend to fall into the categories of active and passive users. Active users do not need encouraging to report problems. They will usually report even the smallest of problems; sometimes they will even determine the cause and report a fix. While they can often be wearisome in a stressful situation, power users of this type are our friends and go a long way towards spreading the burden of problem solving. Remember the community principle of delegation: if we cannot make good use of resources, then the community is not working.

Power users are sometimes more enthusiastic than they are experienced, however, so the system administrator's job is not simply to accept on trust what they say. Their claims need to be verified and perhaps improved upon. Sometimes, users' proposed solutions cannot be implemented because they are in conflict with the system policy, or because the solution would break something else. Only the system administrator has that kind of bird's-eye view of the system to make the judgement.

In contrast to active users, passive users normally have to be encouraged to report errors. They will fumble around trying to make something work, without understanding that there is necessarily a problem. Help desk systems such as Rust, Gnats, Nearnet, Netlog, PTS, QueueMH and REQ [222] can help in this way, but they also tend to encourage reports of problems which are only misunderstandings.

Suggestion 11 (FAQs) *Providing users with a roadmap for solving problems, starting with Frequently Asked Questions and ending with an error report, can help to rationalize error reporting.*

7.6.2 A Diagnostic Principle

Once an error has been reported, we must determine its cause. A good principle of diagnostics comes from an old medical addage:

Principle 36 (Diagnostics) *When you hear the sound of distant hooves, think horses not zebras, i.e. always eliminate the obvious first.*

What this means is that we should always look for the most likely explanation before toying with exotic ideas. It is embarassing to admit how many times apparently impossible problems have resulted from a cable coming out, or forgetting to put in a plug after being distracted in the middle of a job. If the screen is dark, is it plugged in, is the brightness turned up, is the picture centred? Power failures, loose connections, and accidentally touching an important switch can all confuse us. Since these kinds of accident are common, it is logical to begin here. Nothing is too simple or menial to check. A systematic approach, starting with simple things and progressing through the numbers, often makes light work of many problems. The urge to panic is often strong in novices, when there is no apparent explanation; with experience,

however, we can quell the desire to run for help. A few tests will almost always reveal a problem. Experience allows us to expand our repertoire and recognize clues, but there is no reason why cold logic should not bring us home in every case.

Having eliminated the obvious avenues of error, we are led into murkier waters of fault diagnosis. When a situation is confusing, it is of paramount importance to keep a clear head. Writing down a log of what we try and the effect it has on the problem prevents a forgetful mind from losing its way. Drawing a conceptual map of the problem, as a picture, is also a powerful way of persuading the human mind to do its magic.

One of the most powerful features of the human mind (the thing which makes it, by far, the most powerful pattern recognition agent in existence) is its ability to associate information input with conceptual models from previous experience. Even the most tenuous of connections can lead us to be amused at a likeness. We recognize human faces in clouds and old cars; we recognize a song from just a few notes. The ability to make connections leads us in circles of thought which sooner or later lead to 'inspiration'. As most professionals know, however, inspiration is seldom worth waiting for. A competent person knows how to work through these mental contortions systematically to come up with the same answer. While this might be a less romantic notion than waiting for inspired enlightenment, it is usually more efficient.

7.6.3 Establishing Cause and Effect

If a problem has arisen, then something in the system is different than it was before the error occurred. Our task then is to determine the source of that change, and identify a chain of events which resulted in the unfortunate effect. The hope is that this will tell us whether we can prevent the problem from recurring, and perhaps also whether we can fix it. It is not merely so that we can fill out a report in triplicate that we need to debug errors.

Problem diagnosis is one of the hardest problems in any field, be it system administration, medicine or anything else. Once a cause has been found, a cure can be simple, but finding the problem itself often requires experience, a large knowledge base and an active imagination. There is a three stage process:

- Gather evidence from users and from other tests.
- Make an informed guess as to probable cause.
- Try to reproduce (or perhaps just fix) the error.

There is no particular order in which these pieces of the puzzle must be executed. Normally, they will all be repeated until a satsifactory explanation has been uncovered. It is only when we have shown that a particular change can switch the error on or off that we can say with certainty what the cause of the error was.

Sometimes it is not possible to directly identify the causal chain which led to an error with certainty. Trying to reproduce a problem on an unimportant host is one way of verifying a theory, but this will not always work. Computers are complex systems which are affected by the behaviour of users, interactions between subsystems, network traffic, and any combination of these things. Any one of these factors can have changed in the meantime. Sometimes it can be a chance event which creates a unique set of conditions for an error to occur[2].

[2] I tend to classify all such inexplicable occurrances under the heading 'cosmic ray'.

Usually this is not the case, though; most problems are reproducible with sufficient time and imagination.

Trying to establish probable cause in such a web of intrigue as a computer system is enough to task the best detectives. Indeed, we shall return to this point in Chapter 11, and consider the nature of the problems in more detail. To employ a tried and tested stategy, in the spirit of Sherlock Holmes, we can gradually eliminate possibilities and therefore isolate the problem, little by little. This requires a certain inspiration for hypothesizing causes which can be found from any number of sources:

- One should pay attention to all the facts available about the problem. If users have reported it, then one should take seriously what they have to say, but always attempt to verify the facts before taking too much on trust.
- Reading documentation can sometimes reveal simple misunderstandings in configuration which would lead to the problem.
- Talking to others who might have seen the problem before can provide a short cut to the truth. They might have done the hard work of diagnosis before. Again, their solutions need to be verified before taking them on trust.
- Reading old bug and problem reports can provide important clues.
- Examining system log files will sometimes provide answers.
- Performing simple tests and experiments, based on a best guess scanario, sharpens the perception of the problem, and can even allow cause to be pinpointed.
- If the system is merely running slower than it should, then some part of it is struggling to allocate resources. Is the disk nearing full, or the memory, or even the process table? Entertain the idea that it is choking in garbage. For instance, deleted files take up space on systems like Novell, since the files are stored in such a way that they can be undeleted. One needs to purge the file system every so often to remove these, otherwise the system will spend much longer than it should looking for free blocks. Unix systems thrash when processes build up to unreasonable levels. Garbage collection is a powerful tool in system maintenance. Imagine how human health would suffer if we could never relieve ourselves of dead cells or the biproducts of a healthy consumption. All machines need to do this.

7.6.4 Gathering Evidence

From best guess to verification of fault can be a puzzling time in which one grapples with the possible explanations and seeks tests which can confirm or deny their plausibility. One could easily write a whole book exemplifying techniques for troubleshooting, but that would take us beyond the limits set for this book. Let us just provide two examples of real cases which help to illustrate how the process of detection can proceed.

- *Network services become unavailable*: a common scenario is the sudden disappearence of a network service, like, say, the WWW. If a network service fails to respond it can only be due to a few possibilties:
  - The service has died on the server host.

– The line of communication has been broken.

– The latency of the connection is so long that the service has timed-out.

A natural first step is to try to send a sonar ping to server-host:

```
ping www.domain.country
```

to see whether it is alive. A ping signal will normally return with an answer within a couple of seconds, even for a machine halfway across the planet. If the request responds with

```
www.domain.country is alive
```

then we know immediately that there is an active line of communication between the our host and the server hosts, and we can eliminate the second possibility. If the ping request does not return, then there are two further possibilities:

– The line of communication is broken.

– The DNS lookup service is not responding.

The DNS service can hang a request for a long period of time if a DNS server is not responding. A simple way to check whether the DNS server is at fault or not is to bypass it, by typing the IP address of the WWW server directly:

```
ping 128.39.74.4
```

If this fails to respond then we know that the fault was not primarily due to the name service. It tends to suggest a broken line of communication. The `traceroute` command on Unix-like operating systems, or `tracert` on NT, can be used to follow a net connection through various routers to its destination. This often allows us to narrow down the point of failure to a particular group of cables in the network. If a network break has persisted for more than a few minutes, a ping or traceroute will normally respond with the message

```
ICMP error: No route to host
```

and this tells us immediately that there is a network connectivity problem.

But what if there is no DNS problem and ping tells us that the host is alive? Then the natural next step is to verify that the WWW service is actually running on the server host. On a Unix-like OS we can simply log onto the server host (assuming it is ours) and check the process table for the `httpd` daemon which mediates the WWW service.

```
ps aux | grep httpd BSD
ps -ef | grep httpd Sys V
```

On an NT machine, we would have to go to the host physically and check its status. If the WWW service is not running, then we would like to know why it stopped working. Checking log files to see what the server was doing when it stopped working can provide clues or even an answer. Sometimes a server will die because of a bug in the program. It is a simple matter to start the service again. If it starts and seems to work normally afterwards, then the problem was almost certainly a bug in the program. If the service fails to start, then it will log an error message of some kind which will tell us more. One possibility is that someone has changed something in the WWW service's

configuration file and has left an error behind. The server can no longer make sense of its configuration and it gives up. The error can be rectified and the server can be restarted.

What if the server process has not died? What if we cannot even log onto the server host? The latter would be a clear indication that there was something more fundamentally wrong with the server host. Resisting the temptation to simply reboot it, we could then try to test other services on the server host to see if they respond. We already know that the ping service is responding, so the host is not completely dead. There are therefore several things which could be wrong:

- The host is unable to respond (e.g. it is overloaded).
- The host is unwilling to respond (e.g. a security check denying access to our host).

We can check that the host is overloaded by looking at the process table, to see what is running. If there is nothing to see there, the host might be undergoing a denial of service attack (see Chapter 9). A look at `netstat` will show how many external connections are directed towards to host and their nature. This might show something that would confirm or deny the attack theory. An effective attack would be difficult to prevent, so this could be the end of the line for this particular investigation and the start of a new one, to determine the attacker. If there is no attack, we could check that the DNS name service is working on the server-host. This could cause the server to hang for long periods of time. Finally, there are lots of reasons why the kernel itself might prevent the server from working correctly: the TCP connection close time in the kernel might be too long, leading to blocked connections; the kernel itself might have gone amok; a full disk might be causing errors which have a knock-on effect (the log files from the server might have filled up the disk), in which case the disk problem will have to be solved first. Notice how the DNS and disk problems are problems of *dependency*: a problem in one service having a knock-on effect in another.

- *Disks suddenly become full*: a second example, with a slightly surprising conclusion, begins with an error message from a program telling us that the system disk of a particular host has become full. The nature of this particular problem is not particularly ambiguous. A full disk is a disk with no space left on it. Our aim is to try to clear enough space to get the system working again, at least until a more permanent solution can be found. To do this, we need to know why the disk became full. Was it for legtimate reasons, or because of a lack of preventative garbage collection, or in this case a completely different reason? There are many reasons why a disk partition might become full. Here are some obvious ones:

  - A user disk partition can become full if users download huge amounts of data from the Internet, or if they generate large amounts of temporary files. User disks can become full both for valid reasons and for mischievous reasons.
  - A system disk does not normally change but for two reasons: log files which record system activity can grow and fill up a disk; temporary files written to public directories can grow and fill a disk.

  If a user disk becomes full, it is usually possible to find some unnecessary files which can be deleted to make space temporarily. The files we deem as unnecessary have to be defined as such as a matter of *policy*. It would be questionable ethically to make a habit if

deleting files which users did not know could be removed, in advance. Some administrators follow the practice of keeping a large file on every disk partition which can be removed to make space. This is a somewhat strange practice which only delays the inevitable. Of course, if we have done our preventative maintenance, then there should not be any junk files taking up space on the system. In the end, all user disks grow monotonically and new disks have to be bought, users can be moved to new disks to spread the load, and so on.

If a system disk becomes full, there are only three things to look for:

- `core` files (Unix).
- Temporary files.
- Log files.

Core files are image files which are dumped when programs crash. They are meant to be used for debugging purposes; in practice they cause more problems than they solve. Core files are very large and one or two can easily fill a tight partition. Preventative maintenance should delete such files regularly. Temporary files `/tmp` and `/var/tmp` in Unix-like systems, or `C:\Temp` + on NT are publicly writeable directories which usually take up space on the system disk. These can be filled up either accidentally or maliciously. Again, these should be cleared regularly. The final source of trouble is log files. Log files need to be *rotated* on a regular basis so that they do not grow too large. Rotation means starting a new log and saving a small number of old log files. This means that old log data eventually get thrown away, rather than keeping it forever.

In all of the above cases, we can identify the recent change in a file system by searching for files which have changed in the last 24 hours. On a Unix-like system, this is easily done by running a command to look at all subdirectories of the current directory:

```
find . -mtime -1 -print -xdev
```

On other systems it is harder and requires special software. A GNU version of the Unix `find` utility is available for NT.

A third reason why a file system can become full is *corruption*. In one instance a Unix disk continued to grow, despite verifying that no new files had been created and after removing all old log files. The Unix `df` disk utility eventually reported that the file system was 130% full (an impossibility) and it continued to grow. The eventual cause of this problem was identified as a fault in the file system structure, or inode corruption. This was brought about by the host concerned overheating and causing memory errors (system log errors confirmed memory write errors). The problem recurred twice before the host was moved to a cooler environment, after which time it righted itself (though the file system had to be repaired with `fsck` on each occasion).

There are many tips for tracing the activity of programs. For instance, to trace what files are read by a program, use `strace` or `truss` to for watch file descriptors

```
truss -t open,close program
```

This runs the program concerned in a monitor which prints out all the listed system calls. This can be a good way of finding out which libraries a program uses (or tries and fails to use), or which configuration files it opens.

Complete your own list of troubleshooting tips. This is a list you will be building for the rest of your life.

7.7 System Performance Tuning

When is a fault not a fault? When it is an inefficiency. Sooner or later, user perception of system performance passes a threshold. Beyond that threshold we deem the performance of a computer to be unacceptably slow and we become irritated. Long before that happens, the system itself recognizes the symptoms of a lack of resources and takes action to try to counter the problem, but not always in the way we would like.

Efficiency and users' perception of efficiency are usually two separate things. The host operating system itself can be timesharing perfectly and performing real work at a break-necked pace, while one user sits and waits for minutes for something as simple as a window to refresh. For anyone who has been in this situation, it is painfully obvious that system performance is a highly subjective issue. If we aim to please one type of user, another will be disappointed. To extract maximal performance from a host, we must focus on specific issues and make particular compromises. Note that the system itself is already well adjusted to share resources: that is what a kernel is designed to do. The point of performance tuning is that what is good for one task is not necessarily good for another. Generic kernel configurations try to walk the line of being adequate for everyone, and in doing so they are not great at doing any of them in particular. The only way we can truly achieve maximal performance is to specialize. Ideally, we would have one host per task and optimize each host for that one task. Of course this is a huge waste of resources, which is why multi-tasking operating systems exist. The inevitability of sharing resources between many tasks is to strike compromise. This is the paradox of multi-tasking.

Whole books have been written on the subject of performance tuning, so we shall hardly be able to explore all of the avenues of the topic in a brief account. See, for instance, refs. [132, 71, 169, 265, 11, 275, 251, 227]. Our modest aim in this book is, as usual, to extract the essense of the topic, pointing fingers at the key performance bottlenecks. If we are to tune a system, we need to identify what it is we wish to optimize, i.e. what is most important to us. We cannot make everything optimal, so we must pick out a few things which are most important to us, and work on those.

System performance tuning is a complex subject, in which no part of the system is sacrosanct. Although it is quite easy to pin-point general performance problems, it is harder to make general recommendations to fix these. Most details are unique to each operating system. A few generic pointers can nonetheless offer the greatest and most obvious gains, while the tweaking of system-dependent parameters will put the icing on the cake.

To identify a problem, we must first measure the performance. Again, there are the two issues: *user perception* of performance (interactive response time); and *system throughput*, and we have to choose the criterion we wish to meet. When the system is running slowly, it is natural to look at what resources are being tested, i.e.

- What processes are running.
- How much available memory the system has.
- Whether disks are being used excessively.

- Whether the network is being used heavily.
- What software dependencies does the system have (e.g. DNS, NFS)?

The last point is easy to overlook. If we make one host dependent on another then the dependant host will always be limited by the host upon which it depends. This is particularly true of file servers (e.g. NFS, DFS, Netware distributed file systems) and of the DNS service.

Principle 37 (Symptoms and cause) *Always try to fix problems at the source, rather than patching symptoms.*

7.7.1 Resources and Dependencies

Since all resources are scheduled by processes, it is natural to check the process table first and then look at resource usage. On NT, one has the process manager and performance monitor for this. On Unix-like systems, we check the process listing with `ps aux`, if a BSD compatible `ps` command exists, or `ps -efl` if the system is derived from System V. If the system has both, or a BSD compatible output mode, as in Solaris and Digital Unix (OSF1), for instance, then the BSD style output is recommended. This provides more useful information and orders the processes so that the heaviest process comes at the top. This saves time. Another useful Unix tool is `top`. A BSD process listing looks like this:

```
nexus % ps aux | more
USER  PID   %CPU %MEM SZ   RSS  TT    S START    TIME   COMMAND
root  3     0.2  0.0  0    0    ?     S Jun 15   55:38  fsflush
root  22112 0.1  0.5  1464 1112 pts/2 O 15:39:54 0:00   ps aux
mark  22113 0.1  0.3  1144 720  pts/2 O 15:39:54 0:00   more
root  340   0.1  0.4  1792 968  ?     S Jun 15   3:13   fingerd
...
```

This one was taken on a quiet system, with no load. The columns show the user ID of the process, the process ID, an indication of the amount of CPU time used in executing the program (the percentage scale can be taken with a pinch of salt, since it means different things for different kernels). An indication of the amount of memory allocated. The `SZ` post is the size of the process in total (code plus data plus stack), while `RSS` is the resident size, or how much of the program code is actually resident in RAM, as opposed to being paged out, or never even loaded. The `TIME` shows the amount of CPU time accumulated by the process, while `START` indicates the amount of clock time which has elapsed since the process started. Problem processes are usually identified by:

- `%CPU` is large. A CPU intensive process, or a process which has gone into an endless loop.
- `TIME` is large. A program which has been CPU intensive, or which has been stuck in a loop for a long period.
- `%MEM` is large. `SZ` is large. A large and steadily growing value can indicate a memory leak.

One thing we notice is that the `ps` command itself uses quite a lot of resources. If the system is low on resources, running constant process monitoring is an expensive intrusion.

Unix-like systems also tell us about memory performance through the virtual memory statistics, e.g. the `vmstat` command. This command gives a different output on each operating system, but summarizes the amount of free memory as well as paging performance, etc. It can be used to get an idea of whether or not the system is paging a lot (a sign that memory is low). Another way of seeing this is to examine the amount of swap space which is in use:

OS	List virtual memory usage
AIX	`lsps -a`
HPUX	`swapinfo -t -a -m`
Digital Unix/OSF1	`swapon -s`
Solaris 1 or SunOS 3/4	`pstat -s`
Solaris 2 or SunOS 5	`swap -l`
GNU/Linux	`free`
NT	Performance manager

Excessive network traffic is also a cause of impaired performance. We should try to eliminate unnecessary network traffic whenever possible. Before any complex analysis of network resources is undertaken, we can make sure that we have covered the basics:

- Make sure that there is a DNS server on each large subnet to avoid sending unnecessary queries through a router. (On small subnets this would be overkill.)
- Make sure that the name servers themselves use the loop back address `127.0.0.1` as primary name server on Unix-like hosts, so that we do not cause collisions by having the name server talk to itself on the public network.
- Try to avoid distributed file accesses on a different subnet. This loads the router. If possible, file servers and clients should be on the same subnet.
- If we are running X-windows, make sure that each workstation uses has its `DISPLAY` variable set to `:0.0` rather than `hostname:0.0`, to avoid sending data out onto the network, only to come back to the same host.

Some operating systems have nice graphical tools for viewing network statistics, while others have only `netstat`, with its varying options. Collision statistics can be seen with `netstat -i` for Unix-like OSes or `netstat /S` on NT. DNS efficiency is an important consideration, since all hosts are more or less completely reliant on this service.

Measuring performance reliably, in a scientifically stringent fashion is a difficult problem (see Chapter 11), but *adequate* measurements can be made, for the purpose of improving efficiency, using the process tables and virtual memory statistics. If we see frantic activity in the virtual memory system, it means that we are suffering from a lack of resources, or that some process has run amok.

Once a problem is identified, we need a strategy for solving it. Performance tuning can involve everything from changing hardware to tweaking software:

- Optimizing choice of hardware.
- Optimizing chosen hardware.
- Optimizing kernel behaviour.

- Optimizing software configurations.
- (Optimizing service availability).

Hardware has *physical limitations*. For instance, the heads of a hard disk can only be in one place at a time. If we want to share a hard disk between two processes, the heads have to be moved around between two regions of the disk, back and forth. Moving the read heads over the disk platter is the slowest operation in disk access, and perhaps the computer as a whole, and unfortunately something we can do nothing about. It is a fundamental limitation. Moreover, to get the data from disk into RAM, it is necessary to interrupt processes and involve the kernel. Time spent executing kernel code is time not spent on executing user code, and so it is a performance burden. Resource sharing is about balancing *overheads*. We must look for the sources of overheads and try to minimize them, or mitigate their effects by cunning.

7.7.2 Hardware

The fundamental principle of any performance analysis, as HiFi buffs know only too well, is:

Principle 38 (Weakest link) *The performance of any system is limited by the weakest link amongst its components. System optimization should begin with the source. If performance is weak at the source, nothing which follows can make it better.*

Obviously, any effect which is introduced after the source will only reduce the performance in a chain of data handling. A later component cannot 'suck' the data out of the source faster than the source wants to deliver it. This tells us that the logical place to begin in with the system hardware. A corollary to this principle follows from a straightforward observation about hardware. As Scotty said, we cannot change the laws of physics:

Corollary 39 (Performance) *A system is limited by its slowest moving parts. Resources with slowly moving parts, like disks, CD-ROMs and tapes, transfer data slowly and delay the system. Resources which work purely with electronics, like RAM memory and CPU calculation, are quick, because electrons are light and move around quickly. However, electronic motion/communication over long distances takes much longer than communication over short distances (internally within a host) because of impedances and switching.*

Already, these principles tell us that RAM is one of the best investments we can make. Why? To avoid mechanical devices like disks as much as possible, we store things in RAM; to avoid sending unnecessary traffic over networks, we cache data in RAM. Hence RAM is the primary workhorse of any computer system. After we have exhausted the possibilities of RAM usage, we can go on to look at disk and network infrastructure.

- *Disks*: when assigning partitions to new disks, it pays to use the fastest disks for the data which are accessed most often, e.g. for user home directories. To improve disk performance, we can do two things: One is to buy faster disks, and the other is to use *parallism* to overcome the time it takes for physical motions to be executed. The mechanical problem which is inherent in disk drives is that the heads which read and write data have to move as a unit. If we need to collect two files concurrently, which lie spread all over the disk, this has to be done *serially*. *Disk striping* is a technique whereby file systems are

spread over several disks. By spreading files over several disks, we have several sets of disk heads which can seek independently of one another, and work in parallel. This does not increase the transfer rate, but it does lower seek times, and thus performance improvement can approach as much as N times with N disks. RAID technologies, employ striping techniques, and are widely available commercially. GNU/Linux also has RAID support. Spreading disks and files across multiple disk controllers will also increase parallelism.

- *Network:* to improve network performance, we need fast interfaces. All interfaces, whether they be Ethernet or some other technology, vary in quality and speed. This is particularly true in the PC world, where the number of competing products is huge. Network interfaces should not be trusted to give the performance they advertise. Some interfaces which are sold at 100 Mbits/sec, Fast Ethernet, manage little more than 40 Mbits/s. Some network interfaces have intelligent behaviour and try to detect the best available transmission rate. For instance, newer Sun machines use the *hme* Fast Ethernet interface. This has the ability to detect the best transmission protocol for the line a host is connected to. The best transmission type is 100 Mbits/sec, full duplex (simultaneous send and receive), but the interface will switch down to 10 Mbits/sec, half duplex (send or receive, one direction at a time) if it detects a problem. This can have a huge performance effect. One problem with auto-detection is that, if both ends of the connection have auto-detection, it can become a timing issue as to which speed we end up with. Sometimes it helps to try setting the rate explicitly, assuming that the network hardware supports that rate. There are other optimizations also for TCP/IP tuning, which we shall return to below. References [254, 270] are excellent references on this topic.

The sharing of resources between many users and processes is what networking is about. The competition for resources between several tasks leads to another performance issue.

Principle 40 (Contention/competition) *When two processes compete for a resource, performance can be dramatically reduced as the processes fight over the right to use the resource. This is called contention. The benefits of sharing have to be weighed against the pitfalls.*

Contention could almost be called a strategy in some situations, since there exist technologies for avoiding contention altogether. For example, Ethernet technology allows contention to take place, whereas Token Ring technology avoids it. We shall not go into the arguments for and against contention. Suffice it to say that many widely used technologies experience this problem.

- *Ethernet collisions*: Ethernet communication is like a television panel of politicians: many parties shouting at random, without waiting for others to finish. The Ethernet cable is a shared bus. When a host wishes to communicate with another host, it simply tries. If another host happens to be using the bus at that time, there is a collision and the host must try again at random until it is heard. This method naturally leads to contention for bandwidth. The system works quite well when traffic is low, but as the number of hosts competing for bandwidth increases, the probability of a collision increases in step. Contention can only be reduced by reducing the amount of traffic on the network segment. The illusion of many collisions can also be caused by incorrect wiring, or

incorrectly terminated cable, which leads to reflections. If collision rates are high, a wiring check might also be in order.

- *Disk thrashing*: thrashing[3] is a problem which occurs because of the slowness of disk head movements, compared to the speed of kernel time-sharing algorithms. If two processes attempt to take control of a resource simultaneously, the kernel and its device drivers attempt to minimize the motion of the heads by queueing requested blocks in a special order. The algorithms really try to make the disks traverse the disk platter uniformly, but the requests do not always come in a predictable or congenial order. The result is that the disk heads can be forced back and forth across the disk, driven by different processes and slowing the system to a virtual standstill. The time for disk heads to move is an eternity to the kernel, some hundreds of times slower than context switching times.

  An even worse situation can arise with the virtual memory system. If a host begins paging to disk, because it is low on memory, then there can be simultaneous contention both for memory and for disk. Imagine, for instance, that there are many processes, each loading files into memory, when there is no free RAM. To use RAM, some has to be freed by paging to disk; but the disk is already busy seeking files. To load a file, memory has to be freed, but memory can't be freed until the disk is free to page, this drags the heads to another partition, then back again... and so on. This nightmare brings the system to a virtual standstill as it fights both over free RAM and disk head placement. The system spends more time juggling its resources than it does performing real work, i.e. the overhead to work ratio blows up. The only cure for thrashing is to increase memory, or reduce the number of processes contending for resources.

A final point to mention in connection with disks is to do with standards. Disk transfer rates are limited by the protocols and hardware of the disk interfaces. This applies to the interfaces in the computer and to the interfaces in the disks. Most serious performance systems will use SCSI disks, for their speed (see section 2.7). However, there are many versions of the SCSI disk design. If we mix version numbers, the SCSI bus will default to the lowest common denominator, i.e. if we mix slow disks with fast disks, then all the disks will work slowly. If one needs to support legacy disks together with new disks, then it pays to collect like disks with a special host for each type, or alternatively, buy a second disk controller rather than to mix disks on the same controller.

7.7.3 Software Tuning and Kernel Configuration

It is true that software is constrained by the hardware on which is runs, but it is equally true that hardware can only follow the instructions it has received from software. If software asks hardware to be inefficient, hardware will be inefficient. Software introduces many inefficiencies of its own. Hardware and software tuning are inextricably intertwined.

Software performance tuning is a more complex problem than hardware performance tuning, simply because the options we have for tuning software depend upon what the software is, how it is written, and whether or not the designer made it easy for us to tune

[3] For non-native English speakers, note the difference between *thrash* and *trash*. Thrashing refers to a beating, or the futile fight for survival, e.g. when drowning.

its performance. Some software is designed to be stable rather than efficient. Efficiency is not a fundamental requirement; there are other priorities, such as simplicity and robustness.

In software the potential number of variables is much greater than in hardware tuning. Some software systems can be tuned individually. For instance, high availability server software such as WWW servers and SMTP (e-mail) servers be be tuned to handle traffic optimally for heavy loads. See, for instance, tips on tuning sendmail [156], and other general tuning tips [265, 169, 262].

More often than not, performance tuning is related to the availability or sharing of system resources. This requires tuning the system kernel. The most configurable piece of software on the system is the kernel. All Unix-like systems kernel parameters need to be altered and tuned. The most elegant approach to this is taken by Unix SVR4 and Solaris. Here, many kernel parameters can be set at run time using the kernel module configuration command `ndd`, while others can be configured in a single file `/etc/system`. The parameters in this file can be set with a reboot of the kernel, using the reconfigure flag

```
reboot -- -r
```

For instance, on a heavily loaded system which allows many users to run external logins, terminals, or X terminal software, we need to increase many of the default system parameters. The `maxusers` parameter (actually in most Unix-like systems) is used as a guide to estimating the size of many tables and limits on resources. Its default value is based on the amount of available RAM, so one should be careful about changing its value in Solaris, though other OSes are less intelligent. Solaris also has a separate parameter `pt\_cnt` for extending the number of virtual terminals (pty's). It is possible to run out if many users are logged in to the same host simultaneously. Many graphics intensive programs use shared memory in large blocks. The default limit for shared memory segments is only a megabyte, so it can be increased to optimize for intensive graphics use, but should not be increased on heavily loaded file servers, where memory for caching is more important. The file `/etc/system`, then looks like this:

```
set maxusers=100
set shmsys:shminfo_shmmax = 0x10000000
set pt_cnt=128
```

After a reboot, these parameters will be set. Some caution is needed in editing this file. If it is non-existent or unparsable, the host will not be able to boot (a questionable design feature). The `ndd` command in Solaris can be chosen to optimize its over-safe defaults set on TCP/IP connections.

For busy servers which handle many TCP connections, the time it takes an operating system to open and close connections is important. There is a limit on the number of available connections and open sockets (see Chapter 8); if finished socket connections are not purged quickly from the kernel tables, new connections cannot be opened in their place. On non-tuned hosts, used sockets can hang around for five minutes or longer, for example on a Solaris host. On a heavily loaded server this is unacceptable. The close time on sockets can be shortened to half a minute so as to allow newer sockets to be opened sooner. The parameters can be set when the system boots, or patched at any later time. The times are measured in milliseconds. See refs. [254, 270] for excellent discussions of these values.

```
/usr/sbin/ndd -set /dev/tcp tcp_keepalive_interval 900000
/usr/sbin/ndd -set /dev/tcp tcp_time_wait_interval 30000
```

Prior to Solaris 2.7 (SunOS 5.7), the latter line would have read:

```
/usr/sbin/ndd -set /dev/tcp tcp_close_wait_interval 30000
```

which illusrates the futility of documenting these fickle parameters in a static medium like a book. Note that setting these parameters to ultra-short values could cause file transmissions to be terminated incorrectly. This might lead to a corruption of data. On a web server, this is a nuisance for the client, but it is not mission-critical data. For security, longer close times are desirable to ensure correct closure of sockets. After setting these values, the network interface needs to be restarted by taking it down and up with `ifconfig`. Alternatively, the values can be configured in a start-up script which is executed before the interface is brought up at boot time.

Most Unix-like operating systems do not permit runtime configuration; new kernels have to be compiled and the values hard-coded into the kernel. This requires not just a reboot, but a recompilation of the kernel in order to make a change. This is not an optimal way to experiment with parameters. Modularity in kernel design can save us memory, since it means that static code does not have to take up valuable memory space. However, the downside of this is that modules take time to load from disk, on demand. Thus, a modular kernel can be slower than a statically compiled kernel. For frequently used hardware, static compilation is a must, since it eliminates the load-time for the module, at the expense of extra memory consumption.

The GNU/Linux system kernel is a modular kernel, which can load drivers for special hardware at run time, in order to remain small in the memory. When we build a kernel, we have the option to compile in modules statically (see section 4.7). Tips for Linux kernel configuration can readily be found by searching the Internet, so we shall not reproduce these tips here, where they would quickly become stale (see, for instance ref. [71]).

NT performance tuning can be undertaken by perusing the multitudinous screens in the graphical performance monitor and editing the values. For once, this useful tool is a standard part of the NT system.

7.7.4 Data Efficiency

Efficiency of storage and transmission depends upon the configuration parameters used to confure disks and networks, and also on the amount of traffic the devices see. We have already mentioned the problem of contention.

Some file system formatting programs on Unix-like systems allow us to reserve a certain percentage of disk space for privileged users. For instance, the default for BSD is to reserve 10% of the size of a partition for use by privileged processes only. The idea here is to prevent the operating system from choking, due to the activities of users. This practice goes back to the early times when disks were small and expensive and partition numbers were limited. Today, these limits are somewhat in appropriate. Of a one bigabyte disk, 10% is a huge amount of space, which many users could live happily with for many weeks. If we have partitioned a host so as to separate users from operating system, then there is no need to reserve space on user disks. Better to let users utilize the existing space until a real problem

occurs. Preventitve tidying helps to avoid full disks. Whether one regards this as maintenance or performance tuning is a moot point. The effect is to save us time and loss of resource availability. See section about making file systems.

Another issue with disk efficiency is the configuration of block sizes. This is a technical issue which one probably does not want to play with too liberally. Briefly, the standard unit of space which is allocated on a file system is a *block*. Blocks are quite large, usually around 8 kilobytes. Even if we allocate a file which is one byte long, it will be stored as a separate unit, in a block by itself, or in a *fragment*. Fragments are usually around 1 kilobyte. If we have many small files, this can clearly lead to a large wastage of space, and it might be prudent to decrease the file system block size. If, conversely, we deal with mostly large files, then the block size could be increased to improve transfer efficiency. The file system parameters can, in other words, be tuned to balance file size and transfer rate efficiency. Normally, the default settings a good compromise.

Tuning the network is a complex subject and few operating systems allow us to do it at all. Solaris' `ndd` command can be used to configure TCP/IP parameters which can lead to noticable performance improvements. See the excellent discussion in ref. []. As far as software tuning is concerned, we have few options. The time we wait for a service to reply to a query is called the *latency*. Latency clearly depends upon many factors, so it is difficult to pin down, but it is a useful concept since it reflects users' perceptions of performance. Network performance can degrade for a variety of reasons. Latency can increase as a result of network collisions, making traffic congested, and it can be increased due to server load, making the server slow to respond. Network latencies clearly increase with distance from the server: the more routers, switches and cables a signal has to travel through, the slower it will be. Our options are to reduce traffic congestion, increase server performance, or increase parallelism (if possible) with failover servers [111]. Some network services are multithreaded (using either light or heavyweight processes), and can be configured to spawn more server threads to handle a greater number of simultaneous connections (e.g. nfsd, httpd, cfd). If traffic congestion is not the problem, then too small a number of servers might help in expediting multiple connections (many multithreaded servers set limits on the number of threads allowed, so as not to run a machine into the ground in the event of spamming). These measures help to reduce the need for retransmission of TCP segments and timeouts on connection. Assuming that the network interface is working as fast as it can (see the previous section), a server will then respond as quickly as it can.

Exercises

Exercise 7.1 Discuss why system homogeneity is a desirable feature of network infrastructure models. How does homogeneity simplify the issues of configuration and maintenance? What limits have to be placed on homogeneity, i.e. why can't hosts all be exactly identical?

Exercise 7.2 Find out about process priorities. How are process priorities changed on the computer systems on your network? Formulate a policy for handling processes which load the system heavily. Should they be left alone, killed, re-scheduled, etc?

Exercise 7.3 Describe the process you would use to troubleshoot a slowly running host.

Exercise 7.4 Suppose you are performance tuning, trying to find out why one host is slower than another. Write a program which tests the efficiency of *CPU intensive work only.* Write programs which test the speed of *memory intensive work* and *disk intensive* work. Would comparing the time it takes to compile a program on the hosts be a good way of comparing them?

Exercise 7.5 Determine the network transmission speed on the servers on your network. Are they as high as possible? Do they have auto-detection? If not, how are they configured?

Exercise 7.6 Review the role of cfengine in system administration. What is it used for? What are its special strengths? Review also the role of Perl. What are its special strengths? Is there any overlap between Perl and cfengine? Do the languages compete or supplement one another?

Exercise 7.7 Collect and compile cfengine. Set up a simple cfengine script which you can build on. Make it run hourly.

Exercise 7.8 Why are Unix shell scripts not portable? Is Perl portable? How can cfengine help in the issue of script portability?

Exercise 7.9 Discuss the advantages of having all scripts which perform configuration and maintenance in one place, and of spreading them around the network on the hosts to which they apply.

Exercise 7.10 Discuss the ethical issues associated with garbage collection of files and processes. Is it right to delete users' files? How would a garbage collection policy at a research laboratory differ from a policy at a high-school?

Exercise 7.11 What is meant by an Ethernet collision? How might doubling the speed of all hosts on an Ethernet segment, make the total system slower?

Chapter 8

Services

In previous chapters we have spent a lot of time discussing what cooperation between hosts means for networked communities. Now it is time to address the practical issues of setting up basic services.

Network services are the crux of network cooperation (see section 3.3). They distinguish a cooperative network from a loose association of hosts. A community is bound together by a web of delegation and sharing. We give this job to *A* and that job to *B*, and they carry out their specialized tasks, making the whole function. In a computer network, we assign specific functions to specific hosts, thereby consolidating effort while distributing functionality.

The way in which services are handled by most operating systems is to use the socket abstraction. A socket is, loosely speaking, a file-like interface with an IP address plus a TCP or UDP port number [138], where some kind of data are communicated. A server has a listening socket which responds to client requests by opening a new temporary socket at a random port number. Information is exchanged and then any connection is terminated.

The system administrator has the task of organizing and configuring network services. That includes installing, planning and implementing the daemons which carry out the services.

For definiteness, the examples discussed in this chapter are based on Unix-like operating systems. In a network of Unix-like hosts, we have complete freedom to locate a server on whatever host we wish. Although some services (e.g. remote login) run on every host, most are confined to one or two hosts, whose special function it is to perform the tasks on behalf of the network community.

8.1 High Level Services

Internet networks use many high level protocols to provide the distributed services which most users take for granted:

- FTP. The File Transfer Protocol. Passwords are sent in clear text.
- HTTP. The hypertext transfer protocol for the transmission of data on the world wide web. All data are sent in clear text.
- S-HTTP is a superset of HTTP, which allows messages to be encapsulated for increased security. Encapsulations include encryption, signing and MAC-based authentication. An S-HTTP message can have several security transformations applied to it. S-HTTP also

includes support for key transfer certificate transfer and similar administrative functions. It is generally regarded as being superior to HTTPS, a pure SSL encryption.

- HTTPS. The secure world wide web protocol for exchanging hypertext and multimedia data. All data are encrypting using Netscape's Secure Socket Layer (SSL).
- SSH. The secure shell. A replacement for the remote shell (rsh) Unix protocol. The secure shell provides full encryption and forwarding of X11 display data through a secure pipe.
- LDAP. The *Lightweight Directory Access Protocol* is a generalized protocol for looking up data in simple databases. It is a lightweight version of the Director Access Protocol originally written for X.500, and is currently at Version 3. LDAP can be used to register user information, passwords, telephone numbers, etc., and interfaces through gateways to the NDS (Novell Directory Service), Microsoft's Exchange server and NIS (Sun's Network Information Service). The advantage of LDAP will be a uniform protocol for accessing table lookups. Currently, the spread of LDAP is hindered by few up-to-date implementations of the protocol.
- NTP is the Network Time Protocol, used for synchronizing clocks throughout the network.
- IMAP. *Internet Mail Access Protocol* provides a modern mailbox format and a number of network services for reading and transferring mail over the network.
- RPC (Remote Procedure Call) is not one service, but a whole group of services with a common mode of communication. It is an extra layer of functionality built on top of TCP/IP sockets. Normally when we speak of RPC, we are referring to Sun's RPC, which has been widely adopted on Unix-like systems, but there are several methods for remote procedure call. A whole class of daemons make use of RPC services. These are often (but not always) prefixed by 'rpc', e.g. `rpc.mountd`, `rpcbind`. RPC services are assigned port numbers dynamically by a service known as `portmap`, or `rpcbind`. If the portmapper fails or dies, all RPC services have to be restarted from scratch.

There is an almost endless list of services which are registered by the `/etc/services` file. These named services perform a wide range of functions.

8.2 Proxies and Agents

A proxy is an agent which works on behalf of another. Proxies are used for two reasons: security and caching. Some proxy agents collect information and cache it locally so that traffic over a slow network can be minimized. Web proxies can perform this kind of function. Rather than sending WWW requests out directly, they are sent to a proxy server which registers the requests and builds a list of popular requests. These requests are collected by the proxy and copied into local storage so that the next time the request is made, the data can be served from local storage. This improves both speed and traffic load, in principle. The proxy's agents make sure that its cached copies are up to date.

Another type of proxy is the firewall type. One of the advantages of asking another to do a job, is that the original agent doesn't need to get its hands dirty. It is a little bit like the robots which bomb squads use to defuse bombs: better to send in a robot than get blown to bits

yourself. Firewall proxies exist for most services to avoid handling potentially dangerous network connections directly. We shall return to the issue of proxy services in the discussion of firewalls in section?

8.3 Installing a New Service

We need to configure the system to accept a new service by editing the file `/etc/services`. This file contains the names of services and their protocol types and port numbers. The format of entries is like this:

```
service    portnumber/protocol    aliases

pop3         110/tcp           postoffice
bootp        67/udp
cfinger      2003/tcp
```

There are two ways in which a service can run under Unix: one is that a daemon runs all the time in the background, handling connections. This method is used for services which are used often. Another way is to start the daemon only when an outside connection wishes to use it; this method is used for less frequently used services. In the second case, a master Internet daemon is used, which listens for connections for several services at once and starts the correct daemon only long enough to handle one connection. The aim is to save the overhead of running many daemons.

If we want to run a daemon all the time, then we just need to make sure that it is started in the appropriate `rc` start-up files for the system. To add the service to the internet daemon, on the other hand, we need to add a line of the following form the the configuration file `/etc/inetd.conf`:

```
# service type protocol threading user-id server-program
  server-command

pop3      stream tcp nowait root /local/etc/pop3d          pop3d
cfengine stream tcp nowait root /local/iu/bin/cfpush     cfpush -x -c
```

The software installation instructions tell us what we should add to this file.

Once we have configured a new service, it must be started by running the appropriate daemon (see section 8.3).

8.4 Summoning Daemons

Network services are run by daemons. Having done the deed of configuring a network service, (see section 8.3) by editing text files and ritually sacrificing a few doughnuts, we reach the point where we have to actually start the daemon in order to see the fruits of those labours. There are two ways to start network daemons:

- When the system boots, by adding an appropriate shell command to one of the system's startup scripts. When we use this method, the daemon hangs around in the background all the time waiting for connections.

- On demand: that is, only when a network request arrives. We use the `inetd` daemon to monitor requests for a new service. It starts the daemon to handle requests on a one-off basis. Not all services should be started in this way. One should normally follow the guidelines in the documentation for the service concerned.

The behaviour of Unix-like systems at boot-time is very far from being standard. Older systems use a series of scripts called `/etc/rc*` (short for 'read commands'). On such a system one normally finds a file called `/etc/rc.local`, where it is possible to add our own commands. Newer operating systems use a program called `init` and a series of run-levels to control what happens when the machine boots.

8.4.1 System 5 `init`

The SVR4 version of the `init` program is attaining some popularity, and is used by several GNU/Linux distributions. The idea with this program is to start the system in one of a number of run-levels. Run levels decide how many services will be started when the system boots. The minimum level of operation is single user mode, or run level 's'. Full operation is usually run level 2 or 3, depending on the type of system. (NB: be sure to check this!) When entering a run level, `init` looks in a directory called `/etc/rc?.d` and executes scripts in this directory. For instance, if we are entering run-level 2, `init` would look in the directory `/etc/rc2.d` and execute scripts lying there in order to start necessary services for this run-level. All one has to do to add a new service is to make a new file here which conforms to `init`'s simple rules. The files in these directories are usually labelled according to the following pattern:

S *number-function*

K *number-function*

Files beginning with S are for starting services and files beginning with K are for killing them again when the system is halted. The number is used to determine the order in which the scripts are read. It does not matter if two scripts have the same number. Finally, the function tells us what the script does.

Each script is supposed to accept a single argument, the word 'start' or the word 'stop'. Let's consider an example of how we might start the `httpd` daemon using `init`. Here is a checklist:

1 Determine the correct run-level for the service. Let us suppose that it is run level 2.

2 Choose an unused file name, say `S99http`.

3 Create a script accepting a single argument:

```
#!/bin/sh
case 1 in
  start) /usr/local/bin/httpd -d /usr/local/lib/httpd ;;
  stop) kill 'cat /usr/local/lib/httpd/logs/httpd.pid' ;;
  *) echo Syntax error starting http
esac
```

The advantage of this system is that software packages can be added and removed transparently just by adding or removing a file. No special editing is required, as is the case for BSD unix.

8.4.2 BSD `init`

The BSD style `init` program is quite simple. It starts executing a shell script called `/etc/rc` which then generally calls other child-scripts. These scripts start important daemons and configure the system.

To add our own local modifications, we have to edit the file `/etc/rc.local`. This is a Bourne shell script.

The BSD approach has a simpler structure than the System 5 inittab directories, but it is harder to manipulate package-wise.

8.4.3 `inetd` Configuration

The internet daemon is a *service demultiplexer*. In English, that means that it is a daemon which listens on the network for messages to several services simultaneously. When it receives a message intended for a specific port, it starts the relevant daemon to handle the request just long enough to handle one request. `inetd` saves the system some resources by starting daemons only when they are required, rather than having the clutter up the process table all the time.

The format this file can differ slightly on older systems. The best way to gleaning its format is to look at the entries which are already there. Here is a common example for the format:

```
#
# Service|type|protocol|wait|user|daemon-file|command line
#
# NB wu-ftpd needs -a now
#
ftp     stream tcp nowait root   /usr/sbin/in.ftpd      in.ftpd -a
telnet  stream tcp nowait root   /usr/sbin/in.telnetd   in.telnetd
finger  stream tcp nowait finger /local/etc/in.fingerd  in.fingerd
cfinger stream tcp nowait finger /local/etc/in.cfingerd in.cfingerd
```

The first column is the name of the service from `/etc/services`. The next column is the type of connection (stream or dgram or tli), then comes the protocol type (tcp/udp, etc.). The wait column indicates whether the service is to be single or multi-threaded, i.e. whether new requests should wait for an existing request to complete or whether a new daemon should be started in parallel. The last two columns contain the location of the program which should handle the request and the actual command line (including options) which should be executed. Notice that the finger daemon runs as a special user with no privileges.

To add a new service, we edit the file `/etc/inetd.conf` and then send the `inetd` process the HUP signal. To do this, we find the process id:

```
ps aux | grep inetd
```

Then we type:

```
kill -HUP process-id
```

8.4.4 Binding to Sockets

When a daemon is started, it creates a listening socket or *port* with a specific port number, which then gets 'bound' to the host running the service concerned. The act of binding a socket to a host's IP address identifies a fixed port service with that host. This has a specific consequence. It is only possible to bind a socket port to an address once. If we try to start another daemon, we will often see the error message

```
host: Couldn't bind to socket
bind: Address already in use
```

This means that another daemon is already running. This error can occur if two copies of `inetd` are started, or if we try to start a daemon twice, or indeed if we try to place a service in `inetd` and start a daemon at the same time. The error can also occur within a finite time-window after a service has crashed, but the problem should right itself within a few minutes.

8.4.5 TCP Wrapper Security

One of the problems with `inetd` is that it accepts connections from any host and passes them to services registered in its configuration file without question. In today's network climate this is a dangerous step, and it is usually desirable to limit the availability of certain services. For instance, services which are purely local (like RPC) should never be left open so that outside users could try to exploit them. In short, services should only be made available to those who need them. If they are left open to those who do not need them, we invite attacks on the system.

TCP wrappers is a solution to this problem. In short, it gives us the possibility of adding Access Control Lists (ACLs) to network services. TCP wrappers exists in two forms: as the `tcpd` daemon, and as a library which standalone programs can link to, called `libwrap.a`. Services which are not explicitly compiled with the library can use the daemon as a wrapper, if the services can be started from `inetd`. See section 10.3.5 for more details. TCP wrapper expects to find the daemons it proxies for in a special directory. It requires two configuration files, one which grants access to services and one which denies access. If services are not listed explicitly, TCP wrappers does nothing to prevent connection. The file to allow access to a service overrides the file to deny access, thus one normally denies access to all services as a default measure, and opens specific services one by one (see below). The `hosts.allow` file contains the names of daemons followed by a list of hosts or IP addresses, or domains or network series. The word `LOCAL` matches any host which has an unqualified host name. If we are opening a service to our local domain, it is often necessary to have both the domain suffix and the word `LOCAL`, since different operating systems employ different name services in different ways.

```
# hosts.allow
in.fingerd: .domain.country LOCAL
in.cfingerd: .domain.country LOCAL
sendmail: ALL
cfd: .domain.country LOCAL
in.rlogin: ALL in.telnetd: ALL
sshd: ALL sshdfwd-X11: .domain.country
# Portmapper doesn't understand DNS for security
```

```
portmap: 192.0.2.
rpc.mountd: 192.0.2.
rpc.nfsd: 192.0.2.
```

The TCP wrapper service works mainly for plain TCP streams, but in some operating systems (notably GNU/Linux) RPC services can also be placed under its umbrella. The portmapper and NFS mount daemons are also subject to TCP wrapper access controls. Note that we have to use IP addresses here. Host names are not accepted.

Apart from those explicitly mentioned above, all other services are denied access by adding:

```
ALL: ALL
```

in `/etc/hosts.deny`.

8.5 Setting up the DNS Name Service

The Domain Name Service (DNS/BIND) is that most important of Internet services which converts Fully Qualified host Names (FQHN) like `host.domain.country` into IP addresses like `192.0.2.10`, and vice versa. A FQHN includes the full name of the host and the domain in which it is registered. The name service consists of a database for local addresses together with a daemon `named` or `in.named` which handles look-ups in the database. Recently, the BIND software has been rewritten to solve a number of pressing problems. The resulting version is called BIND 8 [25]. Most vendor releases do not incorporate this new BIND as of 1999, but the BIND software can be fetched and installed freely from the network. Anyone running a name server ought to do this.

Establishing a name service is not difficult, but BIND is complex and we shall only skim the surface in this book. More detailed accounts of DNS configuration can be found in refs. [5, 187]. A tool for managing domain naming and electronic mail has been described in ref. [228].

8.5.1 Primary and Secondary Servers (Master and Slave)

Each domain which is responsible for its own host registration requires at least one primary nameserver. A primary nameserver (or master) is a nameserver whose data lie in authoritative source files on the server-host, maintained by a local system administrator. A domain can also have a number of secondary nameservers which mirror the primary data. A secondary nameserver (or slave) does not use source file data, but downloads its data second-hand from a primary server at regular intervals. The purpose of a secondary nameserver is to function as a backup to the primary server, or to spread the load of serving all the hosts in the domain. The only difference in setting up primary and secondary servers is one word in a configuration file.

If their primary source becomes unavailable, secondary nameservers are unable to download their data and they can hang. An alternative to secondary servers, preferred at some sites, is to set up several primary servers by mirroring the source files across several server-

[1] Traditionally, `rdist` has been used for this. However, `rdist` requires us to establish a potentially dangerous trust relationship between the hosts which these days must be regarded as a security risk; `cfengine` and `rsync` can grant limited access privileges with far less risk.

hosts. Duplicating primary servers in this fashion can be achieved easily using `cfengine` or `rsync`[1], and avoids problems which can arise with secondary servers. The disadvantage with this approach is that it is harder to administrate. The master-slave relationship exists precisely to make the mirroring of data as straightforward as possible.

In practice, master and slave servers are identical, as seen from the outside. The only difference is in the source of the data: a primary or master server knows that it should propagate changes to all secondary servers, while a secondary or slave server knows that it should accept updates.

The nameserver daemon is started once by `root`, since the DNS port is a privileged port. To function, the daemon needs to be told about its status within the DNS hierarchy, and it needs to be told where to find the files of domain data. This requires us to set up a number of configuration files. The files can seem cryptic at first, but they are easy to maintain once one has a working configuration.

8.5.2 File Structure on the Primary

Since the mapping of (even fully qualified) host names to IP addresses is not one-to-one (a host can have several aliases and network interfaces), the DNS database needs information about conversion both from the FQHN to the IP address, and the other way around. That requires two sets of data. To set up a primary nameserver, we need to complete a checklist:

- We need to make a directory in our local or site-dependent files where the DNS domain data can be installed, for instance called `dns` or `named`, and change to this directory.
- We then make subdirectories `pz` and `sz` for primary (master) and secondary (slave) data. We might not need both on the same host, but some servers can be master servers for a zone and slave servers for another zone. We shall only refer to the primary data in this book, but we might want to add secondary servers later, for whatever reason. Secondary data are cached files which can be placed in this directory.
- Assume that our domain name is called *domain.country*. We create a file `pz/`*`domain.country`*. We shall worry about its contents shortly. This file will contain data for converting names into addresses.
- Now we need files which will perform the reverse translation. It is convenient, but not essential, to keep different subnet addresses separate, for clarity. This is easy if we have a netmask which gives the subnets in our domain easily separatable addresses.

  The domain `iu.hioslo.no`, for instance, has four networks: `128.39.89.0`, `128.39.73.0`, `128.39.74.0` which includes `128.38.75.*`. So we would create files `pz/128.39.89`, `pz/128.39.73`, etc., one for each network. These files will contain data for converting addresses into 'canonical' names, or official host names (as opposed to aliases). We shall call these network files generically `pz/subnet`. Of course, we can call any of the files anything we like, since the file names must be declared in the configuration boot file.
- Dealing with the Unix loopback address requires some special attention. We handle this by creating a file for the loopback pseudo-network `pz/127.0.0`.
- Create a cache file `named.cache` which will contain the names of the Internet's primary (root) nameservers.

- (Old BIND v 4.x) Create a boot configuration file for the name-service daemon `named.boot`. We shall later link this file to `/etc/named.boot` where the daemon expects to find it. We place it here, however, so that it doesn't get lost or destroyed if we should choose to upgrade the operating system.
- (New BIND v 8.x) Create a configuration file `named.conf`. We shall later link or synchronize this file to `/etc/named.conf` where the daemon expects to find it. We place it here, however, so that it doesn't get lost or destroyed if we should choose to upgrade the operating system.
- Link the boot/conf file to the `/etc` directory, so that it appears to be at the location `/etc/named.boot`. Start the name-service daemon by typing `in.named`.

At this stage, one should have the following directory structure in site dependent files:

Names	Examples
`named/named.boot`	`dns/named.boot` (old BIND)
`named/named.conf`	`dns/named.conf` (new BIND)
`named/named.cache`	`dns/named.cache`
`named/pz/domain.country`	`dns/pz/domain.country`
`named/pz/subnet`	`dns/pz/192.0.2`
	`dns/pz/128.39.73`
	`dns/pz/128.39.89`

8.5.3 `Sample named.boot` for BIND v 4.x

The boot file tells the name-service daemon which files provide information for which networks. The syntax of this file is somewhat bizarre, and has its roots in history. It begins with the name of the directory in which we have chosen to store our data. Next it contains a line telling the daemon the name of the cache file. Finally, we list all primary and secondary domains and networks. In our case we have only primary data, no mirrored data from other servers. Suffice it to say that the network addresses have to be written backwards for reverse lookup (i.e. IP address to FQHN resolution), and the string `in-addr.arpa` is appended[2]. Here is an example file from the primary server at `domain.country`:

```
;
; Boot file for primary nameserver
; Note that there should be one primary entry for each SOA record.
;
; type domain source file or host
;
directory /usr/local/site/dns ; running directory for named
;
; cache (mandatory). Needed to "prime" nameserver with
  startup info so it
; can reach the root nameservers.
;
```

[2] Every subnet domain is a direct subdomain of the ARPA net, since IP subnet numbers do not form any numerical hierarchy.

```
cache . named.cache
;
; Primary and secondary name/address zone files follow.
;
primary 0.0.127.in-addr.arpa pz/127.0.0
primary 2.0.192.in-addr.arpa pz/192.0.2
primary domain.country                    pz/domain.country
```

Note that comments are written after a semi-colon in DNS files.

8.5.4 Sample `named.conf` for BIND v 8.x

If we are going to use the new BIND software (recommended for any a name server), then we need to replace the `named.boot` file with a new format file called `named.conf`. The information contained in this file is the same as that for `named.boot`, but many more options can be set in the new file, particularly in connection with logging. Here is a translation of the above `named.boot` file into the new format:

```
options
   {
   directory "/local/site/dns";
   check-names master ignore;
   check-names response ignore;
   check-names slave warn;
   named-xfer "/local/site/bind/bin/named-xfer"; /* Location of
     daemon */
   fake-iquery no;                                /* security */
   notify yes;
   };

zone "."
   {
   type hint;
   file "named.cache";
   };

//
// Primary and secondary name/address zone files follow.
//
zone "0.0.127.in-addr.arpa"
   {
   type master;
   file "pz/127.0.0";
   };

zone "2.0.192.in-addr.arpa"
   {
   type master;
   file "pz/192.0.2";
   };

acl trustedhosts
   {
   ! 192.0.2.11;  // Not this host!
   192.0.2.0/24;  // Net with 24 bit netmask set. i.e. 255.255.255.0
   192.0.74.0/23; //         23                      . 255.255.254.0
   };
```

```
zone "domain.country"
   {
   type master;
   file "pz/domain.country";

allow-transfer   // Allows ls domain.country in nslookup
   {             // and domina downloads
   trustedhosts; // Access Control List defined above
   }
   };

// dns.domain.country server options

server 192.0.2.11
   {
   transfer-format many-answers;
   };

logging
   {
   channel admin_stuff
      {
      file "/local/site/logs/admin" versions 7;
      severity debug;
      print-time yes;
      print-category yes;
      print-severity yes;
      };

   channel xfers
      {
      file "/local/site/logs/xfer" versions 7;
      severity debug;
      print-time yes;
      print-category yes;
      print-severity yes;
      };

   channel updates
      {
      file "/local/site/logs/updates" versions 10;
      severity debug;
      print-time yes;
      print-category yes;
      print-severity yes;
      };

   channel security
      {
      file "/local/site/logs/security" versions 7;
      severity debug;
      print-time yes;
      print-category yes;
      print-severity yes;
      };

   category config
      {
      admin_stuff;
      };
```

```
category parser
   {
   admin_stuff;
   };
category update
   {
   updates;
   };
category load
   {
   updates;
   };
category security
   {
   security;
   };
category xfer-in
   {
   xfers;
   };
category xfer-out
   {
   xfers;
   };
category db
   {
   updates;
   };
category lame-servers
   {
   null;
   };
category cname
   {
   null;
   };
};
```

Note the `allow-transfer` statement which allows a user of `nslookup` to obtain a dump of the local domain, using the 'ls' command within the `nslookup` shell. If this is not present, version 8 BIND will not allow such a listing. BIND now allows ACLs to control access to these data. In the example we have created an ACL alias for all of the trusted hosts on our network. The ACLs use an increasingly popular, if somewhat obscure, notation for groups of IP addresses. The 'slash' notation is supposed to represent all of the hosts on a subnet. To fully specifiy a subnet (which, in practice, might be part of a class A, B or C network), we need to specify the network address and the subnet mask. The slash notation does this by giving the network address followed by a slash, followed by the number of bits in the netmask which are set to one. So, for example, the address series

```
192.0.2.0/24
```

means all of the addresses from 192.0.2.0 to 192.0.2.255, since the netmask is 255.255.255.0. The example

```
192.0.74.0/23
```

is an example of a doubled-up subnet. This means all of the hosts from 192.0.74.0 to 192.0.74.255 and 192.0.75.0 to 192.0.75.255, since the netmask is 255.255.254.0, i.e. 23 non-zero bits. ACLs can contain any list of hosts. The pling '!' operator negates an address, or entry. The important thing to remember about ACLs in general is that they work taking each entry in turn. As soon as there is a match, the access algorithm quits. So if we were to write

```
acl test
  {
  192.0.2.11;
  !192.0.2.11;
  }
```

the result would always be to grant access to 192.0.2.11. Conversely, if we wrote

```
acl test
  {
  !192.0.2.11;
  192.0.2.11;
  }
```

the result would always be to deny access to this host, since the second instance of the addresses is never reached.

Note that for a secondary, or slave server mirroring a master, we would replace the word master with slave, and `pz` with `sz` for clarity.

8.5.5 Sample `named.cache` or `named.root`

The cache file (now often referred to as the root file) contains the names of root nameservers. The data for the cache file were formerly maintained by the American military at `nic.ddn.mil`. Today they are retrieved by anonymous ftp from the INTERNIC `ftp.rs.internic.net`. The list of Internet root servers (which bind together all Internet domains) are listed in a file called `domain/named.root`. The retrieved data are simply included in a file called `named.cache` or `named.root`.

```
;       This file holds the information on root nameservers needed to
;       initialize cache of Internet domain nameservers
;       (e.g. reference this file in the "cache . <file>"
;       configuration file of BIND domain nameservers).
;
;       This file is made available by InterNIC registration services
;       under anonymous FTP as
;         file                /domain/named.root
;         on server           FTP.RS.INTERNIC.NET
;
.                       3600000   IN   NS   A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET.     3600000        A    198.41.0.4
;
; formerly NS1.ISI.EDU
;
```

```
.                       3600000      NS   B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET.     3600000      A    128.9.0.107
;
; formerly C.PSI.NET
;
.                       3600000      NS   C.ROOT-SERVERS.NET.
C.ROOT-SERVERS.NET.     3600000      A    192.33.4.12
;
; formerly TERP.UMD.EDU
;
.                       3600000      NS   D.ROOT-SERVERS.NET.
D.ROOT-SERVERS.NET.     3600000      A    128.8.10.90
;
; formerly NS.NASA.GOV
;
.                       3600000      NS   E.ROOT-SERVERS.NET.
E.ROOT-SERVERS.NET.     3600000      A    192.203.230.10
;
; formerly NS.ISC.ORG
;
.                       3600000      NS   F.ROOT-SERVERS.NET.
F.ROOT-SERVERS.NET.     3600000      A    192.5.5.241
;
; formerly NS.NIC.DDN.MIL
;
.                       3600000      NS   G.ROOT-SERVERS.NET.
G.ROOT-SERVERS.NET.     3600000      A    192.112.36.4
;
; formerly AOS.ARL.ARMY.MIL
;
.                       3600000      NS   H.ROOT-SERVERS.NET.
H.ROOT-SERVERS.NET.     3600000      A    128.63.2.53
;
; formerly NIC.NORDU.NET
;
.                       3600000      NS   I.ROOT-SERVERS.NET.
I.ROOT-SERVERS.NET.     3600000      A    192.36.148.17
;
; temporarily housed at NSI (InterNIC)
;
.                       3600000      NS   J.ROOT-SERVERS.NET.
J.ROOT-SERVERS.NET.     3600000      A    198.41.0.10
;
; housed in LINX, operated by RIPE NCC
;
.                       3600000      NS   K.ROOT-SERVERS.NET.
K.ROOT-SERVERS.NET.     3600000      A    193.0.14.129
;
; temporarily housed at ISI (IANA)
;
.                       3600000      NS   L.ROOT-SERVERS.NET.
L.ROOT-SERVERS.NET.     3600000      A    198.32.64.12
;
; housed in Japan, operated by WIDE
;
.                       3600000      NS   M.ROOT-SERVERS.NET.
M.ROOT-SERVERS.NET.     3600000      A    202.12.27.33
; End of File
```

8.5.6 Sample pz/*domain.country*

The main domain file contains data identifying the IP addresses of hosts in our domain; it defines possible aliases for those names, and it also identifies special servers such as mail-exchangers which mail-relay programs use to send electronic mail to our domain. *Note that the IP addresses used here are examples. You should replace them with your own valid IP addresses.*

Each host has a *canonical name* or CNAME, which is its official name. We may then define any number of aliases to this canonical name. For instance, it is common to create aliases for hosts which provide well known services, like `www.domain.country` and `ftp.domain.country` so that no-one needs to remember a special host name in order to access these services in our domain. Here is an abbreviated example file. There are several kind of records here:

- `SOA` Indicates the Start Of Authority for this domain (referred to as `@`).
- `NS` Lists a name server for this domain or a sub-domain. NS records are not used for anything other than delegation. They exist only for convenience.
- `MX` Lists a mail exchanger for this domain (with priority).
- `A` Create an A record, i.e. define the canonical name of a host with a given IP address.
- `CNAME` Associate an alias with a canonical name.
- `HINFO` Advertise host information. No longer advisable for security reasons.

DNS database files contain *resource records* which refer to these different elements. The general form of such a record is

```
Name   [ttl] [class]   assoc   data
```

The square brackets imply that the second and third fields are optional. The default class is `IN` for Internet. Other values for class refer to little used name services which will not be considered here. Each resource record is an association of a name with an item of data. The type of association is identified by the *`assoc`* field, which is one of the set `A`, `PTR`, `NS`, etc. For example

```
host1 86400 IN A     10.20.30.40
host1 86400 IN CNAME host2
host1 86400 IN MX    10 mailhost
```

or simply

```
host1 A     10.20.30.40
host1 CNAME host2
host1 MX    10 mailhost
```

Since we need two sets of data, both for name lookups and address lookups, there are records in which *`name`* and *`data`* are host names and host numbers, respectively, and another set in which those roles are reversed, e.g.

```
40 PTR host1
```

In addition to mapping host names to addresses and vice versa, the DNS tables also tell e-mail services how to deliver mail. We will need to have a so-called 'mail-exchanger' record in the

DNS tables in order to tell e-mail which host handles e-mail for the domain. An entry of the form

```
domain-name MX  priority  mailhost
```

tells e-mail services that mail sent to *name domain-name* should be routed to the host *mailhost*. For instance,

```
domain.country.  MX 10 mailhost
                 MX 20 backup
```

tells our server that mail addresses of the form *name*`@domain.country` should be handled by host called mailhost (which is an alias for a host called mercury, as we shall see below). The priority number 10 is chosen at random. Several records can be added with backup servers if the first server does not respond.

Mail records are also possible on a per-host basis. If we want mail sent to host `XXX` handled by host `YYY`, we would add a record,

```
XXX MX 10 YYY
```

This would mean that mail sent to

```
XXX.  domain-name
```

would be handled by `YYY`. For instance, mail addressed to

```
name@XXX.domain.country
```

would actually be sent to

```
name@YYY.domain.country
```

Here is an example file with all the elements in place. Note the meaning of the following special symbols:

;	Comment
@	Stands for the local domain
( )	Continue record over several lines

```
ORIGIN domain.country.    ; @ is an alias for this

@      IN  SOA   mercury.domain.country. sysadm.mercury.
  domain.country.
                 (
                 1996111300 ; Serialnumber
                 3600       ; Refresh, 1 hr
                 600        ; Retry, 10 min
                 604800     ; Expire 1 week
                 86400      ; Minimum TTL, 1 day
                 )

           A     192.0.2.237 ; domain.country points to
                             ; this host by default
```

```
; Name servers:
      IN  NS      mercury.domain.country.
      IN  NS      backup.domain.country.
      IN  NS      dns.parent.co.

; Mail exchangers for whole domain

@             MX    10      mercury

; Common aliases for well known services

www           CNAME mercury    ; aliases
ftp           CNAME mercury
mailhost      CNAME mercury

; Router

domain-gw     A     192.0.2.1
              A     128.39.73.129 ; 2 addresses

iu-gw         CNAME domain-gw

localhost     A     127.0.0.1

; example net

mercury       A     192.0.2.10
thistledown   A     192.0.2.233
jart          A     192.0.2.234
nostromo      A     192.0.2.235
daystrom      A     192.0.2.236
borg          A     192.0.2.237
dax           A     192.0.2.238
axis          A     192.0.2.239
```

Note that, as this file stands, mail exchanger data described by the `MX` record are only described for the domain as a whole. If an external mailer attempts to send directly to a specific host, it is still allowed to do so. We can still override this by adding an explicit `MX` record for each `A` record. For example,

```
mercury       A     192.0.2.10
              MX    10 mailhost
              MX    20 backup
thistledown   A     192.0.2.233
              MX    10 mailhost
              MX    20 mailhost
jart          A     192.0.2.234
              MX    10 mailhost
              MX    20 mailhost
```

This will tell an external mailer to send mail to each of these hosts to the mailhost instead. Normally, this would not be a problem. Once could simply configure the non-mailhosts as so-called null clients, meaning that they would just forward the mail on to the mailhost. However, it can be important to avoid relying on client forwarding if we are using a hub solution inside some kind of filtering router, since the SMTP ports might be blocked to all the

other hosts. Thus, mail would not be send-able to the other hosts unless these MX records were in place see section 8.7 for more details.

8.5.7 Sample pz/*network*

The network files are responsible for producing a fully qualified domain name given an IP address. This is accomplished with so-called PTR records. In other words, these records provide reverse lookup. *Note that the IP addresses used here are examples. You should replace them with your own valid IP addresses.* The reverse lookup-file looks like this:

```
$ORIGIN 89.39.128.in-addr.arpa.
@       IN      SOA mercury.domain.country. sysadm.mercury.
                    domain.country.
                      (
                      1996111300 ; Serialnumber
                      3600       ; Refresh, 1 hr
                      600        ; Retry, 10 min
                      604800     ; Expire 1 week
                      86400 ; Minimum TTL, 1 day
                      )
; Name servers:
        IN      NS      mercury.domain.country.
        IN      NS      dns.parent.co.
        IN      NS      backup.domain.country.
;
; Domain data:
;
1       PTR     domain-gw.domain.country.
; etc
10      PTR     mercury.domain.country.
; etc
233     PTR     thistledown.domain.country.
234     PTR     jart.domain.country.
```

Note carefully how the names end with a full-stop. If we forget this, the nameserver appends the domain name to the end, resulting in something like `lore.domain.country.domain.country`.

8.5.8 Sample pz/127.0.0

To avoid problems with the loop-back address, all domains should define a fake 'loop-back' network simply to register the Unix loop-back address correctly within the DNS. Since 127.0.0 is not a network and the loop-back address doesn't belong to anyone, it is acceptable for everyone to define this as part of the local nameserver. No name collisions will occur as a result.

```
; Zone file for "localhost" entry.
$ORIGIN 0.0.127.IN-ADDR.ARPA.
@    IN       SOA      mercury.domain.country. sysadm.mercury.
                       domain.country.
                         (
                         1995070300 ; Serialnumber
                         3600       ; Refresh
                         300        ; Retry
                         3600000    ; Expire
                         14400      ; Minimum
                         )
     IN       NS       mercury.domain.country.
;
; Domain data
;
1    PTR    localhost.
0.0.127.in-addr.arpa.  IN NS mercury.domain.country.
0.0.127.in-addr.arpa   IN NS backup.domain.country.
1.0.0.127.in-addr.arpa. IN PTR localhost.
```

8.5.9 Zone Transfers

A zone is a portion of a complete domain which is self-contained. Responsibility for a zone is delegated to the administrators of that zone. A zone administrator keeps and maintains the files we have discussed above. When changes are made to the data in a domain, we need to update the serial number of the data in the source files. Secondary nameservers use this serial number to register when changes have occurred in the zone data, i.e. to determine when they should download new data.

8.5.10 Sub-domains and Structure

Suppose we are in a DNS domain `college.edu` and would like to name hosts according to their departmental affiliation. We could use a naming scheme which made each department look like a sub-domain of the true domain `college.edu`. For instance, we might want the following hosts:

```
einstein.phys.college.edu
darwin.bio.college.edu
von-neumann.comp.college.edu
```

We can achieve this very simply, because having the extra 'dot' in the name makes no difference to the name service. We just assign the A record for the host accordingly. In the zone file for `college.edu`

```
$ORIGIN college.edu.    ; @ is an alias for this
@       IN  SOA    chomsky.college.edu.sysadm.chomsky.
                   college.edu
                     (
                     1996111300 ; Serialnumber
                     3600 ; Refresh, 1 hr
                     600 ; Retry, 10 min
                     604800 ; Expire 1 week
                     86400 ; Minimum TTL, 1 day
                     )

; ...

einstein.phys      A 192.0.2.5
darwin.bio         A 192.0.2.6
von-neumann.comp   A 192.0.2.7
```

It does not matter that we have dots in the names on the left-hand side of an A record assignment. DNS does not care about this. It still looks and behaves like a sub-domain. There is no need for an SOA record for these sub-domains, as written, since we are providing authoratative information about them here explicitly. However, we could handle this differently. According the principle of delegation, we would like to empower local units of a network community with the ability to organize their own affairs. Since the computer science department is growing fat on the funds it receives, it has many hosts and it starts to make sense to delegate this sub-domain to a local administrator. The emaciated physics and biology departments don't want this hassle, so we keep them under out wing in the parent zone records.

Delegating the sub-domain `comp.college.edu` means doing the following:

- Setting up a name server which will contain an authoratative SOA database for the sub-domain.
- Delegating responsibility for the sub-domain to that nameserver, using an `IN` record in our parent domain's zone data.
- Informing the parent organization about the changes required in the top-level `in-addr.arpa` domain, required for reverse lookups.

Normally, the `NS` records for a zone are only present to remind local administrators which hosts are name servers. These records do not serve any real function. However, a parent domain seeking to delegate a sub-domain does need these. Suppose we choose the host `markV` to be the name server for the sub-domain. A pair of records like this:

```
comp                    86400 IN  NS m5.comp.college.edu
m5.comp.college.edu     86400 IN  A 192.0.2.200
```

creates a sub-domain called 'comp'. The first line tells us the name of a name server in the sub-domain which will be authoratitive for the sub-domain. We give it a specific time to live, for definiteness. Notice, however, that this is a cyclic definition: we have defined the comp sub-domain using a member of the comp-subdomain. To break the infinite loop, we have to also provide a *glue record* (an A record, borrowed from the sub-domain) which tells the system the actual IP address of the nameserver which will contain information about the remainder of the domain.

To delegate a sub-domain, we also have to delegate the reverse pointer records. To do this we need to contact the parent organization which owns the network above our own. In the 'olden days' in-addr.arpa delegation was performed by the 'nic' military, then the 'internic'. Today, the logistics of this mapping has become too large a job for any one organization. For ISPs who delegate subnets, it is passed back up the line to the network owners. Each organization knows the organization above it, and so we contact these until someone has authority to modify the in-addr.arpa domain. An erroneous delegation is usually referred to as a lame delegation. This means that a name server is listed as being authoratative for a domain, but in fact does not service that domain. Lame delegations and other problems can be diagnosed with programs such as `dnswalk`. See also ref. [22].

8.5.11 Compiling BIND v 8.x

The new BIND software can be collected from the Internet Software Consortium's [25] web site. This software contains everything required to build replacement name server software and complementary `nslookup` and `dig` clients. Unfortunately, this software is rather akward to compile, since it used an antiquated BSD configuration process and some non-standard tools. Here is a brief checklist for building and installing the BIND 8 name server:

- Bearing in mind that we want to separate local patches from the operating system (so that an OS upgrade will not take us back to an inferior version), it is advisable to make a new directory called, say, `bind` in our site-dependent files. Collect the gzipped tar file `bind-8.x-src.tar.gz` from the WWW site [25].

  ```
  host% mkdir bind-src
  host% cd bind-src
  host% get file
  host% tar zxf bind-8.x-src.tar.gz
  ```

  The tar file unpacks into a sub-directory called `src`, since it is part of a larger tree of code, so we must make sure to unpack it in a fresh directory where it will not get lost.

- Decide on a directory where the compiled code will reside. For instance, within site-dependent files `/local/site/bind`. It is crucial to keep this code separate from operating system code. BIND is likely to go through many versions while operating system code stays static. Installing it where the operating system would have installed it is just going to be confusing: we need to distinguish the good BIND from the OS BIND with certainty.

- Although the build procedure tends to encourage it, it is not necessary to be the root user to build BIND, so long as we follow a simple procedure. The build procedure requires an intermediate directory; let's call is `bind-compile`. In what follows, we can assume that we are working as an ordinary user, e.g. mark. Change directory to the sources and follow the instructions there:

  ```
  host% cd src
  host% make DST=HOME/bind-compile SRC='pwd' links
  host% cd HOME/bind-compile
  host% make clean
  host% make depend
  host% make
  ```

- The default installation requires us to have `byacc` installed on the system. Most systems will not have this program, so it is necessary to edit a configuration file in the `ports/platform` directory called `Makefile.set` and replace it with `bison`. Note that, if our system uses bison, we have to use the command `bison -y -d` for compatibility. The default paths for file installation are based on a vendor installation, and break the principle of separating operating system from local modifications. We can fix this by changing a `Makefile.set` file in the `ports` directory before executing `make` above, e.g. for Solaris, installed in `/local/site/bind`

```
host% more Makefile.set
'CC=gcc'
'CDEBUG=-g -O2'
'DESTBIN=/local/site/bind/bin'
'DESTSBIN=/local/site/bind/bin'
'DESTEXEC=/local/site/bind/bin'
'DESTMAN=/usr/local/share/man'
'DESTHELP=/local/site/bind/lib'
'DESTETC=/etc'
'DESTRUN=/local/site/bind/var'
'LDS=:'
'AR=/usr/ccs/bin/ar cru'
'LEX=flex'
'YACC=bison -y -d'
'SYSLIBS=-ll -lnsl -lsocket'
'INSTALL=/usr/ucb/install'
'PIDDIR=/etc'
'MANDIR=man'
'MANROFF=man'
'CATEXT=$ $N'
'PS=ps -p'
'RANLIB=/usr/ccs/bin/ranlib'
```

Once the sources are built, we can move them to a permanent location. The compilation process compiles each separate binary in a subdirectory of its own, so to copy them all to a final location, we first make the target directory and then copy them all:

```
host# mkdir -p /local/site/bind-new/bin
```

```
host# cd $HOME/bind-compile/bin
host# foreach prog ( * )
foreach? cp $prog/$prog /local/site/bind-new/bin
foreach? end
```

To replace an old version of bind we kill the old `named` daemon and move the new binaries into place:

```
host# mv /local/site/bind /local/site/bind-old
host# mv /local/site/bind-new /local/site/bind
```

Finally, we start the new daemon:

```
/local/site/bind/bin/named
```

Remember to set appropriate permissions on the files. None of them need to be setuid root! Remember to set the correct path to the compiled xfer-daemon in the file `named.conf`.

Note: BIND does not allow us to simply configure the location of important files. It takes several decisions without asking, on an operating system specific basis. For instance, on Solaris hosts the `named.conf` file must be available in `/usr/local/etc/named.conf` and the `nslookup` help file needs to be placed in `/usr/local/lib/nslookup.help`. This is an awkward mess, which hopefully will be cleaned up one day. Symbolic links can be used to make a master file (kept at a sensible location) appear to be at the location required by the daemon. Cfengine is a useful tool for managing such links.

Finally, we can remove the source code from the system:

```
rm -r bind-compile
rm -r bind-src
```

8.6 Setting up a WWW Server

The World Wide Web (or W3) service is mediated by the daemon `httpd`. There are several publicly available daemons which mediate the WWW service. This description is based on the freely available Apache daemon which is widely regarded as the best and most up-to-date. It can be obtained from `http://www.apache.org`.

Configuring `httpd` is a relatively simple matter, but it does involve a few subtleties which we need to examine. Most of these have to do with security. Some are linked to minor changes between versions of the server. This discussion is based on Apache version 1.3.*x*.

An `httpd` server can be used for two purposes:

- *Website:* for publishing information which is intended to be open to the world. The more people who see this information, the better.
- *Intranet:* for publishing private information for internal use within an organization.

Unless we are going to operate within a firewall configuration, there is probably no need to separate these two services. They can run on the same server, with access control restricting information, on a need to know basis. Special attention should be given to CGI programs, however. These are particularly insecure and can compromise any access controls which we place on data. If we need restricted information access we should not allow arbitrary users to have accounts or CGI privileges on a server: CGI programs can always be written to circumvent server security.

The WWW service will also publish two different kinds of web pages:

- *Site data:* a site's official welcome page and subsequent official data (access by `http://www.domain.country`).
- *Personal data:* the private web pages of registered, local users (access by `http://www.domain.country`).

The matter of whether to allow local users private web pages at a given organization is a matter of policy.

The WWW is an open service: it gives access to file information, usually without requiring a password. For that reason it has the potential to be a security hazard: not with respect to itself, or the information which one intends to publish, but to the well-being of the host on which it runs. A typical configuration error in a large cooperation's web server, a few

years ago, allowed an attacker to delete all users' home directories from the comfort of his browser.

To start a WWW service we need some `html`-files containing information we wish to publish and a server-daemon. We then need to edit configuration files which tell the daemon where to find the web pages it will be publishing. Finally, we need to tell it what we do *not* want it to tell the outside world. The security of the whole system can depend upon which files and directories outsiders have access to.

8.6.1 Choosing a Server Host

Personal data are accessed from users' home directories, usually under a sub-directory called www. It makes considerable sense for the WWW server to be on the host which has the physically mounted disks. Otherwise, the WWW server would first have to access the files via NFS and then transmit them back to the requester. This would lead to an unecessary doubling of network traffic.

8.6.2 Installation

A survey which was carried out in 1998 revealed that about 70% of all WWW servers in the world were Apache WWW servers running on FreeBSD or GNU/Linux PCs. The Apache server is amongst the most efficient and versatile, and it is Free Software so we shall adopt it as our working model. Apache-`httpd` runs both on Unix-like OSes and NT.

The server is compiled in the usual way by unpacking a `.tar.gz` file and by running `configure` then `make`. Apache has support for many bells and whistles which enhance its performance. For instance, one can run an embedded language called PHP from within HTML pages to create 'active web pages'. The `httpd` daemon must then be configured with special PHP support. Debian GNU/Linux has a ready-made package for the Apache server, but it it old. It is always worth collecting the latest version of the server from an official mirror site (see the Apache web site for a list of mirrors).

Apache uses a GNU autoconf `configure` program to prepare the compilation. As always, we have to choose a prefix for the software installation. If none is specified, the directory `/usr/local` is the default.

There is no particular reason for installing the binaries elsewhere, however Apache does generate a convenient startup/shutdown script which compiles in the location of the configuration files. The configuration files are kept under the installation-prefix, in `etc/apache/*.conf`. On the principle of separating files which we maintain ourselves, from files installed by other sources, we almost certainly do not want to keep the true configuration files there, but rather would like to keep them together with other site-dependent configuration files. We shall bear this in mind below.

To build a basic web server, then, we follow the usual sequence for compiling Free Software:

```
% configure
% make
% make -n install
% su
# make -n install
```

8.6.3 Configuration

Having compiled the daemon, we have to prepare some infrastructure. First we make sure that we have the two lines

```
www 80/tcp http
www 80/udp http
```

in the `/etc/services` file on Unix-like systems. Next we must:

- Create a directory in our site-dependent files called, say, `www`, where HTML documents will be kept. In particular, we will need a mercury file `www/index.html` which will be the root of the web site.
- Edit the files `httpd/conf/*.conf` with a text editor so that we configure in our site-specific data and requirements.
- Create a special user and a special group which will be used to restrict the privilege of the `httpd` service.

The daemon's configuration determines the behaviour of the WWW service. It is decided by the contents of a set of files:

`httpd.conf`	Properties of the daemon itself.
`access.conf`	Access control for documents and CGI.
`srm.conf`	Server resource management.
`mime.types`	Multimedia file extensions.

This breakdown is a matter of convention. There is no real difference between the files as far as the configuration language is concerned.

The Apache server provides example configuration files which we can use as an initial template. In recent versions, Apache has moved away from the idea of using several configuration files, towards keeping everything in one file. You may wish to form your own opinion about what is best policy here. In practice it makes no difference, since the old file-structure is still supported. For clarity, we shall assume the traditional file structure.

The `httpd` is started by root/Administrator, but the daemon immediately relinquishes its special privileges in order to run with the access rights of a `www` user for all operations. The `User` and `Group` directives specify which user the daemon should run as. The default here is usually the user `nobody`. This is the default because it is the only non-privileged user name which most systems already have. However, the `nobody` user was introduced to create a safe mapping of privileges for the Unix NFS (Network File System), so to use it here could lead to confusion and possibly even accidents later. A better approach is to use a completely separate user ID for the service. In fact, in general:

Principle 41 (Separate uids for services) *Each service which does not require privileged access to the system should be given a separate, non-privileged user-ID. This restricts service privileges, preventing any potential abuse should the service be hijacked by system attackers; it also makes clear which service is responsible for which processes in the process table.*

Corollary 42 (Privileged ports) *Services which run on ports 1–256 must started with Administrator privileges in order for the socket to be validated, but can switch internally to a safer level of privilege once communications have been established.*

8.6.4 The `httpd.conf` file

Here is a cut-down example file which points out some important parameters in the configuration of the server. The actual example files distributed with the server are more verbose and contain additional options. You should probably not delete anything from those files unless you have read the documentation carefully, but you will need to give the following points special attention:

```
# httpd.conf

ServerRoot      /local/site/httpd/

ServerAdmin     sysadm@domain.country
User            www
Group           www

ServerType      standalone  # not inetd
HostnameLookups off         # save time

  # Several request-transfers per connection is efficient

KeepAlive            On
MaxKeepAliveRequests 100
KeepAliveTimeout     15

  # Looks like several servers, really only one ..

NameVirtualHost 192.0.2.220

<VirtualHost www.domain.country>
ServerAdmin webmaster@domain.country
DocumentRoot /site/host/local/www-data
ServerName www.domain.country
</VirtualHost>

<VirtualHost project.domain.country>
ServerAdmin webmaster@domain.country
DocumentRoot /site/host/local/project-data
ServerName project.domain.country
</VirtualHost>
```

The `ServerRoot` directive tells the daemon which directory is to be used to look for additional configuration files (see below) and to write logs of transactions. When the daemon is started, it normally has to be told the location of the server root, with the '`-d`' option:

```
httpd -d /local/site/httpd
```

The daemon then looks for configuration files under `conf`, for log files under `logs`, and so on. The location of the server root does not have to have anything to do with the location of the binaries, as installed above. Indeed, since configuration files and log files can both be considered local, site-dependent data, we have placed them here amongst local, site-dependent files.

The `User` and `Group` directives tell the daemon which users' privileges it should use after connecting to the privileged port 80. A special user and group ID should be created for this purpose. The user ID should be an account which it is not possible to log on to, with no valid shell and a barred password (see section 5.2).

The `ServerType` variable indicates whether we are planning to run `httpd` on demand from `inetd`, or whether it should run as a standalone daemon. Running as a standalone daemon can give a considerable saving of overhead in forking new processes. A standalone daemon can organize its own resources, rather than relying on a multiplexer like `inetd`.

`HostnameLookups` determines whether the DNS names of hosts connecting to the server will be looked up and written into the access log. DNS lookups can add a significant amount of delay to a connection, so this should be turned off on busy servers. In a similar efficiency vein, the `KeepAlive` variable tells the server to not close a connection after every transaction, but to allow multiple transactions up to a limit of `MaxKeepAliveRequests` on the same connection. Since the overhead of starting a new connection is quite high, and of shutting one down even higher, a considerable improvement in efficiency can be achieved by allowing persistent connections.

The final part of the file concerns the `VirtualHost` environment. This is a feature of Apache which is very useful. It enables one to maintain the appearance of separate web servers, with just one daemon. For instance, we might want to have a generic point of contact for our domain, called `www.domain.country`, but we might also want to run a special project machine, whose data were maintained by a separate research group, called `project.domain.country`. To do this we need to create a `VirtualHost` structure for each virtual-hostname we would like to attach to the server.

We also need to register these alternative names as DNS aliases so that others will be able to use them in normal URLs in their web browsers. Suppose the actual canonical name of the host we are running on is `workhorse.domain.country`. In the primary zone of the domain `domain.country`, we would make the following aliases:

```
www CNAME workhorse
project CNAME workhorse

workhorse A 192.0.2.220
```

The IP address of `workhorse` must also be declared in `httpd.conf` so that we have a reliable address to bind the socket to. The declarations as shown then create two virtual hosts, `www` and `project`, each of which has a default data root-directory pointed to by the `DocumentRoot` variable.

8.6.5 The `access.conf` File

This file determines what rights various users will have when trying to access data on the server. It also implicitly determines whether `httpd` will search for `.htaccess` files in directories. Such files can be used to override the settings in the `access.conf` file.

The `Directory` structure works like an access control list, granting or denying access to directories (and implicitly all subdirectories). A good place to start is to make a general structure denying access to directories which are not, later, dealt with explicitly.

```
# access.conf

AccessFileName .htaccess

<Directory />
 order allow,deny
 deny from all
 AllowOverride None
</Directory>
```

This initializer tells the daemon that it should neither grant rights to arbitrary directories on the disk, nor allow any overriding of access rights by `.htaccess` files. This simple precaution can yield a performance gain on a web server because the daemon will search for `.htaccess` files in every directory from the top of the file tree to the directory mentioned in a `Directory` structure. This can consume many disk operations which, on a busy server, could waste valuable time. We then need to go through the independent subtrees of the file system which we want to publish:

```
#
# Don't allow users to make symlinks to files
# they don't own, thus circumventing .htaccess
#

<Directory /home>
 order allow,deny
 allow from all
 AllowOverride All
 Options Indexes SymLinksIfOwnerMatch
</Directory>
```

In a `Directory` structure, we express rules which determine how `httpd` evaluates access rights. The ordering `allow` followed by `deny` means that files are allowed unless explicitly denied. The line which follows has the form `allow from all`, meaning that the data in `/home` (users' home directories) are open to every caller. The `Options` directive is quite important. `Indexes` means that a browser will be able to present a directory listing of `.html` files which can be accessed, if a user browses a directory which does not contain a standard `index.html` file. The option `SymLinksIfOwnerMatch` means that `httpd` will follow symbolic links, only if the user who made the symbolic link is also the owner of the file it points to. The point of the conditional is that a local user should not be able to bypass normal access controls simply by creating a symbolic link to a file which is otherwise access restricted. `AllowOverride` means that we can override access controls for specific directories using `.htaccess` files (see section 8.6.9).

```
<Directory /local/site/www>
 order allow,deny
 allow from all
 AllowOverride All
 Options Indexes FollowSymLinks
</Directory>
```

In this stanza, things are almost the same. The files under `/local/site/www` are the site's main web pages. They are available to everyone, and symbolic links are followed regardless of owner. We can afford to be magnanimous here since the site's main pages are controlled

by a trusted user (probably us), whom we assume would not deliberately circumvent any security mechanisms. The story is different for ordinary users, whom we do not necessarily have any reason to trust.

```
<Directory /local/site/www/private>
 order allow,deny deny
 from all
 allow from 192.0.2.
 AllowOverride All
 Options Indexes FollowSymLinks
</Directory>
```

In this example, we restrict access to the sub-directory `private`, to hosts orginating from network addresses `192.0.2.`*x*. This is useful for controlling access to certain documents to within an organization. Another way of doing this would be to write

```
allow from .domain.country
```

but we might have a special reason for restricting on the basis of subnets, or network IP series. This kind of access control is a way of making an *intranet* server.

Finally, as an extra check to prevent ordinary (untrusted) users from making symbolic links to the password file, we can add a `FilesMatch` structure which checks to see whether the actual file pointed to matches a regular expression. In the event of a match, access is denied to everyone.

```
# Don't allow anyone to download a copy of the passwd file
# even by symbolic link

<FilesMatch ".*passwd.*">
order allow,deny
deny from all
</FilesMatch>
```

This is not an absolute security. If local users really want to publish the password file they can simply copy it into an HTML document. However, it does help to close obvious avenues of abuse.

8.6.6 `srm.conf` File

The `srm.conf` file is the file where we define the remaining behaviour of the server in response to requests from clients. The first issue to deal with is that of users' private web pages. These are searched for in a subdirectory of each user's home directory which we must specify. Normally, this is called `www` or `www-data`. The `UserDir` directive is used to set this. Using this directive, it is also possible to say that certain users will not have web pages. The obvious contender here is the administrator account root.

```
# srm.conf

UserDir www
UserDir disabled root

DirectoryIndex index.html
FancyIndexing on
```

The `DirectoryIndex` directive determines the default filename which `httpd` looks for, if the URL provided by the client is the name of a directory. Using this arrangement the start home-page for a user becomes

```
~user/www/index.html
```

which, using the file scheme */site/host/contents, becomes*

```
/site/server/home/user/www/index.html
```

Next, it useful to be able to specify the way in which the server responds to errors. The default behaviour is to simply send the client a rather dull text string indicating the number and nature of the error. We can alter this by asking the daemon to send a specially customized page, tailored to our own special environment, perhaps with personal logo, etc. The `ErrorDocument` directive is used for this. It traps error numbers and maps them to special pages. For example, to map to a standard local file in the root of the server's pages one would add

```
ErrorDocument 500 /errorpage.html
ErrorDocument 401 /errorpage.html
ErrorDocument 402 /errorpage.html
```

Another possibility is to make a generic CGI script for handling error conditions. An example script is provided in section 8.6.8. In that case, we declare all error-codes to point to the generic CGI-script.

```
# Customizable error response (Apache style)

ErrorDocument 500 /cgi-bin/error.pl
ErrorDocument 404 /cgi-bin/error.pl
ErrorDocument 401 /cgi-bin/error.pl
ErrorDocument 403 /cgi-bin/error.pl
ErrorDocument 407 /cgi-bin/error.pl
```

The final issue to mention about the `srm.conf` file is that of script aliases. For `httpd` to allow the execution of a CGI script on the server, it must be referred to with the help of a `ScriptAlias`. There are two purposes to this. The script alias points to a single directory, usually a directory called `cgi-bin` which lies under the user's own www directory. The script alias means that only programs placed in this directory will be executed. This helps to prevent the execution of arbitrary programs which were not intended for web use; it also hides the actual directory structure on the server host. It is necessary to add one script alias entry for each directory that we want to execute CGI programs from. That usually means at least one directory for each user, plus one for general site scripts. Here are two examples:

```
ScriptAlias /cgi-bin/        /local/site/www/cgi-bin
ScriptAlias /cgi-bin-mark/   /home/mark/www/cgi-bin
```

The script alias is used in all references to the CGI programs. For instance, in an HTML form, we refer to

```
<FORM method="POST" action="/cgi-bin-mark/script.pl">
```

8.6.7 Perl Script for Generating Script Aliases

A convenient way of generating script aliases for all users is to write a short Perl script which rewrites the `srm.conf` file by looking through the password file and adding a `ScriptAlias` entry for every user. In addition, a general `cgi-bin` directory is often desirable, where it is possible to place scripts which anyone can use. In the example below, we call this alias `cgi-bin-public`. Each user has a script alias called `cgi-bin-`*`username`*.

```
#!/local/bin/perl
#
# Build script aliases from password file
#

# Path to the template and real srm.conf files

$srmconf = "/local/httpd/conf/srm.conf";
$srmbase = "/local/httpd/conf/srm.conf.in";

open (OUT,">$srmconf") || die;

open (BASE,"$srm") || die;

  while (<BASE>)    # Copy base file to output
    {
    print OUT $_;
    }

close (BASE);

setpwent();

while (($name,$pw,$uid,$gid,$qu,$com,$full,$dir) = getpwent)
  {
   # SKip system accounts

  next if ($uid < 100);

  print OUT "ScriptAlias /cgi-bin-$name $dir/www cgi-bin\n";

  last if ($name eq "");
  }

close OUT;
```

8.6.8 Perl Script for Handling Errors

This Perl script can be used to generate customized or intelligent responses to error conditions.

```
#!/local/bin/perl

 #
 # Error handler
 #

# Environment variables set
#
# REDIRECT_STATUS contains the error type
```

```
# REDIRECT_URL contains the requested URL
# REDIRECT_REQUEST_METHOD e.g. GET
# REMOTE_ADDR : 192.0.2.238
# HTTP_USER_AGENT : Mozilla/4.05 [en] (X11; I; SunOS 5.6 sun4m)

if ($ENV{"REDIRECT_STATUS"} == 500)
  {
    $color = "#ff0000";
    $error_type = "Server error";
    $error_message = "An error occurred in the configuration of
     the server.<br>";
  }
elsif ($ENV{"REDIRECT_STATUS"} == 403)
  {
    $color = "#ffff67";
    $error_type = "Access restricted";
    $error_message = "Sorry, that file is not available to you.";
  }
elsif ($ENV{"REDIRECT_STATUS"} == 404)
  {
    $color = "#ffff67";
    $error_type = "File request error";
    $error_message = "The file which you accessed was not found
     or was<br>not available to you.";
  }
else
  {
    $color = "#ffff67";
    $error_type = "Unknown error";
    $error_message = "Please try again";
  }

#
# Spit out a standard format page
#
print "Content-type: text/html\n\n";
print <<END;

<html>
  <head>
    <title>$error_type</title>
  </head>

  <body bgcolor="#eeeeee">
  <img src="image.gif">

<blockquote>
<br>
<h1>$error_type</h1>

<br>
<table>
```

```
<tr>
<td>
 <table border="0" cellpadding=4>
  <tr><td bgcolor=$color>
  <br>
   $error_message
  <br><br>
</td></tr>
</table>
</td></tr>
</table>
<br><br><br><br>
Make sure that the error is not a mistake on your part. If you
  continue to have<br>
trouble, please contact the <a href="mailto:webmaster
  \@domain">Webmaster\@domain</a>.
</blockquote>

<br>

END

print "<br></body></html>";
```

The error codes are in `http_protocol.c` of the Apache distribution.

8.6.9 `mime.types` File

This file tells the server how to respond to file requests containing special data. It consists of a list of protocol names followed by a list of file extensions. Unrecognized files are displayed in a browser as text/ASCII files. If we see graphics files (like VRML files) displayed as text, then we need to add a line here to inform the server about the existence of such files. Here is a brief excerpt:

```
video/mpeg          mpeg mpg mpe
video/quicktime     qt mov
video/x-msvideo     avi
video/x-sgi-movie   movie
x-world/x-vrml      wrl
```

8.6.10 Private Directories

In some cases one requires certain information to be made available to local users of our domain but not to general outside users. This can be accomplished by using a `.htaccess` file to override the default access rights set in the server configuration files. The assumes that we have set `AllowOverride All` in an appropriate `<Directory>` structure.

Creating a directory that is only available from the local domain is a simple matter of creating the directory and creating a `.htaccess` file owned by the 'www' user (i.e. the user running the daemon) with read permission for 'www', containing the lines:

```
order deny,allow
deny from all
allow from   .local.domain
```

8.6.11 SQL/PHP

PHP is a script language rather like Perl, but it has special functions which are well suited to web programming. Amongst other things, it can be used to query an SQL database like Oracle, Sybase or MySQL. PHP has to be compiled into the WWW server at the same time that we build `httpd`.

For many sites the possibility to combine a database of information with its web pages is a powerful one. The language PHP is designed for this purpose. To use PHP it has to be coupled to a local WWW server and to a specific database. The order in which these three components is configured is important. Here is an example using MySQL, PHP and the Apache web server. This is a powerful and popular combination of free components.

For some of its graphical functions, PHP uses a library called gd. This can be obtained from ref. [164] and added with a configuration option as described below.

The first step is to compile the database engine. Note that the source tree has to remain on the system, so it should be unpacked in its final location. The other software packages expect to find the msql server in the directory `/usr/local/mysql`. One should make a special user for the mysql daemon so that it does not need to run as root.

```
cd /usr/local
tar zxf mysql-xxx.tar.gz

cd /local/mysql-xxx
configure --with-pthread \
--with-unix-socket-path=/home/mysql.socket \
--with-mysqld-user=mysql --prefix=/usr/local/mysql
make
make install
```

Note that an option is used to specify the location of a socket. The SQL database uses a Unix domain socket for internal communication. The default location for this is in the `/tmp` directory. However, placing private objects in public directories can be a security issue, so this should be placed in a private directory for the daemon. To complete the installation, we run a script which sets up a test database.

```
mysql_install_db
```

This script also sets up a 'root' account and appropriate privileges for a get-you-started database. After making and intalling the MySQL daemon, one should check that the script which starts it, called `bin/safe_mysqld` actually does the job of starting the daemon with the correct user ID. In my experience, it does not and the script has to be edited manually, inserting the option `--user=mysql` into the script at the appropriate place:

```
if test "$#" -eq 0
then
  nohup $ledir/mysqld --user=mysql \
--basedir=$MY_BASEDIR_VERSION \
      --datadir=$DATADIR >> $err_log 2>&1
else
  nohup $ledir/mysqld --user=mysql \
--basedir=$MY_BASEDIR_VERSION \
      --datadir=$DATADIR "$@" >> $err_log 2>&1
fi
```

After this the daemon may be started with the command

```
cd /local/mysql
bin/safe_mysqld +
```

Next it is necessary to unpack the WWW server. After this is done, PHP can be compiled and installed. Then one must go back and make a proper build of the WWW server. These can be kept in `/usr/local` or in site-dependent files `/local/site`, as one sees fit. Again, the source trees have to remain on the system.

```
cd /local/site
tar zxf apache_1.3.3.tar.gz
cd /local/site/apache_1.3.3

./configure --prefix=/local/site

cd /local/site
tar zxf php3.tar.gz

cd /local/site/php3
./configure --prefix=/local/site --with-apache=/local/site
 apache_1.3.3 \
            --with-mysql=/local/mysql --with-zlib \
            --with-gd=/usr/local/lib
make install

cd /local/site/apache_1.3.3
./configure --prefix=/local/site \
--activate-module=src/odules/php3/libmodphp3.a
make install
```

The name of the PHP library has changed from `libphp.a` to `libmodphp.a` through various versions of the language. There are lots of things which can go wrong with options here. These options work for building a PHP interpreter into the Apache server. We might also want to use PHP as a CGI scripting language. In that case, we need to build a binary separately. To do this one configures as follows:

```
cd /local/site/php3
./configure --prefix=/local/site \
--with-mysql=/local/mysql--with-zlib \
            --with-gd=/usr/local/lib --enable-discard-path
make install
```

8.7 E-mail Configuration

Configuration of e-mail is one of the most complex issues for the system administrator, because it involves both nagging policy decisions and technical acrobatics. For many system administrators, the phrase *Nightmares on ELM[3] street* does not conjure up a vision of Freddie Kruger, but of dark nights spent with e-mail configuration. E-mail is used for so many crucial

[3] ELM is a free mail reader written by an employee of Hewlett-Packard which redefined the standard for e-mail interfaces in the 1980s.

purposes; it is the *de facto* form of communication in a network environment [171, 6, 149, 58, 63].

Why should e-mail be so complex? Part of the trouble is that, in the past, there were many different kinds of network and many different ways of connecting up to different hosts. This made it quite a complex issue to relay messages all over the world. Today things are much simpler: most sites use the Internet protocols and some of the technical aspects of mail configuration can be simplified. Some operating systems like GNU/Linux provide a program which automatically helps set up e-mail for simple host configurations, but these are no substitute for a carefully considered e-mail system.

In this chapter we shall consider only the popular mail transfer agent `sendmail`. Sendmail changes so often that anything specific written about it is likely to be out of date by the time you read this, so this section will necessarily be of a schematic nature. The source code and documentation for `sendmail` are available from ref. [233]. No matter whether the majority of local users read mail on a PC or on a Unix workstation, every site requires a mail transfer agent like sendmail in order to handle incoming and outgoing transfers. Because of the way Unix-like operating systems multitask, and because their behaviour is well known in a network context, it is highly recommended that e-mail traffic be handled by a Unix-like operating system.

8.7.1 Models of Mail Delivery

E-mail can be delivered in two ways: either locally (where the sending host is the same as the destination host); or across a network (where the destination host is different from the sending host). Local delivery is an almost trivial problem and requires no special configuration. It applies almost exclusively to Unix-like operating systems, since it assumes that multiple users will have simultaneous access to the same host. The alternative to local delivery is to transmit mail across a network using the SMTP (actually ESMTP) mail protocol. Regardless of whether local or network delivery is used, e-mail has to end up in a *mailbox system*. For Unix-like operating systems, there are two actual choices:

- *Traditional Berkeley mailbox format*: each user has one mailbox file. New messages are appended to the end of mailbox file. Although tried and trusted, this system has certain inadequacies:
  - The mailbox is easily corrupted, leaving users in a fix.
  - The file has to be locked during mail delivery. This often results in problems.
  - NFS sharing of the mailboxes can lead to locking problems, if configured badly.
- *IMAP mailbox format*: each user has their own spool directory. New messages are written to new files, which simplifies the problem of locking during delivery. This is not strictly a mailbox format, it is a storage method. IMAP mailboxes are also available via a private network protocol (similar though superior to POP) which does not rely on NFS, so IMAP mailboxes are available from any host with an IMAP-enabled mail reader. This is a useful way of integrating e-mail across many different OS platforms.

As soon as a network is involved in e-mail transmission, there are many choices to be made. Some of the basic choices involve deciding a logistic topology for the e-mail service: should we consolidate mail services to one host, or should we deliver mail to every host

independently? The consequences of the latter are that users will have different e-mail on every host they have an account on. Usually, users require and desire only one mailbox.

One way to avoid having different e-mail on every host, is to share the mailbox file system between all hosts, using NFS. The Berkeley mail spool system is kept in one of the dirctories

```
/var/spool/mail
/usr/spool/mail
/var/mail
/usr/mail
```

depending on the flavour of operating system. To do this, we pick a special host which has the physical disk, and we force every other host to mount that disk so that users see the same mailbox, independently of which host they happen to be logged onto. This lends itself to a non-distributed solution to e-mail, however: if all mail has to end up on one disk, then the host with the disk should get the mail. If independent hosts try to perform local mail delivery to the same NFS mounted file system there can be mailbox corruption due to locking contentions across many hosts. Some sites report that this is not a problem, but it is not generally advisable to use NFS in this way. A centralized solution is preferable. For a discussion of scalable sendmail configurations see ref. [109].

Another issue which has attracted focus in recent times is whether or not a site should relay mail from other hosts, and if so which hosts. To build a flexible local mail solution, we usually need to relay mail between machines within our local domain. However, relaying of e-mail from other sites has become a security and ethical issue in recent times, with the explosion of mail 'spamming'. Hostile senders often attempt to cover their tracks by relaying e-mail via an intermediate domain. This has led the latest revisions of sendmail to revise policy on relaying. Whereas mail relaying was allowed by default, it is now denied by default. In most cases this is correct and safe behaviour; however, some sites, within particularly complex organizations, might find the need to relay e-mail from a limited number of other additional sites.

8.7.2 Consolidated and Distributed Mail Solutions

There are two main models for handling electronic mail at a domain. One is that every host receives mail independently. Since users normally have the same password and account on all of the hosts on a network, this is not usually appropriate.

The second approach is to have a mail 'hub', or central mail processor. In this model, all incoming mail is diverted to the hub and all outgoing mail is sent via the hub. With this approach, we focus all our effort into optimzing e-mail configuration on the hub, and all other machines have a simple configuration which simply collects or forwards mail on to the hub.

For mail to be diverted to a hub, we have to arrange for the mail exchanger data in DNS to point to the hub, for every system, i.e. for every host in DNS we should have an MX record accompanying the A record:

```
hostname   A    xxx.yyy.zzx.mmm
           MX   mailhub.domain
```

Without such an MX record, mail which is addressed to

```
user@hostname.domain
```

will be sent directly to hostname. With such a record the mail for hostname is sent to `mailhub.domain` instead. It can later be forwarded to hostname if desired using a mailertable. This has several advantages: first it means that mail configuration can be centralized, spam filtering can be performed even for dumb hosts and aliases can be expanded here without the need for a distributed alias database like NIS. The second advantage concerns security. If all mail is forced to pass through this hub, then a secure set-up here will prevent SMTP attacks on weaker hosts, so this also simplifies the security administration of mail. One may concentrate most of one's effort on securing this hub, knowing that nothing very bad will happen. A further precaution is then to configure the site router to accept SMTP traffic only for the mailhub, since it is supposed to go there anyway. In that way, if one forgets an MX record in DNS there will be no back-doors for would-be attackers.

8.7.3 Compiling and Installing Sendmail

In this section we shall look only at the mail agent called *Berkeley sendmail*. This is the most up to date version of sendmail, and all sites should strongly consider using this in favour of older vendor versions of the program[4]. Sendmail is very susceptible to attack from the network, and only the Berkeley version is well enough equipped with deal with this threat.

Information out about sendmail and the latest version can be obtained from ref. [233]. After unpacking the distribution, we need to compile it. Sendmail is one of those old dinosaur programs which does not follow modern standards of configuration and installation. It retains its historical baggage as a part of the BSD Unix source tree, and so it is advisable to perform the configuration and installation largely by hand. Note that a new release of the GNU make program is required in order to parse the Makefiles in the `sendmail` source code.

Before doing this, we have to make sure that we have all of the libraries needed to compile. Sendmail uses BIND and TCP-wrappers libraries. Consider searching for the latest versions of these libraries on the Internet before compiling. BIND is the resolver library. The official place to get BIND is ref. [25]. This also contains a library `lib44bsd.a` which might be necessary. The latest version of TCP wrappers may be obtained from ref. [269]. Many of the database-lookup features require the Berkeley `db` package. This is obtainable from ref. [64].

Here is an example for sendmail-8.9.3:

```
host# cd sendmail-8.9.3
host# ls

FAQ         READ_ME         cf.tar   mail.local praliases src
KNOWNBUGS   RELEASE_NOTES   contrib  mailstats  rmail     test
Makefile    cf              doc      makemap    smrsh

host# cd src
host# sh makesendmail
Configuration: os=SunOS, rel=5.5.1, rbase=5, rroot=5.5,
 arch=sun4, sfx=
Making in obj.SunOS.5.5.1.sun4
...
```

[4] Some newer alternatives to sendmail now exist, such as `smail` and `exim`, but I am not in a position to evaluate their merits or problems.

The script `makesendmail` selects the operating system type and compiles the program. In the process it creates a directory for the compilation. In the example above, it creates

```
obj.SunOS.5.5.1.sun4
```

We might still have to edit the Makefile in this new directory, so do a `CTRL-C` to stop the compilation and edit the file which corresponds to

```
obj.SunOS.5.5.1.sun4/Makefile
```

in the example above.

In the Makefile, we can switch on and off several features. The first thing to do is to switch off NIS alias lookups, since these can cause a large overhead. The best way to do alias lookups is to use *only* the Berkeley `db` database. That means editing out the line beginning DBMDEF=, as in the example below. Use of NIS, NIS+ or other databases is not recommended for several reasons, the main one being that it not required as long as mail passes through a central hub, which is a good approach to mail configuration for small to medium sized sites.

We must also decide where we want the distribution to be placed. A good place is `/usr/local/mail` or `/usr/local/site/mail`, in order to separate our local modifications from the operating system. Create this directory now and make a subdirectory `bin` where executable files will be kept.

Here is an example Makefile for a system that will not use NIS or NISPLUS, but will make use of the Berkeley database. This is well-suited to a hub model of mail configuration, where aliases and other databases only need to be on the hub.

```
# This Makefile is designed to work on the old "make" program.
# It does not use the obj subdirectory. It also does not install
# documentation automatically -- think of it as a quick start
# for sites that have the old make program (I recommend that you
# get and port the new make if you are going to be doing any
# significant work on sendmail).
#
# This has been tested on Solaris 2.5.
#
#      @(#)Makefile.SunOS.5.5 8.10 (Berkeley) 4/13/97
#

# use O=-O (usual) or O=-g (debugging)
# warning: do not use -O with versions of gcc prior to 2.6
O=      -O

CC=     gcc
DESTDIR = /local/mail

# define the database mechanism used for alias lookups:
#       -DNDBM -- use new DBM
#       -DNEWDB -- use new Berkeley DB
#       -DNIS -- include NIS support
# The really old (V7) DBM library is no longer supported.
# See READ_ME for a description of how these flags interact.
#
```

```
# DBMDEF= -DNDBM -DNIS -DNISPLUS
DBMDEF= -DNEWDB

# environment definitions (e.g., -D_AIX3)
ENVDEF= -DSOLARIS=20500

# see also conf.h for additional compilation flags

# include directories
INCDIRS=-I/usr/sww/include

# library directories
LIBDIRS=-L/usr/sww/lib -L/local/lib

# libraries required on your system
# delete -l44bsd if you are not running BIND 4.9.x
# add -ldb if you add -DNEWDB above (in DBMDEF)
# LIBS= -lresolv -l44bsd -lsocket -lnsl -lkstat
LIBS= -lwrap -lresolv -l44bsd -lsocket -lnsl -lkstat -ldb

# location of sendmail binary (usually /usr/sbin or /usr/lib)
BINDIR=$DESTDIR/lib

# location of sendmail.st file (usually /var/log or /usr/lib)
STDIR=$DESTDIR/log

# location of sendmail.hf file (usually /usr/share/misc or /
usr/lib) HFDIR=$DESTDIR/etc

...
```

When we have compiled the program successfully, the finished executable files must be installed. The files are best copied manually like this:

```
cp obj.SunOS.5.5.1.sun4/sendmail /usr/local/mail/bin/sendmail
cp obj.SunOS.5.5.1.sun4/makemap /usr/local/mail/bin/makemap
```

Our operating system most likely expects to find the sendmail executable file in either the `/usr/lib/` directory, or the `/usr/sbin` directory on newer systems. We must replace the old executable in these directories by making a link to the new executable. For example:

```
mv /usr/lib/sendmail /usr/lib/sendmail.org
ln -s /usr/local/mail/bin/sendmail /usr/lib/sendmail
```

8.7.4 Configuring Sendmail

To finish off the installation, we need to create configuration files for our mail domain. Begin by going back to the sendmail distribution and copying the `cf` directory to the mail directory, like this:

```
cp -r sendmail-8.9.3/cf /usr/local/mail
```

Next make a `lib` directory:

```
mkdir /usr/local/mail/lib
```

To create a `sendmail.cf` file, we need to create a so-called macro file `/usr/local/mail/lib/domain.mc`. Here is an example file for domain `domain.country`. We should only need to change the domain name and the OS name of the mailhost in the first three lines. Using this file we will be able to build the sendmail configuration more or less automatically. This example is for sendmail-8.9.3 for a mail hub:

```
divert(-1)
include('/local/site/mail/cf/m4/cf.m4')

VERSIONID('$Id: mercury.mc,v 1.1 1997/04/08 08:52:28 mroot Exp mroot $')
OSTYPE(solaris2)dnl
DOMAIN(domain.country)dnl

MASQUERADE_AS(domain.country)
MASQUERADE_DOMAIN(sub.domain.country)

FEATURE(use_cw_file)
FEATURE(use_ct_file)
FEATURE(redirect)
FEATURE(relay_entire_domain)
FEATURE(always_add_domain)
FEATURE(allmasquerade)
FEATURE(masquerade_envelope)
FEATURE(domaintable, 'hash -o /local/site/mail/lib/domaintable')
FEATURE(mailertable, 'hash -o /local/site/mail/lib/mailertable')
FEATURE(access_db, 'hash -o /local/site/mail/lib/access_db')
FEATURE(genericstable, 'hash -o /local/site/mail/lib/genericstable')
FEATURE(virtusertable, 'hash -o /local/site/mail/lib/virtusertable')

FEATURE(local_procmail,'/local/bin/procmail')

GENERICS_DOMAIN_FILE(/local/site/mail/lib/sendmail.cG)

EXPOSED_USER(root)

define('ALIAS_FILE', /local/site/mail/lib/aliases)dnl
define('HELP_FILE', /local/site/mail/lib/sendmail.hf)dnl
define('STATUS_FILE', /local/site/mail/etc/sendmail.st)dnl
define('QUEUE_DIR', /var/spool/mqueue)
define('LOCAL_MAILER_CHARSET', iso-8859-1)dnl
define('SMTP_MAIL_CHARSET', iso-8859-1)dnl
define('SMTP_MAIL_MAX','2000000')
define('confMAX_MESSAGE_SIZE', '20000000')
define('confHOST_STATUS_DIRECTORY', '.hoststat')
define('confPRIVACY_FLAGS', 'authwarnings,noexpn,novrfy')
define('confME_TOO', 'True')
define('confMIME_FORMAT_ERRORS', 'False')
define('confTIME_ZONE', 'MET-1METDST')
define('confDEF_CHAR_SET', 'iso-8859-1')
define('confEIGHT_BIT_HANDLING', 'm')
define('confSMTP_MAILER', 'esmtp')
define('confCW_FILE', '/local/site/mail/lib/sendmail.cw')
define('confCT_FILE', '/local/site/mail/lib/sendmail.ct')
define('confUSERDB_SPEC', '/local/site/mail/lib/userdb.db')
define('LOCAL_SHELL_PATH','/local/site/mail/bin/smrsh')

MAILER(local)
MAILER(smtp)

FEATURE(rbl) dnl vixie's black hole database for spammers
```

Next create /usr/local/mail/Makefile which will build the configuration for us:

```
MAKEMAP=        bin/makemap
SENDMAIL=       bin/sendmail
PIDFILE=        /etc/mail/sendmail.pid
MAILTABLE=      lib/mailertable
MCFILE=         lib/nexus.mc
ALIASES=        lib/aliases
USERDB=         lib/userdb
GENERICS=       lib/genericstable
ACCESSDB=       lib/access_db
CF_DIR=         cf/

all: nullclient.cf sendmail.cf $(ALIASES).db $(GENERICS) \
.db $(MAILTABLE).db $(USERDB).db $(ACCESSDB).db .restart

$(ALIASES).db: $(ALIASES)
        $(SENDMAIL) -bi

$(USERDB).db: $(USERDB)
        $(MAKEMAP) btree $(USERDB) < $(USERDB)

$(ACCESSDB).db: $(ACCESSDB)
        $(MAKEMAP) hash $(ACCESSDB) < $(ACCESSDB)

$(MAILTABLE).db: $(MAILTABLE)
        $(MAKEMAP) hash $(MAILTABLE) < $(MAILTABLE)

$(GENERICS).db: $(GENERICS)
        $(MAKEMAP) hash $(GENERICS) < $(GENERICS)

sendmail.cf: $(MCFILE)
        m4 -D_CF_DIR_=$(CF_DIR) cf/m4/cf.m4 $(MCFILE) >
          sendmail.cf


nullclient.cf: lib/nullclient.mc
        m4 -D_CF_DIR_=$(CF_DIR) cf/m4/cf.m4 lib/nullclient/
          .mc > nullclient.cf

.restart: sendmail.cf lib/sendmail.cw lib/access_db.db lib/
 mailertable.db
        kill -1 'head -1 (PIDFILE)
        touch .restart
```

This is a shorter example for a system attached to a mail hub, whose only function is to send the mail to the hub for processing:

```
divert(0)

OSTYPE(solaris) dnl
FEATURE(nullclient, mailhost.domain.country)dnl
MASQUERADE_AS(domain.country)dnl
define('MAIL_HUB','mailhost.domain.country')
define('SMART_HOST','mailhost.domain.country')
```

Typing `make` in the `/usr/local/mail` directory should now result in a configuration file `/usr/local/mail/sendmail.cf`. Wait until you have read the next section before doing this.

We will need to create a file `lib/sendmail.cw` which contains a list of possible machines or domains for which the sendmail program will accept mail. It is, amongst other things, this file which allows us to send mail of the form `mark@domain.country`, i.e. to an entire domain, without specifying a particular machine. This file should contain a list of all the valid addresses, like this:

```
domain.country
mailhost.domain.country
www.domain.country
mercury.domain.country
dax.domain.country
borg.domain.country
worf.domain.country
daystrom.domain.country
regula.domain.country
ferengi.domain.country
lore.domain.country
```

Finally, we need to make the key files readable for normal users. There is no harm in giving everyone read access to all the files and directories.

8.7.5 Rewriting Outgoing Addresses

Some organizations like to make their outgoing e-mail addresses look professional by using full name addresses rather than user names, which can look unfriendly. There are arguments for and against using full names. One argument against is that full names are not necessarily unique, so some kind of fudging mechanism has to be used to avoid name collisions in reply addresses. However, full name mail addresses do look nice, even if they take twice as long to type. Some universities use the practice of using student numbers as user names, or class registration, which leads to extremely difficult-to-remember mail addresses. In that case, a full name alias would be welcome.

The `lib/userdb` file and `lib/aliases` file are used by sendmail for resolving incoming aliases. It is common for sites to create mail aliases for all their users, by taking the full name of each user and joining the pieces with dots. For example, user `mark` with full name 'Mark Burgess' would map to an alias `Mark.Burgess`.

Aliases of this kind look nice and are user friendly to outsiders, but they do not work unless we configure this specially. To do this we need to create files `lib/aliases` and `lib/userdb`.

For each user, the `userdb` file should contain a line of the form

```
Mark.Burgess@domain.country:maildrop mark@mailhost.domain \
.country
```

which tells sendmail where to deliver incoming mail that is sent to the user alias 'Mark.Burgess' at the mail hub. Outgoing messages are processed if we have a line of the form

```
mark:mailname Mark.Burgess@domain.country
```

This means that mail messages which originate from 'mark' will be rewritten so that they look as though they orginate from 'Mark.Burgess@domain.country'. This form usually only applies to mail which originates from the local mail host, since that is the only case where the outgoing name is just the short user-name 'mark'. Mail which is sent from other hosts to the mail hub for outgoing processing generally produces from-information in the form 'mark@-domain.country'. For this to be rewritten, we need a line of the form as well:

```
mark@domain.country:mailname Mark.Burgess@domain.country
```

The `aliases` file is set up as in section 8.7.10. Rewriting of outgoing mail should be handled adequately by these two mechanisms alone, but sendmail is nothing if not inconsistent, and it will fail in some circumstances when mail is routed through a mail hub from client hosts. If the client hosts do not masquerade their outgoing addresses, we have to fix the problem manually. There is another table introduced into Berkeley sendmail, called `genericstable` with the format:

```
mark@host1.college.edu Mark.Burgess@college.edu
mark@host2.college.edu Mark.Burgess@college.edu
```

If rewriting fails for mail sent from special hosts, through the mail hub, this file seems to fix the problem.

8.7.6 `smrsh`

The sendmail remote shell is a security measure to prevent system crackers from executing an arbitrary program on the system. The `smrsh` program is contained in the sendmail distribution and is configured by using the FEATURE.

8.7.7 Spam and Junk Mail

Spam mail is e-mail which is sent repetetively as a would-be denial of service attack. The word comes from the Monty Python spam song sketch. Junk mail is unwanted mail, often advertisements about financial opportunities or pornography, sometimes hoaxes. Often these two kinds of unwanted mail are lumped together and called collectively 'spam'.

Spam has become a major problem, since it is very easy to send e-mail and very hard to pick out what is useful from what is useless. There are two approaches to the filtering of spam, both of which are needed together:

- Site rules for rejecting mail (ACLs).
- Private user-rules for rejecting mail.

The reason why both of these is needed is that what one user wants to reject, another user might be glad to receive. Users prospecting for finanical opportunities or collecting the latest 'artwork' might live for the messages which most of us get annoyed with.

Sendmail has rules for filtering mail at the site level. These include the ability to deny access to connecting mailers from certain domains. At the time of writing they seem to be only partially successful in practice [115].

At the user level, users of `procmail` can use `junkfilter` to create their own rules for rejecting spam (see ref. [207]).

8.7.8 Policy Decisions

To protect our site from e-mail attacks, even ones made in innocence, we might want to restrict mail by other criteria too. For example, multimedia attachments can now allow users to send huge files by e-mail. This is a very inefficient way of sending large amounts of data, and it causes problems for mailbox storage space. A possibility is to limit the size of mail messages handled by sendmail so that mail which is too large will be rejected with an error message. For example, the following rules limit e-mail to approximately 20MB. Even with such a large reject size, a handful of messages per month are rejected on the basis of this rule:

```
define('SMTP_MAIL_MAX','2000000')
define('confMAX_MESSAGE_SIZE', '20000000')
```

Again, this must be a policy decision like garbage collection of users' files. It is never desirable to restrict the personal freedom of users, but it becomes a matter of survival. If one provides an opening, it will be exploited either through ignorance or malice.

8.7.9 Filtering Outgoing Mail

An organization might want to prevent certain types of e-mail from being sent. For example, mail generated by CGI scripts is impossible to trace to a specific user, but is stamped with the domain name of the WWW server which sent it. CGI mail is therefore readily abused, and many institutions would therefore disallow it. If ordinary users are allowed to write their own CGI scripts, however, this can be a difficult problem to contain. Unfortunately, sendmail does not currently permit true filtering of outgoing e-mail, only filtering of *relayed* mail, so if the web server and mail filter are on the same host, this is not possible. One can *discard* such mail, however, with a local rule of the form:

```
HReturn-Path: $>local_ret_path
D{SpamMessage}"553 You are a spammer. Go away."

Slocal_ret_path
R<www>                  $#discard $: discard
R<www>                  $#error $@ $: ${SpamMessage}
```

This is not terribly sociable, since no-one will be informed that the mail was discarded. The error message above does not work in sendmail 8.9.3.

8.7.10 The Mail Queue

When mail cannot be delivered immediately it is placed in the mail queue. To see what mail is waiting in the mail queue, we must log onto the mail server (i.e. the host which handles outgoing mail) on the network. We can see what is in the queue by typing

```
mercury% mailq
```

or

```
mercury% sendmail -bp
```

These two forms are equivalent. We can force sendmail to process the mail queue by typing:

```
mercury% sendmail -q -v
```

8.7.11 Mail Aliases

One of the first things to locate on a system is the `sendmail` alias file. This is a file which contains e-mail aliases for users and system services. Common locations for this file are `/etc/aliases` and `/etc/mail/aliases`. On some systems, the mail aliases are in the NIS network database.

If this file actually lies in the `/etc` directory, or some other place amongst the system files, then we should move it to your special area for site-dependent files and make a symbolic link to `/etc/aliases` instead. Mail aliases are valuable and we want to make sure that nothing happens to them if we reinstall the OS.

The format of the mail aliases file is as follows:

```
# Alias for mailer daemon; returned messages from our MAILER-
  DAEMON
# should be routed to our local Postmaster.

postmaster: mark, toreo

MAILER-DAEMON: postmaster

nobody: /dev/null

#
# alias: list of addresses
#

sysadm:mark@domain.country,toreo@domain.country
root:sysadm

#
# Alias for distribution list, members specified elsewhere:
#  alias: :include:file of names
#

maillist: :include:/iu/mercury/local/maillists

#
# Dump mail to a file
#

archive: /iu/mercury/local/archive/email.archive
```

8.8 Mounting NFS Disks

The sharing of disks over the network is the province of the NFS (Network File System). Unix disks on one host may be accessed across the network by other UNIX hosts, or by PCs running PC-NFS. A disk attached physically to a host called a *server* is said to be *mounted* on

a client host. To maintain a certain level of security, the server must give other hosts permission to mount disks. This is called *exporting* or *sharing* disks.

8.8.1 Server Side Exporting

To mount a disk on a server we must export the disk to the client (this is done on the server), and we must tell the client to mount the disk. Permission to mount disks is given on the server in a file which is called `/etc/exports` or on recent SVR4 hosts `/etc/dfs/dfstab`. The format for information in these files differs from system to system, so one should always begin by looking at the manual page for these files. Here are two examples. The first is from GNU/Linux:

```
# See exports(5) for a description.
# This file contains a list of all dirs exported to other
   computers.
# It is used by rpc.nfsd and rpc.mountd.

/iu/borg/local daystrom(rw) worf(rw) nanite(rw) *.domain/
   .country(ro)
```

In this example, a file system called `/iu/borg/local` is exported read-write explicitly to the client hosts `daystrom`, `worf` and `nanite`. It is also exported read-only to any host in the domain `domain.country`. This last feature is not available on most types of Unix.

On some brands of Unix (such as SunOS 4.1.*), one must run a command after editing this file in order to register the changes. The command is `exportfs -a` to export all file systems. The command `exportfs` alone shows which file systems are currently exported, and to whom.

Our second example is from Solaris (SVR4). The file is called `/etc/dfs/dfstab`. Under Solaris, one can use the `share` command to export file systems manually from the shell, using a command line of the form

```
share -F nfs -o rw=hostname file system
```

The `/etc/dfs/dfstab` file is in fact a shell script which simply executes such a command for each file system of interest. This has several advantages over traditional export files, since one may define variables, as in the example below.

```
# place share(1M) commands here for automatic execution
# on entering init state 3.
#
# share [-F fstype] [ -o options] [-d "<text>"] <pathname> \
   [resource]
# .e.g,
# share -F nfs -o rw=engineering -d "home dirs" /export/home2

hostlist=starfleet:axis:ferengi:borg:worf:daystrom:worf \
.domain.country:daystrom.domain.country:nostromo:voyager \
:aud4:aud4.domain.country:aud1:aud1.domain.country:aud2 \
:bajor:nostromo:galron:ds9:thistledown:rama

share -F nfs -o rw=$hostlist /iu/mercury/local
share -F nfs -o rw=$hostlist,root=starfleet /iu/mercury/u1
```

```
share -F nfs -o rw=$hostlist,root=starfleet /iu/mercury/u2
share -F nfs -o rw=$hostlist,root=starfleet /iu/mercury/u3
share -F nfs -o rw=$hostlist,root=starfleet /iu/mercury/u4
share -F nfs -o rw=$hostlist /var/mail
```

This script exports the six named file systems, read-write to the entire list of hosts named in the variable `hostlist`. The command `shareall` runs this script, or it can be run manually by typing `sh /etc/dfs/dfstab`. The command `share` without arguments shows the currently exported file systems. Notice that the host name `daystrom` is repeated, once unqualified, and again with a fully qualified host name. This is sometimes necessary in order to make the entry recognized. The mount daemon is not particularly intelligent when it verifies host names. Some systems send the fully qualified name to verify, and others send the unqualified name. If in doubt, list both.

8.8.2 Client Side Mounting

Clients may mount any subdirectory of the exported directory onto any local directory by becoming `root` and either executing a shell command of the form

```
mount server:remote-directory local-directory
```

or by adding a line to the file system table file, usually called `/etc/fstab`. On some brands of Unix, this file has been renamed as `/etc/checklist` or `/etc/file systems`. On Solaris systems it is called `/etc/vfstab`. The advantage of writing the disks in the file system table is that the mount commands will not be lost when we reboot our system. The file systems in the file system table file are mounted automatically when the system is booted. All the file systems in this file are mounted with the simple command `mount -a`.

We begin by looking at the manual page on the appropriate file for the system, or better still, looking at examples which are already in the file. The form of a typical file system table is as below[5]:

```
/dev/sda2                 swap             swap  rw,bg    1  1
/dev/sda1                 /                ext2  rw,bg    1  1
/dev/sda3                 /iu/borg/local   ext2  rw,bg    1  1
mercury:/iu/mercury/u1     /iu/mercury/u1     nfs      rw,bg
mercury:/iu/mercury/u2     /iu/mercury/u2     nfs      rw,bg
mercury:/iu/mercury/u3     /iu/mercury/u3     nfs      rw,bg
mercury:/iu/mercury/local /iu/mercury/local nfs      rw,bg
```

This example is from GNU/Linux. Notice the left-hand column. These are disks which are to be mounted. The first disks which begin with `/dev` are local disks, physically attached to the host concerned. Those which begin with a host name followed by a colon (in this case host `mercury`) are NFS file systems which lie physically on the named host. The second column in this table is the name of a directory on which the disk or remote file system is to be mounted, i.e. where the files are to appear in the local host's file-tree. The remaining columns are options and file system types: `rw` means mount for read and write access, and `bg` means

[5] On older HPUX systems, there is a bug which causes mysterious numbers to appear in the `/etc/checklists` file. These have no meaning.

'background' which tells `mount` to continue trying to mount a file system in the background if it fails on a first attempt.

Editing the `/etc/fstab` (or equivalent) file is a process which can be automated very nicely with the help of the system administration tool cfengine. We shall discuss this in the next chapter.

8.8.3 Troubleshooting NFS

If you get a message telling you 'Permission denied' when you try to mount a remote file system, you may like to check the following:

- Did you remember to add the name of the client to the `export` or `dfstab` file on the server?
- Some systems require a fully qualified host name (i.e. host name with domain name appended) in the export file. Try using this.
- Did you mis-spell the name of the client or the server?
- Are the correct network daemons running which support NFS? On the server side, you must be running `mountd` or `rpc.mountd`. This is an authentication daemon. The actual transfer of data is performed by `nfsd` or `rpc.nfsd`. On older systems there should be at least four of these daemons running to handle multiple requests. Modern systems use a multi-threaded version of the program, so that only one daemon is required.

  On the client side, some systems use the binary input/output daemon to make transfers more efficient. This is not strictly necessary to get NFS working. This daemon is called `biod` on older systems and `nfsiod` on newer systems like FreeBSD. Solaris no longer makes use of this daemon, its activities are now integrated into a kernel thread.
- The portmapper (`portmap` or `rpcbind`) is a strange creature. On some Unix-like systems, particularly GNU/Linux, the portmapper requires an entry in the TCP wrapper file `/etc/hosts.allow` in order for it to accept connections. Otherwise, you might see the error

  ```
  RPC service not registered.
  ```

  The portmapper requires numerical IP addresses in the TCP wrapper configuration. Host names will not do, for security reasons (see section 8.4.5).
- The exports file on GNU/Linux hosts is also somewhat unusual. If you are using a non-standard netmask, it is necessary to tell the mount daemon:

  ```
  # /etc/exports: the access control list for file systems which
  #               may be exported # to NFS clients. See exports(5).
  /site/cube/local *.college.edu/255.255.255.0(rw)
  /site/cube/local 192.0.2./255.255.255.0(rw)
  ```

8.9 The Printer Service

Printing services vary from single printers coupled to private workstations to huge consolidated spooling services serving large organizations [283, 212]. Host print services need to be

told about available printers by registering the printers in a local database. In BSD-like print servers this database is kept in a flat file called `/etc/printcap`. In System V print servers, a program called `lpadmin` is used to register printers, and it's anyone's guess what happens to that information.

The way in which we register printers thus depends upon

- What kind of Unix we are using.
- Whether we are running any special network printer software [205, 90].

The main difference is between BSD-like systems and System V. Recently, a replacement print service was introduced for a generic heterogeneous network. Called LPRng, this package preserves the simplicity of the BSD system while providing superior functionality to both [205].

To register a printer with a BSD-like printer service, we do the following:

- Think of a name for the printer.
- Decide whether it is going to be connected directly to a host or standalone on the network.
- Register the printer with the printing system so that the daemons which provide the print service know how to talk to it. This can inlcude manually making a 'spool' directory for its queue files. This normally lies under `var/spool` or `/usr/spool`.

```
mkdir /var/spool/printer-name
```

- Most Unix systems assume the existence of a *default* printer which is referred to by the name 'lp'. If one does not specify a particular printer when printing, your data are sent to the default printer. It is up to us to name or alias one of our printers 'lp'. Each printer may have several names or aliases.

With some print spoolers, we also need to decide whether to send all data to a common central server, or whether to let each host handle its own negotiations for printing. If we are interested in maintaining a record of how many pages each user has printed, then a centralized solution is a much simpler option. The downside of this is that, if there is a large user base, the traffic might present a considerable load for one host. A central print spooler must have sufficient disk space to temporarily store all the incoming print jobs.

8.9.1 BSD Printer with `lpd`

The file `/etc/printcap` is used to register a printer with a BSD system. If we are lucky, there might be a script or a user interface for helping to register a printer; if not follow the recipe below.

The format of the `/etc/printcap` file can be quite simple in most cases. The manual page for `printcap` contains a description of the file format. This file consist of a list of entries (each on a single line, or split over several lines using the continuation character \). A simple template entry looks like this:

```
printer-name-1|printer-alias-1|printer-\alias-2
:lp=:\
```

```
:sd=spool-directory:\
:rm=remote machine or IP address of printer:\
:rp=name of remote printer on remote machine:
```

This file should be installed on all hosts which need to access the printer, regardless of whether the printer is physically attached to them or not. Here is an example which registers two printers. The first is called 'myprinter', and is connected physically to the remote host `mercury`. The second is a standalone printer which we have named 'diff-engine' and which has IP address `192.0.2.99`[6].

```
#
# /etc/printcap
#
myprinter|lp|default|SPARCprinter, a Sun SPARCprinter: \
:lp=:\
:lf=/var/adm/lpd-errs:\
:sd=/var/spool/VirtualLight:\
:rm=mercury:\
:rp=myprinter:

diff-engine|HP laser stand-alone:\
        :lp=:\
        :lf=/var/adm/lpd-errs:\
        :sd=/var/spool/otherprint:\
        :rm=192.0.2.99:\
        :rp=(none):
```

Note that the `rp` field exists in case a given printer has a different name on the remote host to the one we have given it locally on our machine. On a standalone printer this name is irrelevant.

8.9.2 System V

The `lpadmin` command is used to install printers under System 5. This commmand is complex and has many command line options to add and remove printers. If you have System 5 systems, consult the manual page on your system. Sun Microsystems provide a script front-end to help simplify this procedure called `AdminTool`.

8.9.2 LPRng

A recent and welcome addition to the printer debate is the Next Generation LPR package by Patrick Powell [205]. LPRng is a drop in replacement for both BSD and System 5 print systems. It is configured quite simply in a manner very similar (but not identical) to the Berkeley printcap system. LPRng can be obtained from `http://www.astart.com/lprng/LPRng.html` and it is a real god-send if one has System 5 (e.g. Solaris) hosts to grapple with.

Suggestion 12 (Unix printing) *Install LPRng on all hosts in the network. Forget about trying to understand and manage the native printing systems on System V and BSD hosts. LPRng can replace them all with a system which is at least as good.*

[6] This address is set up on the printer when one installs it.

If one follows this suggestion there is only a single printer system to worry about. Note that some GNU/Linux distributions (e.g. Debian) have adopted this system, so it will not need to be installed from scratch.

The software uses a printcap file and two other optional files called `lpd.conf` and `lpd.perms`. The printcap file is like a regular printcap file but without the backslash continuation characters. LPRng provides effectively both `lpr`, `lpd`, `lpq` and `lprm` commands from Berkeley, and `lp`, `lpstat` and `cancel` commands from System 5. The daemon reads the three configuration files and handles spooling. The configuration is challenging but straightforward, and there is extensive documentation. Here is a simple example for a network printer (with its own IP address) which allows logged on users to start and delete their own printjobs:

```
# /etc/printcap (lprng)

myprinter|lp
        :if=/local/bin/lpf              # LF/CR filter
        :af=/var/spool/lpd/acctfil
        :lf=/var/spool/lpd/printlog
        :sd=/var/spool/myprinter
        :lp=xxx.yyy.zzz.mmm%9100
        :rw
        :sh
```

The IP address of the printer is *xxx.yyy.zzz.mmm* and it must be written in numerical form. The percent symbol marks the standard port 9100. The `lpd.conf` file is slightly mysterious, but has a number of useful options. Most, if not all, of these can also be set in the printcap file, but options set here apply for all printers. One nice feature, for instance, is the ability to reject printouts of binary (non-printable) files. This can save a few rain forests if someone is kind enough to dump `/bin/ls` to the printer.

```
#
# lpd.conf
#

# Purpose: name of accounting file (see also la, ar)
af=/var/spool/lpd/acctfil

# Purpose: accounting at start (see also af, la, ar)
as=jobstart $H $n $P $k $b $t

# Purpose: check for nonprintable file
check_for_nonprintable

# Purpose: default printer
default_printer=local

# Purpose: error log file (servers, filters and prefilters)
lf=/var/adm/printlog

# Purpose: lpd lock file
lockfile=/var/spool/lpd/lpd.lock.%h

# Purpose: lpd log file
logfile=/var/spool/lpd/lpd.log.%h
```

```
# Purpose: /etc/printcap files
printcap_path=/etc/printcap

# Purpose: suppress headers and/or banner page
sh
```

The `lpd.perms` file sets limits on who can access the printers and from where, unlike the traditional services which are open to everyone.

```
#
# lpd.perms
#
# allow root on server to control jobs
ACCEPT SERVICE=C SERVER REMOTEUSER=root
# allow anybody to get status
ACCEPT SERVICE=S
# reject all others, including lpc commands permitted by
  user_lpc
REJECT SERVICE=CSU
#
# allow same user on originating host to remove a job
ACCEPT SERVICE=M SAMEHOST SAMEUSER
# allow root on server to remove a job
ACCEPT SERVICE=M SERVER REMOTEUSER=root
REJECT SERVICE=M
# All other operations disallowed
DEFAULT REJECT    # orACCEPT
```

LPRng claims to support Berkeley printcap files directly. In trials its behaviour has been quirky, however, with some things working and others not. In any event, LPRng is a highly welcome piece of software which works supremely well, once configured.

8.9.4 Environment Variable `PRINTER`

The BSD print command and some application programs read the environment variable `PRINTER` to determine which printer destination to send data to. The System V the print command `lp` does not.

8.9.5 BSD Print Queue

- `lpr -p` *printer* `file` Send file to named print queue.
- `lpq` Show the printer queue for the default printer, or the printer specified in the environment variable `PRINTER` if this is set. This lists the queue-ids.
- `lprm` *queue-id* Remove a job from the print queue. Get the queue id using `lpq`.
- `lpd` Start the print service. (Must be killed to stop again.)
- `lpc` An incredibly stupid user interface for print administration. This program tells lies.

8.9.6 System V Print Queue

- `lp -d` *printer file* Send a file to the named print queue.

- `lpstat -o all` Show the printer queue for the default printer. This lists the queue-ids.
- `lpstat -a` Tells lies about when the print service was started.
- `lpsched` Start the print service.
- `lpshut` Stop the print service.
- `cancel` *queue-id* Remove a job from the print queue. Get the queue id using `lpstat`.

The Solaris operating system has an optional printing system called Newsprint in addition to the SVR4 printing commands.

Exercises

Exercise 8.1 Set up an Apache web server.

Exercise 8.2 Build a tree of documents, where some files are public and others are restricted to access by your local organization, using the `.htaccess` file capability.

Exercise 8.3 Show that a CGI script can always be written which reveals all of the files restricted using `.htaccess`. This shows that untrusted CGI scripts are a security risk.

Exercise 8.4 Write a Perl script for handling WWW errors at your site.

Exercise 8.5 Estimate the number of megabytes transferred per week by the file servers at you domain. Could any of this traffic be avoided by reorganizing the network?

Exercise 8.6 Where are the default name servers placed around your network? Is there a name server on each subnet, i.e. does DNS lookup traffic have to pass through a router?

Exercise 8.7 Set up TCP wrappers on your system (Unix-like OSes only).

Chapter 9

Principles of Security

Computer security is about protecting the data and availability of computing systems. In order to have security, we must sacrifice a certain level of convenience [144]. The key words are *access*, *privacy*, *integrity* and *trust*. To understand computer security we have to understand the inter-relationships between all of the hosts and services on our networks, as well as the ways in which those hosts can be accessed. A system can be compromised by:

- Malicious attacks.
- Accidental erasure of data.
- Disk crashes.
- User ignorance.

Protecting against these issues requires both pro-active (preventative) measures and damage control after breaches.

Security is an increasingly important problem. Just in the last few years the number of attacks and break-ins to computer systems has risen to millions of cases a year. Crackers[1] have found their way inside the computers of the Pentagon, the world's security services, warships, fighter plane command computers, banks and major services such as electrical power grids. With this kind of access the potential for causing damage is great. Computer warfare is the next major battlefield we have to conquer. It is happening now, as you read these words. It is here, like it or not. Moreover, it is estimated that the banks lose millions of dollars a year to computer crime.

Security can also embrace other issues such as reliability. For instance, many computers are used in mission-critical systems, such as aircraft controls and machinery, where human lives are at stake. Thus, reliability and safety are also concerns. *Real-time systems* are computer systems which are guaranteed to respond in real-time to every request which is made of them. That means that a real-time system must always be fast enough to cope with any demand which is made of it. Real time systems are required in cases where human lives and huge sums of money are involved. For instance, in a flight control system it would be unacceptable to give a command 'Oh my goodness, we're going to crash, flaps NOW!' and have the computer reply with 'Processing, please wait...'.

Security is a huge subject, because modern computer systems are complex and the connectivity of the Internet means that millions of people can try to break into networked

[1] It is incorrect to call intruders hackers: hackers are legitimate programmers.

systems. In this chapter we consider the basic principles of security. Having studied this, you might wish to read more about security in refs. [103, 98, 44, 39, 237].

9.1 Physical Security

For a computer to be secure it must be physically secure. If we can get our hands on a host then we are never more than a screwdriver away from all of its assets. Disks can be removed. Sophisticated users can tap network lines and listen to traffic. The radiation from monitor screens can be captured and recorded, showing an exact image of what a user is looking at on his/her screen. Or one can simply look over the shoulder of a colleague while he or she types a password. The level of physical security one requires depends upon the sophistication of the potential intruder, and therefore in the value of the assets which one is protecting.

Assuming that hosts are physically secure, we then still have to deal with the issues of software security, which is a much more difficult topic. Software security is about access control and software reliability. No single tool can make computer systems secure. Major blunders have been made out of the belief that a single product (e.g. a 'firewall') would solve the security problem. The bottom line is that there is no such thing as a secure operating system, firewall or no firewall. What is required is a persistent mixture of vigilance and adaptability.

9.2 Four Independent Issues

For many, security is perceived as being synonymous with network privacy or network intrusion. Privacy is one aspect of security, but the network is not our particular enemy. Many breaches of security happen from within. In reality, there is little difference between the dangers of remote access by network or direct access from a console: privacy is about access control, no matter where a hostile user might be. If we focus exclusively on network connectivity we ignore a possible threat from internal employees (e.g. the janitor who is a computer expert and has an axe to grind, or the mischievous son of the director who was left waiting to play in mom's office, or perhaps the unthinkable: a disgruntled employee who feels as though his/her talents go unappreciated). Software security is a vast subject, because modern computer systems are complex. It is only exacerbated by the connectivity of the Internet which allows millions of people to have a go at breaking into networked systems. What this points to is the fact that a secure environment requires a tight control of access control on every host individually, not merely at specific points such as firewalls.

If we stretch our powers of abstraction even to include loss by natural disaster, then system security can be summarized by a basic principle:

Principle 43 (Security) *The fundamental requirement for security is the ability to restrict access and privilege to data.*

The word privilege does not apply well to loss by accident or natural disaster, but the word access does. If accidental actions or natural disasters do not have access to data, they they cannot cause them any harm. Any attempt to run a secure system where restriction of access is not possible is fundamentally flawed.

There are four basic elements in security:

- *Privacy*: restriction of access.
- *Authentication*: verification of identity.
- *Trust*: trusting the source.
- *Integrity*: protection against corruption or loss (redundancy).

9.3 Trust Relationships

There are many implicit trust relationships in computer systems: it is crucial to understand them. If we do not understand where we are placing our trust, that trust can be exploited by attackers who have thought more carefully than we have.

For example, any Unix NFS server which shares users' home directories trusts the root user on the hosts which mount those directories. Some bad accidents are prevented by mapping root to the user 'nobody' on remote systems, but this is not security, only convenience. The root user can always use 'su' to becomes any user in its password file and access/change any data within those file systems. The `.rlogin` and `hosts.equiv` files on Unix machines grant root (or other user) privileges to other hosts without the need for authentication.

When collecting software from remote servers, we should make sure that they come from a machine that is trustworthy, particularly if the files could lead to privileged access to the system. For example, it would be an extremely foolish idea to copy a binary program such as the Unix program `/bin/ps` from a host one knows nothing about. This program runs with root privileges. If someone were to replace that version of `ps` with a Trojan horse command, the system would have effectively been opened to attack.

Most users trust anonymous FTP servers where they collect free software. In any remote copy we are setting up an implicit trust relationship. First of all, we trust the integrity of the host we are collecting files from. Secondly, we trust that they have the same user name database with regard to access control. The root user on the collecting host has the same rights to read files as the root user on the server. The same applies to any matched user name.

In any remote file transfer one is also forced to trust the integrity of the data received. No matter how hard a program may work to authenticate the identity of the host, even once the host's identity is verified, the accuracy or trustworthiness of unknown data is still in doubt. This has nothing to do with encryption, as users sometimes believe: encrypted connections do not change these trust relationships: they improve the privacy of the data being transmitted, but neither their accuracy nor trustworthiness.

Implicit trust relationships lie at the heart of so many software systems which grant access to services or resources that it would be impossible to list them all here. Trust relationships are important to grasp because they can lead to security holes.

9.4 Security Policy

Security only has meaning when we have defined what we mean by it. Defining what the local community, means by security is essential. Only then will we know when security has been breached, and what to do about it. Some sites which contain sensitive data require strict

security and spend a lot of time enforcing it, others do not particularly care about their data, and would rather not waste their time on pointless measures to protect them. Security must be balanced against convenience [144]. How secure must we be

- From outside the organization?
- From inside the organization (different host)?
- From inside the organization (same host)?
- Against the interruption of services?
- From user error?

Finally, how much inconvenience are the users of the system willing to endure in order to uphold this level of security? This point should not be under-estimated: if users consider security to be a nuisance, they try to circumvent it.

Principle 44 (Work defensively) *Expect the worst, do your best, preferably in advance of a problem.*

Visible security can be a problem in itself. Systems which do not implement high level security tend to attract only low-level crackers – and those who manage to break in tend to use the systems only as a springboard to go other places. The more security one implements, and the more visible it is, the more of a challenge it is for a cracker. So spending a lot of time on security might only have the effect of asking for trouble.

Principle 45 (Network Security) *Extremely sensitive data should not be placed on a computer which is attached in any way to a public network.*

What resources are we trying to protect?

- *Secrets*: some sites have secrets they wish to protect. They might be government or trade secrets or the solutions to a college exam.
- *Personell data:* in your country there are probably rules about what you must do to safeguard sensitive personal information. This goes for any information about employees, patients, customers or anyone else we deal with. Information about people is private.
- *CPU usage/System downtime:* we might not have any data that we are afraid will fall into the wrong hands. It might simply be that the system is so important to that we cannot afford the loss of time incurred by having someone screw it up. If the system is down, everything stops.
- *Abuse of the system:* it might simply be that we do not want anyone using our system to do something for which they are not authorized, like breaking into other systems.

Who are we trying to protect them from?

- *Competitors*, who might gain an advantage by learning your secrets.
- *Malicious intruders*. Note that people with malicious intent might come from inside or outside our organization. It is wrong to think that the enemy is simply everyone outside of our domain. Too many organizations think 'inside/outside' instead of dealing with

proper access control. If one *always ensures that systems and data are protected* on a need-to-know basis, then there is no reason to discriminate between inside or outside of an organization.

- *Old employees with a grudge* against the organization.

Next: what will happen if the system is compromised?

- Loss of money.
- Threat of legal action against you.
- Missed deadlines.
- Loss of reputation.

How much work will we need to put into protecting the system? Who are the people trying to break in?

- Sophisticated spies.
- Tourists, just poking around.
- Braggers, trying to impress.

Finally: *what risk is acceptable?* If we have a secret which is worth 4 Lira, would we be interested in spending 5 Lira to secure it? Where does one draw the line? How much is security worth?

The social term in the security equation should never be forgotten. One can spend a hundred thousand dollars on the top of the range firewall to protect data from network intrusion, but someone could walk into the building and look over an unsuspecting shoulder to obtain it instead, or use a receiver to collect the stray radiation from your monitors. Are employees leaving sensitive printouts lying around? Are we willing to place our entire building in a Faraday cage to avoid remote detection of the radiation expelled by monitors? In the final instance, someone could just point a gun at someone's head and ask nicely for their secrets. Some examples of security policies can be found at refs. [1, 202, 203, 17].

9.5 Protecting from Loss

Prevention of loss is better than recovery, after the fact. Any preventative measures we can take are worth the investment.

9.5.1 Loss of Data: Backup

The data collected and produced by an organization are usually the primary reason for them owning a computer installation. The loss of those data, for whatever reason, would be a catastrophe, second to none.

Data can be lost by accident, by fire or natural catastrophe, by disk failure, or even vandalism. If you live in a war-zone or police state, you might also have to protect data from bombs or brutal incursions onto your premises. Once destroyed, data cannot be recovered. The laws of thermodynamics dictate this. So, to avoid complete data loss, you need to employ a policy of *redundancy*, i.e. you need to make several copies of data,

and make sure that they do not befall the same fate. Of course, no matter how many copies of data you make, it is possible that they might all be destroyed simultaneously, no matter what you do to protect them, but we are aiming to minimize the likelihood of that occurrence.

Principle 46 (Data invulnerability) *The purpose of a backup copy is to provide an image of data which is unlikely to be destroyed by the same act that destroys the original.*

There is an obvious corollary:

Corollary 47 *Backup copies should be stored at a different physical location to the originals.*

The economics of backup has changed in recent times for several reasons: first, storage media are far more reliable than they once were. If a disk does not show signs of a problem within a few months then it will probably never fail of its own accord, before you change the whole machine on other grounds. Disks tolerate continuous usage for perhaps five years, after which time you will almost certainly want to replace them for other reasons, e.g. performance. The other important change is the almost universal access to networks. Networks can be used to transport data simply and cheaply from one physical location to another.

Traditionally, backups have been made to tape, since tape is relatively cheap and mobile. This is still the case at many sites, particularly larger ones; but tapes usually need to be dealt with manually, by a human or by an expensive robot. This adds a price tag to tape-backup which smaller institutions can find difficult to manage. By way of contrast, the price of disks and networking has fallen dramatically. For an organization with few resources, a cheap solution to the backup problem is to mirror disks across a network [206], using well-known tools like `rdump`, `rdist` or `cfengine`. This solves the problems of redundancy and location; and, for what it costs to employ a human or tape robot, one can purchase quite a lot of disk space.

Another change is the development of fast, reliable media like CD-ROM. In earlier times, it was normal to back up the operating system partitions of hosts to tape. Today that practice is nonsense: the operating system is readily available on a CD-ROM which is at least as fast as a tape streamer and consumes a fraction of the space. It is only necessary to make backups of whatever special configuration files have been modified locally. Sites which use cfengine can simply allow cfengine to reconstruct local modifications after an OS installation. In any event, if we have followed the principle of separating operating system from local modifications, this is no problem at all.

Similar remarks can be made about other software. Commercial software is now sold on CD-ROM and is trivial to re-install (remember only to keep a backup of license keys). For freely available software, there are already many copies and mirrors at remote locations by virtue of the Internet. For convenience, a local source repository can also be kept, to speed up recovery in the case of an accident. In the unlikely event of every host being destroyed simultaneously, downloading the software again from the network is the least of your worries!

Reconstructing a system from source rather than from backup has never been easier than now. Moreover, a policy of not backing up software which is easily accessible from source,

can make a considerable saving in the volume of backup space required, at the price of more work in the event of accident. In the end this is a matter of policy.

It should be clear that user data must have maximum priority for backup. This is where local creativity manifests itself; these are the data which form your assets.

9.5.2 Loss of Service

Loss of service might be less permanent than the loss of data, but it can be just as debilitating. Downtime costs money for businesses and wastes valuable time in academia.

The basic source of all computing power is electricity. Loss of electrical power can be protected against, to a limited extent, with an *Uninterruptible Power Supply* (UPS). This is not an infallible security, but it helps to avoid problems due to short breaks in the power. UPS solutions use a battery backup to keep the power going for a few hours when power has failed. When the battery begins to run down, they can signal the host so as to take it down in a controlled fashion, thus minimizing damage to disks and data. Investing in a UPS for an important server could be the best thing one ever does. Electrical spike protectors are another important accessory for anyone living in a region where lightning strikes are frequent, or where the power supply is of variable quality. No fuse will protect a computer from a surge of electricity: microelectronics burn out much quicker than any fuse.

Service can also be interrupted by a breach of the network infrastructure: a failed router or broken cable, or even a blown fuse. It can be interrupted by cleaning staff, and carelessness. A backup or stand-by replacement is the only option for hardware failure. It helps to have the telephone number of those responsible for network hardware when physical breaches occur.

Software can be abused in a *denial of service attack*. Denial of service attacks are usually initiated by sending information to a host which confuses it into inactivity. There are as many variations on this theme as there are vandals on the network. Some attacks exploit bugs, while others are simply spamming episodes, repeatedly sending a deluge of service requests to the host, so that it spends all of its resources on handling the attack.

9.6 System and Network Security

Since the explosion of interest in the Internet, the possibility of hosts being attacked from outside sources has become a significant problem. With literally millions of users on the net, the tiny percentage of malicious users becomes a large number.

9.6.1 Security through Obscurity

There is a commonly held belief that, if one makes it difficult for intruders to find out information, they will not bother to try. This is completely wrong. Often the reverse is true. If attackers can see that there is nothing worth finding they will leave systems alone. If everything is concealed they will assume that there must be something interesting worth breaking in for.

Security through obscurity is a naive form of security which at best delays break-ins. Shadow passwords are an example of this. By making the encrypted password list

inaccessible to normal users, one makes it harder for them to automate the search for poor passwords, but one does not prevent it! It is still possible to guess passwords in exactly the same way as before, but it takes much longer. The NT password database is not in a readable format. Some people have claimed that this makes it more secure than the Unix password file. Since then tools have been written which rewrite the NT password file in Unix format with visible encrypted passwords. In other words, making it difficult for people to break in does not make it impossible.

Clearly there is no need to give away information to potential intruders. Information should be available to *everyone* on a need-to-know basis, whether they be local users or people from outside the organization. But at the same time, obscurity is no real protection. Even the invisible man could get shot. Time spent securing systems is better than time spent obscuring them. Obscurity might attract more attention than we want and make the system as obscure to us as to a potential intruder.

9.6.2 Honey-pots and Sacrificial Lambs

A *honey pot* is a host which is made to look attractive to attackers. It is usually placed on a network with the intention of catching an intruder or distracting them from more important systems. A *sacrificial lamb* host is one which is not considered to be particularly important to the domain. If it is compromised by an attacker then that is an acceptable loss and no real harm is done.

Some network administrators believe that the use of such machines contributes to security. For example, WWW servers are often placed on sacrificial lamb machines which are placed outside firewalls. If the machine is compromised then it can simply be reinstalled and the data reloaded from a secure backup. This practice might seem rather dubious. There is certainly no evidence to support the idea that either honey pot havens or sacrificial lamb chops actually improve security.

9.6.3 Security Holes

One way that outside users can attack a system is by exploiting security holes in software. Classic examples usually involve *setuid-root* programs, which give normal users temporary superuser access to the system. Typical examples are programs like `sendmail` and `finger`. These programs are constantly being fixed, but even so, new security holes are found with alarming regularity. Faults in software leave back-doors open to intruders. The only effective way of eliminating such attacks is to build a so-called *firewall* around your network (see section 10.6).

The computer emergency response team (CERT) was established in the wake of the Internet Worm incident to monitor potential security threats. CERT publish warnings to a mailing list about known security holes. This is also available on the newsgroup *comp.security.announce*. Several other organizations, often run by staff who work as security consultants, are now involved in computer security monitoring. For instance, the SANS organization performs an admirable job of keeping the community informed about security developments, both technical and political. Moreover, old phreaker organizations like Phrack and the l0pht (pronounced loft) now apply their extensive knowledge of system vulnerabilities for the good of the network community. See refs. [4, 259, 50, 237, 199, 160].

9.6.4 System Homogeneity

In a site with a lot of different kinds of platforms, perhaps several Unix variants, NT and Windows 9x, the job of closing security holes is much harder. Inhomogeneity often provides the determined intruder with more possibilities for bug-finding. You might ask yourself whether you need so many different kinds of platform. If you do, then perhaps a firewall solution would provide an extra level of protection, giving you a better chance of being able to upgrade your systems before something serious happens.

9.6.5 Modem Pools

Some companies expend considerable effort to secure their network connections, but forget that they have dial-in modems. Modem pools are a prime target for attackers because they are often easy targets. There are many problems associated with modem pools. Sometimes they are quite unexpected. For example, if one has network access to Windows systems using the same modem, then those systems are automatically on a shared segment and can use one another's resources, regardless of whether they have any logical relationship. Modems can also succumb to denial of service attacks by repetitive dialling.

Modems should never allow users to gain access to a part of the network which needs to be secure. Modems should never be back-doors into firewalled networks.

9.6.6 Laptops

Laptop computers are increasingly popular and they are popular targets for thieves. There have been cases of laptop computers being stolen containing sensitive information, often enough to give crackers access to further systems, or simply to give competitors the information they wanted!

9.6.7 Backups

If you make backups of important data (private data), then we must take steps to secure the backups also. If an intruder can steal your backups, then he/she doesn't need to steal the originals.

9.7 Social Engineering

Network attackers (i.e. system crackers) are people. It is easy to become so consumed by a fascination of the network, that we forget that people can just walk into a building and steal something *in the real world.* If one can avoid complex technical expertise in order to break into a system, then why not do it? There is more than one way to crack a system.

The only secure computer is a computer which is locked into a room, not connected to a network, shielded from all electromagnetic radiation. In social studies of large companies, it has been demonstrated that – in spite of expensive firewall software and sophisticated anti-cracking technology – all most crackers had to do to break into the system was to make a

phone call to an unwary employee of the company and ask for their user name and password [281]. Some crackers posed as system administrators trying to fix a bug, others simply questioned them as in a marketing survey until they gave away information which allowed the crackers to guess their passwords. Some crackers will go one step further and visit the building they are trying to break into, going through the garbage/refuse to look for documents which would give clues about security. Most people do not understand the lengths that people will go to to break into systems if they really want to.

Another social phenomenon which motivates break-ins is bragging. Crackers who have broken into the system like to tell people that they have been there and done that. They sometimes try to scare administrators by telling them how much damage they have caused. Here the trick is not to panic and do something hasty, but to try to verify what the crackers claim. In many cases it is nonsense, empty words.

There is a few things which can be done to counteract social threats.

- Never disclose information over the telephone, especially through a voice-mail system. Phone calls can be spoofed.
- Examine system logs, check the system regularly, run cfengine to ensure consistency.
- Make hard-copies of messages sent with all the headers printed out. Most people don't know how to hide their true identity on the Internet.
- Make proper backups of the system regularly.
- Inform the whole system of attacks so that anyone who knows something can help you.
- Do not assume that what crackers tell you is true. Make a judgment as to whether you want to act or ignore the event.
- Have a clear security policy. Death threats, serious or not, should probably be reported to the company responsible for sending the message, and perhaps even the police, but not the person sending the message.

9.8 TCP/IP Security

On top of the hardware, there are many levels of protocol which make network communication work. Many of these layers are invisible or are irrelevant to us, but there are two layers in the protocol stack which are particularly relevant to network security, namely the IP layer and the TCP/UDP layer.

9.8.1 The Internet Protocol (IPv4)

The Internet Protocol was conceived in the 1970s as a military project. The aim was to produce a routable network protocol. The version of this protocol in use today is version 4, with a few patches. Let's revise some of the basics of IPv4, which we discussed earlier in the operating systems course. TCP/IP is a transmission protocol which builds on lower level protocols like Ethernet and gives it extra features like 'streams' or virtual circuits, with automatic handshaking. UDP is a cheaper version of this protocol which is used for services that do not require connection-based communication. The TCP/IP protocol stack consists of several layers (see Figure 9.1).

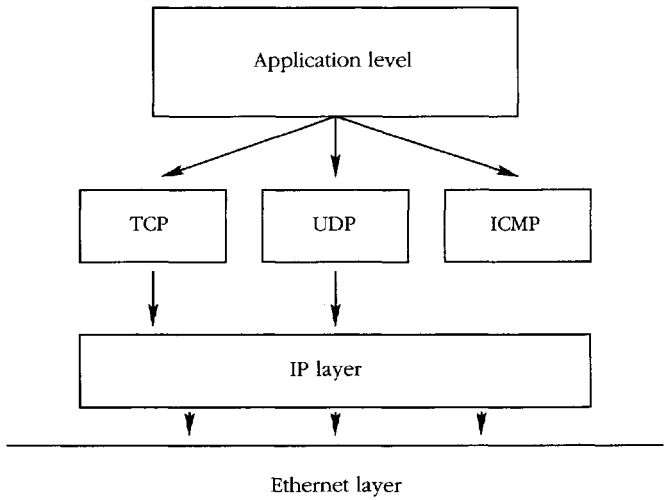


Figure 9.1 The Internet Protocol stack

At the application level we have text-based protocols like Telnet and FTP, etc. Under these lies the TCP (Transmission Control Protocol), which provides reliable connection based handshaking, in a virtual circuit. TCP and UDP introduce the concept of the port and the socket (=port+IP address). We base our communications on these, so we also base the security of our communications on these. Under TCP/UDP is the IP transport layer, then Ethernet or token ring, etc. ICMP is a small protocol used by network hardware to send control and error messages as a part of the IP protocol set, e.g. ping uses ICMP.

With all of its encapsulation packaging, a TCP/IP packet looks as in Figure 9.2. TCP packets are reliable connection oriented data. They form *streams* or continuous data-flows with handshaking. This is accomplished by using a three-way handshake based on so-called SYN (synchronize) and ACK (acknowledge) bits in the TCP header. Suppose *host A* wishes to set up a connection with *host B*. *Host A* sends a TCP segment to *host B* with its SYN bit set and a sequence number X which will be used to keep track of the order of the segments. *Host B* replies to this with its SYN and ACK bits set, with Acknowledgement=X+1 and a new sequence number Y. *Host A* then replies to *host B* with the first data and the Acknowledge field=Y+1. The reason why each side acknowledges every segment with a sequence number which is one greater than the previous number sent is that the Acknowledgement field actually determines the next sequence number expected. This sequence is a weakness which network attackers have been able to exploit through different connections, in

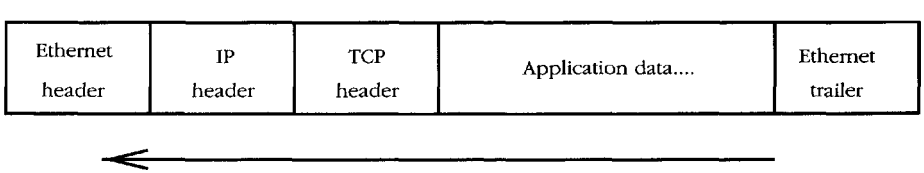


Figure 9.2 Encapsulation with Ethernet and TCP/IP

'sequence number guessing' attacks. Now many implementations of TCP allow random initial sequence numbers.

The purpose of this circuit connection is to ensure that both hosts know about every packet which is sent from source to destination. Because TCP guarantees delivery, it retransmits any segment for which is has not received an ACK after a certain period of time (the TCP timeout).

At the end of a transmission the sender sends a FIN (finished) bit, which is replied to with FIN/ACK. In fact, closing connections is quite complex since both sides must close their end of the connection reliably. See the reference literature for further details of this.

Let us consider Telnet as an example, and see how the Telnet connection looks at the TCP level (see Figure 9.3). Telnet opens a socket from a random port address (e.g. 54657) to a standard well-known port (23) where the Telnet service lives. The combination of a port number at an IP address, over a communication channel is called a socket. The only security in the Telnet service lies in the fact that port 23 is a reserved port which only root can use. (Ports 0–1023 are reserved.)

The TCP protocol guarantees to deliver data to their destination in the right order, without losing anything. To do this it breaks up a message into segments and numbers the parts of the message according to a sequence. It then confirms that every part of that sequence has been received. If no confirmation of receipt is received, the source retransmits the data after a timeout. The TCP header contains handshaking bits. Reliable delivery is achieved through a three-way handshake. Host A begins by sending host B a packet with a SYN (synchronize) bit set and a sequence number. This provides a starting reference for the sequence of communication. Host B replies to this message with a SYN,ACK which confirms receipt of an open-connection request and provides a new sequence number which confirms identity. Host A acknowledges this. Then B replies with actual data. We can see this in an actual example (see Figure 9.4). This handshaking method of sending sequence numbers with the acknowledgement allows the TCP protocol to guarantee and order every piece of a transmission. The ACK return values are incremented by one because in earlier implementations this would be the next packet required in the sequence. This predictability in the sequence is unfortunately a weakness which can be exploited by so-called sequence guessing attacks. Today, in modern implementations, sequence numbers are randomized to avoid this form of attack. Older operating systems still suffer from this problem. Future implementations of TCP/IP will be able to solve this problem by obscuring the sequence numbers entirely through encryption.

The TCP handshake is useful for filtering traffic at the router level, since it gives us something concrete to latch on to. TCP would rather drop a connection than break one of

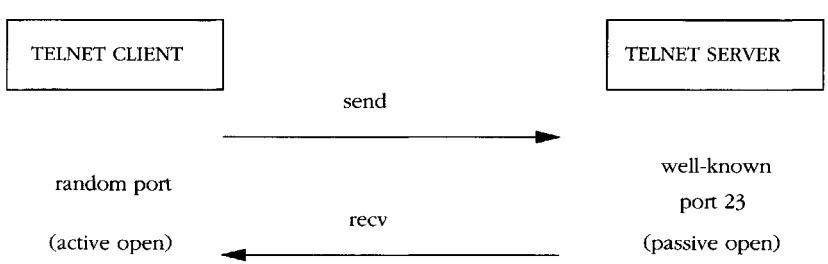


Figure 9.3 A Telnet connection

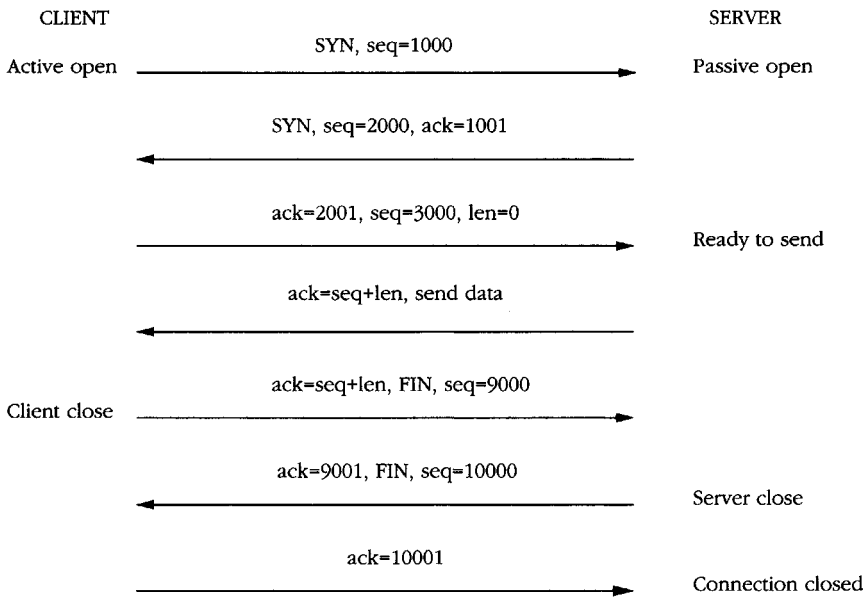


Figure 9.4 The TCP three-way handshake

its promises about data integrity, so if we want to block Telnet connections, say, we only have to break one part of this fragile loop. The usual strategy is to filter all incoming connections which do not have their ACK bit set, using router filtering rules. This will prevent any new connections from being established with the 'outside'. We can, on the other hand, allow packets which come from inside the local network. This provides a simple router-level firewall protection. It is useful for stopping IP spoofing attempts. The UDP protocol does not have SYN,ACK bits, and so it is more difficult to filter.

9.8.2 Example Telnet Session

Aside from the theory, it is helpful to see a real example. Although slightly cumbersome, it is very informative to see how the communication actually takes place. The first thing we see is how inefficient the Telnet protocol is, how passwords are transmitted in clear text over the network, and how fragmentation and retransmission of IP fragments is performed to guarantee transmission. Notice also how the sequence numbers are randomized.

```
from% telnet to.domain.country
Trying 192.0.2.238...
Connected to to.domain.country
Escape character is '^]'.

SunOS 5.6

login: mark
Password:
SunOS Release 5.6 Version Generic [UNIX(R) System V Release 4.0]
```

```
  [/etc/motd]
to% echo hei
to% exit

-------------------------------------------

  Send Syn to establish connection, + random Seq

    from -> to ETHER Type=0800 (IP), size = 58 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=44, ID=53498
    from -> to TCP D=23 S=54657 Syn Seq=4095044366 Len=0
              Win=8760
    from -> to TELNET C port=54657

-------------------------------------------

  Reply with Syn,Ack and Ack=prev Seq+1

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=44, ID=43390
to -> from TCP D=54657 S=23 Syn  Ack=4095044367 Seq=826419455
          Len=0 Win=8760
to -> from TELNET R port=54657

-------------------------------------------

  Reply with Ack = prev Seq+1

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53499
    from -> to TCP D=23 S=54657  Ack=826419456 Seq=4095044367
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 81 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=67, ID=53500
    from -> to TCP D=23 S=54657  Ack=826419456 Seq=4095044367
              Len=27 Win=8760
    from -> to TELNET C port=54657

-------------------------------------------

  Now send data: ack = seq + Len each time until Fin

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43391
to -> from TCP D=54657 S=23  Ack=4095044394 Seq=826419456
          Len=0 Win=8760
to -> from TELNET R port=54657
 (retrans)
to -> from ETHER Type=0800 (IP), size = 69 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=55, ID=43396
to -> from TCP D=54657 S=23  Ack=4095044394 Seq=826419456
          Len=15 Win=8760
to -> from TELNET R port=54657

-------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53504
```

```
    from -> to TCP D=23 S=54657  Ack=826419471 Seq=4095044394
              Len=0 Win=8760
    from -> to TELNET C port=54657

 (retrans with different Len! = fragmentation, same Ack)

    from -> to ETHER Type=0800 (IP), size = 66 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=52, ID=53505
    from -> to TCP D=23 S=54657  Ack=826419471 Seq=4095044394
              Len=12 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 69 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=55, ID=43397
to -> from TCP D=54657 S=23    Ack=4095044394 Seq=826419471
          Len=15 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53506
    from -> to TCP D=23 S=54657  Ack=826419486 Seq=4095044406
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 75 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=61, ID=43398
to -> from TCP D=54657 S=23  Ack=4095044406 Seq=826419486
          Len=21 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 120 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=106, ID=53507
    from -> to TCP D=23 S=54657  Ack=826419507 Seq=4095044406
              Len=66 Win=8760
    from -> to TELNET C port=54657 \377\372\30\0VT100\377\360
              \377\372#\0from

(Transfers TERM variable - VT100)
```

```
to -> from ETHER Type=0800 (IP), size = 75 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=61, ID=43399
to -> from TCP D=54657 S=23  Ack=4095044472 Seq=826419507
          Len=21 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53508
    from -> to TCP D=23 S=54657  Ack=826419528 Seq=4095044472
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=46, ID=43400
to -> from TCP D=54657 S=23  Ack=4095044472 Seq=826419528
          Len=6 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 60 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=46, ID=53509
    from -> to TCP D=23 S=54657  Ack=826419534 Seq=4095044472
              Len=6 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=43, ID=43401
to -> from TCP D=54657 S=23  Ack=4095044478 Seq=826419534
          Len=3 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53510
    from -> to TCP D=23 S=54657  Ack=826419537 Seq=4095044478
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 61 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=47, ID=43402
to -> from TCP D=54657 S=23  Ack=4095044478 Seq=826419537
          Len=7 Win=8760
to -> from TELNET R port=54657 login:
```

```
  Here comes the login name
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53511
    from -> to TCP D=23 S=54657  Ack=826419544 Seq=4095044478
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans, bad Len)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53512
    from -> to TCP D=23 S=54657  Ack=826419544 Seq=4095044478
              Len=1 Win=8760
    from -> to TELNET C port=54657 m
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43403
to -> from TCP D=54657 S=23    Ack=4095044479 Seq=826419544
          Len=1 Win=8760
to -> from TELNET R port=54657 m
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53513
    from -> to TCP D=23 S=54657  Ack=826419545 Seq=4095044479
              Len=0 Win=8760
    from -> to TELNET C port=54657
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53514
    from -> to TCP D=23 S=54657  Ack=826419545 Seq=4095044479
              Len=1 Win=8760
    from -> to TELNET C port=54657 a
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43404
to -> from TCP D=54657 S=23  Ack=4095044480 Seq=826419545
          Len=1 Win=8760
to -> from TELNET R port=54657 a
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53515
    from -> to TCP D=23 S=54657  Ack=826419546 Seq=4095044480
              Len=0 Win=8760
    from -> to TELNET C port=54657
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53516
    from -> to TCP D=23 S=54657  Ack=826419546 Seq=4095044480
              Len=1 Win=8760
    from -> to TELNET C port=54657 r
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43405
to -> from TCP D=54657 S=23  Ack=4095044481 Seq=826419546
          Len=1 Win=8760
to -> from TELNET R port=54657 r
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53517
    from -> to TCP D=23 S=54657  Ack=826419547 Seq=4095044481
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53518
    from -> to TCP D=23 S=54657  Ack=826419547 Seq=4095044481
              Len=1 Win=8760
    from -> to TELNET C port=54657 k
-------------------------------------------------
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43406
to -> from TCP D=54657 S=23  Ack=4095044482 Seq=826419547
          Len=1 Win=8760
to -> from TELNET R port=54657 k
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53519
    from -> to TCP D=23 S=54657  Ack=826419548 Seq=4095044482
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 56 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=42, ID=53520
    from -> to TCP D=23 S=54657  Ack=826419548 Seq=4095044482
              Len=2 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=42, ID=43407
to -> from TCP D=54657 S=23  Ack=4095044484 Seq=826419548
          Len=2 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53521
    from -> to TCP D=23 S=54657  Ack=826419550 Seq=4095044484
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 64 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=50, ID=43408
to -> from TCP D=54657 S=23  Ack=4095044484 Seq=826419550
          Len=10 Win=8760
to -> from TELNET R port=54657 Password:
```

```
  Here comes the password, in plain text, for all to see!
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53522
    from -> to TCP D=23 S=54657  Ack=826419560 Seq=4095044484
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53523
    from -> to TCP D=23 S=54657  Ack=826419560 Seq=4095044484
              Len=1 Win=8760
    from -> to TELNET C port=54657 p
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43409
to -> from TCP D=54657 S=23  Ack=4095044485 Seq=826419560
          Len=0 Win=8760
to -> from TELNET R port=54657 p
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53524
    from -> to TCP D=23 S=54657  Ack=826419560 Seq=4095044485
              Len=1 Win=8760
    from -> to TELNET C port=54657 a
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43410
to -> from TCP D=54657 S=23  Ack=4095044486 Seq=826419560
          Len=0 Win=8760
to -> from TELNET R port=54657 a
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53525
    from -> to TCP D=23 S=54657  Ack=826419560 Seq=4095044486
              Len=1 Win=8760
    from -> to TELNET C port=54657 s
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43411
to -> from TCP D=54657 S=23  Ack=4095044487 Seq=826419560
          Len=0 Win=8760
to -> from TELNET R port=54657 s
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53526
    from -> to TCP D=23 S=54657  Ack=826419560 Seq=4095044487
              Len=1 Win=8760
    from -> to TELNET C port=54657 w
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43412
to -> from TCP D=54657 S=23  Ack=4095044488 Seq=826419560
          Len=0 Win=8760
to -> from TELNET R port=54657 w
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=53530
    from -> to TCP D=23 S=54657  Ack=826419560 Seq=4095044491
              Len=1 Win=8760
    from -> to TELNET C port=54657 d
-------------------------------------------------
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43416
to -> from TCP D=54657 S=23  Ack=4095044492 Seq=826419560
          Len=0 Win=8760
to -> from TELNET R port=54657 d
```

```
    from -> to ETHER Type=0800 (IP), size = 56 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=42, ID=53531
    from -> to TCP D=23 S=54657   Ack=826419560 Seq=4095044492
              Len=2 Win=8760
    from -> to TELNET C port=54657 \n
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=42, ID=43417
to -> from TCP D=54657 S=23  Ack=4095044494 Seq=826419560
          Len=2 Win=8760
to -> from TELNET R port=54657
  (fragment)
to -> from ETHER Type=0800 (IP), size = 357 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=343, ID=43484
to -> from TCP D=54657 S=23    Ack=4095044494 Seq=826419562
          Len=303 Win=8760
to -> from TELNET R port=54657 SunOS Release 5.6 Ve
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=53599
    from -> to TCP D=23 S=54657   Ack=826419865 Seq=4095044494
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 130 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=116, ID=43487
to -> from TCP D=54657 S=23  Ack=4095044494 Seq=826419865
          Len=76 Win=8760
to -> from TELNET R port=54657 1:33pm up 2 day(s
  (fragment)
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=42, ID=43882
to -> from TCP D=54657 S=23  Ack=4095044494 Seq=826419941
          Len=2 Win=8760
to -> from TELNET R port=54657
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54316
    from -> to TCP D=23 S=54657   Ack=826419943 Seq=4095044494
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
to -> from ETHER Type=0800 (IP), size = 101 bytes
```

```
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=87, ID=43887
to -> from TCP D=54657 S=23  Ack=4095044494 Seq=826419943
          Len=47 Win=8760
to -> from TELNET R port=54657 You have mail (total

-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54319
    from -> to TCP D=23 S=54657  Ack=826419990 Seq=4095044494
              Len=0 Win=8760
    from -> to TELNET C port=54657

-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=45, ID=43890
to -> from TCP D=54657 S=23  Ack=4095044494 Seq=826419990
          Len=5 Win=8760
to -> from TELNET R port=54657 prompt\%

-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43891
to -> from TCP D=2049 S=1023  Ack=4258218482 Seq=1642166507
          Len=0 Win=8760

-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54320
    from -> to TCP D=23 S=54657  Ack=826419995 Seq=4095044494
              Len=0 Win=8760
    from -> to TELNET C port=54657

-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=54321
    from -> to TCP D=23 S=54657  Ack=826419995 Seq=4095044494
              Len=1 Win=8760
    from -> to TELNET C port=54657 e

-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43892
to -> from TCP D=54657 S=23  Ack=4095044495 Seq=826419995
          Len=1 Win=8760
to -> from TELNET R port=54657 e

-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54322
    from -> to TCP D=23 S=54657  Ack=826419996 Seq=4095044495
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=54323
```

```
    from -> to TCP D=23 S=54657  Ack=826419996 Seq=4095044495
               Len=1 Win=8760
    from -> to TELNET C port=54657 c
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43893
to -> from TCP D=54657 S=23  Ack=4095044496 Seq=826419996
           Len=1 Win=8760
to -> from TELNET R port=54657 c
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54324
    from -> to TCP D=23 S=54657  Ack=826419997 Seq=4095044496
               Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=54325
    from -> to TCP D=23 S=54657  Ack=826419997 Seq=4095044496
               Len=1 Win=8760
    from -> to TELNET C port=54657 h
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43894
to -> from TCP D=54657 S=23  Ack=4095044497 Seq=826419997
           Len=1 Win=8760
to -> from TELNET R port=54657 h
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54326
    from -> to TCP D=23 S=54657  Ack=826419998 Seq=4095044497
               Len=0 Win=8760
    from -> to TELNET C port=54657
  (frag)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=54327
    from -> to TCP D=23 S=54657  Ack=826419998 Seq=4095044497
               Len=1 Win=8760
    from -> to TELNET C port=54657 o
-------------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43895
to -> from TCP D=54657 S=23  Ack=4095044498 Seq=826419998
           Len=1 Win=8760
to -> from TELNET R port=54657 o
-------------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54328
    from -> to TCP D=23 S=54657  Ack=826419999 Seq=4095044498
               Len=0 Win=8760
```

```
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=54329
    from -> to TCP D=23 S=54657  Ack=826419999 Seq=4095044498
               Len=1 Win=8760
    from -> to TELNET C port=54657
------------------------------------------------------------
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43896
to -> from TCP D=54657 S=23  Ack=4095044499 Seq=826419999
          Len=1 Win=8760
to -> from TELNET R port=54657
------------------------------------------------------------
    from -> to ETHER Type=0800 (IP), size = 56 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=42, ID=54333
    from -> to TCP D=23 S=54657  Ack=826420001 Seq=4095044500
               Len=2 Win=8760
    from -> to TELNET C port=54657 ei
------------------------------------------------------------
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43898
to -> from TCP D=54657 S=23  Ack=4095044502 Seq=826420001
          Len=1 Win=8760
to -> from TELNET R port=54657 e
------------------------------------------------------------
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54334
    from -> to TCP D=23 S=54657  Ack=826420002 Seq=4095044502
               Len=0 Win=8760
    from -> to TELNET C port=54657
------------------------------------------------------------
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=41, ID=43899
to -> from TCP D=54657 S=23  Ack=4095044502 Seq=826420002
          Len=1 Win=8760
to -> from TELNET R port=54657 i
------------------------------------------------------------
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54335
    from -> to TCP D=23 S=54657  Ack=826420003 Seq=4095044502
               Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 56 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=42, ID=54336
    from -> to TCP D=23 S=54657  Ack=826420003 Seq=4095044502
               Len=2 Win=8760
    from -> to TELNET C port=54657
------------------------------------------------------------
```

```
to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=44, ID=43900
to -> from TCP D=54657 S=23 Ack=4095044504 Seq=826420003
          Len=4 Win=8760
to -> from TELNET R port=54657
-------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54337
    from -> to TCP D=23 S=54657  Ack=826420007 Seq=4095044504
              Len=0 Win=8760
    from -> to TELNET C port=54657
-------------------------------------------

to -> from ETHER Type=0800 (IP), size = 64 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=50, ID=43901
to -> from TCP D=54657 S=23 Ack=4095044504 Seq=826420007
          Len=10 Win=8760
to -> from TELNET R port=54657 hei\r\nprompt\%
-------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54338
    from -> to TCP D=23 S=54657  Ack=826420017 Seq=4095044504
              Len=0 Win=8760
    from -> to TELNET C port=54657
  (retrans)
    from -> to ETHER Type=0800 (IP), size = 55 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=41, ID=54339
    from -> to TCP D=23 S=54657  Ack=826420017 Seq=4095044504
              Len=1 Win=8760
    from -> to TELNET C port=54657
-------------------------------------------

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=44, ID=43902
to -> from TCP D=54657 S=23 Ack=4095044505 Seq=826420017
          Len=4 Win=8760
to -> from TELNET R port=54657
-------------------------------------------

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54343
    from -> to TCP D=23 S=54657  Ack=826420021 Seq=4095044505
              Len=0 Win=8760
    from -> to TELNET C port=54657
-------------------------------------------

to -> from ETHER Type=0800 (IP), size = 62 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=48, ID=43907
to -> from TCP D=54657 S=23 Ack=4095044505 Seq=826420021
          Len=8 Win=8760
to -> from TELNET R port=54657 logout\r\n
-------------------------------------------
```

```
    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54348
    from -> to TCP D=23 S=54657  Ack=826420029 Seq=4095044505
              Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
  Send Fin, end of connection

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43911
to -> from TCP D=54657 S=23 Fin    Ack=4095044505
Seq=826420029 Len=0 Win=8760
to -> from TELNET R port=54657
```

```
  Send Fin,Ack with Ack=previous Seq+1

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54349
    from -> to TCP D=23 S=54657  Ack=826420030 Seq=4095044505
              Len=0 Win=8760
    from -> to TELNET C port=54657

    from -> to ETHER Type=0800 (IP), size = 54 bytes
    from -> to IP D=192.0.2.238 S=192.0.2.10 LEN=40, ID=54350
    from -> to TCP D=23 S=54657 Fin Ack=826420030
              Seq=4095044505 Len=0 Win=8760
    from -> to TELNET C port=54657
```

```
  Send Ack+1 to end

to -> from ETHER Type=0800 (IP), size = 60 bytes
to -> from IP D=192.0.2.10 S=192.0.2.238 LEN=40, ID=43912
to -> from TCP D=54657 S=23  Ack=4095044506 Seq=826420030
              Len=0 Win=8760
to -> from TELNET R port=54657
```

9.9 Attacks

There are many ways to attack a networked computer in order to gain access to it, or simply disable it. Some well known examples are listed below. The actual attack mechanisms used by attackers are often intricate and ingenious, but the common theme in all of them is to exploit naive limitations in the way network services are implemented. Time and again, one sees crackers make use of software systems which were written in good faith, by forcing them into unnatural situations where the software fails through inadequate checking.

9.9.1 Ping Attacks

The RFC 791 specifies that Internet datagrams shall not exceed 64kb. Some implementations of the protocol can send packets which are larger than this, but not all implementations can receive them.

```
ping -s 65510 targethost
```

Some older network interfaces can be made to crash certain operating systems by sending them a 'ping' request like this with a very large packet size. Most modern operating systems are now immune to this problem (e.g. NT 3.51 is vulnerable, but NT 4 is not). If not, it can be combatted with a packet filtering router. See http://www.sophist.demon.co.uk/ping/ .

9.9.2 Denial of Service (DoS) Attacks

Another type of attack is to overload a system with so many service requests that it grinds to a halt. One example is mail spamming[2], in which an attacker sends large numbers of repetitive e-mail messages, filling up the server's disk and causing the `sendmail` daemon to spawn rapidly and slow the system to a stand-still.

Newer versions of Berkeley sendmail have built in anti-spamming mechanisms to help protect from this problem. Vendors' sendmails are less advanced.

Denial of service attacks are almost impossible to protect against. It is the responsibility of local administrators to prevent their users from initiating such attacks wherever possible.

Denial of service attacks on NT have been embarrassingly simple

```
%myhost telnet ne-host 1028
Trying ???.???.???.???...
Connected to nt-host
Escape character is '^]'.
Hello there.

^]quit

myhost%
```

This sends NT4 services into a loop which can stop all services. Service Pack 2 fixes this. There are many other examples. Starting full system auditing can also bring a respectable Pentium system to a virtual standstill.

Cfengine employs a system of anti-spamming locks [38] to limit denial of service attacks.

9.9.3 TCP/IP Spoofing

Most network resources are protected on the basis of the host IP addresses of those resources. Access is granted by a server to a client if the IP address is contained in an Access Control List (ACL). Since the operating system kernel itself declares its own identity when packets are sent, it has not been common to verify whether packets actually do arrive from the hosts which they claim to arrive from. Ordinary users have not traditionally had access to privileges which allow them to alter network protocols. Today everyone can run a PC with privileged access to the networking hardware.

Normally an IP datagram passing from *host A* to *host B* has a destination address 'host B' and source address 'host A' (see Figure 9.5). IP spoofing is the act of forging IP datagrams in such a way that they appear to come from a third party host, i.e. an attacker at *host A* creates a packet with destination address 'host B' and source address 'host C'. The reasons for this are varied. Sometimes an attacker wants to appear to be *host C* in order to gain access to a special

[2] From the Monty Python song "Spam spam spam spam..."

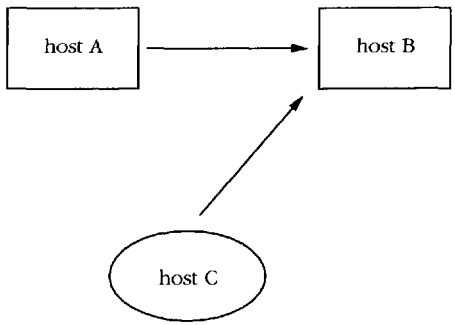


Figure 9.5 IP spoofing. A third party host C assumes the role of host A

resource which *host C* has privileged access to. Another reason might be to attack *host C* as part of a more elaborate attack. Usually it is not quite this simple, however, since the forgery is quickly detected. The TCP handshake is such that *host A* sends a packet to *host B* and then replies to the source address with a sequence number which has to match the next number of an agreed sequence. If another packet is not received with an agreed sequence number, the connection will be reset and abandoned. Indeed, if *host C* received the confirmation reply for a message which it never sent, it would send a reset signal back immediately, saying effectively 'I know nothing about this'. To prevent this from happening it is common to take out *host C* first by attacking it with some kind of Denial of Service method, or simply choosing an address which is not used by any host. This prevents it from sending a reset message. The advantage of choosing a real *host C* is that the blame for the attack is placed on *host C*.

9.9.4 SYN Flooding

IP spoofing can also be used as a denial of service attack. By choosing an address for *host C* which is not in use so that it cannot reply with a reset, *host A* can send SYN packets (new connections) on the same and other ports repeatedly. The *RECV* queue quickly fills up and cannot be emptied since the connections cannot be completed. Because the queues are filled the services are effectively cut off.

These attacks could be prevented if routers were configured so as to disallow packets with forged source addresses.

9.9.5 TCP Sequence Guessing

This attack allows an attacker to make a TCP connection to a host by guessing the initial TCP sequence number used by the other end of the connection. This is a form of IP spoofing. The attack was first described in the references below. It was made famous by the break in into Tsutomo Shinomrua's computers which led to the arrest of Kevin Mitnick. This attack is used to impersonate other hosts for trusted access [185, 23, 252]. This approach can now be combatted by using random sequence numbers.

9.9.6 IP/UDP Fragmentation (Teardrop)

The Teardrop attack was responsible for the now famous twelve hour attack which 'blue-screened' thousands of NT machines all over the world. This attack uses the idea of datagram fragmentation. Fragmentation is something which happens as a datagram passes through a router from one network to another network where the transmission rate is lower. Large packets can be split up into smaller packets for more efficient network performance. In the Teardrop attack, the attacker forges two UDP datagrams which appear to be fragments of a larger packet, but with data offsets which overlap.

When fragmentation occurs it is always the end host which reassembles the packets. To allocate memory for the data, the kernel calculates the difference between the end of the datagram and the offset at which the datagram fragment started. In a normal situation that would look like that in Figure 9.6. In a Teardrop attack the packets are forged so that they overlap like this: The assumption that the next fragment would follow on from the previous one leads to a negative number for the size of the fragment. As the kernel tries to allocate memory for this it calls `malloc(`*`size`*`)` where the size is now a negative number. The kernel panics and the system crashes on implementations which did not properly check the bounds.

9.9.7 ICMP Flooding (Smurf)

ICMP flooding is another denial of service attack. The ICMP protocol is the part of TCP/IP which is used to transmit error messages and control information between hosts. Well known services like `ping` and `echo` use ICMP. Normally all hosts respond to ping and echo requests without question, since they are useful for debugging. In an ICMP flooding attack, the attacker sends a spoofed ICMP packet to the broadcast address of a large network. The source address of the packet is forged so that it appears to come from the host which the attacker wishes to attack. Every host on the large network receives the ping/echo request and replies to the same host simultaneously. The host is then flooded with requests. The requests consume all the system resources.

9.9.8 DNS Cache Poisoning

This attack is an example of the exploitation of a trusted service in order to gain access to a foreign host. Again it uses a common theme, that of forging a network service request. This time, however, the idea is to ask a server to cache some information which is incorrect so that

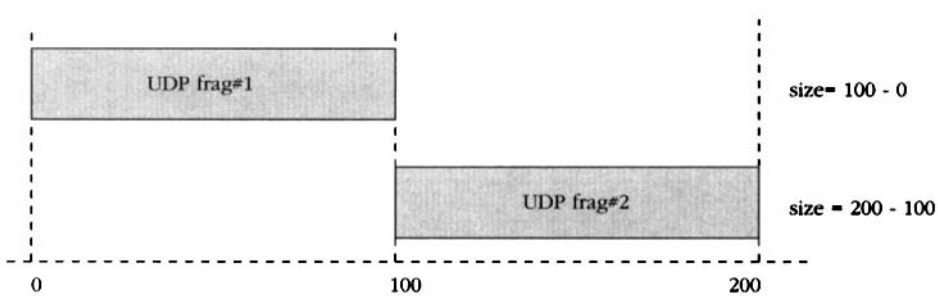


Figure 9.6 Normal UDP fragmentation

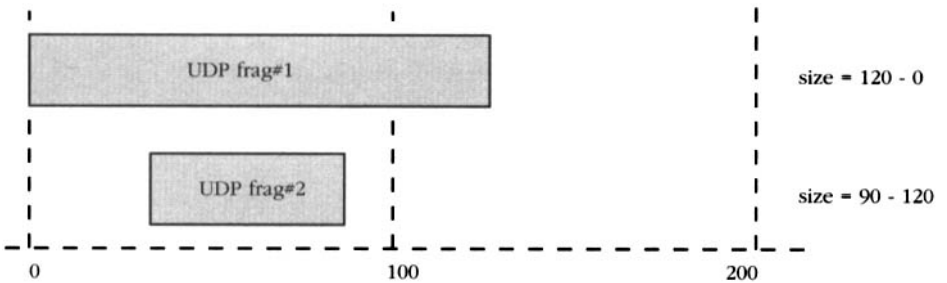


Figure 9.7 Spoofed UDP fragmentation, generates a negative size

future look-ups will result in incorrect information being given instead of the correct information [24].

DNS is a hierarchical service which attempts to answer queries about IP names and addresses locally. If a local server does not have the information requested, it asks an authoritative server for that information. Having received the information from the authoritative server it caches it locally to avoid having to contact the other server again; after all, since the information was required once, it is likely that the same information will be required again soon. The information is thus placed in the cache for a period of time called the TTL (*Time To Live*). After that time has expired it has to be obtained again from the authoritative server.

In a cache poisoning attack, the aim is to insert incorrect information into the cache of a server. Once it is there it will be there for the TTL period. To arrange this an attacker does the following:

1 The attacker launches his/her attack from the authoritative name server for his/her network. This gives him/her the chance to send information to another name server which will be trusted.

2 The attacker sends a query for the IP address of the victim host to the victim's default DNS server in order to obtain a DNS query ID. This provides a point of reference for guessing, i.e. forging the next few query IDs from that server.

3 The attacker then sends a query asking for the address of a host which the victim machine trusts, i.e. the host which the attacker would like to impersonate.

4 The attacker hopes that the victim host will soon need to look up the IP address of the host it trusts; he/she sends a fake 'reply' to such a DNS lookup request, forged with the query ID to look as though it comes from a lookup of the trusted host's address. The answer for the IP address of the trusted host is altered so that it is the IP address of the attacker's host.

5 Later, when the victim host actually sends such a DNS request, it finds that it has already received a UDP reply to that request (this is the nature of UDP), and it ignores the real reply because it arrives later. Now the victim's DNS cache has been poisoned.

6 The attacker now attempts to connect directly to the victim host, posing as the trusted host. The victim host tries to verify the IP address of the host by looking up the address in its DNS server. This now responds from its cache with the forged address.

7 The attacker's system is accepted.

This kind of attack requires the notion of external login based on trust, e.g. with Unix `.rhosts` files. This doesn't help with NT because NT doesn't have trusted hosts in the same way. On the other hand, NT is much easier to gain access to through NULL sessions.

Exercises

Exercise 9.1 What are the basic requirements for computer security? Look around your network. Which hosts satisfy these basic requirements?

Exercise 9.2 Devise a checklist for securing a PC attached to a network in your organization. How would you secure a PC in a bank? Are there any differences in security requirements between your organization and a bank? If so, what are they and how do you justify them?

Exercise 9.3 Determine what password format is used on your own system. Are shadow password files used? Does you site use NIS (i.e. can you see the password database by typing `ypcat passwd)?`

Exercise 9.4 Assume that passwords may consist of only the 26 letters of the alphabet. How many different passwords can be constructed if the number of characters in the password is 1, 2, 3, 4, 5, 6, 7 or 8 characters?

Exercise 9.5 Suppose a password has four characters, and it takes approximately a millisecond (10^{-3} *s*) to check a password. How long would a brute force attack take to determine the password?

Exercise 9.6 Discuss how you can really determine the identity of another person. Is it enough to see the person? Is a DNA test sufficient? How do you know that a person's body has not been taken over by aliens, or brainwashed by a mad scientist? This problem is meant to make you think carefully about the problem of *authentication.*

Exercise 9.7 Password authentication works by knowing a shared secret. What other methods of authentication are used?

Exercise 9.8 The secure shell uses a Virtual Private Network (VPN) or encrypted channel between hosts to transfer data. Does this offer complete security? What does encryption not protect against?

Exercise 9.9 Explain the significance of redundancy in a secure environment.

Exercise 9.10 When the current TCP/IP technology was devised, ordinary users did not have personal computers or access to network listening devices. Explain how encryption of TCP/IP links can help to restore the security of the TCP/IP protocol.

Exercise 9.11 Explain the purpose of a sacrificial lamb. Would sacrificing small children, or users help?

Exercise 9.12 Discuss the point of making a honey pot. Would this attract anyone other than bears of little brain?

Chapter 10

Security Implementation

In the previous chapter, we looked the meaning of security in the context of a computer system. Now we apply the basic principles, and consider what practical steps can be taken to provide a basic level of security. See RFC 1244 about how to devise a site security plan.

10.1 The Recovery Plan

When devising a security scheme, think of the post-disaster scenario. When disaster strikes, how will the recovery proceed? How long is this likely to take? How much money or time will be lost as a result?

The network is a jigsaw puzzle in which every piece has its place and plays its part. Recall the principle of redundancy: the more dependent we are on one particular piece of the puzzle, the more fragile the set up. Recovery will occur more quickly if we have backups of all key hardware, software and data.

In formulating a recovery plan, then, we need a scheme for replacing key components either temporarily or permanently, and we should also bear in mind that we do rely on many things which are outside of our immediate control. What happens, for instance, if a digger (back-hoe) goes thought the net cable, our only link to the outside world? Who should we call? Less fundamental but more insidious, what if the network managers above us decide to decouple us from the network without informing us in advance? In a large organization, different people have responsibility for different maintenance tasks. It has happened on more than one occasion that the power has been shut down without warning—a potentially lethal act for a hard disk.

10.2 Data Integrity

As part of any infrastructure, we need to apply the principle of redundancy to the network's data. Although backup copies will not protect us against loss, they do provide minimal insurance against accidents, intentional damage and natural disasters, and make the business of *recovery* less painful.

10.2.1 Preventing Loss

Once a file is deleted in a Unix-like operating system, it is not possible to get it back. Unlike DOS and its successors, there is no way to undelete a file. Some system administrators like to protect inexperienced users by making an alias (in the C shell)

```
alias rm rm -i
```

which causes the `rm` command to ask whether it really should delete files before actually doing so. This is a simple idea and it is not fool-proof, but it is an example of the kind of small details which make systems safer for the inexperienced. The only real security against deletion is to keep extensive backups of user disks. In Windows environments, it is not uncommon to hear screams of anguish as users lose two hours work because they didn't save before the system crashed, or reformatted their text according to some arbitrary template. Sensible software defaults can go a long way towards preventing loss of data.

Loss against physical disk failure can be mitigated by using RAID (Redundant Array of Inexpensive Disks) solutions which offer real redundancy. The idea is that, since disks are relatively cheap compared to human time and labour, we can build a system which uses extra disks in order to secure increased performance and redundancy. RAID disks systems are sold my most manufacturers and come in a variety of levels. Not all of the RAID levels have anything at all to do with redundancy. Indeed, some are more concerned with striping disks to increase performance and are more insecure than using single disks. There are currently seven levels of RAID:

1 *Disk striping:* this is a reorganization of the file system structure on a group of disks. Data are spread across disks, using parallelism to increase data throughput and improve search rate. This can improve performance dramatically, but reduces security by an equal amount, since if one disk fails, all the data are lost from the other disks.

2 *Real-time mirroring:* when data are written to one disk, they are simultaneously written to a second disk, rather than mirroring as a batch job performed once per day (see the next section). This increases security. This protects against random disk failure, but not necessarily against power failures, natural disasters, etc., since RAID disks are usually located all in one place.

3 *Disk striping with parity:* data are split across several disks to utilize parallelism, and a special parity disk enables data to be reconstructed provided no more than one disk fails randomly. Again, this does not help us against loss due to outside influences like power failure or natural disaster.

4 *As 3 except for striping algorithm:* allows single threaded writes but parallel reads. Each drive holds an entire data word.

5 *As 3 except for parity storage:* parity data are spread across all disks, to avoid a single point of failure.

6 *Enhanced raid 5:* two drives can fail randomly, and data can still be recovered.

7 *Disk striping again:* experimental high throughput system with dedicated driver, can connect to more than one host.

New RAID solutions appear frequently. RAID provides enhancements for performance and fault tolerance, but it cannot protect us against deliberate vandalism or widespread failure.

Last but far from least, network services are an important source of loss. They open a host to outside attack. Network services, indeed any daemons, which are not explicitly required on a given host (e.g. snmpd, powerd, nfsd) should be disabled.

10.2.2 Backup Schemes

Information can be lost in many ways: by accident, technical failure, natural disaster or even sabotage. We must make sure that there are several copies of the data so that everything may be recovered from a secure backup. Backups are one of the favourite topics of the system administration community. Everyone has their own local tricks. Many schemes for backup have been described; most of them resemble one another apart from cosmetic differences. Descriptions of backup schemes are manifold. Regular incremental style backups with site customizations can be found in refs. [268, 131, 141, 201, 120, 194, 287, 184, 216, 179]. A forward looking backup scheme with a broad generality in its ability to use different services and devices for remote backups was described in ref. [243], and backup to optical disks is discussed in refs. [46, 276]. Automated tape backup and restore was discussed in ref. [155] and in the Amanda system [241]; the AFS backup system is discussed in ref. [123]. A review of how well backup systems deal with special Unix sparse files was conducted in ref. [290].

Backup applies to individual changes, to system setup and to user data alike. In backing up data according to a regular pattern, we are assuming that no major changes occur in the structure of the data [239]. If major changes occur, we need to start backups afresh. The network has completely changed the way we have to think about backup. Transmitting copies of files to secondary locations is now much simpler. The basics of backup are these:

- *Physical location:* a backup needs to be kept at a different physical location than the original. If data were lost because of fire or natural disaster, then copies will also be lost if they are stored nearby.
- *How often?* How often do the data change significantly, i.e. how often do we need to make a backup? Every day? Do you need to archive several different versions of files, or just the latest version? The cost of making a backup is a relevant factor here.
- *Relevant and irrelevant files:* there is no longer any point in making a backup of parts of the operating system distribution itself. Today it is just as quick to reinstall the operating system from source, using the original CD-ROM. If we have followed the principle of separating local modifications from the system files, then it should be trivial to back up only the files which cannot be recovered from the CD-ROM, without having to backup everything.
- *Backup policy:* some sites might have rules for defining what is regarded as valid information, i.e. what it is worth making a backup of. Files like `prog.tar.gz` might not need to be kept on backup media since they can be recovered from the network just as easily. Also, one might not want to make backups of teen 'artwork' which certain users collect from the network, nor temporary data, such as browser cache files.

Medium

Traditionally, backups have been made from disk to tape (which is relatively cheap and mobile), but tape backup is awkward and difficult to automate unless one can afford a specialized robot to change and manage the tapes. For small sites it is also possible to

perform disk mirroring. Disk is cheap, while human operators are expensive. Many modern file systems (e.g. DFS) are capable of automatic disk mirroring in real-time. A cheap approach to mirroring is to use cfengine:

```
# cfengine.conf on backup host

copy:

  /home dest=/backup/home
        recurse=inf
        server=myhost
        exclude=core
```

When run on the backup host, this makes a backup of all the files under the directory `/home` on the host `myhost`, apart from core files. RAID disks also have inbuilt redundancy which allows data to be recovered in the event of a single disk crash. Another advantage with a simple mirroring scheme is that users can recover their files themselves, immediately, without having to bother a system administrator.

Of course, as the size of an institution grows, the economics of backup change. If one part of an organization has the responsibility for making backups for the entire remainder, then disk mirroring suddenly looks expensive. Of course, if each department of the organization invests in its own mirror disks, then the cost is spread. Economics has a lot to do with appearance as well as reality.

One criticism of disk mirroring is that it is not always possible to keep the disk mirrors far enough away from the original to be completely safe. An additional tape backup as a last resort is probably a good idea anyway.

A Backup Schedule

How often we need to make backups depends upon two competing rates of change:

- The rate at which new data are produced.
- The expected rate of loss or failure.

For most sites, a daily backup is sufficient. In a war zone, where risk of bombing is a threat at any moment, it might be necessary to back up more often. Most organizations do not produce huge amounts of data every day; there are limits to human creativity. However, other organizations, such as research laboratories, collect data automatically from instruments which would be impractically expensive to re-acquire. In that case, the importance of backup would be even greater.

Suggestion 13 (Static data) *When new data are acquired and do not change, they should be backed up to write only media at once. CD-ROM is an excellent medium for storing permanent data.*

For a single, un-networked host used only occasionally, the need for backup might be as little as once per week or less.

The options we have for creating backup schemes depend upon the tools we have available for the job. On NT we have NTBackup. On Unix-like systems there is a variety tools which can be used to copy files and file systems.

Backup	Restore
`cp -ar`	`cp -ar`
`tar cf`	`tar xpf`
GNU `tar zcf`	`tar zxpf`
`dd`	`dd`
`cpio`	`cpio`
`dump`	`restore`
`ufsdump`	`restore`
`rdump`	`rrestore`
NTBackup	NTBackup

Of course, commercial backup solutions exist for all operating systems, but they are often costly.

On both Unix and NT, it is possible to back up file systems either *fully* or *differentially*, also called *incrementally*. A full dump is a copy of every file. An incremental backup is a copy of only those files which have changed since the last backup was taken. Incremental backups rely on dump timestamps and a consistent and reliable system clock to ISO/IEC 9798 to avoid files being missed. For instance, the Unix dump utility records the dates of its dumps in a file `/etc/dumpdates`. Incremental dumps work on a scheme of levels, as we shall see in the examples below.

There are many schemes for performing system dumps:

- *Mirroring:* by far the simplest backup scheme is to mirror data on a daily basis. A tool like `cfengine` or `rsync` (Unix) can be used for this, copying only the files which have changed since the previous backup. Cfengine is capable of retaining the last two versions of a file, if disk space permits. A disadvantage with this approach is that it places the onus of keeping old versions of files on the user. Old versions will be mercilessly overwritten by new ones.
- *Simple tape backup:* tape backups are made at different *levels*. A level 0 dump is a complete dump of a file system. A level 1 dump is a dump of only those files which have changed since the last level 0 dump; a level 2 dump backs up files which have changed since the last level 1 dump, and so on, *incrementally*. There are commonly nine levels of dumps using the Unix dump commands. NTBackup also allows incremental dumps.

  The point of making incremental backups is that they allow us to capture changes in rapidly changing files without having to copy an entire file system every time. The vast majority of files on a file system do not change appreciably over the space of a few weeks, but the few files which we are working on specifically do change often. By pinpointing these for special treatment we save both time and tapes.

So how do we choose a backup scheme? There are many approaches, but the key principle to have in mind is that of *redundancy*. The more copies of a file we have, the less likely we are to lose the file. A dump sequence should always begin with a level 0 dump, i.e. the whole file system. This initializes the sequence of incremental dumps. Monday evening, Tuesday morning or Saturday are good days to make a level 0 dump, since that will capture most large changes to the file system in the level zero dump rather than in the subsequent incremental

ones. Studies show that users download large amounts of data on Mondays (after the weekend break), and it stands to reason that after a week of work, large changes will have taken place by Saturday. So we can take our pick. Here is a simple backup sequence for user home-directories, then, assuming that the backups are taken at the end of each day:

Day	Dump Level
Mon	0
Tue	1
Wed	2
Thu	3
Fri	4
Sat	1

Notice how this sequence works. We start with a full dump on Monday evening, collecting all files on the file system. Then on subsequent days we add only those files which have changed since the previous day. Finally, on Saturday we go back to a level 1 dump which captures all the changes from the whole week (since the Monday dump) in one go. By doing this, we have two backups of the changes, not just one. If we do not expect much to happen over the weekend, we might want to drop the dump on Saturday.

A variation on this scheme, which captures several copies of every file over multiple tapes, is the so-called *Towers of Hanoi* sequence. The idea here is to switch the order of the dump levels every other day. This has the effect of capturing not only the files which have changed since the last dump, but also all of the files from the previous dump as well. Here is a sample for Monday to Saturday:

Towers of Hanoi sequence over four weeks
$0 \rightarrow 3 \rightarrow 2 \rightarrow 5 \rightarrow 4 \rightarrow 6$
$1 \rightarrow 3 \rightarrow 2 \rightarrow 5 \rightarrow 4 \rightarrow 6$
$1 \rightarrow 3 \rightarrow 2 \rightarrow 5 \rightarrow 4 \rightarrow 6$
$1 \rightarrow 3 \rightarrow 2 \rightarrow 5 \rightarrow 4 \rightarrow 6$

There are several things to notice here. First, we begin with a level zero dump at the beginning of the month. This captures primarily all of the static files. Next we begin our first week with a level 3 dump, which captures all changes since the level 0 dump. Then, instead of stepping up, we step down and capture all of the changes since the level zero dump again (since 3 is higher than 2). This means that we get everything from the level 3 dump and all the changes since then too. On day 4 we go for a level 5 dump, which captures everything since the last level 3, and so on. At every stage, each backup captures not only new changes, but all of the previous backup also. This provides double the amount of redundancy as would be gained by a simple incremental sequence. When it comes to Monday again, we begin with a level one backup which grabs the changes from the whole of the previous week. Then once a month, a level zero backup grabs the whole thing again.

The Towers of Hanoi sequence is clever and very secure, in the sense that it provides a high level of redundancy, but it is also expensive since it requires a lot of tapes and time. The level of redundancy which is appropriate for a given site has to be a question of economics based on four factors:

1 The cost of the backup (time and media).
2 The expected rate of loss.
3 The rate of data production.
4 Media reliability.

These factors vary for different kinds of data, so the calculation needs to be thought out for each file system independently. The final point can hardly be emphasized enough. It helps us nothing to make ten copies of a file, if none of those copies are readable when we need them.

Suggestion 14 (Tape backup) *Tapes are notoriously unreliable media, and tape streamers are mechanical nightmares, with complex moving parts which frequently go wrong. Verify the integrity of each substantial backup tape backup once you have made it. Never trust a tape. If the tape streamer gets serviced or repaired, check old tapes again afterwards. Head alignment changes can make old tapes unreadable.*

Needless to say, backups should be made when the system is virtually quiescent: at night, usually. The most obvious reason for this is that, if files are being changed while the backup is progressing, then data can be corrupted or backed up incorrectly. The other reason is one of load: traversing a file system is a highly disk intensive operation. If the disk is being used extensively for other purposes at the same time, both backup and system will proceed at a snail's pace.

File Separation

The principle of keeping independent files separate was not merely to satisfy any high-flying academic aesthetic, it also has a concrete practical advantage, particularly when it comes to backing up the system. There is little sense in backing up the static operating system distribution. It can be reinstalled just as quickly from the original CD-ROM (a non-perishable medium). However, changing files such as `/etc/passwd` or `/etc/shadow` which need to be at special locations should be copied to another file system which is backed up often. This follows automatically from the principle of keeping local changes separate from the OS files. The same thing applies to other files like `/etc/fstab` or `/etc/group` and `/etc/system` which have been modified since the operating system was installed. However, here one can reverse the policy for the sake of a rational approach. While the password and shadow files have to be at a fixed place, so that they will be correctly modified when users change their passwords, none of the other files have to be kept in their operating system recommended locations.

Suggestion 15 (OS configuration files) *Keep master versions of all configuration files like* `/etc/fstab`, `/etc/group` *or* `/etc/system` *in a directory under site-dependent files, and use a tool which synchronizes the contents of the master files with the operating*

system files (e.g. cfengine). This also allows the files to be distributed easily to other hosts which share a common configuration, and provide us with one place to make modifications, rather than having to hunt around the system for long-forgotten modifications. Site-dependent files should be on a partition which is backed up. Do not use symbolic links for synchronizing master files with the OS: only the root file system is mounted when the system boots, and cross-partition links will be invalid. You might render the system unbootable.

10.2.3 Recovery from Loss

The ability to recover from loss presupposes that we have enough of the pieces of the system from which to reconstruct it, should disaster strike. This is where the principle of redundancy comes in. If we have done an adequate job of backing up the system, then we will not lose data, but we can still lose valuable time.

Recovery plans can be useful provided they are not merely bureaucratic exercises. Usually a checklist is sufficient, provided the system administration team is all familiar with the details of the local configuration. A common mistake in a large organization, which is guaranteed to lead to friction, is to make unwarranted assumptions about a local department. Delegation can be a valuable strategy in the fight against time. If there are sufficient local system administrators who know the details of each part of the network, then it will take such a person less time to make the appropriate decisions and implement the recovery plan. If a higher authority comes crashing down from too high a level, then it usually only aggravates the situation. Higher level authorities tend to think in terms of generalities rather than realities.

When loss occurs, we have to recover files from the backups. One of the great advantages of a disk mirroring scheme is that users can find backups of their own files without having to involve an administrator. For larger file recoveries, it is more efficient for a system administrator to deal with the task. Restoring from tape backup is a much more involved task. Unfortunately, it is not merely a matter of monkey work. First of all, we have to locate the correct tape (or tapes) which contain the appropriate versions of backed up files. This involves having a system for storage, reading labels and understanding any incremental sequence which was used to perform the dump. It is a time-consuming business. One of the awkwardnesses of incremental backups is that backing up files can involve changing several tapes to gather all of the files. Also, imagine what would happen if the tapes were not properly labelled.

Suggestion 16 (URL file system names) *Use a global URL naming scheme for all file systems and you will never lose a file on a tape, even if the label falls off (see section 3.9.2). Each file will be sufficiently labelled by its time-stamp and its name.*

We have two choices in recovery: reconstruction from backup or from source. Recovery from source is not an attractive option for local data. It would involve typing in every document from scratch. For software which is imported from external sources (CD-ROMs or FTP repositories), it is possible to reconstruct software repositories like `/usr/local` or NT's software directories. Whether or not this is a realistic option depends upon how much money one has to spend. Companies tend to have more money to throw at security issues

than universities or research labs. For a particularly impoverished department, reconstruction from source is a cheap option.

ACLs present an awkward problem for NT file systems. Whereas Unix's root account always has permission to change the ownership and access rights of a file, NT's Administrator account does not. On NT systems, it is important not to reinstate files with permissions intact if there is a risk of them belonging to a foreign domain. If we did that, the files would be unreadable to everyone, with no possibility of changing their permissions.

Data directory loss is one thing, but what if the system disk becomes corrupted? Then it might not even be possible to start the system. In that case it is necessary to boot from floppy disk or CD-ROM. For instance, a PC with GNU/Linux can be booted from a 'rescue disk' or boot disk, in single user mode (see section 4.2.1), just by inserting a disk into the floppy drive. This will allow full access to the system disk by mounting it on a spare directory:

```
mount /dev/hda1 /mnt
```

On Sun Sparc systems, we can boot from CD-ROM, from the monitor:

```
boot cdrom
```

or `boot sd(0,6,2)` with very old PROMs[1]. Then, assuming we know which is the root partition, it can be mounted and examined:

```
mount /dev/dsk/c0t0d0t3 /mnt
```

Recovery also involves some soul searching. We have to consider the reason for the loss of the data. Could the loss of data have been prevented? Could it be prevented at a later time? If the loss was due to a security breach or some other form of vandalism, then it is prudent to consider other security measures at the same time as we reconstruct the system from the pieces.

10.2.4 Checksum Verification

Every time we use the privileged system account, we are at risk of installing a virus or a Trojan horse, or of editing the contents of important files which define system security. The list of ingenious ploys for tricking root privileged processes into working on behalf of attackers makes an impressive ream. The virtual inevitability of it, sooner or later, implores us to verify the integrity of programs and data by comparing them to a trusted source. A popular way to do this is to use a checksum comparison. To all intents and purposes, an MD5 checksum cannot be forged by any known procedure. An MD5 checksum is a numerical value which summarizes the contents of a file. Any small change in a file changes its checksum. A checksum can therefore be used to determine whether a file has changed. First we must compile a database of checksums for all important files on the system, in a trusted state. Then we check the actual files against this database over time. Assuming that the database itself is secure, this enables us to detect changes in the files and programs. The Tripwire program was the original program written to perform this function. Tripwire can be configured to cross check several types of checksum, just on the off-chance that someone manages to find a way to forge an MD5 checksum. Cfengine can also perform this task

[1] The SunOS CD player has to be on controller 0 with SCSI id 6.

```
control :
  actionsequence = ( files )
files :
  /usr  owner=root, bin mode=o-w checksum=md5 recurse=inf
```

Figure 10.1 A cfengine program to gather and check MD5 checksums of the `/usr` file tree

routinely, while doing other file operations. Cfengine currently uses only MD5 checksums (see Figure 10.1).

10.3 Analysing Network Security

To assess the potential risks to a site, we must gain some kind of overview of how the site works. We have to place ourselves in the role of outsider: how would someone approach the network from outside? Then we have to consider the system from the viewpoint of an insider: how do local users approach the system? To begin the analysis, we form a list:

- What hosts exist on our site?
- What OS types are used?
- What services are running?
- What bug patches are installed?
- Run special tools, nmap, satan, saint, titan to automate the examination procedure and find obvious holes.
- Examine trust relationships between hosts.

This list is hardly a trivial undertaking. Simply building the list can be a lesson to many administrators. It is so easy to lose control over a computer network, so difficult to keep track of changes and the work of others in a team, that one can easily find oneself surprised by the results of such a survey. Having made the list, it should become clear as to where potential security weaknesses lie. Network services are a common target for exploitation. FTP servers and NT's commercial WWW servers have had a particularly hard time with bugs which have been exploited by attackers. At this stage it might be prudent to revise the organization of the above items in the network in order to tighten the rein on things.

Correct host configuration is one of the prerequisites for network security. Even if we have a firewall shielding us from outside intrusion, an incorrectly configured host is a security risk. Firewalls do not protect us from the contents of data which are relayed to a host. If a bug can be exploited by sending a hidden message, then it will get through a firewall. Some form of automated configuration checking should be installed on hosts. Manual checking of hosts is impractical even with a single host; a site which has hundreds of hosts requires an automated procedure for integrity checking. On Unix and NT one has cfengine and Perl for these tasks.

Trust relationships are amongst the hardest issues to debug. A trust relationship is an implicit *dependency*. Any host which relies on a network service implicitly trusts that service to be reliable and correct. This can be the cause of many stumbling blocks. The complexity of interactions between host services makes many trust relationships opaque. Trust relation-

ships occur in any instance in which there is an external source of information: remote copying, hostname lookup, directory services, etc. The most important trust relationship of all is the Domain Name Service (DNS). Many access control systems rely on an accurate identification of the host name. If the DNS service is compromised, hosts can be persuaded to do almost anything. For instance, access controls which assign special privileges to a named host can be be spoofed if the DNS lookups are corrupted or intercepted. DNS servers are therefore a very important pit-stop in a security analysis.

Access control is the fundamental requirement for security. Without access controls there can be no security. Access controls apply to files on a file system and to services provided by remote servers. Access should be provided on a need-to-know basis. If we are too lax in our treatment of access rights, we can fall foul of intrusion. For example: a common error in the configuration of Unix file servers is to grant arbitrary hosts the right to mount file systems which contain the personal files of users. If one exports file systems which contain users' personal data to Unix-like hosts, it should be done on a host-by host basis, with strict controls. If a user, who is root on their own host (e.g. a portable PC running GNU/Linux), can mount a user file system (with files belonging to a non-root user), that person owns the data there. The privileged account can read any file on a mounted file system by changing its user ID to whatever it likes. That means that anyone with a laptop could read any user's mail or change any user's files. This is a huge security problem. Hosts which are allowed to mount NFS file systems containing users' private data should be secured and should be active at all times to prevent IP spoofing; otherwise it is trivial to gain access to a user's files.

There are many tools written for Unix-like operating systems which can check the security of a site, literally by trying every conceivable security exploit. Tools like `SPY [250]`, `COPS`, `SATAN`, `SAINT` and `TITAN` are aimed at Unix-like hosts. Port scanners such as nmap will detect services on any host with any operating system. These tools can be instrumental in finding problems. Recent and frightening statistics from the Computer Emergency Response Team indicated that only a pitiful number of sites actually upgrade or install patches and review their security, even after successful network intrusions [133].

Having mapped out an overview of a network site, and used the opportunity both to learn more about the specifics of the system, as well as fix any obvious flaws, we can turn our attention to more specific issues at the level of hosts.

10.3.1 Password Security

Perhaps the most important issue for network security, beyond the realm of accidents, is the consistent use of strong passwords. Unix-like operating systems which allow remote logins from the network are particularly vulnerable to password attacks. The `.rhosts` and `hosts.equiv` files which allowed login without password challenge via `rsh` and `rlogin` were acceptable risks in bygone times, but these days one cannot afford to be lax about security. The problem with this mechanism is that `.rhosts` and `hosts.equiv` use hostnames as effective passwords. This mechanism trusts DNS name service lookups which can be spoofed in elaborate attacks. Moreover, if a cracker gets into one host, he/she will then be able to log in on every host in these files without a password. This greatly broadens the possibilities for effective attack. Typing a password is not such a hardship for users, and there are alternative ways of performing remote execution for administrators, without giving up password protection (e.g. use of cfengine).

Password security is the first line of defence against intruders. Once a malicious user has gained access to an account, it is very much easier to exploit other weaknesses in security. Experience shows that many users have little or no idea about the importance of using a good password. Consider some examples from a survey of passwords at a university. About 40 physicists had the password 'Einstein', around 10 had 'Newton' and several had 'Kepler'. Hundreds of users used their login-name as their password, some of them really went to town and added '123' to the end. Many girls chose 'horse' as their password. Even after extensive campaigns encouraging good passwords, users have a shocking tendency to trivialize this matter. User education is clearly an important weapon against weak passwords.

Some sites use schemes such as password aging to force users to change passwords regularly. This helps to combat password familiarity gained over time by local peer users, but it has an unfortunate side-effect. Users who tend to set poor passwords will not appreciate having to change their passwords repeatedly, and will tend to rebel by setting trivial passwords if they can. Once a user has a good password, it is often advantageous to leave it alone. The problems of password aging are insignificant compared to the problem of weak passwords.

Passwords are not visible to ordinary users, but their encrypted form is often visible. Even on NT systems, where a binary file format is used, a freely available program like `PwDump` can be used to decode the binary format into ASCII. There are many publicly available programs which can guess passwords and compare them with the encrypted forms, e.g. `crack`, which is available both for Unix and for NT. No one with an easy password is safe. Passwords should never be any word in a dictionary or a simple variation of such a word or name. It takes just a few seconds to guess these.

Newer operating systems like FreeBSD, NetBSD and Solaris and GNU/Linux have *shadow password files* which are not readable by normal users. The regular password file contains an 'x' instead of a password, and the encrypted password is kept in an unreadable file. This makes it much harder to scan the password file for weak passwords.

Suggestion 17 (Passwords) *A useful hint in choosing a password is to incorporate the PIN code from a little-used credit card as a part of the password. This helps users to remember both – and it means that there will be secret numbers in the password.*

Tools for password cracking (e.g. Alec Muffet's `crack` program) can help administrators find weak passwords before crackers do. Other tools can be obtained from security sites to prevent users from typing in weak passwords. See refs. [259, 50, 4, 125].

10.3.2 Password Sniffing

Many communication protocols (Telnet, FTP, etc.) were introduced before security was a concern amongst those on the Internet. So many of these protocols are very insecure. Passwords are often sent over the network as plain text. This means that a sophisticated cracker could find out passwords simply by listening to everything happening on the network and waiting for passwords to go by. If a cracker has privileged access to at least one machine with a network interface on the same subnet, he/she could use `tcpdump` to capture all network traffic. Normal users do not have this privilege for precisely this reason.

These days, however, anyone with a laptop, an Ethernet card and a GNU/Linux installation could do this. Switched networks are immune to this problem since traffic is routed directly from host to host.

Programs which dump all network traffic include `tcpdump`, `etherfind` and `snoop`. Here is a sample of the output from Solaris' `snoop` program showing the Ethernet traffic on a segment of cable. Snoop recognizes common high level protocols (SMTP/FTP/ARP, etc.), and lists them explicitly. Unknown protocol types (in this case IPX) are simply listed as ETHER. In the right-hand column is the information which an intruder would try to use to sniff passwords.

```
Using device /dev/hme (promiscuous mode)
   post.eet.no -> nexus        SMTP C port=4552 oJyhnJycoZyhnKCcnGCc
torget.drammensnett.no -> nexus        SMTP C port=54621 AGoHRPVU9VT3
        nexus -> torget.drammensnett.no SMTP R port=54621
pc111-75.iu.hioslo.no -> nexus         FTP C port=1093
        nexus -> pc111-75 FTP R port=1093 226 Transfer complet
        nexus -> post.eet.no  SMTP R port=4552
   post.eet.no -> nexus       SMTP C port=4546 UHAQcBB/UB9QcBBwH1AQ
        nexus -> post.eet.no  SMTP R port=4546
   post.eet.no -> nexus       SMTP C port=4546 H2AQcBBwH1AfYBAQH1Af
      fw.nki.no -> nexus       SMTP C port=11424 O3Jw+XF7cMFCCweEQ/
        nexus -> fw.nki.no    SMTP R port=11424
   post.eet.no -> nexus       SMTP C port=4552 niYmJgomChomChoaChoK

        nexus -> post.eet.no  SMTP R port=4546
        nexus -> (broadcast)  ARP C Who is 128.39.89.230, takpeh ?
        nexus -> post.eet.no  SMTP R port=4552

          ? -> *              ETHER Type=0000 (LLC/802.3), size = 86 bytes
          ? -> *              ETHER Type=0000 (LLC/802.3), size = 128 bytes
          ? -> *              ETHER Type=0000 (LLC/802.3), size = 80 bytes
```

One way to avoid the problem of password sniffing is to use fully encrypted links such as `ssh` [285] and SSL (Secure Socket Layer) enabled services which replace the standard services. Another is to use a system of one-time passwords. One time passwords are designed to eliminate the need for users to send their passwords over the network at all. Instead of typing an actual password, one types the remote password for a host into a program on a local machine, in order to generate a sequence of throw-away passwords which can be used in place of the actual remote password. The passwords are used only once, so even if someone gets to overhear them, it will already be too late: the password will have expired. Also, the system is ingeniously designed so that the actual remote password (which is used to generate the one-time passwords) never gets sent over the network at all. S/KEY is such a system. Here is an example of how it works:

1 We want to make a connection from host A to host B.

2 We have earlier set a password on host B.

3 We Telnet to host B from host A.

4 Host B prompts us with a code string: 659 ta55095 and asks for our user name. We type the user name and host B asks for the one-time password.

5 We now need to find the one-time password by running a local program on host A with the code string as an argument:

```
key 659 ta55095
passwd: *******
```

The key program prompts us for the secret password on host B. When we type this it does not go across the network. The key program returns a clear text, one-time password valid for one session: 'EASE FREY WRY NUN ANTE POT'

6 We type 'EASE FREY WRY NUN ANTE POT'on host B (sent over the network) and the password is accepted.

7 Next time we follow the same procedure and get a different password.

10.3.3 Network Services

When installing a new service which is available to more than one user, it is appropriate to ask the questions:

- Do I need this service?
- Whom or what information do I have to trust in order to use this?
- What will happen if someone abuses that trust?

For example, the `rlogin` feature of Unix has a file called `.rhosts` in which a user can add a list of trusted hosts. That user can log into the host with the `.rhosts` file from any one of those trusted hosts without giving a password. The user is clearly willing to trust this list of hosts. But that is not the only trust relationship here. Unix uses DNS (the Domain Name Service) to verify the identity of connecting machines, so the `rlogin` service *implicitly* trusts the DNS service. If someone could corrupt that service, there would be a potential security problem (see section 9.9.8).

Another example is in software distribution, both for Unix and NT. To distribute software from a central server to many clients, the clients have to trust the information being sent to them from the server. They have to give the server permission to install unknown files which might be security hazards.

SNMP control systems accept information from a controller, based only on a fairly weak password (community string). The password has a default value of 'public' which many sites forget to change (a potentially huge security risk). This information can be used to change control functions of key network components, and is even used for performing remote system administration in certain products.

Cfengine places all of its trust in the correctness of its input file; it does not accept input from the network at all. In software distribution it will trust files from a software server of its own choosing, but arbitrary servers cannot send data to it uninvited.

10.3.4 Protecting Against Attacks

- Look out for users with weak passwords. This is the easiest way for an attacker to enter the system.
- Do not give trusted access to other hosts unless absolutely necessary. Make sure there are no NIS wildcards + in `/etc/hosts.equiv`. Avoid using `.rhosts` files altogether.

- Disable unused services in `/etc/inetd.conf` which might contain security leaks, like UUCP or TFTP.
- Make sure the router filters all unnecessary traffic. Usually there is no reason to permit RPC traffic outside of the local domain, for instance.
- Make sure that the latest security patches are installed on all systems.
- Monitor connections using `netstat -a` to show all listening connections. Use `tcpd` logging.
- Monitor processes running on the system. How many copies of important processes are running? How many should be running? Often it is possible to see that one is under attack by looking at what processes are running and who is running them. For instance, an attempt at port sniffing or spamming might be seen with a bunch of processes like this:

```
nobody .... /usr/sbin/inetd
nobody .... /usr/sbin/inetd
nobody .... /usr/sbin/inetd
nobody .... /usr/sbin/inetd
nobody .... /usr/sbin/inetd
```

  Attempts at ping attack have been identified by large numbers of persistent ping processes. `inetd` is a multiplexer which starts Internet services on many ports. Normally it is only root who runs this. The above indicates that a user is trying to use the well-known account `nobody` to start services, or to overload the system with requests.
- Check file systems for suspicious looking hidden files, i.e. files with names like `. . .` These are often used to hide dangerous programs or shells which users can use to gain root privileges. Cfengine performs this task automatically when it examines file systems.
- Make sure that `.` is not in the root's path. It is possible to inadvertently execute a Trojan horse program.
- Make sure that files like `/var/adm/utmp` are not world writable, allowing crackers to cover their tracks.

Cfengine can be used to automate many of these issues.

10.3.5 Access Control with TCP Wrappers

Even without a firewall, one can go half way by verifying the authenticity of packets coming into a Unix host. A *wrapper* is a front line against 'spoofing', i.e. Internet impersonation. The `tcpd` program is a wrapper for internet services which provides host-based access control. Instead of starting services directly from the `/etc/inetd.conf` file, we start the tcp-daemon with instructions to start the named service by proxy (see section 8.4.5). The `tcpd` daemon checks where the request comes from and only starts it if it comes from a trusted host. A trusted host is one which is listed in `/etc/hosts.allow`. Another file `/etc/hosts.deny` lists services which are to be denied to non-trusted hosts.

replace:

```
finger stream tcp nowait nobody /usr/etc/in.fingerd in.fingerd
```

with:

```
finger stream tcp nowait nobody /local/sbin/tcpd   in.fingerd
```

TCP wrappers does not exact a performance price, since it is only operative during the establishment of the connection. During actual transmission, it does nothing. TCP wrappers logs connections to the syslog service, so we can see which hosts are connecting to which service.

10.3.6 `pidentd` Authentication Server

One of the problems with socket-based communication is authentication. How do we determine the identity of the user making the connection? Normally the only certain way is to require a password to be given, i.e. some kind of shared secret. The RFC documents RFC981 and RFC1413 specify a standard protocol for identifying the user *name* of a connecting user. For this to work, a server which is being connected to has to contact the host which is attempting to connect and check who owns the transmitting socket. This is done with the help of the `identd` daemon, or its implementation `pidentd` (the portable identity daemon).

The identity or authentication service (port 113) is alas not installed as standard by most vendors. You need to get this daemon and install it yourself. This is quite straightforward:

```
ftp ftp.lysator.liu.se/pub/ident/servers/

tar zxf pidentd-3.0.4.tar.gz
cd pidentd-3.0.4
configure --with-des=no
make
make install
```

The latest versions of the server are multithreaded and are most easily run outside of the inetd service. The daemon is started by

```
identd
```

Although it can be started from inetd, one should never invoke this daemon with TCP wrappers, since TCP wrappers attempts to contact with daemon. This will result in an infinite loop. You might have to register the service in `/etc/services`

```
auth 113/tcp    authentication tap ident
```

The identification service is a public service which a site provides for its own benefit and for others' benefit. It allows the authentication of connections with greater confidence. It is also worth noting, for the record, that the ident service is another network service which can be exploited. Port scanners can use ident to obtain information about a host. Every service is a two-edged sword.

10.3.7 Port Scanning

To find back-doors into vulnerable systems, many network attackers scan ports on network hosts in order to find out which services are running on them. Programs for performing such

scans (e.g. `nmap`, or `queso`) can be obtained freely from the network, as can many other intrusion tools, so crackers require little or no intelligence to carry out these simple attacks these days. Most intrusion tools can also be used to help secure systems.

In a poorly configured system a cracker might find active services which even the system owner did not realize were running. UUCP and TFTP are typical examples. These services can often be exploited to install files in illegal places. Known faults in services can be exploited if one knows about the services which are running. TCP/IP fingerprinting is the process by which port scanners determine the type of operating system from the quirks of a host's TCP stack. If a Telnet to a host does not immediately reveal a banner with the OS type (it usually does on any operating system):

```
nomad% telnet 127.0.0.1
Trying 127.0.0.1...
Connected to 127.0.0.1.
Escape character is '^]'.

Red Hat Linux release 4.2 (Biltmore)
Kernel 2.0.30 on an i586
login:
```

then more intricate signatures can be combed for tell-tale signs.

Primitive port scanning attempts are detectable if network activity is followed closely. Strings of attempted 'connect' requests to one port after the other are easily spotted. Recently, however, the trend has expanded to include 'stealth scanning' in which scans are performed at random over long periods of time to avoid attracting attention. Port scanning is only dangerous if there are poorly configured hosts on the network. Perhaps the most important issue is the consistent use of strong passwords. The `.rhosts` and `hosts.equiv` files which allow login without password challenge via `rsh/rlogin` were okay in bygone times, but these days we cannot afford to be lax about security. The problem with this mechanism is that `.rhosts` and `hosts.equiv` use host names as effective passwords. This mechanism trusts DNS lookups which can be spoofed in elaborate attacks in order to mislead a host about the identity of a connecting host. Moreover, if a hacker gets into one host, he/she will then be able to log in on every host in these files without a password. This greatly broadens the possibilities for effective attack.

10.3.8 X11 Security

Although many users are not aware of it, it is often possible to download the screen image of another user who is using the X-windows system. While this might occasionally be a useful opportunity to help remote users with a specific problem [289], in general it must be considered a grave security risk. It is equally possible to 'bug' the keyboard and listen to all the key-presses. The problem is an out-dated security mechanism which has long since been replaced, but which is still used by very many users. The problem is the `xhost` program. This is used to grant other hosts permission to draw on your X server – in other words, if you are remotely logged on to a host other than the one you are using as a display, you must grant the remote host access to write on your screen.

In the old X windows system, prior to release 5, one had to grant access to a particular host. Once this was done, *anyone* on that host had access to your server, not just you. This

was later replaced by the `xauth` magic-cookie mechanism which works on a user basis. Some users still insist on using `xhost`, however, with a command like this:

```
xhost +
```

Any user writing this opens their display to everyone in the world. The antidote, of course is the command `xhost -`. Users of the secure shell `ssh` (see section 10.4) can now have automatic X11 forwarding with authentication cookies. Everyone should therefore execute `xhost -` once and never use the xhost mechanism again.

10.4 VPNs: Secure Shell and FreeS/WAN

VPN stands for a Virtual Private Network. This is simply an armoured pipe connecting two locations: a line of communication which is reinforced by encryption and authentication. Privacy is obtained through encryption and the line is virtual because it sits on top of regular TCP/IP communication. Secure shell software can be used to build VPNs or most popular services.

The secure shell [285] is a secure replacement for the `rsh` commands. It protects against IP spoofing where a remote host pretends to be trusted host by faking IP datagrams; DNS spoofing where an attacker forges name entries in the name-service; the interception of passwords in network packets; and several other kinds of attack. Note that the secure shell is *not* free software. It is only freely usable in academic contexts for Unix-like systems. Anything else you have to pay for.

To install this, collect it from an FTP site and perform the usual steps:

```
(get ftp file ssh-2.0.9.tar.gz)

host% tar zxf ssh-2.0.9.tat.gz
host% cd ssh-2.0.9
host%./configure
host% make
host% su
Passwd: ********
host# make install
```

You then need to start the daemon by adding a command of the form

```
/usr/local/sbin/sshd
```

to the startup scripts on your system.

FreeS/WAN is another project [208], started for GNU/Linux systems, which will provide encrypted tunnels. See also the Virtual Private Network Consortium [52].

10.5 WWW Security

The concept of WWW security sounds like a contradiction in terms. The WWW is designed to publish information to the masses. Security has to do with restricting access. What has the WWW got to do with security? Web security has to do with

- Protecting the published data from corruption.
- Granting access only to those files we wish to publish.
- Preventing users from tricking the WWW server into executing unauthorized commands on the server host.

Although there have been many security problems with the feature over-leaden Internet Information Server for NT [259], there is nothing principally insecure about the WWW service. Any file server can, in principle, compromise the security of a host by making information about that host available to others. If a server provides access to unauthorized files, this will clearly be the case. All we need to do is to ensure that proper access controls are maintained.

The Free Apache WWW server (see section 8.6) has all of the features one requires to operate a secure web service. It can be run without special privilege, and it has quite sophisticated mechanisms for restricting access to data. It is nevertheless possible to configure the server in an insecure fashion, so one needs to be cautious. There are three distinct categories for web use:

- External web service for organization.
- Internal web service for organization.
- Private user web pages.

The last of these is the greatest potential security risk for the web: we usually trust the files and programs which we write ourselves in the name of our organization, but we have no reason to trust the integrity of private users. There are two areas where a security breach can occur:

- File ownership and access rights.
- CGI scripts.

CGI scripts can be used to execute commands on the server-host with the user-privileges of the WWW user. Although the WWW user is introduced precisely to isolate the powers of the WWW service, we can still do quite a bit of damage – not to the host directly, but to other users and to the web server access controls. It is an inevitable consequence of running a public service with a private ID, that any file which gets written by a CGI script can also be overwritten by another CGI script, regardless of which user is responsible for that script. Thus, users could wage war on one another with CGI scripts such as guestbooks, corrupting or even deleting one another's data freely. This is a fundamental weakness in the WWW service: if we allow the existence of arbitrary CGI scripts on the system, then we can carry out arbitrary operations with the privileges of the WWW user. Users can:

- Send anonymous, untraceable mail which appears to come from the WWW user at the organization hosting the CGI program.
- Circumvent `.htaccess` access controls to certain files on most types of operating system, by executing the command `/bin/cat filename` as part of a CGI script.

The first principle of server security is thus:

Principle 48 (WWW corruption) *If a web server runs with the privileges of user* www, *then none of the data files should be owned, or be writable by, the* www *user, otherwise it is trivial to alter the contents of the data with a CGI script.*

If we violate this principle, any local user can overwrite and corrupt web pages simply by writing a CGI script. Of course, the WWW server does not have any special privileges. It is just an ordinary, non-privileged user who has to obey normal file permissions, yet this is not enough to prevent a few accidents, nevertheless. This brings us to the fundamental flaw in WWW security.

Any files which are to be served by a WWW server have to be readable by the WWW user. All CGI scripts run with the rights of the WWW user. It therefore follows that any CGI script can read any file which is capable of being served by the daemon. To put it bluntly: *any user can write a CGI script to circumvent* `.htaccess` *security barriers.* The solution to this problem is to either disallow CGI scripts, or to move sensitive (non-public) documents to a separate host, which regular users do not have access to.

CGI scripts which send mail are a conundrum. If a user decides to write a CGI script which sends e-mail, it executes the mail program with the user identity of the WWW user. The identity of the true sender is irrelevant, since the actual sender is the WWW server. This could be an unfortunate situation for an organization. If private users can send e-mail anonymously, but which can be traced back to the WWW server of our organization, we clearly stand in the firing line for all kinds of trouble. A user could harass anyone with impunity, and only the organization would be responsible. At the time of writing, the sendmail mailer places restrictions only on mail which is relayed, not on mail which originates from the localhost. This presents a problem if our WWW server runs on the same host as the mail hub. It means, in particular, that we are not able to prevent normal users from sending e-mail from CGI scripts, using sendmail's standard access control mechanisms. This is a big drawback, and the only solution is to keep the mail exchanger on a different host to untrusted user accounts and the web server.

10.6 Firewalls

A firewall is a network configuration which isolates some machines from the rest of the network. It is a gate-keeper which limits access to and from a network. Our human bodies are relatively immune to attack by bacteria and viruses because we have a barrier: skin. The skin contains layers of various fatty acids in which bacteria and viruses cannot normally survive. If we lose the skin from a part of the body, wounds become quickly infected; indeed, prior to antibiotics, many people died from infected wounds. A firewall is like a skin for a local area network.

The idea is this: if we can make a barrier between our local network and the Internet which is impenetrable, then we would be safe from network attacks. But if there is an impenetrable barrier so that no one can get into the network, then no one can get out either. Why pay for a firewall when we could just pull out the network cable? Think of the body again: we have to put food and air into our bodies and we have to let stuff out, so we need a hole in the skin (preferably several). We do not usually die of the food we eat because the body has filters which screen out and break down dangerous organisms (stomach acid and layers of mucus

etc). These then hand us the 'input' by proxy. We do the same thing with computer networks. A firewall is not an impenetrable barrier: it has holes in it with passport checks. We demand that only network data with appropriate credentials should be allowed to pass.

10.6.1 A Concept

A firewall is a concept. It is not a thing; there is no single firewall solution. The name 'firewall' is a collective description for a variety of methods which restrict access to a network. They all involve placing restrictions on the way in which network packets are routed. A firewall might be a computer which is programmed to act like a router, or it might be a dedicated router or combination of routers and software systems. The idea with a firewall is to keep important data behind a barrier which has some kind of passport-control and can examine and restrict network packets, allowing only 'harmless' packets to pass.

- All traffic from inside to outside (or vice versa) must pass through the firewall.
- Only authorized traffic is allowed to pass.
- Potentially risky network services (like mail) are rendered safer using intermediary systems.
- The firewall itself should be immune to attack.

A firewall cannot help with the following:

- Badly configured hosts or misconfigured networks.
- Data based attacks (where the attack involves sending some harmful information, like the code word which makes you take your own life, or an e-mail which bolts a Trojan horse).

There are two firewall philosophies: *block everything unless we make an explicit exception* and *pass everything unless we make a specific exception.* The first of these is clearly the most secure or at least the most paranoid of the two.

Here's a few concepts which get bandied around in firewall-speak:

- *Screening router/Choke* A router which can be programmed to filter or reject packets directed at certain IP ports.
- *Bastion host* A computer, specially modified to be secure.
- *Dual-homed host* A computer with two net-cards, which can be used to link an isolated network to a larger network.
- *Application gateway* A filter, usually run on a bastion host, which has the ability to reject or forward packets at a high level (i.e. at the application level).
- *Screened subnet/DMZ* An isolated subnet, between the Internet and the private network. Also called a DMZ (de-militarized zone). A DMZ is the bit between a screening router and the firewall, bastion host. This is a good place for external WWW services.

The firewall philosophy builds on the idea that it is easier to secure one host (the bastion host) than it is to secure hundreds or thousands of hosts on a local network. One focuses on a single machine, and ensures that it is the only one effectively coupled directly to the network.

One forces all network traffic to stop at the bastion host, so if someone tries to attack the system by sending some kind of IP attack there can be little damage to the rest of the network because the private network will never see the attack. This is, of course, a simplification. It is important to realize that installing a firewall does not give absolute protection, and it does not remove the importance of configuring and securing the hosts on the inside of the firewall.

10.6.2 Firewall Proxies

Of course we do not want all traffic to stop; some services like e-mail and maybe HTTP should be able to pass through. To allow this, one uses a so-called 'proxy' service or a 'gateway'[2]. A common solution is to give the bastion host two network interfaces (one is then connected to the unsafe part of the network and the other is connected to the safe part), though the same effect can be obtained with a single interface. A service is said to be *proxied* if the bastion host forwards the packets from the unsafe network to the safe one. It only does this for packets which meet the requirements of your security policy. For instance, you might decide that the services you require to cross the firewall are inbound/outbound Telnet, inbound/outbound SMTP (mail), DNS, HTTP and FTP, but no others.

Principle 49 (Community borders) *Proxying is about protecting against breaches to the fundamental principle of communities. A firewall proxy provides us with a buffer against violations of our own community rights from outside, and also provides others with a buffer against what we choose to do in our own home.*

Proxying requires some special software, often at the level of the kernel where the validity of connections can be established. For instance, packets with forged addresses can be blocked. Data arriving at ports where there was no registered connection can be discarded. Connections can be discarded if they do not relate to a known user-account.

10.6.3 Example: Dual-homed Bastion Host

A simple firewall configuration is shown in Figure 10.2. In this example we effectively have two routers, a DMZ and a protected network. The first packet filtering router will route packets between the Internet and one of three hosts. FTP is routed directly to a special FTP server. The same applies to HTTP packets. These services are dealt with by separate hosts, so that (if something should go wrong and the machines are broken into) it is no worse than having to restore these single hosts from backup. None of the servers in the DMZ have user accounts, so there would be no help to crackers trying to crack password files there, if they managed to break in. The bastion host gets all packets which are not for the other services. The bastion host forwards okay-looking packets to the internal router which is really just a further packet filter (a backup in case of failure of the bastion host). The internal router accepts only packets passing between the safe network and the bastion host; all others are rejected. The bastion host proxies all of the appropriate protocols, including FTP and HTTP, between the safe network and the DMZ.

[2] This should not be confused with a WWW proxy, which is a kind of cache for frequently used HTML pages.

10.6.4 Example: Two Routers

A second example, in which there is no dual-homed host, is shown in Figure 10.3. In this configuration we use two routers to allow increased protection. An exterior router connects the site to the Internet, or the untrusted 'outside' network. The interior router protects the internal network from the bastion host and from the DMZ. Although the bastion host does

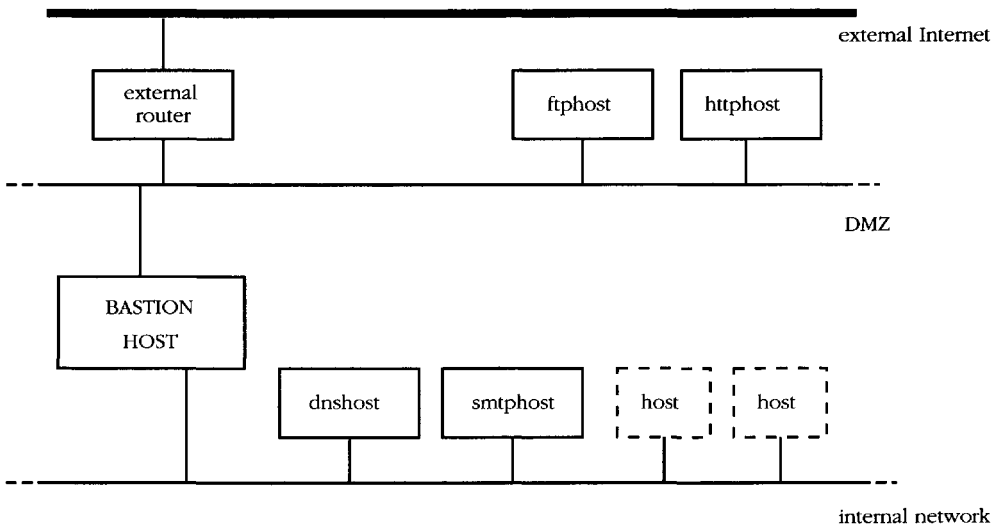


Figure 10.2 A simple firewall with a dual-homed bastion host

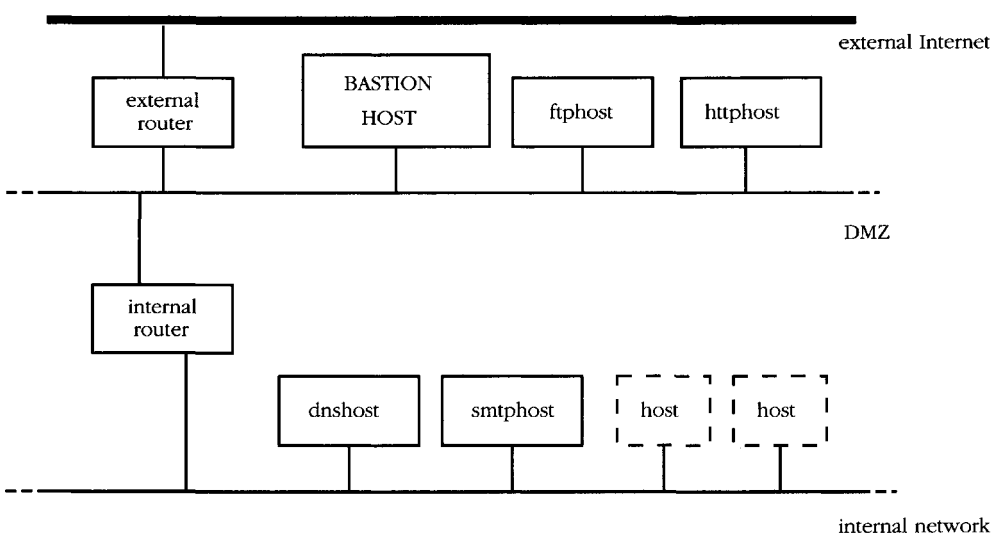


Figure 10.3 Firewall with no dual-homed host, but two routers

not physically separate the exterior and interior networks, it still separates them through proxy software, by forcing packets to be routed through the bastion host's proxy services.

To illustrate router filtering tables more explicitly, let us assume that we have a WWW server and FTP server in the DMZ, and an SMTP server on the internal network. The DNS service is split. Data pertaining to the outside network are kept in an authoritative server on the bastion host itself. DNS data about the internal network are not visible to the outside world; they are kept on the internal network, on a separate internal DNS server. Local machines are clients of the internal DNS server, so that DNS data are maximally protected.

The router filtering tables are shown and explained in Figures 10.4 and 10.5 (see also ref. [44] for an excellent discussion of filtering and firewalls in general). They are designed to

Rule	I/O	Src Addr	Dest Addr	Proto	Scr Port	Scr Port	ACK set	Action
spoof	in	intern	any	any	any	any	any	deny
telnet1	in	bastion	intern	TCP	23	>1023	yes	permit
telnet2	in	bastion	intern	TCP	>1023	23	any	permit
telnet3	out	intern	bastion	TCP	>1023	23	any	permit
telnet4	out	intern	bastion	TCP	23	>1023	yes	permit
ftp1	out	intern	bastion	TCP	>1023	21	any	permit
ftp2	in	bastion	intern	TCP	21	>1023	yes	permit
ftp3	in	bastion	intern	TCP	20	>1023	any	permit
ftp4	out	intern	bastion	TCP	>1023	20	yes	permit
smtp1	out	intern	bastion	TCP	>1023	25	any	permit
smtp2	in	bastion	smtphost	TCP	25	>1023	any	permit
smtp3	out	smtphost	bastion	TCP	25	>1023	yes	permit
smtp4	in	bastion	intern	TCP	25	>1023	yes	permit
http1	out	intern	bastion	TCP	>1023	80	any	permit
http2	in	bastion	intern	TCP	80	>1023	yes	permit
dns1	out	dnshost	bastion	UDP	53	53	N/A	permit
dns2	in	bastion	dnshost	UDP	53	53	N/A	permit
dns3	out	dnshost	bastion	TCP	>1023	53	any	permit
dns4	in	bastion	dnshost	TCP	53	>1023	yes	permit
dns5	in	bastion	dnshost	TCP	>1023	53	any	permit
dns6	out	dnshost	bastion	TCP	53	>1023	yes	permit
default1	out	any	any	any	any	any	any	deny
default1	in	any	any	any	any	any	any	deny

Figure 10.4 *Firewall example 2: Internal router filter table.* Incoming Telnet traffic is allowed from the bastion host Telnet proxy to the internal hosts, but not to the DMZ. Outgoing traffic is channeled through the bastion host proxy, which avoids the origin IP address being seen by outsiders. FTP, HTTP and SMTP traffic is allowed between the respective server-hosts and the bastion hosts' proxies. Note how WWW and FTP servers are on special 'sacrificial lamb' hosts in the DMZ, with data backed up on internal hosts. Note that FTP uses two channels: a transmission channel and a control channel on ports 20 and 21. An SMTP mail hub is used. DNS MX records should be set to point to the bastion host proxy. DNS filters are a little complex, since the DNS services use both UDP for lookup and TCP for bulk transfer and forwarding

Rule	I/O	Src Addr	Dest Addr	Proto	Scr Port	Dest Port	ACK set	Action
spoof1	in	intern	any	any	any	any	any	deny
spoof2	in	outside	any	any	any	any	any	deny
telnet1	in	any	any	TCP	23	>1023	yes	permit
telnet2	in	any	any	TCP	>1023	23	any	permit
telnet3	out	bastion	bastion	TCP	>1023	23	any	permit
telnet4	out	bastion	bastion	TCP	23	>1023	yes	permit
ftp1	out	bastion	any	TCP	>1023	21	any	permit
ftp2	in	any	bastion	TCP	21	>1023	yes	permit
ftp3	in	any	bastion	TCP	20	>1023	any	permit
ftp4	out	bastion	any	TCP	>1023	20	yes	permit
smtp1	out	bastion	any	TCP	>1023	25	any	permit
smtp2	in	any	bastion	TCP	25	>1023	yes	permit
smtp3	in	any	bastion	TCP	>1023	25	any	permit
smtp4	out	bastion	any	TCP	25	>1023	yes	permit
http1	out	bastion	any	TCP	>1023	80	any	permit
http2	in	any	bastion	TCP	80	>1023	yes	permit
http3	in	any	httphost	TCP	>1023	80	any	permit
http4	out	httphost	any	TCP	80	>1023	yes	permit
dns1	out	bastion	any	UDP	53	53	N/A	permit
dns2	in	any	bastion	UDP	53	any	N/A	permit
dns3	in	any	bastion	UDP	any	53	N/A	permit
dns4	out	bastion	any	UDP	53	any	N/A	permit
dns5	out	bastion	any	TCP	53	>1023	any	permit
dns6	in	any	bastion	TCP	53	>1023	yes	permit
dns7	in	any	bation	TCP	>1023	53	any	permit
dns8	out	bastion	any	TCP	53	>1023	yes	permit
default1	out	any	any	any	any	any	any	deny
default1	in	any	any	any	any	any	any	deny

Figure 10.5 *Firewall example 2: External router filter table.* External connections are forced to go through the bastion host proxies

route traffic through the proxy servers on the bastion host, and direct special services (SMTP, HTTP, etc.) only to the hosts which need to receive them. The bastion host, as usual, has a stripped down operating system, to remove as many potential exploits from the reach of potential intruders. The filter rules distinguish between traffic which is incoming and traffic which is outgoing. Note that TCP and UDP traffic differs here. Whereas TCP traffic generally involves fixed port addresses on servers and random ports (with port numbers greater than 1023, actually usually much higher than this) on clients, we have to be careful about filtering possible traffic based on port numbers. In practice, 1023 is probably far too low a port number to set here, but it is difficult to make generic rules for random port numbers, so we use this number as a mnemonic. Some spoofing attempts are prevented by requiring the ACK bit to be set on TCP connection requests. The ACK bit is not set on SYN packets which initiate connections, only on replies, so requiring the ACK bit to be set is a way of saying that these

rules require traffic to be part of an already established dialogue between legitimate ports. This prevents a remote attacker from using a well-known port externally to attempt to bypass the filter rules to attack a server living at a port number over 1023[3]. The corresponding outgoing rule can be considered a service to other sites, which stifles local spoofing attempts.

10.6.5 A Warning

These are just examples. In practice we might not have all the hardware we need to separate things as cleanly as shown here. Although there is a public domain firewall toolkit [87], most firewall software is commercial in nature because it needs to live in the kernel and make use of code which is proprietary.

Firewall management is a complex issue. We cannot set up a firewall and then forget about it. Firewalls need constant maintenance and they are susceptible to bugs just like any other software. It is best to build up a firewall system slowly, understanding each step. A good place to start is with packet filtering routers to eliminate the most offensive or least secure service requests from outside your local network. These include NFS (RPC), IRC, ping, finger, etc.

10.7 Intrusion Detection and Forensics

In the last few years the reality of network intrusion has led to attempts to build systems which can detect break-ins, either while they are in progress or afterwards.

There are several ways in which we can gather evidence about intrusions. Evidence can be direct and indirect. Direct evidence might come from audits and log files, smoking guns, user observations, records of actions conducted by intruders, and so on. Checksums of important files can detect unauthorized changes, for instance. Indirect evidence can be obtained by looking at system activity and trying to infer unusual activity. Changes in the behaviour of programs can signal changes in the patterns of usage of a system, perhaps flagging the exploitation of a vulnerability in software.

Intrusion detection by process monitoring is a relatively new idea. The idea is to gather a profile of what is normal and compare it to software behaviour over time. This idea is a little like the idea of an immune system which tolerates 'self' and reacts against 'non-self'. Forrest *et al.* have pioneered system call profiling, inspired by vertebrate immune systems [92, 129] in order to detect hostile patterns of activity in special software processes. The system builds a database of short patterns of system call usage, and then performs direct pattern search on subsequent data to detect anomalous patterns. The rationale for this approach is that intrusions are often caused by exploitations of system calls which do not follow intended patterns. The beauty of this approach is its natural simplicity; its disadvantage is that it incurs a high overhead in resources to implement in pattern searching in real-time; also, the system needs to be taught what is normal in advance. Unfortunately, 'normal' is a rather fickle concept [36], so in spite of its appealing simplicity, this is unlikely to be a complete, workable solution to the problem.

[3] Attackers are devious. We should not imagine that, simply because a filtering rule was intended for, say, SMTP traffic, that it could not be manipulated for some other purpose.

Another approach is to go to the network level and examine the totality of traffic arriving at a host. To detect an intrusion in progress, programs like *Network Flight Recorder* [75] (NFR) and *Big Brother* [196] (Bro) attempt to examine every packet on the network to look for telltale signatures of network break-in activity. This is an extremely resource consuming task, and it is best with a number of problems. Network monitors look for packets containing data which might represent an attack, as they arrive. Network monitoring has its problems, however. One problem is that of *fragmentation.* Fragmentation is something which occurs to IP datagrams which pass between networks with different transmission rates. Larger packets can be broken up into smaller packets in order to optimize transmission. These fragments are reassembled at the final destination. This presents a problem for intrusion detection systems because the fragmented packets might not contain enough data to identify them as hostile. This would allow them to get past the detection system. An intruder might be able to generate packets which were fragmented in such as way as to confound the attempts at detection. Another problem is that switches and routers limit the spread of traffic to specific cables. An intrusion detection system needs to see all packets in order to cover every attack. In spite of the difficulties, network intrusion detection is a hot research topic. A number of conferences on intrusion detection methods has sprung up to explore this problem in depth.

Network forensics is what one does after an intrusion. The idea is to examine logs and system audits in order to name the intruder and determine the damage. Network forensics is perhaps most important for the purpose of possible legal action against intruders. The cost of keeping the necessary logs and audits is very great, and the work required after a break-in is far from trivial.

Exercises

Exercise 10.1 Research the appropriate commands for making file system backups at your site. Consider backups to disk and backups to tape.

Exercise 10.2 Determine how many copies of each file are made in the Towers of Hanoi backup sequence.

Exercise 10.3 Collect and compile the secure shell. (Note that this software is a commercial product. You are allowed to download for strictly educational purposes, but commercial organizations must pay.)

Exercise 10.4 Explain why a switched network reduces the risk of password sniffing.

Chapter 11

Analytical System Administration

System administration has always involved a high degree of experimentation. Inadequate documentation, combined with a steep learning curve, had made that a necessity. As the curve continues to steepen and the scope of the problem only increases, the belief has gradually deepened that system administration is not merely a mechanic's job, but a scientific discipline.

A research community has grown up, led by a mixture of academics and working administrators, encouraged by organizations such as USENIX and SAGE, mainly in the US though increasingly in Europe and Australasia. The work has often been dominated by the development of software tools, since tools for the trade have been most desperately required. Now that many good tools exist, at least for Unix-based networks, the focus is changing towards more careful analyses of system administration [32, 37, 81, 263, 35], with case studies and simple experiments.

The purpose of this chapter is to stimulate objective discussions about system administration by discussing some of the possibilities for collecting scientific evidence to support hypotheses about system administration[1]. In short, we must establish a scientific basis for system administration.

11.1 Science vs Technology

Most of the research which is presently undertaken in system administration is of an applied nature. In most cases, it involves the construction of a tool which solves a specific local problem, a one-off solution to a general problem, i.e. a demonstration of possibility. A minority of authors has attempted to collate the lessons learned from these pursuits and distill their essence into a general technology of more permanent value. This is partly the nature of technological research. Science, on the other hand, deals in abstraction. The aim of science is to regard the full horror of reality and condense it into a few themes which capture its essence, without undue complication. We say that scientific knowledge has increased if

[1] Many system administrators originate from scientific disciplines and are already well-versed in scientific methods, but this could change as the future washes over us... and it never hurts to be reminded of basic techniques.

we are able to perform this extraction of the foundations some study, and if that knowledge empowers us with some increased understanding of the problem.

In science, knowledge advances by undertaking a series of studies, in order to either verify or falsify a hypothesis. Sometimes these studies are theoretical, sometimes they are empirical, and frequently they are a mixture of the two. The aim of a study is to contribute to a larger discussion, which will eventually lead to progress in the field. A single piece of work is rarely, if ever, an end in itself. Once a piece of work is published, it needs to be verified or shown to be false by others also. *Reproducibility* is an important criterion for any result, otherwise it is worthless.

How we measure progress in a field is often a contentious issue, but it can involve several themes. To test an idea it is often necessary to develop a suitable 'technology' for the investigation. That technology might be mathematical, computational or mechanical. It does not relate directly to the study itself, but it makes it possible for the study to take place. In system administration, software tools form this technology. For example, the author's management system cfengine [32] is a tool which was created to implement and refine a conceptual scheme, namely the immunity model of system maintenance [35]. There is a distinction between the tool which makes the idea possible, and the idea itself.

Having produced the tool, it is still necessary to test whether or not the original idea was a good one, better or worse than other ideas, or simply unworkable in practice. Scientific progress is made with the assistance of a tool only if the results of previous work can be improved upon, or if an increased understanding of the problem can be achieved, perhaps leading to greater predictive power or a more efficient solution to the original problem.

All problems are pieces of a larger puzzle. A complete scientific study begins with a *motivation*, followed by an *appraisal* of the problems, the construction of a *theoretical model* for understanding or solving the problems, and finally, an *evaluation* or *verification* of the *approach used* and the *results obtained*. Recently, much discussion has been directed towards finding suitable methods for evaluating technological innovations in computer science, as well as to encouraging researchers to use them. Nowadays, many computing systems are of comparable complexity to phenomena found in the natural world, and our understanding of them is not always complete, in spite of the fact that they were designed to fulfil a specific task. In short, technology might not be completely predictable, hence there is a need for experimental verification.

11.2 Studying Complex Systems

There are many issues to be studied in system administration. Some issues are of a technical nature, while others are of a human nature. System administration confronts the human-machine interaction as few other branches of computer science do. Here are some examples:

- *Reliability studies* (e.g. failure rate of hardware/software, evaluation of policies and strategies).
- *Determining and evaluating methods for ensuring system integrity* (e.g. automation, cooperation between humans, formalization of policy, etc.).
- *Observations which reveal aspects of system behaviour which are difficult to predict* (e.g. strange phenomena, periodic cycles).

- *Issues of strategy and planning.*

Science proceeds as a dialogue between theory and by experiment. We need theory to interpret results of observations, and we need observations to back up theory. Any conclusions must be a consistent mixture of the two.

To-date, very little theory has been applied to the problems of system administration. Most studies have been empirical, or anecdotal. Very few of the studies made, in the references of this book, attempt to quantify their findings. In a subject which is complex, like system administration, it is easy to fall back on *qualitative* claims. This is dangerous, however, since one is more easily fooled by qualitative descriptions than by hard numbers. At the same time, one must not believe that it is sensible to demand hard-nosed Popper-like falsification of claims in such a complex environment. Any numbers which we can measure must be considered valuable, provided they actually have a sensible interpretation.

> Computers are *complex systems*. Complexity in a system means that there is a large number of variables to be considered, probably too many to deal with in detail. Many issues are hidden directly from view and have to be discovered with some ingenuity.

A liberal attitude is usually the most constructive in making the best of a difficult lot. Any study will be worthwhile if it has something to tell us, however little. However, it is preferable if studies are authoratative, i.e. if they are able to tell us something of deeper value than mere heresay. Still, we have to judge studies for what they are worth, and no more. Authors should try to avoid marketing language which is prevalent in the commerical world, and also pointless tool-building without regard for any well thought-out model. The following questions are useful:

- What am I trying to study?
- Has it been done before? Can it be improved?
- What are the criteria for improvement?
- Can I formulate my study as a hypothesis which can be verified or falsified to some degree?
- If not, how can I clearly state the aims of my work? What are the methods available for gauging success/failure?
- How general is my study? What is the scope of its validity?
- How can my study be generalized?
- How can I ensure objectivity?

Then afterwards check:

- Is my result unambiguously true or merely a matter of interpretation?
- Are there alternative viewpoints which lead to the same conclusion?
- Is the result worth reporting to others?

Case studies are often used in fields of research where metrics are few and far between. Case studies, or anecdotal evidence, are a poor man's approach to the truth, but in system

administration we suffer from a general poverty of avenues available for investigation. Case studies, made as objectively as possible, are often the best one can do.

11.3 The Purpose of Observation

In technology the act of observation has two objective goals: (i) to gather information about a problem in order to motivate the design and construction of a technology which solves it; and (ii) to determine whether or not the resulting technology fulfils its design goals. If the latter is not fulfilled in a technological context, the system may be described as faulty, whereas in natural science there is no right or wrong. In between these two empirical bookmarks lies a theoretical model which hopefully connects the two.

The problem with technological disciplines is that what constitutes an evaluation of success or failure is often far from clear. This is because both goals and assisting technologies can be dominated by invested interests and dogged by the difficulty of constructing objective experiments with clear metrics. System administration is an example where these problems are particularly acute.

System administration is a mixture of technology and sociology. The users of computer systems are constantly changing the conditions for observations. If the conditions under which observations are made are not constant, then the data lose their meaning: the message we are trying to extract from the data is supplemented by several other messages which are difficult to separate from one another. Let us call the message we are trying to extract *signal* and the other messages which we are not interested in *noise*. Complex systems are often characterized by very noisy environments.

In most disciplines, one would attempt to reduce or eliminate the noise in order to isolate the signal. However, in system administration, it would be no good to eliminate the users from an experiment, since it is they who cause most of the problems which one is trying to solve. In principle, this kind of noise in data could be eliminated by statistical sampling over very long periods of time, but in the case of real computer systems this might not be possible, since seasonal variations in patterns of use often lead to several qualitatively different types of behaviour which should not be mixed. The collection of reliable data might therefore take many years, even if one can agree on what constitutes a reasonable experiment. This is often impractical, given the pace of technological change in the field.

11.4 Evaluation Methods and Problems

The simplest and potentially most objective way to test a model of system administration is to combine heuristic experience with repeatable simulations. Experienced system administrators have the pulse of their system and can evaluate their performance in a way that only humans can. Their knowledge can be used to define *repeatable* benchmarks or criteria for different aspects of the problem. Even so, this approach is not without its difficulties. Many of the administrators' impressions would be very difficult to gauge numerically. For example, a common theme is research which is designed to relieve administrators of tedious work, leaving them to work on more important tasks. Can such a claim be verified? Here are some of the difficulties:

Measure the time spent working on the system.	The administrator has so much to do that he/she can work full time no matter how much one automates 'tedious tasks'.
Record the actions taken by the automatic system, which a human administrator would have been required to do by hand, and compare.	There is no unique way to solve a problem. Some administrators fix problems by hand, while others will write a script for each new problem. The time/approach taken depends upon the person.

In this case, the issue was too broad to be able to quantify. Choosing the appropriate question to ask is often the most difficult aspect of an experimental study. If we restrict the scope of the question to a very specific point, we can end up with an artificial study; if the question is too broad in its scope, we risk not being able to test it convincingly.

To further clarify this point, it is useful to refer to an analogy. Imagine two researchers who create vehicles for the future, one based on renewable solar power and another based on coal. The two vehicles have identical functionality; the solar powered vehicle seems cleaner than the coal powered one, but in fact the level of pollution required to make the solar cells equals the harmful output of the coal vehicle throughout its lifetime. The laws of thermodynamics tell us that there is potential for improving this situation for the electric car but probably not for the coal powered one. The solar vehicle is lighter and more efficient, but it cannot do anything that the coal powered car cannot. All in all, one suspects that the solar powered system is a better solution, since one does not have to refuel it frequently and it is based on a technology which is universally useful, whereas the coal system is quite restricted. So what are the numbers which we should measure to distinguish one from the other, to verify the hypothesis that the solar powered vehicle is better? Is one solution really better than the other? Regardless of whether either solution is optimal, is one of them going in a sustainable direction for future development? It might seem clear that the electric vehicle is a sounder technology, since it is both sustainable in its power source and in its potential for future development, whereas the coal vehicle is something of a dead end. The solution can be ideologically correct, but this is a matter of opinion. Anyone can claim to prefer the coal powered vehicle, whether others would deem that belief to be rational or not. One can attempt to evaluate their basic principles on the basis of anecdotal evidence. One can produce numbers for many small contributing factors (such as the weight of the vehicles, or their power efficiency), but when it comes down to it, anyone can claim that those numbers do not matter because both vehicles fulfil their purpose identically.

This example is not entirely contrived. System administration requires tools. Often such tools acquire a following of users who grow to like them, regardless of what the tools allow them to achieve. Also, the marketing skills of one software producer might be better than those of another. Thus, one cannot rely on counting the numbers of users of a specific tool as an indication of its power or usefulness. On the other hand, one has to rely on the evaluations of the tools by their users.

In some cases, one technology might be better than another only in a certain context. There might be room for several different solutions. For example, are transistors better than thermionic valve devices for building computers? Most people think so, because valve technology is large and cumbersome. But advances in Russian military aerospace technology developed miniature valves because they were robust against electromagnetic pulse inter-

ference. One can think of many examples of technologies which have clear advantages, but which cannot be proven numerically, because it boils down to what people prefer to believe about them. This last case also indicates that there is not necessarily a single universal solution to a problem.

Although questionnaires and verbal evaluations which examine experienced users' impressions can be amongst the best methods of evaluating a hypothesis with many interacting components, the problems in making such a study objective are great. Questionnaires, in particular, can give misleading results, since they are often only returned by users who are already basically satisfied. Completely dissatisfied users will usually waste no time on what they consider to be a worthless pursuit, by filling out a questionnaire.

11.5 Evaluating a Hierarchical System

Evaluating a model of system administration is a little bit like evaluating the concept of a bridge. Clearly, a bridge is a structure with many components, each of which contributes to the whole. The bridge either fulfils its purpose in carrying traffic past obstacles or it does not. In evaluating the bridge, should one then consider the performance of each brick and wire individually? Should one consider the aesthetic qualities of the bridge? There might be many different designs each with slightly different goals. Can one bridge be deemed better than another on the basis of objective measurement? Perhaps only the bridge's maintainer is in a position to gain a feeling for which bridge is the most successful, but the success criterion might be rather vague: a collection of small differences which make the perceptible performance of the bridge optimal, but with no measurably significant data to support the conclusion. These are the dilemmas of evaluating a complex technology.

In refs. [48, 286] and many others, it is clear that computer scientists are embarrassed by this difficulty in bringing respectability to the field of study. In fact the difficulty is general to all fields of technology. To evaluate an approach to the solution of a problem it is helpful to create a model. A model is comprised of a principle of operation, a collection of rules and the implementation of these rules through specific algorithms. It involves a conceptual decomposition of the problem and a number of assertions or hypotheses. System administration is full of intangibles; this restricts model building to those aspects of the problem which can be addressed in schematic terms. It is also sufficiently complex that it must be addressed at several different levels in an approximately hierarchical fashion.

In brief, the options we have for performing experimental studies are

- measurements,
- simulations,
- case studies,
- user surveys,

with all of the incumbent difficulties which these entail.

11.5.1 Evaluation of the Conceptual Decomposition

It is a general principle in analysis that the details of lower level structure, insofar as they function correctly, do not change the structure of higher levels. In physics this is called the

separation of scales; in computer science it is called procedural structure or object orientation. The structure of lower levels does not affect the optimal structure of higher levels, for example. An important part of a meaningful evaluation is to sort out the conceptual hierarchy. Is the separation between high level abstractions and low level primitives sufficient, flexible, restrictive, etc.

11.5.2 Simplicity

Conceptual and practical simplicity are often deemed to be positive attributes of software systems and procedures. User surveys can be used to collect evidence of what users believe about this. The system designer's belief about the relative simplicity of his/her creation is a scientific irrelevancy.

11.5.3 Efficiency

The efficiency of a program or procedure might be an interesting way to evaluate it. Efficiency can mean many things, so the first step is to establish precisely what is meant by efficiency in context.

Most system administration tasks are not resource intensive for individual hosts. The efficiency with which they are carried out is less important than the care with which they are carried out. The reason is simple: the time required to complete most system administration tasks is very short compared to the time most users are prepared to wait.

Efficiency in terms of the consumption of human time is a much more pertinent factor. An automatic system which aims to avoid human interaction is by definition more efficient in man hours than one which places humans in the driving seat. This presupposes, of course, that the setup and maintenance of the automatic system is not so time-consuming in itself as to outweigh the advantages provided by such an approach.

11.6 Faults

The IEEE classification of software anomalies is [139]

- Operating system crash.
- Program hang-up.
- Program crash.
- Input problem.
- Output problem.
- Failed required performance.
- Perceived total failure.
- System error message.
- Service degraded.
- Wrong output.
- No output.

This classification touches on a variety of themes, all of which might plague the interaction between users and an operating system. Some of these issues encroach on the area of performance tuning, e.g. service degraded. Performance tuning is certainly related to the issue of availability of network services, and thus this is a part of system administration. However, performance tuning is of only peripheral importance compared to the matter of possible complete failure. Most of the problems associated with system administration can be attributed to input problems (incorrect or inappropriate configuration) and failed performance through loss of resources. Unlike many software situations, these are not problems which can be eliminated by re-evaluating individual software components. In system administration the problems are partly social and partly due to the cooperative nature of the many interaction software components. The unpredictability of operating systems is dominated by these issues.

11.6.1 How are Faults Corrected?

Faults occur for a plethora of reasons, too complex to present in any summarial fashion. Sometimes diagnosing a fault can take days or even weeks. In spite of this, a working solution to the fault is often extremely simple. It might be as simple as restarting a process, killing a process, editing a file, changing the access rights (permissions) to a file object, and so on. The complexity of fault diagnosis lies in the same place as the complexity of the system, i.e. that operating systems are cooperative systems with intricate causal relationships. It is usually these causal relationships which are difficult to diagnose, not the measurable effects which they have on the system. Such causal relationships make useful studies to publish in journals, since they document important experience.

The root cause of a fault is often not important to the running of the system in practice. One may complain about buggy software, but system administrators are not usually in a position to fix the software. While everyone agrees that the fault needs to be fixed at source, the system must continue to function in lieu of that time. Once a fault has been successfully diagnosed it is usually a straightforward matter to find a recipe for preventing the problem, or for curing it, if it should occur again. Problem diagnosis is way beyond the abilities of current software systems except in the simplest cases, so the best one could do would be to capture the experience of a human administrator using a knowledge based expert system. In artificial intelligence studies expert systems are not the only approach to diagnosis. Another approach, for instance, is the use of genetic algorithms. Such algorithms can be fruitful when looking for trends in statistical data, but statistically meaningful data are seldom available in system administration. The nature of most problems is direct cause and effect, perhaps with a cascade or domino effect. That is not to say that statistical data cannot be used in the future. However, at present no such data exist, and no-one knows what such data are capable of revealing about system behaviour.

11.6.2 Primitives

Suppose we abstract an operating system by considering it as the sum of its interfaces and resources. There is a only handful of operations which can be performed on this collection of objects, and so this set of basic primitives is the complete toolbox of a system administrator. One can provide helpful user interfaces to execute these primitives more easily, but no greater functionality is possible. The basic primitives are

- Examining files.
- Creating files.
- Aliasing files.
- Replacing files.
- Renaming files.
- Removing files.
- Editing files.
- Changing access rights on files.
- Starting and stopping processes or threads.
- Signalling processes or threads.
- Examining and configuring hardware devices.

From these primitives one may build more complex operations such as frequently required tasks for sharing resources. Note that the difference between a thread and a process is not usually relevant for the system administrator, so we shall speak mainly of processes and ignore the concept of a thread. The reason for this is that kernel level threads are usually transparent or invisible to processes, and user level threads cannot normally be killed or restarted without restarting an entire process.

11.6.3 Evaluation of System Administration as a Collective Effort

Few system administrators work alone. In most cases they are part of a team who all need to keep abreast of the behaviour of the system and the changes made in administration policy. Automation of system administration issues does not alter this. One issue for human administrators is how well a model for administration allows them to achieve this cooperation in practice. Does the automatic system make it easier for them to follow the development of the system in (i) theory and (ii) practice? Here theory refers to the conceptual design of the system as a whole, and practice refers to the extent to which the theoretical design has been implemented in practice. How is the task distributed between people, systems, procedures and tools? How is responsibility delegated and how does this affect individuals? Is time saved, are accuracy and consistency improved? These issues can be evaluated in a heuristic way from the experiences of administrators. Longer term, more objective studies could also be performed by analysing the behaviour of system administrators in action. Such studies will not be performed here.

11.6.4 Cooperative Software: Dependency

The fragile tower of components in any functional system is the fundament of its operation. If one component fails, how resilient is the remainder of the system to this failure? This is a relevant question to pose in the evaluation of a system administration model. How do software systems depend upon one another for their operation? If one system fails, will this have a knock-on effect on other systems? What are the core systems which form the basis of system operation? In the present work it is relevant to ask how the model continues to work in the event of the failure of DNS, NFS and other network services

which provide infrastructure. Is it possible to immobilize an automatic system administration model?

11.6.5 Evaluation of Individual Mechanisms

For individual pieces of software, it is sometimes possible to evaluate the efficiency and correctness of the components. Efficiency is a relative concept and, if used, it must be placed in a context. For example, efficiency of low level algorithms is conceptually irrelevant to the higher levels of a program, but it might be practically relevant, i.e. one must say what is meant by efficiency before quoting results. The correctness of the results yielded by a mechanism/algorithm can be measured in relation to its design specifications. Without a clear mapping of input/output, the correctness of any result produced by a mechanism is a heuristic quality. Heuristics can only be evaluated by experienced users expressing their informed opinions.

11.6.6 Evidence of Bugs in the Software

Occasionally, bugs significantly affect the performance of software. Strictly speaking, an evaluation of bugs is not part of the software evaluation itself, but of the process of software development, so while bugs should probably be mentioned they may or may not be relevant to the issues surrounding the software itself. In this work, software bugs have not played any appreciable role in either the development or the effectiveness of the results, so they will not be discussed in any detail.

11.6.7 Evidence of Design Faults

In the course of developing a program, one occasionally discovers faults which are of a fundamental nature, faults which cause one to rethink the whole operation of the program. Sometimes these are fatal flaws, but that need not be the case. Cataloguing design faults is important for future reference to avoid making similar mistakes again. Design faults may be caused by faults in the model itself or merely in its implementation. Legacy issues might also be relevant here: how do outdated features or methods affect software by placing demands on onward compatibility, or by restricting optimal design or performance?

11.6.8 Evaluation of System Policies

System administration does not exist without human attitudes, behaviours and policies. These three fit together inseparably. Policies are adjusted to fit behavioural patterns; behavioural patterns are local phenomena. The evaluation of a system policy has only limited relevance for the wider community, then: normally only relative changes are of interest, i.e. how changes in policy can move one closer to a desirable solution.

Evaluating the effectiveness of a policy in relation to the applicable social boundary conditions presents practical problems which sociologists have wrestled with for decades. The problems lie in obtaining statistically significant samples of data to support or refute the policy. Controlled experiments are not usually feasible since they would tie up resources over long periods. No-one can afford this in practice. To test a policy in a real situation the

best one can do is to rely on heuristic information from an experienced observer (in this case the system administrator). Only an experienced observer would be able to judge the value of a policy on the basis of incomplete data. Such information is difficult to trust, however, unless it comes from several independent sources. A better approach might be to test the policy with simulated data spanning the range from best to worst case. The advantage with simulated data is that the results are reproducible from those data, and thus one has something concrete to show for the effort.

11.6.9 Reliability

Reliability cannot be measured until we define what we mean by it. One common definition uses the *average (mean) time before failure* as a measure of system reliability. This is quite simply the average amount of time we expect to elapse between serious failures of the system. Another way of expressing this is to use the *average uptime*, or the amount of time for which the system is responsive (waiting no more than a fixed length of time for a response). Another complementary figure is, then, the *average downtime*, which is the average amount of time the system is unavailable for work (a kind of informational entropy). We can define the reliability as the probability that the system is available:

$$R = \frac{\text{Mean uptime}}{\text{Total elapsed time}}$$

This is clearly a number between 0 and 1. The meaning of parallelism, or redundancy, can be evaluated with an equal disregard for reality, if we treat the system as a simple linear facsimile of the Ohm's law problem:

$$\text{Rate of service (delivery)} = \text{change in information/rate of failure}$$

This is directly analogous to Ohm's law for the flow of current through a resistance:

$$I = V/R$$

The analogy is captured in this table:

Potential difference V	Change in information
Current I	Rate of service (flow of information)
Resistance R	Rate of failure

This relation is simplistic. For one thing it does not take into account variable latencies (although these could be defined as failure to respond). It should be clear that this simplistic equation is full of unwarranted assumptions, and yet its simplicity justifies its use for simple hand-waving. If we consider Figure 6.1, it is clear that a flow of service can continue, when servers work in parallel, even if one or more of them fails. In Figure 6.2 it is clear that systems which are dependent on other series are coupled in series, and a failure prevents the flow of service. Because of the linear relationship, we can use the usual Ohm's law expressions for combining failure rates:

$$R_{\text{series}} = R_1 + R_2 + R_3 + \ldots$$

and

$$\frac{1}{R_{\text{Parallel}}} = \frac{1}{R_1} + \frac{1}{R_2} + \frac{1}{R_3} \cdots$$

These simple expressions can be used to hand-wave about the reliability of combinations of hosts. For instance, let us define the rate of failure to be a probability of failure with a value between 0 and 1. Suppose we find that the rate of failure of a particular kind of server is 0.1. If we couple two in parallel (a double redundancy) then we obtain an effective failure rate of

$$\frac{1}{R} = \frac{1}{0.1} + \frac{1}{0.1}$$

i.e. $R = 0.05$, the failure rate is halved. This estimate is clearly naive. It assumes, for instance, that both servers work all the time in parallel. This is seldom the case. If we run parallel servers, normally a default server will be tried first, and if there is no response, only then will the second backup server be contacted. Thus, in a fail-over model, this is not really applicable. Still, we use this picture for what it is worth, as a crude hand-waving tool.

The Mean Time Before Failure (MTBF) is used by electrical engineers, who find that its values for the failures of many similar components (say light bulbs) has an exponential distribution. In other words, over large numbers of similar component failures, it is found that the probability of failure has the form

$$P(t) = \exp(-t/\tau)$$

or that the probability of a component lasting time t is the exponential, where τ is the mean time before failure and t is the failure time of a given component. There are many reasons why a computer system would not be expected to have this simple form. One is *dependency*. Computer systems are formed from many interacting components. The interactions with third party components mean that the environmental factors are always different. Again, the issue of fail-over and service latencies arises, spoiling the simple independent component picture. Mean time before failure doesn't mean anything unless we define the conditions under which the quantity was measured. In one test at Oslo College, the following values were measured for various operating systems, averaged over several hosts of the same type:

Solaris 2.5	86 days
GNU/Linux	36 days
Windows 95	0.5 days

While we might feel that these numbers agree with our general intuition of how these operating systems perform in practice, this is not a fair comparison since the patterns of *usage* are different in each case. An insider could tell us that the users treat the PCs with a casual disregard, switching them on and off at will: and in spite of efforts to prevent it, the same users tend to pull the plug on GNU/Linux hosts also. The Solaris hosts, on the other hand, live in glass cages where prying fingers cannot reach. Of course, we then need to ask: what is the reason why users reboot and pull the plug on the PCs? The numbers above cannot have any meaning until this has been determined. To say all of this in a single phrase: the software components of a computer system are not atomic; they are composed of many parts whose behaviour is difficult to catalogue.

Thus the problem with these measures of system reliability is that they are almost impossible to measure, and assigning any real meaning to them is fraught with subtlety. Unless the system fails regularly, the number of points over which it is possible to average is rather small. Moreover, the number of external factors which can lead to failure makes the comparison of any two values at different sites meaningless. In short, this quantity cannot be used for anything other than illustrative purposes. Changes in the reliability, for constant external conditions, can be used as a measure to show the effect of a single parameter from the environment. This is perhaps the only instance in which this can be made meaningful, i.e. as a means of quantitative comparison within a single experiment.

Another point is this: failure probabilities are almost irrelevant in today's computer systems. All of the components are so reliable that simple duplication is enough to guarantee a continuance of service with virtually unit probability. This hardly requires any lofty calculation.

11.6.10 Metrics Generally

The quantifiers which can be usefully measured or recorded on operating systems are the variables which can be used to provide quantitative support for or against a hypothesis about system behaviour. System auditing functionality can be used to record just about every operation which passes through the kernel of an operating system, but most hosts do not perform system auditing because of the huge negative effect it has on performance. Here we consider only metrics which do not require extensive auditing beyond what is normally available.

Operating system metrics are normally used for operating system performance tuning. System performance tuning requires data about the efficiency of an operating system. This is not necessarily compatible with the kinds of measurement required for evaluating the effectiveness of a system administration model. System administration is concerned with maintaining resource availability over time in a secure and fair manner. It is not about optimizing specific performance criteria.

Operating system metrics fall into two main classes: current values and average values for stable and drifting variables, respectively. Current (immediate) values are not usually directly useful, unless the values are basically constant, since they seldom accurately reflect any changing property of an operating system adequately. They can be used for fluctuation analysis, however, over some coarse-graining period. An averaging procedure over some time interval is the main approach of interest. The Nyquist law for sampling of a continuous signal is that the sampling rate needs to be twice the rate of the fastest peak cycle in the data if one is to resolve the data accurately. This includes data which are intended for averaging since this rule is not about accuracy of resolution, but about the possible complete loss of data. The granularity required for measurement in current operating systems is summarized in the following table:

0 – 5 secs	Fine grain work
10 – 30 secs	For peak measurement
10 – 30 mins	For coarse grain work
Hourly average	Software activity
Daily average	User activity
Weekly average	User activity

Although kernel switching times are of the order of microseconds, this timescale is not relevant to users' perceptions of the system. Inter-system cooperating requires many context switch cycles and I/O waits. These compound themselves into intervals of the order of seconds in practice. Users themselves spend long periods of time idle, i.e. not interacting with the system on an immediate basis. An interval of seconds is therefore sufficient. Peaks of activity can happen quickly by user perceptions, but they often last for protracted periods, thus ten to thirty seconds is appropriate here. Coarse grained behaviour requires lower resolution, but as long as one is looking for peaks, a faster rate of sampling will always include the lower rate. There is also the issue of how quickly the data can be collected. Since the measurement process itself affects the performance of the system and uses its resources, measurement needs to be kept to a level where it does not play a significant role in loading the system or consuming disk and memory resources.

The variables which characterize resource usage fall into various categories. Some variables are devoid of any apparent periodicity, while others are strongly periodic in the daily and weekly rhythms of the system. The amount of periodicity in a variable depends upon how strongly it is coupled to a periodic driving force, such as the user community's daily and weekly rhythms, and also how strong that driving force is (users' behaviour also has seasonal variations, vacations and deadlines, etc.). Since our aim is to find a sufficiently complete set of variables which characterize a macrostate of the system, we must be aware of which variables are ignorable, which variables are periodic (and can therefore be averaged over a periodic interval), and which variables are not periodic (and therefore have no unique average).

We rate periodicity as weak, strong, undetermined or fractal. Studies of total network traffic have shown an apparently self-similar (fractal) structure to network traffic when viewed in its entirety [278]. This is in contrast to telephonic voice traffic on traditional phone networks which is bursty, the bursts following a random (Poisson) distribution in arrival time. This almost certainly precludes total network traffic from a characterization of host state, but it does not preclude the use of numbers of connections/conversations between different protocols, which one would still expect to have a Poissonian profile. A value of none means that any apparent peak is much smaller than the error bars (standard deviation of the mean) of the measurements when averaged over the presumed trial period. The periodic quantities are plotted on a periodic time-scale, with each covering adding to the averages and variances. Non-periodic data are plotted on a straightforward, unbounded real line as an absolute value. A running average can also be computed, and an entropy, if a suitable division of the vertical axis into cells is defined [33]. We shall return to the definition of entropy later.

The *average type* referred to below divides into two categories: pseudo-continuous and discrete. In fact, virtually all of the measurements made have discrete results (excepting only those which are already system averages). This categorization refers to the extent to which it is sensible to treat the average value of the variable as a continuous quantity. In some cases, it is utterly meaningless. For the reasons already indicated, there are advantages to treating measured values as continuous, so it is with this motivation that we claim a pseudo-continuity to the averaged data.

In this initial instance, the data are all collected from Oslo College's own computer network, which is an academic environment with moderate resources. Once might expect our data to lie somewhere in the middle of the extreme cases which might be found amongst

the sites of the world, but one should be cognizant of the limited validity of a single set of such data. We re-emphasize that the purpose of the present work is to gauge *possibilities* rather than to extract actualities.

Net

- *Total number of packets:* characterizes the totality of traffic, incoming and outgoing on the subnet. This could have a bearing on latencies, and thus influence all hosts on a local subnet.

Weekly period	Weak/fractal
Daily period	Weak/fractal
Average type	Continuous
Expected entropy	High

- *Amount of IP fragmentation:* this is a function of the protocols in use in the local environment. It should be fairly constant, unless packets are being fragmented for scurrilous reasons.

Weekly period	None expected
Daily period	None expected
Average type	Discrete
Expected entropy	Low

- *Density of broadcast messages:* this is a function of local network services. This would not be expected to have a direct bearing on the state of a host (other than the host transmitting the broadcast), unless it became to high as to cause a traffic problem.

Weekly period	None or weak
Daily period	None or weak
Average type	Continuous
Expected entropy	Low

- *Number of collisions:* this is a function of the network community traffic. Collision numbers can significantly affect the performance of hosts wishing to communicate, thus adding to latencies. It can be brought on by the sheer amount of traffic, i.e. a threshold transition, and by errors in the physical network, or in software. In a well configured site, the number of collisions should be random. A strong periodic signal would tend to indicate a burdened network with too low a capacity for its users.

Weekly period	Undetermined
Daily period	Undetermined
Average type	Continuous
Expected entropy	Low

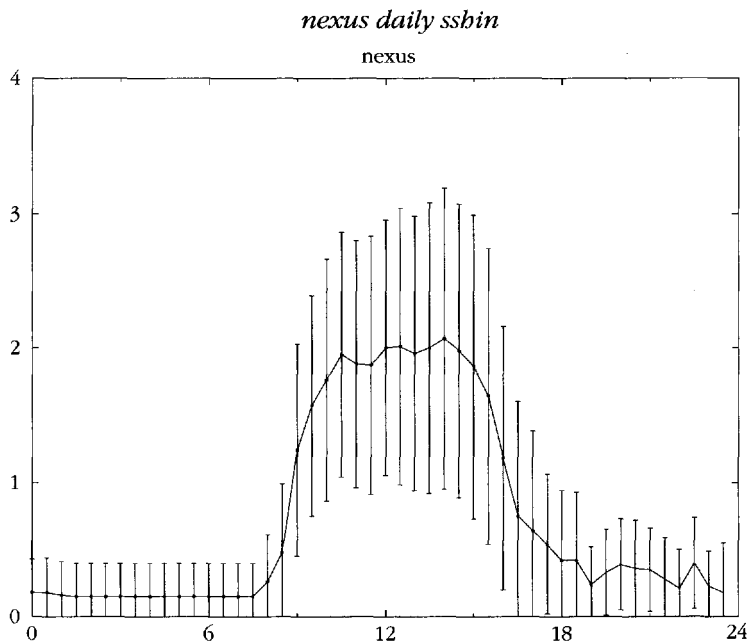


Figure 11.1 The daily rhythm of the external logins shows a strong unambiguous peak during work hours

- *Number of sockets (TCP) in and out:* this gives an indication of service usage. Measurements should be separated so as to distinguish incoming and outgoing connections. We would expect outgoing connections to follow the periodicities of the local site, whereas incoming connections would be a superposition of weak periodicities from many sites, with no net result. See Figure 11.1.

Weekly period	Strong (out)
Daily period	Strong (out)
Average type	Continuous
Expected entropy	Undetermined

- *Number of malformed packets:* this should be zero, i.e. a non-zero value here specifies a problem in some networked host, or an attack on the system.

Weekly period	None
Daily period	None
Average type	Discrete
Expected entropy	Minimal

Storage

- *Disk usage in bytes:* this indicates the actual amount of data generated and downloaded by users, or the system. Periodicities here will be affected by whatever policy one has for garbage collection. Assuming that users do not produce only garbage, there should be a periodicity superposed on top of a steady rise.

Weekly period	Yes/undetermined
Daily period	Yes/undetermined
Average type	Continuous
Expected entropy	Undetermined

- *Disk operations per second:* this is an indication of the physical activity of the disk on the local host. It is a measure of load and a significant contribution to latency both locally and for remote hosts. The level of periodicity in this signal must depend upon the relative magnitude of forces driving the host. If a host runs no network services, then it is driven mainly by users, yielding a strong periodicity. If system services dominate,

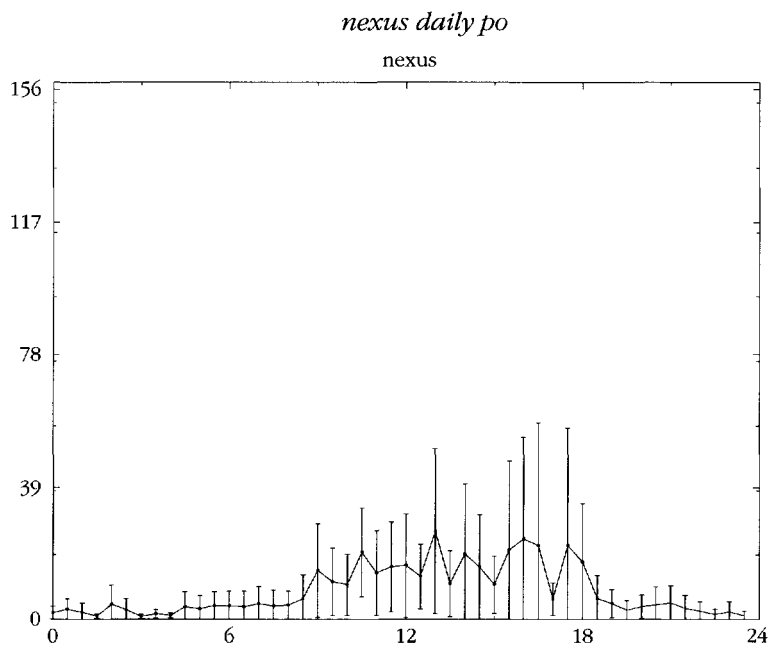


Figure 11.2 The daily rhythm of the paging data illustrates the problems one faces in attaching meaning directly to measurements. Here we see that the error bars (signifying the standard deviation) are much larger than the variation of the graph itself. Nonetheless, there is a marginal rise in the paging activity during daytime hours, and a corresponding increase in the error bars, indicating that there is a real effect, albeit of little analytical value

these could be either random or periodic. The values are thus likely to be periodic, but not necessarily strong.

Weekly period	Yes/undetermined
Daily period	Yes/undetermined
Average type	Continuous
Expected entropy	Undetermined

- *Paging (out) rate (free memory and thrashing):* these variables measure the activity of the virtual memory subsystem. In principle they can reveal problems with load. In our tests, they have proven singularly irrelevant, though we realize that we might be spoiled with the quality of our resources here. See Figures 11.2 and 11.3.

Weekly period	Insignificant
Daily period	Insignificant
Average type	Continuous
Expected entropy	Low

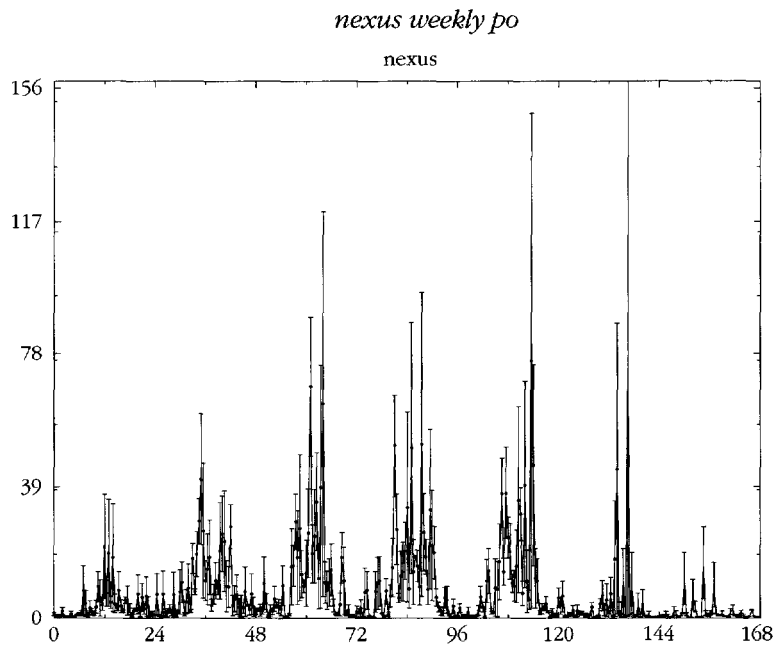


Figure 11.3 The weekly rhythm of the paging data show that there is a definite daily rhythm, but again, it is drowned in the huge variances due to random influences on the system, and is therefore of no use in an analytical context

Processes

- *Number of privileged processes:* the number of processes running system provides an indication of the number of forked processes or active threads which are carrying out the work of the system. This should be relatively constant, with a weak periodicity indicating responses to local users' requests. This is separated from the processes of ordinary users, since one expects the behaviour of privileged (root/Administrator) processes to follow a different pattern. See Figure 11.4.

Weekly period	Weak
Daily period	Weak
Average type	Discrete
Expected entropy	Low

- *Number of non-privileged processes:* this measure counts not only the number of processes but provides an indication of the range of tasks being performed by users, the number of users by implication. This measure has a strong periodic quality, relatively quiescent during weekends, rising sharply on Monday to a peak on Tuesday, followed by a gradual decline towards the weekend again. See Figures 11.5 and 11.6.

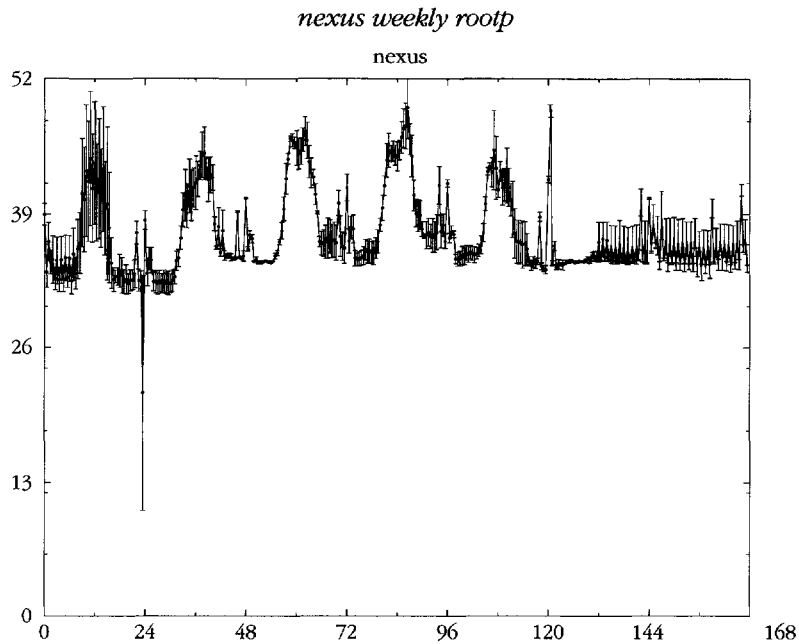


Figure 11.4 The weekly average of privileged (root) processes shows a constant daily pulse, steadily on week days. During weekends there is far less activity, but wider variance. This might be explained by assuming that root process activity is dominated by service requests from users

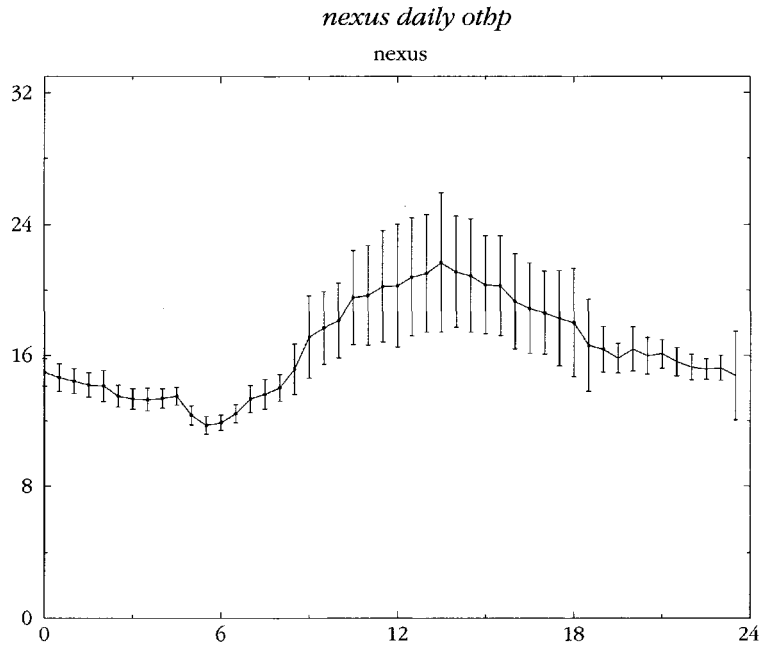


Figure 11.5 The daily average of non-privileged (user) processes shows an indisputable, strong daily rhythm. The variation of the graph is now greater than the uncertainty reflected in the error bars

Weekly period	Strong
Daily period	Strong
Average type	Continuous
Expected entropy	Undetermined

- *Maximum percentage CPU used in processes:* this is an experimental measure which characterizes the most CPU expensive process running on the host at a given moment. The significance of this result is not clear. It seems to have a marginally periodic behaviour, but basically inconclusive. The error bars are much larger than the variation of the average, but the magnitude of the errors also increases with the increasing average, thus; while for all intents and purposes this measure's average must be considered irrelevant, a weak signal can be surmised. The peak value of the data might be important however, since a high max-cpu task will significantly load the system. See Figures 11.7 and 11.8.

Weekly period	Marginal/none
Daily period	Marginal/none
Average type	Peak value only
Expected entropy	High

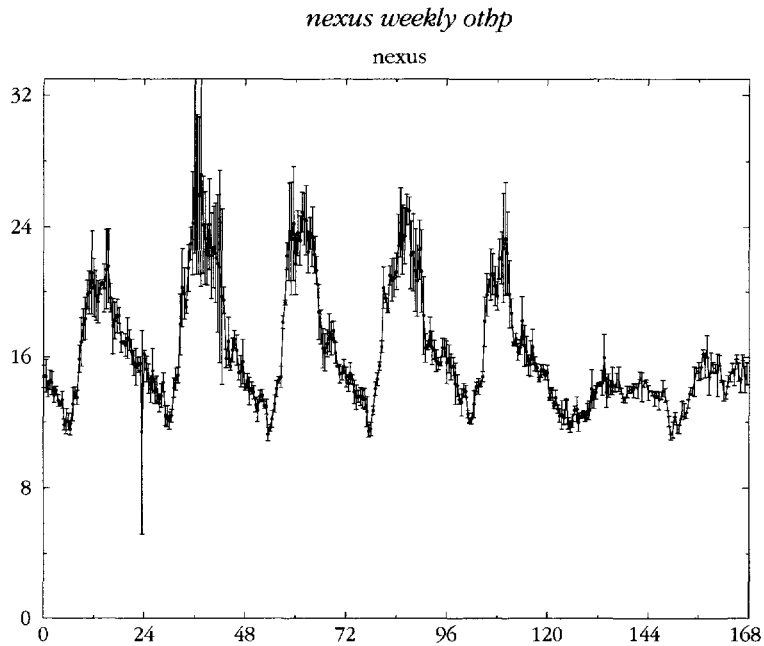


Figure 11.6 The weekly average of non-privileged (user) processes shows a constant daily pulse, quiet at the weekends, strong on Monday, rising to a peak on Tuesday and falling off again towards the weekend

Users

- *Number logged on:* this follows the classic pattern of low activity during the weekends, followed by a sharp rise on Monday, peaking on Tuesday and declining steadily towards the weekend again.

Weekly period	Strong
Daily period	Strong
Average type	Continuous
Expected entropy	Undetermined

- *Total number:* this value should clearly be constant except when new user accounts are added. The average value has no meaning, but any change in this value can be significant from a security perspective.

Weekly period	Irrelevant
Daily period	Irrelevant
Average type	Absolute value
Expected entropy	Minimal

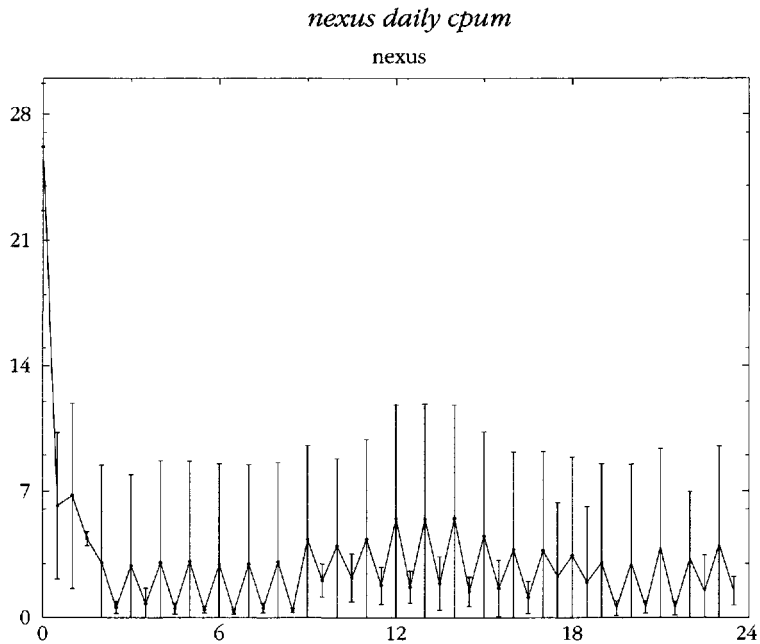


Figure 11.7 The daily average of maximal CPU percentage shows no visible rhythm, if we remove the initial anomalous point then there is no variation, either in the average or its standard deviation (error bars) which justifies the claim of a periodicity

- *Average time spent logged on per user:* can signify patterns of behaviour, but has a questionable relevance to the behaviour of the system.

Weekly period	Undetermined
Daily period	Undetermined
Average type	Continuous
Expected entropy	Undetermined

- *Load average:* this is the system's own back-of-the-envelope calculation of resource usage. It provides a continuous indication of load, but on an exaggerated scale. It remains to be seen whether any useful information can be obtained from this value; its value can be quite disordered (high entropy).

Weekly period	Strong
Daily period	Strong
Average type	Continuous
Expected entropy	High

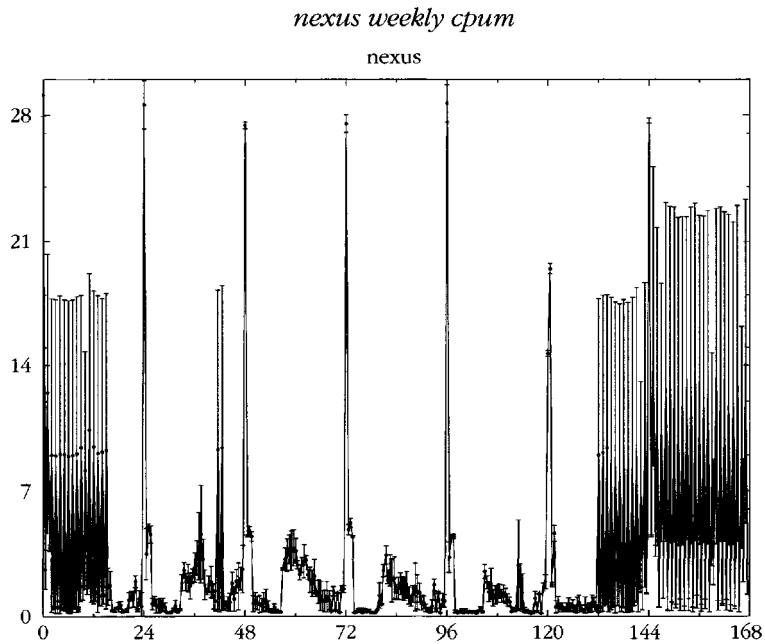


Figure11.8 The weekly average of maximal CPU percentage does appear to indicate the daily rhythm and shows how the huge variances over certain times near the beginning and end of the graph have probably obscured the signal in the daily graph. This indicates the way in which long term measurements must be combined with averages detect the true behaviour and thus be able to detect anomalies

- *Disk usage rise per session per user per hour:* the average amount of increase of disk space per user per session, indicates the way in which the system is becoming loaded. This can be used to diagnose problems caused by a single user downloading a huge amount of data from the network. During normal behaviour, if users have an even productivity, this might be periodic.

Weekly period	Undetermined
Daily period	Undetermined
Average type	Continuous
Expected entropy	Low

- *Latency of services:* the latency is the amount of time we wait for an answer to a specific request. This value only becomes significant when the system passes a certain threshold (a kind of phase transition). Once latency begins to restrict the practices of users, we can expect it to feed back and exacerbate latencies. Thus the periodicity of latencies would only be expected in a phase of the system in which user activity was in competition with the cause of the latency itself.

Weekly period	Strong above threshold
Daily period	Strong above threshold
Average type	Continuous
Expected entropy	Undetermined

Part of what one wishes to identify in looking at such variables is patterns of change. These are classifiable but not usually quantifiable. They can be relevant to policy decisions as well as in fine tuning of the parameters of an automatic response. Patterns of behaviour include

- Social patterns of the users.
- Systematic patterns caused by software systems.

Identifying such patterns in the variation of the metrics listed above is not an easy task, but it is the closest one can expect to come to a measurable effect in the a system administration context.

In addition to measurable quantities, humans have the ability to form value judgments in a way that formal statistical analyses cannot. Human judgement is based on compounded experience and associative thinking, and while it lacks scientific rigour, it can be intuitively correct in a way that is difficult to quantify. The down-side of human perception is that prejudice is also a factor which is difficult to eliminate. Also, not everyone is in a position to offer useful evidence in every judgement:

- *User satisfaction*: software, system-availability, personal freedom
- *Sysadmin satisfaction*: time-saving, accuracy, simplifying, power, ease of use, utility of tools, security, adaptability.

Other heuristic impressions include the 'amount of' dependency of a software component on other software systems, hosts or processes; also, the dependency of a software system on the presence of a human being. Kubicki [157] discusses metrics for measuring customer satisfaction. These involve validated questionnaires, system availability, system response time, availability of tools, failure analysis, and time before reboot measurements.

11.7 Deterministic and Stochastic Behaviour

In this section we turn to a more abstract view of a computer system: to think of it as a generalized dynamical system, i.e. a mathematical model which develops according in time, according to certain rules.

Abstraction is one of the most valuable assets of the human mind: it enables us to build simple models of complex phenomena, eliminating details which are only of peripheral or dubious importance. But abstraction is a double edged sword: on the one hand, abstracting a problem can show us how that problem is really the same as a lot of other problems which we know more about; conversely, unless done with a certain clarity, it can merely plant a veil of fog over our senses, obscuring rather than assisting the truth. Our aim in this section is to think of computers as abstract dynamical systems, such as those which are routinely analysed in physics and statistical analysis. Although this will not be to every working system admin-

istrator's taste, it is an important viewpoint in the pursuit of system administration as a scientific discipline.

11.7.1 Scales and Fluctuations

Complex systems are characterized by behaviour at many levels or scales. To extract information from a complex system it is necessary to focus on the appropriate scale for that information. In physics, three scales are usually distinguished in many-component systems: the *microscopic*, *mesoscopic* and *macroscopic* scales. We can borrow this terminology for convenience.

- *Microscopic* behaviour details exact mechanisms at the level of atomic operations.
- *Mesoscopic* behaviour looks as small clusters of microscopic processes and examines them in isolation.
- *Macroscopic* processes concern the long-term average behaviour of the whole system.

These three scales can also be discerned in operating systems, and they must usually be considered separately. At the microscopic level we have individual system calls and other atomic transactions (on the order of microseconds to milliseconds). At the mesoscopic level we have clusters and patterns of system calls and other process behaviour, including algorithms, procedures, possibly arising from single processes or groups of processes. Finally, there is the macroscopic level at which one views all the activities of all the users over scales at which they typically work and consume resources (minutes, hours, days, weeks). There is clearly a measure of arbitrariness in drawing these distinctions. The point is that there are typically three scales which can be usefully distinguished in a relatively stable dynamical system.

11.7.2 Principle of Superposition

In any dynamical system where several microscopic processes can coexist, there are two possible scenarios:

- Every process is completely independent of every other. System resources change linearly (additively) in response to new processes.
- The addition of each new process affects the behaviour of the others in a non-additive (non-linear) fashion.

The first of these is called the principle of superposition. It is a generic property of *linear* systems (actually this is a defining tautology). In the second case, the system is said to be non-linear because the result of adding lots of processes is not merely the sum of those processes: the processes interact and complicate matters. Owing to the complexity of interactions between subsystems in a network, it is likely that there is at least some degree of non-linearity in the measurements we are looking for. That means that a change in one part of the system will have communicable, knock-on effects on another part of the system, with possible feedback, and so on.

This is one of the things which needs to be examined, since it has a bearing on the shape of the distribution one can expect to find. Empirically one often finds, in non-linear systems, that the probability of a deviation Δx from the expected behaviour is [102]

$$P(\Delta x) = \frac{1}{2\sigma} \exp\left(-\frac{|\Delta x|}{\sigma}\right)$$

for large jumps. This can be contrasted with a Gaussian measure for a random sample

$$P(\Delta x) = \frac{1}{(2\pi)^{1/2}\sigma} \exp\left(-\frac{\Delta x^2}{2\sigma^2}\right)$$

which one might normally expect. It is interesting to determine the extent of non-linearity in the behaviour of computer systems.

11.7.3 The Idea of Convergence

In order to converge to a stable equilibrium one needs to provide countermeasures to change which are switched off when the system has reached its desired state. For this to happen, a policy of checking-before-doing is required. This is actually a difficult issue which becomes increasingly difficult with the complexity of the task involved. Fortunately, most system configuration issues are solved by simple means (file permissions, missing files, etc.), and thus, in practice, it can be a simple matter to test whether the system is in its desired state before modifying it.

In mathematics a random perturbation in time is represented by Gaussian noise, or a function whose expectation value, averaged over a representative time interval, is zero

$$\langle f \rangle = \frac{1}{T}\int_0^T dt\, f(t) = 0.$$

The simplest model of random change is the driven harmonic oscillator.

$$\frac{d^2s}{dt^2} + \gamma\frac{ds}{dt} + \omega_0^2 = f(t),$$

where s is the state of the system and γ is the rate at which it converges to a steady state. To make oscillations converge, they are damped by a frictional or counter force γ (in the present case the immune system is the frictional force which will damp down unwanted changes). To have any chance of stopping the oscillations the counterforce must be able to change direction in time with the oscillations so that it is always opposing the changes at the same rate as the changes themselves. Formally, this is ensured by having the frictional force proportional to the rate of change of the system, as in the differential representation above. The solutions to this kind of motion are *damped oscillations* of the form

$$s(t) \sim e^{-\gamma t} \sin(\omega t + \phi),$$

for some frequency ω and damping rate γ. In the theory of harmonic motion, three cases are distinguished: under-damped motion, damped and over-damped motion. In under-damped motion $\gamma \ll \omega$, there is never sufficient counterforce to make the oscillations converge to any

degree. In damped motion the oscillations do converge quite quickly $\gamma \sim \omega$. Finally, with over-damped motion $\gamma \gg \omega$ the counterforce is so strong as to never allow any change at all.

Under-damped	Inefficient: the system can never quite keep errors in check.
Damped	System converges in a time scale of the order the rate of fluctuation.
Over-damped	Too Draconian: processes killed frequently while still in use.

Clearly, an over-damped solution to system management is unacceptable. This would mean that the system could not change at all. If one does not want any changes then it is easy place the machine in a museum and switch it off. Also, an under-damped solution will not be able to keep up with the changes to the system made by users or attackers.

The *slew rate* is the rate at which a device can dissipate changes in order to keep them in check. If the immune response ran continuously, then the rate at which it completed its tasks would be the approximate slew rate. In the body it takes two or three days to develop an immune response, approximately the length of time it takes to become infected, so that minor episodes last about a week. In a computer system there are many mechanisms which work at different time scales and need to be treated with greater or lesser haste. What is of central importance here is the underlying assumption that an immune response will be timely. The time scales for perturbation and response must match. Convergence is not a useful concept in itself, unless it is a dynamical one. Systems must be allowed to change, but they must not be allowed to become damaged. Presently there are few objective criteria for making this judgement, so it falls to humans to define such criteria, often arbitrarily.

In addition to random changes, there is also the possibility of systematic error. Systematic change would lead to a constant unidirectional drift (clock drift, disk space usage, etc.). These changes must be cropped sufficiently frequently (producing a sawtooth pattern) to prevent serious problems from occurring. A serious problem would be defined as a problem which prevented the system from functioning effectively. In the case of disk usage, there is a clear limit beyond which the system cannot add more files, thus corrective systems need to be invoked more frequently when this limit is approached, but also in advance of this limit with less frequency to slow the drift to a minimum. In the case of clock drift, the effects are more subtle.

11.7.4 Parameterizing a Dynamical System

If we wish to describe the behaviour of a computer system from an analytical viewpoint, we need to be able to write down a number of variables which capture its behaviour. Ideally, this characterization would be numerical since quantitative descriptions are more reliable than qualitative ones, though this might not always be feasible. To properly characterize a system, we need a theoretical understanding of the system or sub-system which we intend to describe. There is a few important points to be clear about.

Dynamical systems fall into two categories, depending on how we choose our problem to analyse. These are called *open systems* and *closed systems*:

- *Open system:* this is a *sub-system* of some greater whole. An open system can be thought of as a black box which takes in input and generates output, i.e. it communicates with its

environment. The names *source* and *sink* are traditionally used for the input and output routes. What happens in the black box depends upon the state of the environment around it. The system is open because input changes the state of the system's internal variables and output changes the state of the environment. Every piece of computer software is an open system. Even an isolated total computer system is an open system as long as any user is using it. If we wish to describe what happens inside the black box, then the source and the sink must be modelled by two variables which represent the essential behaviour of the environment. Since one cannot normally predict the exact behaviour of what goes on outside of a black box (it might itself depend upon many complicated variables), any study of an open system tends to be incomplete. The source and sink are essentially unknown quantities. Normally one would choose to analyse such a system by choosing some special input and consider a number of special cases. An open system is internally *deterministic*, meaning that it follows strict rules and algorithms, but its behaviour is not necessarily determined, since the environment is an unknown.

- *Closed system:* this is a system which is complete, in the sense of being isolated from its environment. A closed system receives no input and normally produces no output. Computer systems can only be approximately closed for short periods of time. The essential point is that a closed system is neither affected by, nor affects, its environment. In thermodynamics, a closed system always tends to a steady state. Over short periods, under controlled conditions, this might be a useful concept in analysing computer subsystems, but only as an idealization. To speak of a closed system, we have to know the behaviour of all the variables which characterize the system. A closed system is said to be completely *determined*[2].

An important difference between an open system and a closed system is that an open system is not always in a steady state. New input changes the system. The internal variables in the open system are altered by external perturbations from the source, and the sum state of all the internal variables (which can be called the system's *macrostate*) reflect the history of changes which have occurred from outside. For example, suppose we are analysing a word processor. This is clearly an open system: it receives input and its output is simply a window on its data to the user. The buffer containing the text reflects the history of all that was input by the user, and the output causes the user to think and change the input again. If we were to characterize the behaviour of a word processor, we would describe it by its internal variables: the text buffer, any special control modes or switches, etc.

Normally we are interested in components of the operating system which have more to do with the overall functioning of the machine, but the principle is the same. The difficulty with such a characterization is that there is no unique way of keeping track of a system's history over time, quantitatively. That is not to say that no such measures exist. Let us consider one simple cumulative quantifier of the system's history, which was introduced by Burgess [33], namely its entropy or disorder. Entropy has certain qualitative, intuitive features which are easily understood. Disorder in a system measures the extent to which it is occupied by files and processes which prevent useful work. If there is a high level of disorder, then –

[2] This does not mean that it is exactly calculable. Non-linear, chaotic systems are deterministic but inevitably inexact over any length of time.

depending on the context – one might either feel satisfied that the system is being used to the full, or one might be worried that it's capacity is nearing saturation.

There are many definitions of entropy in statistical studies. Let us choose Shannon's traditional informational entropy as an example [235]. For the informational entropy to work usefully as a measure, we need to be selective in the type of data which are collected.

In ref. [33], the concept of an informational entropy was used to gauge the stability of a system over time. In any feedback system there is the possibility of instability: of either wild oscillation or exponential growth. Stability can only be achieved if the state of the system is checked often enough to adequately detect the resolution of the changes taking place. If the checking rate is too slow, or the response to a given problem is not strong enough to contain it, then control is lost.

To define an entropy we must change from dealing with a continuous measurement to a classification of ranges. Instead of measuring a value exactly, we count the amount of time a value lies within a certain range and say that all of those values represent a single state. Entropy is closely associated with the amount of granularity or roughness in our perception of information, since it depends upon how we group the values into classes or states. Indeed, all statistical quantifiers are related to some procedure for coarse-graining information, or eliminating detail. To define an entropy one needs, essentially, to distinguish between signal and noise. This is done by blurring the criteria for the system to be in a certain state. As Shannon put it, we introduce redundancy into the states so that a range of input values (rather than a unique value) triggers a particular state. If we consider every single jitter of the system to be an important quantity, to be distinguished by a separate state, then nothing is defined as noise, and chaos must be embraced as the natural law. However, if one decides that certain changes in the system are too insignificant to distinguish between, such that they can be lumped together and categorized as a single state, then one immediately has a distinction between useful signal and error margins for useless noise. In physics, this distinction is thought of in terms of order and disorder.

Let us represent a single quantifier of system resources as a function of time $f(t)$. This function could be the amount of CPU usage, or the changing capacity of system disks, or some other variable. We wish to analyse the behaviour of system resources by computing the amount of entropy in the signal $f(t)$. This can be done by coarse-graining the range of $f(t)$ into N cells:

$$F_-^i < f(t) < F_+^i,$$

where $i = 1..N$,

$$F_+^i = F_-^{i+1}$$

and the constants $F_\pm^i$ are the boundaries of the ranges (imagine drawing horizontal threshold lines from the tick marks of Figure 11.9, thus dividing up the graph into horizontal slices). The probability that the signal lies in the cell i, during the time interval from zero to T, is the fraction of time the function spends in each cell i:

$$p_i(T) = 1T \int_0^T dt \left[\theta\left(f(t) - F_-^i\right) - \theta\left(f(t) - F_+^i\right)\right],$$

where $\theta(t)$ is the step function, defined by

$$\theta(t-t') = \begin{cases} 1 & t-t' > 0 \\ \frac{1}{2} & t = t' \\ 0 & t-t' < 0 \end{cases}$$

Now, let the statistical degradation of the system then be given by the Shannon entropy [235]

$$E(T) = -\sum_{i=1}^{N} p_i(T) \log p_i(T),$$

where p_i is the probability of seeing event i on average. i runs over an alphabet of all possible events from 1 to N, which is the number of independent cells in which we have chosen to coarse-grain the range of the function $f(t)$. The entropy, as defined, is always a positive quantity, since p_i is a number between 0 and 1.

Entropy is lowest if the signal spends most of its time in the same cell $F_{\pm}^{i}$. This means that the system is in a relatively quiescent state, and it is therefore easy to predict the probability that it will remain in that state, based on past behaviour. Other conclusions can be drawn from the entropy of a given quantifier. For example, if the quantifier is disk usage, then a state of low entropy or stable disk usage implies little usage, which in turn implies low power consumption. This might also be useful knowledge for a network; it is easy to forget that computer systems are reliant on physical constraints. If entropy is high it means that the system is being used very fully: files are appearing and disappearing rapidly: this makes it

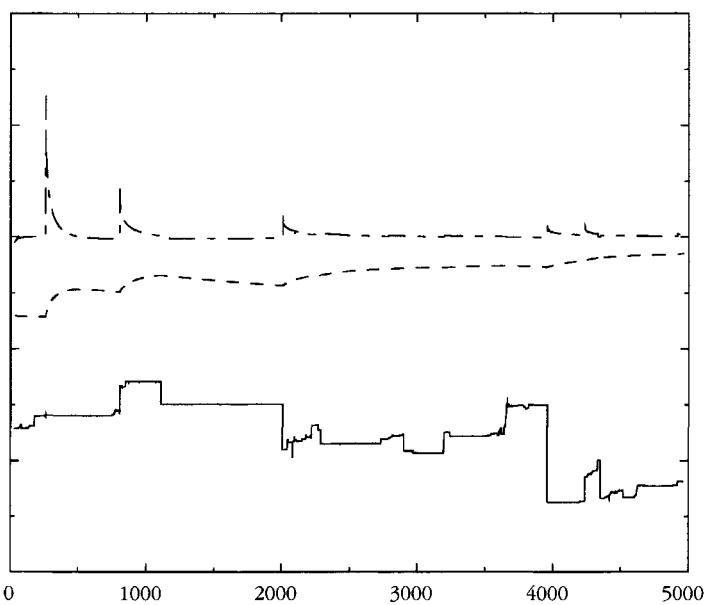


Figure 11.9 Disk usage as a function of time over the course of a week, beginning with Saturday. The lower solid line shows actual disk usage. The middle line shows the calculated entropy of the activity and the top line shows the entropy gradient. Since only relative magnitudes are of interest, the vertical scale has been suppressed. The relatively large spike at the start of the upper line is due mainly to initial transient effects. These even out as the number of measurements increases. Reproduced with permission from ref. [33]

difficult to predict what will happen in the future, and the high activity means that the system is consuming a lot of power. The entropy and entropy gradient of sample disk behaviour is plotted in Figure 11.9.

Another way of thinking about the entropy is that it measures the amount of noise or random activity on the system. If all possibilities occur equally on average, then the entropy is maximal, i.e. there is no pattern to the data. In that case, all of the p_i are equal to $1/N$ and the maximum entropy is $(\log N)$. If every message is of the same type, then the entropy is minimal. Then all the p_i are zero except for one, where $p_x = 1$. Then the entropy is zero. This tells us that, if $f(t)$ lies predominantly in one cell, then the entropy will lie in the lower end of the range $0 < E < \log N$. When the distribution of messages is random, it will be in the higher part of the range.

Entropy can be a useful quantity to plot, in order to gauge the cumulative behaviour of a system within a fixed number of states. It is one of many possibilities for explaining the behaviour of an open system over time, experimentally. Like all cumulative, approximate quantifiers it has a limited value, however, so it needs to be backed up by a description of system behaviour.

11.7.5 Causality and Dependency

We would often like to be able to establish a causal connection between a change of a specific parameter and the resulting change in the system. This can be useful in substantiating claims about the effectiveness of a program or policy, for instance. The principle of causality is simply stated:

Principle 50 (Causality) *Every change or effect happens in response to a cause, which precedes it.*

This principle sounds intuitive and even manifestly obvious, but the way in which cause and effect are related in a dynamical system is not always as clear, as one might imagine. In this section, the aim is to show ways in which we can be deceived as to the true cause of observed behaviour through inadequate analysis.

Suppose we want to consider the behaviour of a small subsystem within the entirety of a networked computer system. First we have to define what we mean by the subsystem we are studying. This might be a straightforward conceptual partitioning of the total system, but conceptual decompositions do not necessarily preserve causal relationships (see Figure 11.10).

In fact we might have to make special allowances for the fact that the subsystem might not be completely described by a closed set of variables. By treating a subsystem as though it were operating in isolation, we might be ignoring important links in the causal web. If we ignore some of the causal influences to the subsystem, its behaviour will seem confusing and unpredictable.

There is a simple mathematical expression of this idea. A total system $S(x_1 \ldots x_n)$ can be treated as two independent subsystems, if and only if the system of variables can be factorized

$$S(x_1 \ldots x_n) \rightarrow s_1(x_1 \ldots x_p) \cdot s_2(x_p \ldots x_n).$$

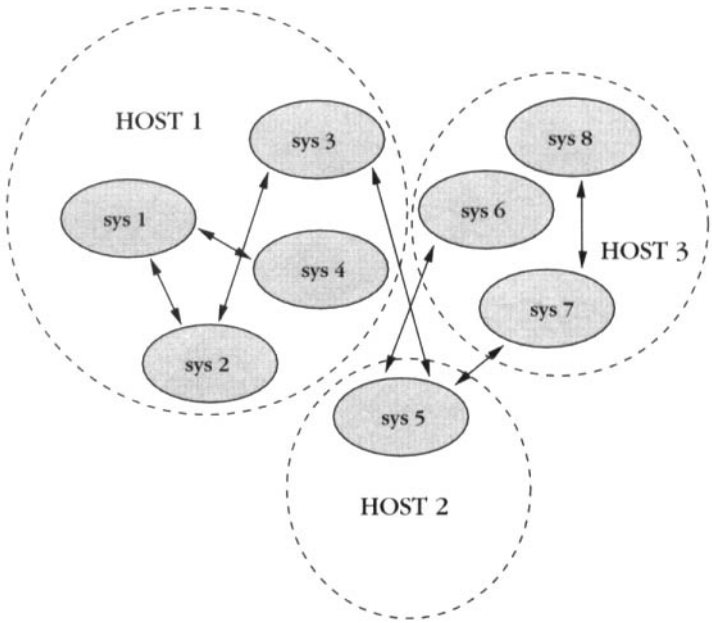


Figure 11.10 A complex system is a causal web or network of intercommunicating parts. It is only possible to truly isolate a subsystem if we can remove a piece of the network from the rest without cutting a connection. If we think of the total system as $S(x_1 \ldots x_n)$, and the individual subsystems as $s_1(x_1 \ldots x_p), s_2(x_p \ldots x_n)$, etc., then one can analyse a subsystem as an open system if the subsystems share any variables, or as a closed system if there are no shared variables

In other words, there has to be a separation of variables. This is a precise statement of something which is intuitively obvious, but which might be practically impossible to achieve. The problem is this: Most of the parts of the causal web in Figure 11.10 are themselves closed to us. We do not know that state of all their internal variables, or how they are affected by other parts of the system. Indeed, the task of knowing all of this information is prohibitively difficult.

Subtle interactions with third-party systems can introduce non-linearities which can lead to behaviour which would appear confusing or even impossible in a closed system. How many times have we cursed the computer for not behaving logically? Of course it always behaves causally and logically, the problem when it seems to make no sense is simply that it is behaving outside of our current conceptual model. That is a failure in our conceptual model. The principle of causality tells us that unpredictable behaviour means that we have an *incomplete description* of the subsystem. Notice that this only applies to a subsystem, since any description of the total system is, by definition, complete. There is another issue, however, by which confusing behaviour can seem to arise. That is by *coarse graining* information. Whenever we simplify data by blurring distinctions, information is lost irretrievably. If we then trust the coarse-grained data, it is possible to obtain the illusion of non-causal behaviour, since the true explanation of the data has been blurred into obscurity.

Causality is a mapping from cause to effect. The point of a complex system with many variables is that this mapping might not be one-to-one or even many-to-one. In general, the mapping is many-to-many. There are knock-on effects. Experimentally, we must have a repeatable demonstration (this establishes a stable context), but we also need a theory about cause and effect (a description of the mapping, sometimes called the *kernel* of the mapping). We need to identify which variables play a relevant role, and we need to factor out any irrelevant variables from the description.

11.7.6 Stochastic (Random) Variables

A stochastic or random variable is a variable whose value depends upon the outcome of some underlying random process. The range of values of the variable is not at issue, but which particular value the variable has at a given moment is random. We say that a stochastic variable X will have a certain value x with a probability $P(x)$:

- Choices made by large numbers of users.
- Measurements collected over long periods of time.
- Cause and effect are not clearly related

Certain measurements can often appear random, because we do not know all of the underlying mechanisms. We say that there are *hidden variables*. If we sample data for long enough, they will fall into a Gaussian type of distribution, by virtue of the *central limit theorem* (see, for instance, ref. [108]).

11.7.7 Probability Distributions and Measurement

Whenever we repeat a measurement and obtain different results, a distribution of different answers is formed. The spread of results needs to be interpreted. There are two possible explanations for a range of values:

- The quantity being measured does not have a fixed value.
- The measurement procedure is imperfect and incurs a range of values due to error or uncertainty.

Often both of these are the case. To give any meaning to a measurement, we have to repeat the measurement a number of times, and show that we obtain approximately the same answer each time. In any complex system, in which there are many things going on which are beyond our control (read: just about anywhere in the real world), we will never obtain exactly the same answer twice. Instead we will get a variety of different answers which we can plot as a graph: on the x-axis, we plot the actual measured value, and on the y-axis we plot the number of times we obtained that measurement divided by a normalizing factor, such as the total number of measurements. By drawing a curve through the points, we obtain an idealized picture which shows the probability of measuring the different values. The normalization factor is usually chosen so that the area under the curve is unity.

There are two extremes of distribution: complete certainty (Figure 11.11) and complete uncertainty (Figure 11.12). If a measurement always gives precisely the same answer, then we say that there is no error. This is never the case in real measurements. Then the

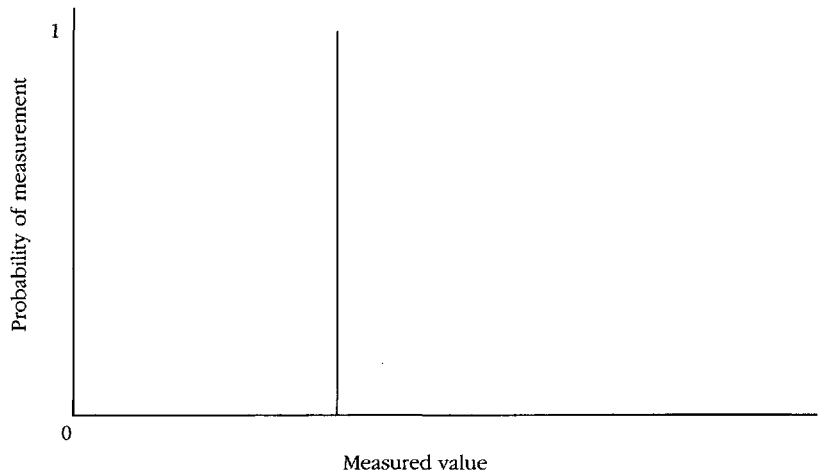


Figure 11.11 The delta distribution represents complete certainty. The distribution has a value of 1 at the measured value

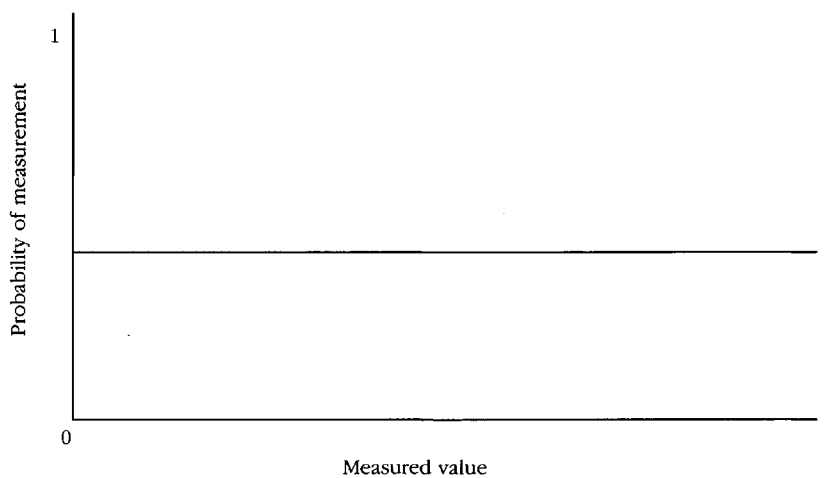


Figure 11.12 The flat distribution is a horizontal line indicating that all measured values, within the shown interval, occur with equal probability

curve is just a sharp spike at the particular measured value. If we obtain a different answer each time we measure a quantity, then there is a spread of results. Normally that spread of results will be concentrated around some more or less stable value (Figure 11.13). This indicates that the probability of measuring that value is biased, or tends to lead to a particular range of values. The smaller the range of values, the closer we approach Figure 11.11. However, the converse might also happen: in a completely random system, there might be no fixed value of the quantity we are measuring. In that case, the measured value is

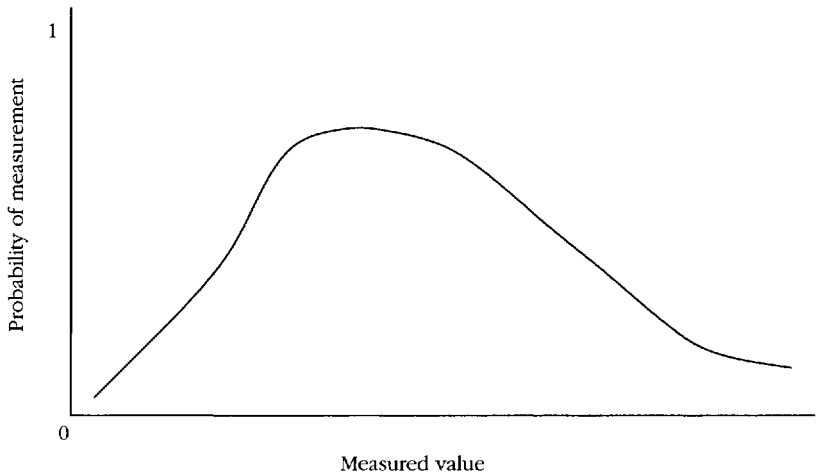


Figure 11.13 Most distributions peak at some value, indicating that there is an expected value (expectation value) which is more probable than all the others

completely uncertain, as in Figure 11.12. To summarize, a flat distribution is unbiased, or completely random. A non-flat distribution is biased, or has an expectation value, or probable outcome. In the limit of complete certainty, the distribution becomes a spike, called the *delta distribution.*

We are interested in determining the shape of the distribution of values on repeated measurement for the following reason. If the variation of the values is symmetrical about some preferred value, (i.e. if the distribution peaks close to its mean value), then we can probably infer that the value of the peak or of the mean is the true value of the measurement and that the variation we measured was due to random external influences. If, on the other hand, we find that the distribution is very asymmetrical, some other explanation is required and we are most likely observing some actual physical phenomenon which requires explanation.

11.8 Observational Errors

All measurements involve certain errors. One might be tempted to believe that, where computers are involved, there would be no error in collecting data, but this is false. Errors are not only a human failing; they occur because of unpredictability in the measurement process, and we have already established throughout this book that computer systems are nothing if not unpredictable. We are thus forced to make estimates of the extent to which our measurements can be in error. This is a difficult matter, but approximate statistical methods are well known in the natural sciences, methods which become increasingly accurate with the amount of data in an experimental sample.

The ability to estimate and treat errors should not be viewed as an excuse for constructing a poor experiment. Errors can only be minimized by design.

11.8.1 Random, Personal and Systematic Errors

There are three distinct types of error in the process of observation. The simplest type of error is called *random error*. Random errors are usually small deviations from the 'true value' of a measurement which occur by accident, by unforeseen jitter in the system, or some other influence. By their nature, we are usually ignorant of the cause of random errors, otherwise it might be possible to eliminate them. The important point about random errors is that they are distributed evenly about the mean value of the observation. Indeed, it is usually assumed that they are distributed with an approximately *normal* or *Gaussian* profile about the mean. This means that there are as many positive as negative deviations, and thus random errors can be averaged out by taking the mean of the observations.

It is tempting to believe that computers would not be susceptible to random errors. After all, computers do not make mistakes. However, this is an erroneous belief. The measurer is not the only source of random errors. A better way of expressing this is to say that random errors are a measure of the unpredictability of the measuring process. Computer systems are also unpredictable, since they are constantly influenced by outside agents such as users and network requests.

The second type of error is a *personal error*. This is an error which a particular experimenter adds to the data unwittingly. There are many instances of this kind of error in the history of science. In a computer controlled measurement process, this corresponds to any particular bias introduced through the use of specific software, or through the interpretation of the measurements.

The final and most insidious type of error is the *systematic error*. This is an error which runs throughout all of the data. It is a systematic shift in the true value of the data, in one direction, and thus it cannot be eliminated by averaging. A systematic error also leads to an error in the mean value of the measurement. The sources of systematic error are often difficult to find, since they are often a result of misunderstandings, or of the specific behaviour of the measuring apparatus.

In a system with finite resources, the act of measurement itself leads to a change in the value of the quantity one is measuring. To measure the CPU usage of a computer system, for instance, we have to start a new program which collects that information, but that program inevitably also uses the CPU, and therefore changes the conditions of the measurement. These issues are well known in the physical sciences and are captured in principles such as Heisenberg's Uncertainty Principle, Schrödinger's cat and the use of infinite idealized heat baths in thermodynamics. We can formulate our own verbal expression of this for computer systems:

Principle 51 (Uncertainty) *The act of measuring a given quantity in a system with finite resources always changes the conditions under which the measurement is made, i.e. the act of measurement changes the system.*

For instance, to measure the pressure in a tyre, you have to let some of the air out, which reduces the pressure slightly. This is not noticeable on a car tyre, but it can be noticeable on a bicycle. The larger the available resources of the system, compared to the resources required to make the measurement, the smaller the effect on the measurement will be.

11.8.2 Adding up Independent Causes

Suppose we want to measure the value of a quantity v whose value has been altered by a series of independent random changes or perturbations Δv_1, Δv_2..., etc. By how much does that series of perturbations alter the value of v? Our first instinct might be to add up the perturbations to get the total:

$$\text{Actual deviation} = \Delta v_1 + \Delta v_2 + \ldots$$

This estimate is not useful, however, because we do not usually know the exact values of Δv_i, we can only guess them. In other words, we are working with a set of guesses Δg_i, whose sign we do not know. Moreover, we do not know the signs of the perturbations, so we do not know whether they add or cancel each other out. In short, we are not in a position to know the actual value of the deviation from the true value. Instead, we have to estimate the limits of the possible deviation from the true value v. To do this, we add the perturbations together as though they were independent vectors.

Independent influences are added together using Pythagoras theorem, because they are independent vectors. This is easy to understand geometrically. If we think of each change as being independent, then one perturbation Δg_1 cannot affect the value of another perturbation Δv_2. But the only way that it is possible to have two changes which do not have any effect on one another is if they are movements at right angles to one another, i.e. they are orthogonal. Another way of saying this is that the independent changes are like the coordinates $x, y, z, \ldots$ of a point which is at a distance from the origin in some set of coordinate axes. The total distance of the point from the origin is, by Pythagoras theorem,

$$d = \sqrt{x^2 + y^2 + z^2 + \ldots}.$$

The formula we are looking for, for any number of independent changes, is just the N-dimensional generalization of this, usually written σ:

$$\sigma = \sqrt{\Delta g_1^2 + \Delta g_2^2 \ldots}.$$

or

$$\sigma^2 = \sum_{i=0}^{N} \Delta g_i^2.$$

This tells us the distance d, or σ by which we can expect the value we are trying to measure to have changed. It does not tell us the sign of the change, so all we can now say is that the true value could be in the range $v \pm \sigma$. To summarize, independent changes in a quantity are like Cartesian coordinates for a vector in an N-dimensional space.

11.8.3 The Mean and Standard Deviation

In the theory of errors, we use the ideas above to define two quantities for a set of data: the mean and the standard deviation. Now the situation is reversed: we have made a number of observations of values v_1, v_2, v_3... which have a certain scatter, and we are trying to find out the actual value v. Assuming that there are no systematic errors, i.e. assuming that all of the

deviations have independent random causes, we define the value $\bar{v}$ to be the arithmetic mean of the data:

$$\bar{v} = \frac{v_1 + v_2 ... v_N}{N} = \frac{1}{N} \sum_{i=1}^{N} v_i.$$

Next we treat the deviations of the actual measurements as our guesses for the error in the measurements:

$$\begin{aligned} \Delta g_1 &= \bar{v} - v_1 \\ \Delta g_2 &= \bar{v} - v_2 \\ &\vdots \\ \Delta g_N &= \bar{v} - v_N \end{aligned}$$

and define the *standard deviation* of the data by

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=0}^{N} \Delta g_i^2}.$$

This is clearly a measure of the scatter in the data due to random influences. σ is the Root Mean Square (RMS) of the assumed errors. These definitions are a way of interpreting measurements, from the assumption that one really is measuring the true value, affected by random interference.

An example of the use of standard deviation can be seen in the error bars of the figures in this chapter. Whenever one quotes an average value, the number of data and the standard deviation should also be quoted in order to give meaning to the value. In system administration, one is interested in the average values of any system metric which fluctuates with time.

11.8.4 The Normal Error Distribution

It has been stated that 'Everyone believes in the exponential law of errors; the experimenters because they think it can be proved by mathematics; and the mathematicians because they believe it has been established by observation' [277]. Some observational data in science closely satisfy the normal law of error, but this is by no means universally true. The main purpose of the normal error law is to provide an adequate idealization of error treatment which is simple to deal with, and which becomes increasingly accurate with the size of the data sample.

The normal distribution was first derived by DeMoivre in 1733, while dealing with problems involving the tossing of coins; the law of errors was deduced theoretically in 1783 by Laplace. He started with the assumption that the total error in an observation was the sum of a large number of independent deviations, which could be either positive or negative with equal probability, and could therefore be added according to the rule explained in the previous sections. Subsequently, Gauss gave a proof of the error law based on the postulate that the most probable value of any number of equally good

observations is their arithmetic mean. The distribution is thus sometimes called the Gaussian distribution, or the bell curve.

The Gaussian normal distribution is a smooth curve which is used to model the distribution of discrete points distributed around a mean. The probability density function $P(x)$ tells us with what probability we would expect measurements to be distributed about the mean value $\bar{x}$ (see Figure 11.14).

$$P(x_i) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp\left(-\frac{(x_i - \bar{x})^2}{2\sigma^2}\right).$$

It is based on the idealized limit of an infinite number of points. No experiments have an infinite number of points, though, so we need to fit a finite number of points to a normal distribution as best we can. It can be shown that the most probable choice is to take the mean of the finite set to be our estimate of the mean of the ideal set. Of course, if we select at random a sample of N values from the idealized infinite set, it is not clear that they will have the same mean as the full set of data. If the number in the sample N is large, the two will not differ by much, but if N is small, they might. In fact, it can be shown that if we take many random samples of the ideal set, each of size N, that they will have mean values which are themselves normally distributed, with a standard deviation equal to $\sigma/\sqrt{N}$. The quantity

$$\alpha = \frac{\sigma}{\sqrt{N}}$$

is therefore called the *standard error of the mean.* This is clearly a measure of the accuracy with which we can claim that our finite sample mean agrees with the actual mean. In quoting a measured value *which we believe has a unique or correct value,* it is therefore normal to write the mean value, plus or minus the standard error of the mean:

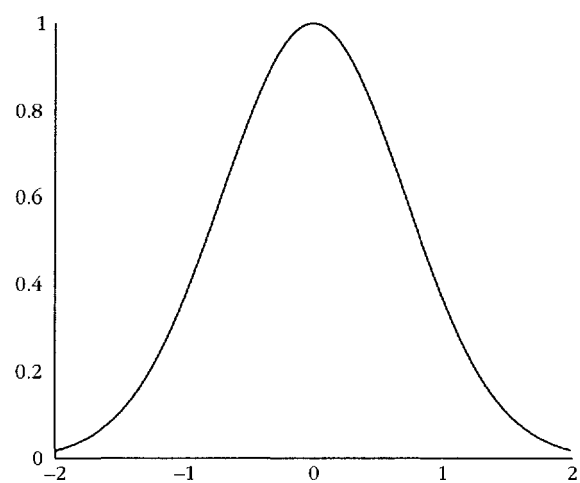


Figure 11.14 The Gaussian normal distribution, or bell curve, peaks at the arithmetic mean. Its width characterizes the standard deviation. It is therefore the generic model for all measurement distributions

$$\text{Result} = \bar{x} \pm \sigma/\sqrt{N} \text{ (for } N \text{ observations)},$$

where N is the number of measurements. Otherwise, if we believe that the measured value should have a distribution of values, one uses the standard deviation as a measure of the error. Many transactional operations in a computer system do not have a fixed value (see the next section).

The law of errors is not universally applicable, but it is still almost universally applied, for it serves as a convenient fiction which is mathematically simple[3].

11.8.5 The Planck Distribution

Another distribution which appears in the periodic rhythms of system behaviour is the Planck radiation distribution, so named for its origins in the physics of blackbody radiation and quantum theory. This distribution can be derived theoretically as the most likely distribution to arise from an assembly of fluctuations in equilibrium with an indefatigable reservoir or source [36]. The precise reason for its appearance in computer systems is subtle, but has to do with the periodicity imposed by users' behaviours, as well as the interpretation of transactions as fluctuations. The distribution has the form

$$D(\lambda) = \frac{\lambda^{-m}}{e^{1/\lambda T} - 1},$$

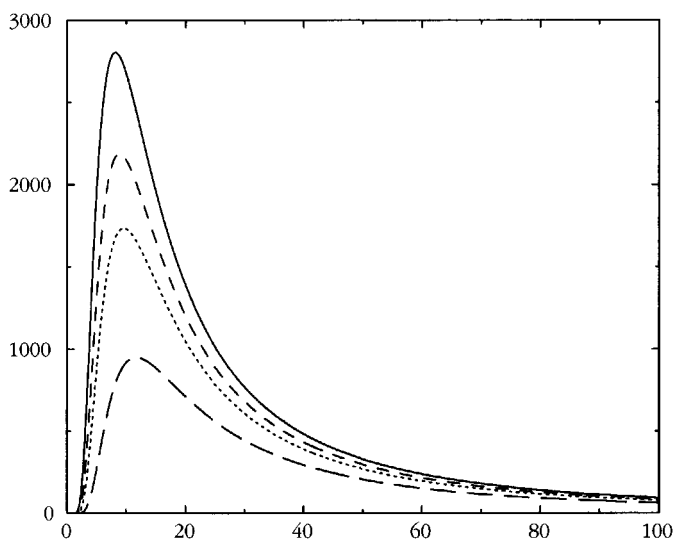


Figure 11.15 The Planck distribution for several temperatures. This distribution is the shape generated by random fluctuations from a source which is unchanged by the fluctuations. Here, a fluctuation is a computing transaction, a service request or new process

[3] The applicability of the normal distribution can, in principle, be tested with a χ^2 test, but this is seldom used in physical sciences, since the number of observations is usually so small as to make it meaningless.

where T is a scale, actually a temperature in the theory of blackbody radiation, and m is a number greater than 2. When $m = 3$, a single degree of freedom is represented. Burgess *et al.* [36] found that a single degree of freedom was sufficient to fit the data measured for a single variable, as one might expect. The shape of the graph is shown in Figure 11.15. Figures 11.16 and 11.17 show fits of real data to Planck distributions.

A remarkable number of transactions take this form. Indeed, it was shown [36] that many transactions on a computing system can be modelled as a linear superposition of a Gaussian distribution and a Planckian distribution, shifted from the origin:

$$D(\lambda) = A\,\mathrm{e}^{-\left(\frac{(\lambda-\bar{\lambda})^2}{4\sigma}\right)} + \frac{B}{(\lambda-\lambda_0)^3(\mathrm{e}^{1/(\lambda-\lambda_0)T}-1)}$$

This is a remarkable result, since it implies the possibility of using methods of statistical physics to analyse the behaviour of computer systems.

11.8.6 Other Distributions

Internet network traffic analysis studies [197, 279] show that the arrival times of data packets within a stream has a long tailed distribution, often modelled as a Pareto distribution (a power law)

$$f(\omega) = \beta\, a^{\beta}\, \omega^{-\beta-1}.$$

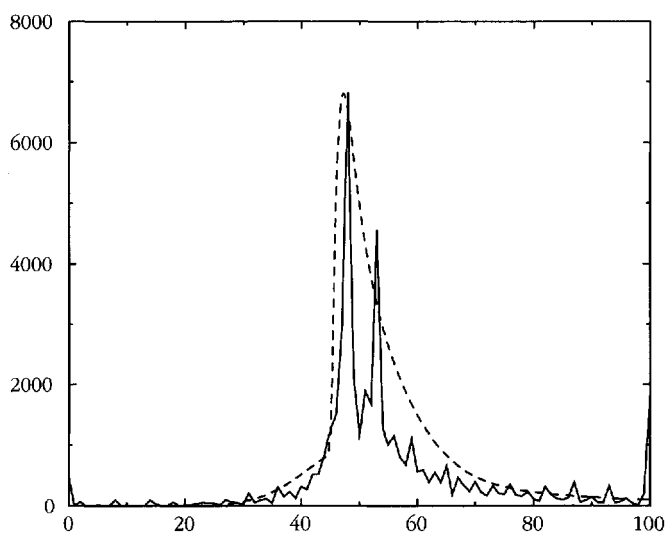


Figure 11.16 The distribution of system processes averaged over a few daily periods. The dotted line shows the theoretical Planck curve, while the solid line shows actual data. The jaggedness comes from the small amount of data (see next graph). The x-axis shows the deviation about the scaled mean value of 50 and the y-axis shows the number of points measured in class intervals of a half σ. The distribution of values about the mean is a mixture of Gaussian noise and a Planckian black-body distribution

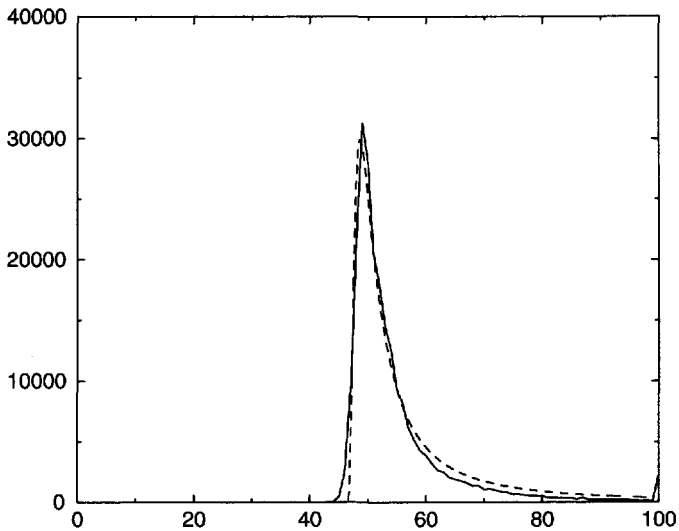


Figure 11.17 The distribution of WWW socket sessions averaged over many daily periods. The dotted line shows the theoretical Planck curve, while the solid line shows actual data. The smooth fit for large numbers of data can be contrasted with the previous graph. The x-axis shows the deviation about the scaled mean value of 50, and the y-axis shows the number of points measured in class intervals of a half σ. The distribution of values about the mean is a pure Planckian black-body distribution

This can be contrasted with the Poissonian arrival times of telephonic data traffic. It is an important consideration to designers of routers and switching hardware. It implies that a fundamental change in the nature of network traffic has taken place. A partial explanation for this behaviour is that packet arrival times consist not only of Poisson random processes for session arrivals, but also of internal correlations within a session. Thus, it is important to distinguish between measurements of packet traffic and measurements of numbers of sockets (or tcp sessions).

11.8.7 Fourier Analysis: Periodic Behaviour

As we have already commented, many aspects of computer system behaviour have a strong periodic quality, driven by the human perturbations introduced by users' daily rhythms. Other natural periods follow from the largest influences on the system from outside. This must be the case since there are no natural periodic sources internal to the system[4]. Apart from the largest sources of perturbation (i.e. the users themselves), there might be other lesser software systems which can generate periodic activity, e.g. hourly updates, or automated backups. The source might not even be known: for instance, a potential network intruder attempting a stealthy port scan might have programmed a script to test the ports periodically, over a length of time. Analysis of system behaviour can sometimes benefit from knowing

[4] Of course, there is the CPU clock cycle and the revolution of disks, but these occur on a time-scale which is smaller than the software operations and so cannot affect system behaviour.

these periods. For example, if one is trying to determine a causal relationship between one part of a system and another, it is sometimes possible to observe the signature of a process which is periodic, and thus obtain direct evidence for its effect on another part of the system.

Periods in data are the realm of Fourier analysis. What a Fourier analysis does is to assume that a data set is built up from the superposition of many periodic processes. This might sound like a strange assumption but, in fact, this is always possible. If we draw any curve, we can always represent it as a sum of sinusoidal-waves with different frequencies and amplitudes. This is the complex Fourier theorem:

$$f(t) = \int d\omega f(\omega) e^{-i\omega t},$$

where $f(\omega)$ is a series of coefficients. For strictly periodic functions, we can represent this as an infinite sum:

$$f(t) = \sum_{n=0}^{\infty} c_n e^{-2\pi i nt/T},$$

where T is some time scale over which the function $f(t)$ is measured. What we are interested in determining is the function $f(\omega)$, or equivalently the set of coefficients c_n which represent the function. These tell us how much of which frequencies are present in the signal $f(t)$, or its *spectrum*. It is a kind of data prism, or spectral analyser, like the graphical displays one finds on some music players. In other words, if we feed in a measured sequence of data and Fourier analyse it, the spectral function shows the frequency content of the data which we have measured.

We shall not go into the whys and wherefores of Fourier analysis, since there are standard programs and techniques for determining the series of coefficients. What is more important is to appreciate its utility. If we are looking for periodic behaviour in system characteristics, we can use Fourier analysis to find it. If we analyse a signal and find a spectrum such as that in Figure 11.18, then the peaks in the spectrum show the strong periodic content of the signal.

To discover these smaller signals, it will be necessary to remove the louder ones (it is difficult to hear a pin drop when a bomb explodes nearby). A word of warning is in order: as hi-fi buffs will know, in finite enclosure (set of data), it is possible to identify false harmonics and sub-harmonics which are fictitious effects of the size of the data sample. Thus, if we find periods which are related to integer multiples of the length of time over which the input signal was analysed, these should be treated with suspicion.

11.9 Strategic Analyses

The use of formal mathematics to analyse system administration has so far been absent from the discussion. There are two reasons why such analyses are of interest: (i) a formal description of a subject often reveals expectations and limitations which were invisible prior to the systematic model; and (ii) optimal solutions to problems can be explored, avoiding unnecessary prejudice.

It is my supposition that the languages of Game Theory and Dynamical Systems will enable us to formulate and model assertions about the behaviour of systems under certain

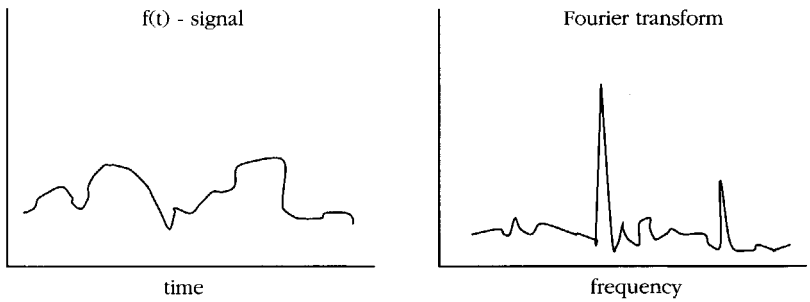


Figure 11.18 Fourier analysis is like a prism, showing us the separate frequencies of which a signal is composed. The sharp peaks in this figure illustrate how we can identify periodic behaviour which might otherwise be difficult to identify. The two peaks show that the input source conceals two periodic signals

adminstrative strategies. At some level, the development of a computer system is a problem in economics: it is a mixed game of opposition and cooperation between users and system. The aims of the game are several: to win resources, to produce work, to gain control of the system, and so on. A proper understanding of the issues should lead to better software and better strategies from human administrators. For instance, is greed a good strategy for a user? How could one optimally counter such a strategy? In some cases, it might even be possible to solve system administration games, determining the maximum possible 'win' available in the conflict between users and administrators.

At the present time, it is too early to discuss this mathematical approach in detail, but we mention it here in passing as a direction for future development.

11.10 Summary

Finding a rigorous experimental and theoretical basis for system administration is not an easy task. It involves many entwined issues, both technological and sociological. The sociological factors in system administration cannot be ignored, since the goals of system administration are, amongst other things, user satisfaction. In this respect, one is forced to pay attention to *heuristic* evidence, as rigorous statistical analysis of a specific effect is not always practical or adequately separable from whatever else is going on in the system. The study of computers is a study of *complexity*. Complexity has only been acknowledged as an object for study, in its own right, for about twenty years.

Exercises

Exercise 11.1 Consider the following data which represent a measurement of CPU usage for a process over time:

```
2.1
2.0
```

```
2.1
2.2
2.2
1.9
2.2
2.2
2.1
2.2
2.2
```

Now answer the following:

(a) To the eye, what appears to be the correct value for the measurement?
(b) Is there a correct value for the measurement?
(c) What is the mean value?
(d) What is the standard deviation?
(e) If you were to quote these data as one value, how would you quote the result of the measurement?

Exercise 11.2 What is meant by random errors? Explain why computers are not immune to random errors.

Exercise 11.3 Explain what is meant by Mean Time Before Failure. How is this quantity measured? Can sufficient measurements be made to make its value credible?

Exercise 11.4 If a piece of software has a MTBF of two hours and an average downtime of 15 seconds, does it matter that it is unstable?

Exercise 11.5 Explain why one would expect measurements of local SMTP traffic to show a strong daily rhythm, while measurements of incoming traffic would not necessarily have such a pronounced daily rhythm.

Exercise 11.6 Discuss whether one would expect to see a daily rhythm in WWW traffic. If such a rhythm were found, what would it tell us about the source of the traffic?

Exercise 11.7 Describe a procedure for determining causality in a computer network. Explain any assumptions and limitations which are relevant to this.

Exercise 11.8 Explain why problems, with quite different causes, lead often to the same symptoms.

Chapter 12

Summary and Outlook

The aim of this book has been to present an overview of the field of system administration for active system administrators, university courses and computer scientists everywhere. For a long time, system administration has been passed on by word of mouth and has resisted formalization. Only in recent times has the need for a formalization of the field been acknowledged, through courses and certifications, determined if not always ideal attempts to crystallize something definite from the fluid and fickle body of knowledge with which system administrators operate.

Compared to many other books on system administration, which are excellent how-to references, this book is quite theoretical. It might disappoint those who hold tradition as an authority. I have gone out of my way to be logical rather than conventional, to ignore redundant quirks where appropriate and to make suggested improvements (with accompanying justifications). History has seldom been the servant for logic, and I believe that it is time to abandon some old practices for the good of the field. That is not to say that I claim to have any ultimate answers, but the main message of this book is to make you, the reader, think and judge for yourself. There are, after all, no questions which should not be asked; there is no authority which should not be questioned.

System administration is about putting together a network of computers (workstations, PCs and supercomputers), getting them running and then *keeping* them running in spite of the activities of *users* who tend to cause the systems to fail. The failure of an operating system can be caused by one of several things. Most operating systems do not fail by themselves: it is users perturbing the system which causes problems to occur. Even in the cases where a problem can be attributed to a bug in a software component, it normally takes a user to provoke the bug. The fact that users play an important role in the behaviour of computer systems is far from doubt. At universities students rush to the terminal rooms to surf on the web during lunch breaks. This can result in the sudden caching of hundreds of megabytes of temporary files which can prevent legitimate work from being carried out. At offices, the workers probably run from their desks giving the opposite pattern of behaviour. The *time-scale* involved here is just a matter of minutes, perhaps an hour. In that short space of time, user behaviour (web surfing) can cause a general failure of the system for all users (disk full). System administration is therefore a mixture of technical expertise and sociology. Patterns of user behaviour need to be taken into account in any serious discussion of this problem. As a consequence, it is necessary to monitor the state of the system and its resources and react swiftly (on the time scale of human behaviour) to correct problems.

12.1 The Next Generation Internet Protocol (IPv6)

The backbone of the network renaissance is the Internet Protocol. As we have seen, the current implementation of the Internet Protocol has a number of problems. It is straightforward to calculate that, because of the structure of the IP addresses, divided into class A, B and C networks, something under 2% of the possible addresses can actually be used in practice. A recent survey from *Unix Review*, March 1998, shows that, of the total numbers of addresses, these are already allocated:

	Max possible	Percent allocated
Class A	127	100%
Class B	16382	62%
Class C	2097150	36%

Of course, this does not mean that all of the allocated addresses are in active use. After all, what organization has 65,535 hosts? In fact the survey showed that under 2% of these addresses were actually in use. This is an enormous wastage of IP addresses. Amongst the class C networks, where smaller companies would like address space, the available addresses are being used up quickly, but amongst the class A networks, the addresses will probably never be used. A new addressing structure is therefore required to solve this problem. Other problems with IPv4 are that it is too easy to take control of a connection by guessing sequence numbers. Moreover, there is no native support for encryption or mobile computing.

To address these issues the IETF (Internet Engineering Task Force) has put together a workgroup to design a new 128-bit protocol which will be called IPv6. Not only will the addressing structure be different, but there will be a considerable number of extra addresses. Even with a certain inefficiency of allocation, it is estimated that there will be enough IPv6 addresses to support a density of at least 10,000 IP addresses per square meter, which ought to be enough for every toaster and wristwatch on the planet and beyond.

12.2 Never-dos in System Administration

One of the main themes of this book has been to see security as an integral part of the system administration *state of mind*. There are literally hundreds of rules of thumb which we could conceivably write down, but let us summarize just a few which are useful to prevent embarrassing accidents:

- The administrator or root account has unlimited privileges. Never log into the system as the root user. Use the su command to gain root privileges when you need them and then quit at once. Never run complex programs with root privileges; it allows the system to be invaded by viruses and could lead to accidental damage to the system.
- Never leave root shells on the console. It is possible to accidentally do something destructive without realizing that one has root privileges: (Put the sledgehammer down when you are not using it, Eugene!)
- Never leave root shells or Administrator logins running so that others might gain access to them in an open room.

- Never leave services running if they are not used for anything. They provide a possible back-door into the system for intruders.
- Never give users physical access to a machine which stores important data. If users can touch the system, it's theirs.
- Operating systems like Windows 95, Windows 98, the MacOS and BeOS are all inherently insecure systems. They *cannot* be secure by virtue of their design (they have no access control of any kind – all access is fully privileged). When setting up a network in a potentially hostile environment, use an operating system (e.g. Unix or NT) which does provide access controls.
- Never give root a shell which is not located on the root disk partition. This will disable the machine. Only the root partition is mounted at boot time.
- Never replace files like `/etc/passwd` or `/etc/system` with links to files on other partitions. You will disable the machine. Only the root partition is mounted at boot time.
- Never make undocumented, non-reproducible changes to system files. They will be destroyed by system upgrades.
- Never make gratuitous changes to hosts by hand. Changes should be automated and all administrators and users should be aware of the changes.

If you need security, work defensively. If you think your machine is secure you are almost certain making a big mistake. Defensiveness, scepticism: treat any new program of subsystem you use with a healthy scepticism. How do you know you are not installing a Trojan horse or a security risk? Check that it is working. Does it really need all those privileges? Check that it is not doing things which you do not want it to. Drive as though the road were slippery! If you believe you are immortal you will take silly risks.

12.3 Information Management in the Future

The future is almost upon us, and no branch of technology has exploded with such a lack of planning and critical review as information technology. The state of our world knowledge is already well beyond our ability to cope with it. We currently have no way of searching and accessing *most* of the scientific and cultural resources which have been produced in the untold years of human endeavour of our history. In short, in our present state, most of our scientific knowledge has gone to waste. This is clearly an unacceptable situation, and it is probably one which will be solved by new information retrieval technology in the future, but the ability to retrieve information is critically dependent on its being organized into an easily parsable structure. This is the basis of programming algorithms in computer software, and the same thing applies to conglomerations of different software systems. The same principle applies to the storage of any kind of information. If information is not organized by a clear principle, it will get lost or muddled.

Structure and organization are the unsung heroes of science and of society. While scientists and computer hackers are frequently portrayed in the popular press as absent-minded muddlers, subject to fits of divine inspiration, the random element plays only a minor role in the true development of knowledge. Contrary to the popular affectation, it is not cool to have a relaxed attitude to organization. Claims to the effect that system administration is a

'dirty' business, not for academics, that we fly by the seats of our pants and so on, only serve to demean the system administration profession. If there is one service we can do for the future, it is to think critically and carefully about the information structures of our network communities.

12.4 Collaboration with Software Engineering

Every computer programmer should have to do service as a network administrator. If computer programs were written together with system administrators they would be efficient at resource usage, they would log useful information, they would be more reliable and more secure. In the future, every piece of software running on a computer system will need to take responsibility for system security and intrusion detection. There is no better way to build reliable and secure software, since every program knows its own internal state better than any external agent can. This is not how software is written today, and we suffer the consequences of this.

12.5 The Future of System Administration

We are approaching a new generation of operating systems, with the capacity for self-analysis and self-correction. It is no longer a question of whether they will arrive, but of when they will arrive. When it happens, the nature of system administration will change.

The day-to-day tasks of system administration change constantly and we pay these changes little attention. However, improvements in technology always lead to changing work practices, as humans are replaced by machinery in those jobs which are menial and repetitive. The core principles of system administration will remain the same, but the job description of the system manager will be rather different. In many ways, the day-to-day business of system administration consists of just a few recipes which slowly evolve over time. However, underneath the veneer of cookery, there is a depth of understanding about computer systems which has a more permanent value. Even when software systems take over many of the tasks which are now performed manually, there will be new challenges to meet.

For understandable reasons, the imaginations and attentions of our college generations have been captured, not by the intrigue of learning machines and intelligent systems, but by the glamour of multimedia. The computer has matured from a mere machine to a creative palette. It is difficult to articulate just why the administration of computer communities is an exciting challenge, but if we are to succeed in pushing through programmes of research which will bring about the level of automation we require, then it will be necessary to attract willing researchers. Fortunately, today there is a high proportion of system administrators with scientific backgrounds with the will and training to undertake such work. However, only the surface has been scratched. The tendency has been to produce tools rather than to investigate concepts, and while the tools are necessary, they must not become an end in themselves. A clearer understanding of the problems we face, looking forward, will only be achieved with more analytical work.

It is on this canvas that we attempt to congeal the discipline of system administration. We began this book by asking whether system administration was indeed a discipline. I hope

that it is now clear that it is – for a long time a diffuse one, but nevertheless real. In many ways system administration is like biology. Animals are machines, just billions of times more complex than our own creations, but the gap is closing and will continue to close as we enter into an era of quantum and biological computing techniques. The essence of experimental observation, and of the complex phenomena and inter-relationships between hosts, is directly analogous to what one does in biology. We may have created computers, but that does not mean that we understand them implicitly. In our field, we are still watching the animals do their thing, trying to learn.

Exercise

Exercise 12.1 Now that we are done, compare your impressions of system administration with those you had at the end of Chapter 1.

Appendix A

Summary

A.1 Summary of Principles

Principle 1 (Privilege) *Restriction of unnecessary privilege protects a system from accidental and malicious damage, infection by viruses and prevents users from concealing their actions with a false identities. It is desirable to restrict users' privileges for the greater good of everyone on the network.*

Corollary 2 (Privilege) *No-one should use a privileged root/Administrator account as a user account. To do so is to place the system in jeopardy.*

Principle 3 (Uniformity) *A uniform configuration minimizes the number of differences and exceptions one has to take into account later. This applies to hardware and software alike.*

Principle 4 (Communities) *What one member of a cooperative community does affects every other member, and vice versa. Each member of the community therefore has a responsibility to consider the well-being of the other members of the community.*

Principle 5 (Multiuser communities) *A multiuser computer system does not belong to any one user. All users must share the resources of the system. What one user does affects all other users, and vice versa. Each user has a responsibility to consider the effect of his/her actions of all the other users.*

Principle 6 (Network communities) *A computer which is plugged into the network is no longer just ours. It is part of a society of machines which shares resources and communicates with the whole. What that machine does affects other machines. What other machines do affects that machine.*

Principle 7 (Delegation I) *Leave experts to do their jobs. Assigning responsibility for a task to a body which specializes in that task is an efficient use of resources.*

Principle 8 (Adaptability) *Optimal structure and performance are usually found only with experience of changing local needs. The need for system revision will always come. Make network solutions which are adaptable.*

Principle 9 (One name for one object I) *Each unique resource should have a unique name which labels it and describes its function.*

Corollary 10 (Aliases) *Sometimes it is advantageous to use aliases or pointers to unique objects so that a generic name can point to a specific resource.*

Principle 11 (Inter-dependency) *Avoid making one service reliant on another. The more independent a service is, the more efficient it will be, and the fewer possibilities there will be for its failure.*

Principle 12 (Separation I) *Data which are separate from the operating system should be kept in a separate directory tree, preferably on a separate disk partition. If they are mixed with the operating system file-tree if makes re-installation or upgrade of the operating system unnecessarily difficult.*

Principle 13 (Separation II) *Data which are logically separate should be kept in separate directory trees, perhaps on separate disk partitions.*

Principle 14 (Separation III) *Independent systems should not interfere with one another, or be confused with one another. Keep them in separate storage areas.*

Principle 15 (Limited privilege) *No process or file should be given more privileges than it needs to do its job. To do so is a security hazard.*

Principle 16 (Temporary files) *Temporary files or sockets which are opened by any program should not be placed in any publicly writable directory like* `/tmp`. *This opens for the possibility of race conditions and symbolic link attacks. If possible, configure them to write to a private directory.*

Principle 17 (Flagging customization) *Customizations and deviations from standards should be made conspicuous to users and administrators. This makes the system easier to understand both for ourselves and our successors.*

Principle 18 (Distributed accounts) *Users move around from host to host, share data and collaborate. They need easy access to data and workstations all over an organization.*

Principle 19 (Environment) *It should always be clear to users which host they are using and what operating system they are working with. Default environments should be kept simple both in appearance (prompts, etc.) and in functionality (specially programmed keys, etc.). Simple environments are easy to understand.*

Principles 20 (Freedom) *Quotas, limits and restrictions tend to antagonize users. Users place a high value on personal freedom. Restrictions should be minimized. Workaround solutions which avoid rigid limits are preferable, if possible.*

Principle 21 (Mind control) *Computers have a perceived authority. We need to be on the look out for abuses of that authority, whether by accident or by design.*

Principle 22 (Homogeneity/Uniformity I) *System homogeneity or uniformity means that all hosts appear to be essentially the same. This makes hosts predictable for users and manageable for administrators. It allows for reuse of hardware in an emergency.*

Principle 23 (Scalability) *Any model of system infrastructure must be able to scale efficiently to large numbers of hosts (and perhaps subnets, depending on the local netmask).*

Principle 24 (Reliability) *Any model of system infrastructure must have reliability as one of its chief goals. Down-time can often be measured in real money.*

Corollary 25 (Redundancy) *Reliability is often safeguarded by redundancy, or backup services running in parallel, ready to take over at a moment's notice [244].*

Principle 26 (Homogeneity/Uniformity II) *A model in which all hosts are basically similar is (i) easier to understand conceptually both for users and administrators, (ii) cheaper to implement and maintain, and (iii) easier to repair and adapt in the event of failure.*

Corollary 27 (Reproducibility) *Avoid improvising system modifications, on the fly, which are not reproducible. It is easy to forget what was done, and this will make the functioning of the system difficult to understand and predict, for you and for others.*

Principle 28 (Abstraction generalizes *Expressing tasks in an operating-system independent language reduces time spent debugging, promotes homogeneity and avoids unnecessary repetition.*

Principle 29 (One name for one object II) *Each user should have the same unique name on every host. Multiple names lead to confusion and mistaken identity. A unique user name makes it clear which user is responsible for which actions.*

Principle 30 (Disorder) *All systems will eventually tend to a state of disorder unless a rigid and automated policy is maintained.*

Principle 31 (Equilibrium) *Deviation from a system's ideal state can be smoothed out by a counteractive response. If these two effects are in balance, the system will stay in equilibrium.*

Principle 32 (Policy) *A clear expression of goals and responses, prepares a site for future trouble and documents intent and procedure.*

Principle 33 (Simplest is best) *Simple rules make system behaviour easy to understand. Users tolerate rules if they understand them.*

Principle 34 (Resource restriction) *Restriction of resources can lead to poor performance and low productivity. Free access to resources prevents bottlenecks.*

Corollary 35 (Resource restriction) *With free access to resources, resource usage needs to be monitored to avoid the problem of runaway consumption, or the exploitation of those resources by malicious users.*

Principle 36 (Diagnostics) *When you hear the sound of distant hooves, think horses not zebras, i.e. always eliminate the obvious first.*

Principle 37 (Symptoms and cause) *Always try to fix problems at the root, rather than patching symptoms.*

Principle 38 (Weakest link) *The performance of any system is limited by the weakest link amongst its components. System optimization should begin with the source. If performance is weak at the source, nothing which follows can make it better.*

Corollary 39 (Performance) *A system is limited by its slowest moving parts. Resources with slowly moving parts, like disks, CD-ROMs and tapes, transfer data slowly and delay the system. Resources which work purely with electronics, like RAM memory and CPU calculation, are quick, because electrons are light and move around quickly. However, electronic motion/communication over long distances takes much longer than communication over short distances (internally within a host) because of impedances and switching.*

Principle 40 (Contention/competition) *When two processes compete for a resource, performance can be dramatically reduced as the processes fight over the right to use the resource. This is called contention. The benefits of sharing have to be weighed against the pitfalls.*

Principle 41 (Separate uids for services) *Each service which does not require privileged access to the system should be given a separate, non-privileged user-ID. This restricts service privileges, preventing any potential abuse should the service be hijacked by system attackers; it also makes clear which service is responsible for which processes in the process table.*

Corollary 42 (Privileged posts) *Services which run on ports 1–256 must started with Administrator privileges in order for the socket to be validated, but can switch internally to a safer level of privilege once communications have been established.*

Principle 43 (Security) *The fundamental requirement for security is the ability to restrict access and privilege to data.*

Principle 44 (Work defensively) *Expect the worst, do your best, preferably in advance of a problem.*

Principle 45 (Network security) *Extremely sensitive data should not be placed on a computer which is attached in any way to a public network.*

Principle 46 (Data invulnerability) *The purpose of a backup copy is to provide an image of data which is unlikely to be destroyed by the same act that destroys the original.*

Corollary 47 *Backup copies should be stored at a different physical location to the originals.*

Principle 48 (WWW corruption) *If a web server runs with the privileges of user* `www`*, then none of the data files should be owned by, or be writable by, the* `www` *user, otherwise it is trivial to alter the contents of the data with a CGI script.*

Principle 49 (Community borders) *Proxying is about protecting against breaches to the fundamental principle of communities. A firewall proxy provides us with a buffer against violations of our own community rights from outside, and also provides others with a buffer against what we choose to do in our own home.*

Principle 50 (Causality) *Every change or effect happens in response to a cause, which precedes it.*

Principle 51 (Uncertainty) *The act of measuring a given quantity in a system with finite resources always changes the conditions under which the measurement is made, i.e. the act of measurement changes the system.*

A.2 Summary of Suggestions

Suggestion 1 (Filer servers with common data) *Place all file servers which serve the same data on a common host, e.g. WWW, FTP and NFS serving user files. Place them on the host, which physically has the disks attached. This will save an unnecessary doubling of network traffic and will speed up services. A fast host with a lot of memory and perhaps several CPUs should be used for this.*

Suggestion 2 (GNU fileutils) *The GNU fileutils programs are superior in functionality than their corresponding vendor versions. Moreover, they work on every platform, bringing a pleasant dose of uniformity to a heterogeneous network. They can be placed in the users'* `PATH` *variable so as to override the vendor commands. In some instances, vendor programs have specially adapted features. One example is the* `ls` *command. Some Unix-like systems have ACLs (Access Control Lists) which give extended file permissions. These are invisible with the GNU version of* `ls`*, but are marked with an additional '*`+`*' to the left of the access bits, when using the vendor* `ls` *command. In the case of* `ls`*, it is probably worth removing or renaming the GNU* `ls` *to, say,* `gls`*.*

Suggestion 3 (Vigilance) *Be on the lookout for software which is configured, by default, to install itself on top of the operating system. Always check the destination using* `make -n install` *before actually committing to an installation. Programs which are replacements for standard operating system components often break the principle of separation*[1]*.*

[1] Software originating in BSD Unix is often an offender, since it is designed to be a part of BSD Unix, rather than an add-on, e.g. sendmail and BIND.

Suggestion 4 (Passwords) *Give users a common user name on all hosts, of no more than eight characters. Give them a common password on all hosts, unless there is a special reason not to do so. Some users never change their passwords unless forced to, and some users never even log in, so it is important to assign good passwords initially. Never assign a simple password and assume that it will be changed.*

Suggestion 5 (Clear prompts) *Try to give users a command prompt which includes the name of the host they are working on. This is important, since different hosts might have different operating systems, or different files. Including the current directory in the prompt, like DOS, is not always a good idea. It uses up half the width of the terminal and can seem confusing. If users want the name of the current directory in the prompt, let them choose that. Don't assign it as a default.*

Suggestion 6 (Unix shell defaults) *Avoid the host-wide files for shell setup in* `/etc`. *They are mixed up in the operating system distribution and changes here will be lost at upgrade time. Use an overridable include strategy in the user's own shell setup to read in global defaults. Do not link a file on a different file system to these in case this causes problems during system boot-up.*

Suggestion 7 (Problem users) *Keep a separate partition for problem users' home directories, so that they only cause trouble for one another, not for more considerate users.*

Suggestion 8 (Delegation II) *For large numbers of hosts, distributed over several locations, consider a policy of delegating responsibility to a local administrators with closer knowledge of the hosts' patterns of usage. Zones of responsibility allow local experts to do their jobs.*

Suggestion 9 (Platform independent languages) *Use languages and tools which are independent of operating system peculiarities, e.g.* `cfengine`,`perl`, `python`. *More importantly, use the right tool for the right job.*

Suggestion 10 (Cron management) *Maintaining cron files on every host individually is awkward. We can use cfengine as a front-end to cron, to give us a global view of the task list (see section 7.4.4).*

Suggestion 11 (FAQs) *Providing users with a road-map for solving problems, starting with Frequently Asked Questions and ending with an error report, can help to rationalize error reporting.*

Suggestion 12 (Unix printing) *Install LPRng on all of hosts in the network. Forget about trying to understand and manage the native printing systems on Sys V and BSD hosts. LPRng can replace them all with a system which is at least as good.*

Suggestion 13 (Static data) *When new data are acquired and do not change, they should be backed up to write only media at once. CD-ROM is an excellent medium for storing permanent data.*

Suggestion 14 (Tape backup) *Tapes are notoriously unreliable media, and tape streamers are mechanical nightmares, with complex moving parts which frequently go wrong. Verify the integrity of each substantial backup tape backup once you have made it. Never trust a tape. If the tape streamer gets serviced or repaired, check old tapes again afterwards. Head alignment changes can make old tapes unreadable.*

Suggestion 15 (OS configuration files) *Keep master versions of all configuration files like* `/etc/fstab`, `/etc/group` *or* `/etc/system` *in a directory under site-dependent files, and use a tool which synchronizes the contents of the master files with the operating system files (e.g. cfengine). This also allows the files to be distributed easily to other hosts which share a common configuration, and provide us with one place to make modifications, rather than having to hunt around the system for long-forgotten modifications. Site-dependent files should be on a partition which is backed up. Do not use symbolic links for synchronizing master files with the OS: only the root file system is mounted when the system boots, and cross-partition links will be invalid. You might render the system unbootable.*

Suggestion 16 (URL file system names) *Use a global URL naming scheme for all file systems and you will never loose a file on a tape, even if the label falls off (see section 3.9.2). Each file will be sufficiently labelled by its time-stamp and its name.*

Suggestion 17 (Passwords) *A useful hint in choosing a password is to incorporate the PIN code from a little-used credit card as a part of the password. This helps users to remember both – and it means that there will be secret numbers in the password.*

Appendix B

Some Useful Unix Commands

Typed commands are infinitely more flexible than graphical (GUI) based programs. You can tell the system what you want to do, rather than having to search through the menus to find out whether or not you are allowed to do what you want. As a system administrator you will find most GUI programs useless for any real tasks which involve looking after more than one host.

Always check the manual page on your local system before trying these commands. Versions, optional and even names differ, especially on older systems.

Who am I?

- `whoami`: prints your user name.
- `who am i`: prints your real and effective user id, and terminal.
- `id`: GNU program which prints all your user ids and groups.

Remote logins

The `telnet` command is the most reliable way of logging onto a remote Unix host. The `rlogin` or `rsh` commands can be used to this effect, but they will sometimes hang without reason, where `telnet` works without problem. The secure shell `ssh` is a secure replacement for the `rsh` command. It is recommended in its place. The `rlogin` command can be used to login without a password using the `.rhosts` authority file for trusted hosts and users. Using secure shell, one may use a public/private key pairs to obtain a much stronger authentication.

Monitoring disk usage

- `df`: display the usage of all mounted disk partitions if no argument is given. If a directory is named, the state of the disk partition on which the given directory resides is displayed. On SVR4 systems the output of this command is hard to understand unless the `-k` option is used.

- `du`: show disk usage on a per-file basis. The file sizes are either in kilobytes or in 512 byte blocks. The `-k` option forces output to be in kilobytes. The `-s` option prevents `du` from outputting information about every file and yields a summary of the named directory instead.
- `swap -s`: System 5 program to show swap space.
- `pstat`: BSD program to show swap space.

Disk backups

- `dump`: raw dump of a disk partition to a file or to tape.
- `rdump`: same as dump, but this can be done over the network, remotely without need for physical contact with the host.
- `ufsdump`: Solaris/SVR4 replaces `dump` with `this command.`
- `restore`: restores a disk partition from a file system dump.
- `cp -r`: copy a directory and all files recursively to a new location. This does not preserve symbolic links but makes multiple copies of the file instead. See `tar below.`
- `tar`: a simple way to copy an entire file system, preserving symbolic links is to do the following:

  ```
  cd source-dir; tar cf - . | (cd destination-dir; tar xf - )
  ```

  This pipes the output directly to the new directory using the streams interface for standard IO.

Mounting file systems

- `mount`: mount a local or remote disk.
- `umount`: unmount a local or remote disk. Note the peculiar spelling.
- `showmount`: show all hosts who are mounting file systems from this server.

Packing and unpacking archives

- `tar cf` *`tarfile`*`.tar` *`source-dir`*: packs all the files and sub-directories in the directory *`source-dir`* into a single 'tape-archive' file. If the `-f` argument is missing, `tar` expects to be able to write data to a default tape-streamer device and will complain with an error message.
- `tar zcf` *`tarfile`*`.tar.gz` *`source-dir`*: same as above, but piped through `gzip` to compress the `data. This only works with` GNU `tar.`
- `tar xf` *`tarfile.tar`*: unpacks the contents of a tar-file into the current directory.
- `tar zxf` *`tarfile`*`.tar.gz`: same as above, but pipes through `gzip` to uncompress data. This only works with GNU tar.

Shared libraries

- `ldd`: display the shared libraries used by a compiled executable file.

- `ldconfig`: some systems require this command to be run after installing or upgrading shared libraries. It updates symbolic links to the latest version of the library and in some cases generates a cache file of library names, especially GNU/Linux and SunOS prior to Solaris.

Handling binaries

- `strings`: this command lists all of the strings in a binary file. It is useful for finding out information which is compiled into software.
- `file`: prints the type of data a file contains.
- `strip`: remove debugging information from a compiled program. This can reduce the size of the program substantially.

Files and databases

- `locate`: GNU fast-find command, part of the GNU `find` package. Locates the names of files matching the argument string in part, by reading from a database. See `updatedb` below.
- `find`: locate by searching through every directory. Slow but powerful search facilities.
- `which`: locate an executable file by searching through directories in the `PATH` or `path` variable lists.
- `whatis`: gives a one-line summary of a command from the manual page (see `catman`).
- `catman -M`: this program builds the `apropos` or `man -k 'whatis' databases`.
- `updatedb`: this shell script updates the `locate` fast-find database.

Process management

- `ps aux`: show all processes on the system (BSD).
- `ps -ef`: show all processes on the system (SysV).
- `kill`: send a signal to the named process (pid), not necessarily to kill it. The process ID is the one listed by the `ps` command. Typical options are `-HUP` to send the hangup signal. This is used by many system daemons like `inetd` and `cron` as a signal which tells them to reread their configuration files. Another option is `-9` which is a non-ignorable kill instruction.
- `nice`: run a program with a non-default scheduling priority. This exists both as a shell command and as a C-shell built-in. The two versions use different syntax. Normal users can only reduce the priority of their processes (make them 'nicer'). Only the superuser can increase the priority of a process. The priority values differ between BSD and SysV systems. Under BSD, the nice values run from -20 (highest priority) to 19 (lowest priority) with 0 being the default. Under SysV, priorities run from 0 to 39, with 20 being the default. The C-shell built-in priorities are always from -20 to 20 for consistency.
- `renice` *`new-priority`* `-p` *`pid`*: resets the scheduling priority of a process to a new value. The priority values used by the system (not C shell) apply here.

- `crontab`: modern releases of Unix use the `crontab` command to schedule commands or scripts which are to be run at a specified time, or at regular intervals. The `crontab -l` command lists currently registered jobs. The `crontab -e` command is used to edit the crontab file. Each user has his or her own `crontab` file on every host. On older BSD systems, only root could alter the `crontab` file, which was typically a single file `/etc/crontab` or `/usr/lib/crontab` containing usernames and jobs to be performed.

Mail management

Sometimes mail gets stuck and cannot be delivered for some reason. This might be because the receiving mailhost is down, or because there is insufficient disk space to transfer the message, or many other reasons. In that case, incoming and outgoing mail gets placed in a queue which usually lies under the Unix directory `/var/spool/mail`, `/var/mail` or one of these with `/var` replaced by `/usr`.

- `mailq`: display any messages waiting in the mail queue. Same as `sendmail -bp`.
- `sendmail -q -v`: manually process the mail queue in verbose mode.

Disk management

- `format`: sun's interactive disk formatting and repair tool.
- `fsck`: the file system check program. A disk doctor. This checks the consistency of the file system (superblock consistency, etc.) and repairs simple problems.
- `newfs`: creates a new file system on a disk partition, erasing any previous data. This is analogous to formatting a diskette.
- `swapon`: this command causes the system to begin using a disk partition or swap file for system swapping/paging. `swapon -a` starts swapping on all devices registered in the file system table `/etc/fstab` or equivalent.
- `mkfile`: creates a special file for swapping inside a file system. The file has a fixed size, it cannot grow or shrink, or be edited directly. Normally swapping should be to a raw partition. Swapping to this kind of file is inefficient, but is used by (for instance) diskless clients.

Name service lookups

- `nslookup`: an interactive query program for reading domain data from the Domain Name Service (DNS/BIND).
- `dnsquery`: a non-interactive query program for reading domain data from the Domain Name Service (DNS/BIND).
- `whois`: displays information about who is responsible for a limited number of domains in the US. For example, the highly irritating domain moneyworld.com can be found with `whois moneyworld.com`.

System statistics

- `iostat`: displays I/O summary from the disks at an interval of *time-in-seconds.*
- `vmstat`: displays virtual-memory summary info at an interval of *time-in-seconds.*
- `netstat`: show all current network socket connections.
- `netstat -i`: show statistics from all network interfaces.
- `netstat -r`: show the static routing table.
- `nfsstat`: show NFS statistics. The `-c` option shows client-side data, while the `-s` option shows server-side data, where appropriate.

Networks

- `ping`: send a sonar 'ping' to see if a host is alive. The `-s` option sends multiple pings on some types of UNIX.
- `traceroute`: show the route, passing through all gateways to the named host. This command normally has to be made setuid-root in order to open the network kernel structures. Here is an example:

```
traceroute to wombat.gnu.ai.mit.edu (128.52.46.26), 30 hops max,
40 byte packets
 1 ca30-gw (128.39.89.1) 3 ms 1 ms 2 ms
 2 hioslo-gw.uninett.no (158.36.84.17) 5 ms 4 ms 5 ms
 3 oslo-gw2.uninett.no (158.36.84.1) 15 ms 15 ms 19 ms
 4 no-gw2.nordu.net (128.39.0.177) 43 ms 34 ms 32 ms
 5 nord-gw.nordu.net (192.36.148.57) 40 ms 31 ms 38 ms
 6 icm-gw.nordu.net (192.36.148.193) 37 ms 21 ms 29 ms
 7 icm-uk-1-H1/0-E3.icp.net (198.67.131.41) 58 ms 57 ms
 8 icm-pen-1-H2/0-T3.icp.net (198.67.131.25) 162 ms 136 ms
 9 icm-pen-10-P4/0-OC3C.icp.net (198.67.142.69) 198 ms 134
10 bbnplanet1.sprintnap.net (192.157.69.51) 146 ms 297 ms
11 * nyc2-br2.bbnplanet.net (4.0.1.25) 144 ms 120 ms
12 nyc1-br1.bbnplanet.net (4.0.1.153) 116 ms 116 ms 123
13 cambridge1-br1.bbnplanet.net (4.0.1.122) 131 ms 136 ms
14 cambridge1-br1.bbnplanet.net (4.0.1.122) 133 ms 124 ms
15 cambridge1-cr1.bbnplanet.net (206.34.78.23) 138 ms 129
16 cambridge2-cr2.bbnplanet.net (192.233.149.202) 128 ms
17 ihtfp.mit.edu (192.233.33.3) 129 ms 170 ms 143 ms
18 B24-RTR-FDDI.MIT.EDU (18.168.0.6) 129 ms 147 ms 148
19 radole.lcs.mit.edu (18.10.0.1) 149 ms * 130 ms
20 net-chex.ai.mit.edu (18.10.0.2) 134 ms 129 ms 134 ms
21 * * *
22 * * * < -- routing problem here
```

- `etherfind`: dump Ethernet packet activity to console, showing traffic, etc., SunOS.
- `snoop`: newer version of etherfind in Solaris.
- `ifconfig`: configure or summarize the setup of the a network interface, e.g. `ifconfig -a` shows all interfaces. Used to set the broadcast address, netmask and Internet address of the host.
- `route`: make an entry in the static routing table. Hosts which do not act as routers need only a default route, e.g.

```
route add default xxx.xxx.xxx.1 1
```

or in GNU/Linux

```
route add default gw xxx.xxx.xxx.1
```

Appendix C

Programming and Compiling

C.1 Make

Make is a *declarative* language which was designed for building software. In fact, its usefulness far outshines this meager goal. Make is, in reality, a generalized hierarchical organizer for instructions which generate file objects.

Nowadays compilers are often sold with fancy user environments driven by menus which make it easier to compile programs. Make was originally written so that Unix programmers could write huge source trees of code, occupying many directories and subdirectories and assemble them efficiently and effortlessly.

Building programs

Typing lines like

```
cc -c file1.c file2.c ...
cc -o target file1.o ....
```

repeatedly to compile a complicated program can be a real nuisance. One possibility would therefore be to keep all the commands in a script. This could waste a lot of time, though. Suppose you are working on a big project which consists of many lines of source code—but are editing only one file. You really only want to recompile the file you are working on and then relink the resulting object file with all of the other object files. Recompiling the other files which hadn't changed would be a waste of time. But that would mean that you would have to change the script each time you change what you need to compile.

A better solution is to use the `make` command. `make` was designed for precisely this purpose. To use `make`, we create a file called `Makefile` in the same directory as our program. `make` is a quite general program for building software. It is not specifically tied to the C programming language – it can be used in any programming language.

A `make` configuration file, called a `Makefile`, contains rules which describe how to compile or build all of the pieces of a program. For example, even without telling it specifically, `make` knows that to go from `prog.c` to `prog.o` the command `cc -c`

`prog.c` must be executed. A Makefile works by making such associations. The Makefile contains a list of all of the files which compose the program and rules as to how to get to the finished product from the source.

The idea is that, to compile a program, we just have to type `make`. The program `make` then reads a configuration file called a `Makefile` and compiles only the parts which need compiling. It does not recompile files which have not changed since the last compilation! How does it do this? `make` works by comparing the time-stamp on the file it needs to create with the time-stamp on the file which is to be compiled. If the compiled version exists and is newer than its source, then the source does not need to be recompiled.

To make this idea work in practice, `make` has to know how to go through the steps of compiling a program. Some default rules are defined in a global configuration file, e.g.

```
/usr/include/make/default.mk
```

Let's consider an example of what happens for the three files `a.c`, `b.c` and `c.c` in the example above—and let's not worry about what the Makefile looks like yet.

The first time we compile, only the '.c' files exist. When we type `make`, the program looks at its rules and finds that it has to make a file called 'myprog'. To make this it needs to execute the command

```
gcc -o myprog a.o b.o c.o
```

So it looks for `'a.o'`, etc., and doesn't find them. It now goes to a kind of subroutine and looks to see if it has any rules for making files called '.o', and it discovers that these are made by compiling with the `gcc -c` option. Since the files do not exist, it does this. Now the files `'a.o b.o c.o'` exist, and it jumps back to the original problem of trying to make 'myprog'. All the files it needs now exist, and so it executes the command and builds 'myprog'.

If we now edit `'a.c'`, and type `make` once again—it goes through the same procedure as before but now it finds all of the files. So it compares the dates on the files—if the source is newer than the result, it recompiles.

By using this recursive method, `make` only compiles those parts of a program which need compiling.

Makefiles

To write a Makefile, we have to tell `make` about *dependencies*. The dependencies of a file are all of those files which are required to build it. In a strong sense, dependencies are like subroutines which are carried out by `make` in the course of building the final target. The dependencies of `myprog` are `a.o`, `b.o` and c.o. The dependencies of `a.o` are simply `a.c`, the dependencies of `b.o` are `b.c`, and so on.

A Makefile consists of rules of the form:

```
target : dependencies
[TAB]                    rule;
```

The target is the thing we eventually want to build, the dependencies are like subroutines to be executed first if they do not exist. Finally, the rule is some code which is to be executed if all if the dependencies exist; it takes the dependencies and turns them into the current target.

Notice how dependencies are like subroutines, so each sub-rule makes a sub-target. In the end, the aim is to combine all of the sub-targets into one final target. There are two important things to remember:

- The file names must start on the first character of a line.
- There must be a TAB character at the beginning of every rule or action. If there are spaces instead of tabs, or no tab at all, `make` will signal an error. This bizarre feature can cause a lot of confusion.

Let's look at an example Makefile for a program which consists of two course files `main.c` and `other.c`, and which makes use of a library called `libdb` which lies in the directory `/usr/local/lib`. Our aim is to build a program called `database`:

```
#
# Simple Makefile for 'database'
#

# First define a macro

OBJ = main.o other.o

CC = gcc
CFLAGS = -I/usr/local/include
LDFLAGS = -L/usr/local/lib -ldb
INSTALLDIR = /usr/local/bin

#
# Rules start here. Note that the $@ variable becomes the name of
# the executable file. In this case it is taken from the $OBJ
# variable
#

database: $OBJ
          $CC -o $@ $OBJ $LDFLAGS

#
# If a header file changes, normally we need to recompile
# everything. There is no way that make can know this unless we
# write a rule which forces it to rebuild all .o files if the
# header file changes...
#

$OBJ: $HEADERS

#
# As well as special rules for special files we can also define a
# "suffix rule". This is a rule which tells us how to build all
# files of a certain type. Here is a rule to get .o files from .c
# files. The $< variable is like $? but is only used in suffix
# rules.
#

.c.o:
          $CC -c $CFLAGS $<
```

```
############################################################
# Clean up
############################################################
#
# Make can also perform ordinary shell command jobs
# "make tidy" here performs a cleanup operation
#
clean:
        rm -f $OBJ
        rm -f y.tab.c lex.yy.c y.tab.h
        rm -f y.tab lex.yy
        rm -f *% *~ *.o
        make tidy

install: $INSTALLDIR/database
        cp database $INSTALLDIR/database
```

The Makefile above can be invoked in several ways:

```
make
make database
make clean
make install
```

If we simple type `make` (i.e. the first of these choices), `make` takes the first of the rules it finds as the object to build. In this case the rule is 'database', so the first two forms above are equivalent. On the other hand, if we type

```
make clean
```

then execution starts at the rule for 'clean', which is normally used to remove all files except the original source code. Make 'install' causes the compiled program to be installed at its intended destination.

`make` uses some special variables (which resemble the special variables used in Perl—but don't confuse them). The most useful one is `$@` which represents the current *target*—or the object which `make` would like to compile, i.e. as `make` checks each file it would like to compile, `$@` is set to the current file name:

- `$@` This evaluates to the current target, i.e. the name of the object you are currently trying to build. It is normal to use this as the final name of the program when compiling.
- `$?` This is used only outside of suffix rules, and means the name of all the files which must be compiled in order to build the current target:

```
target: file1.o file2.o
TAB cc -o $@ $?
```

- `$<` This is only used in suffix rules. It has the same meaning as `$?` but only in suffix rules. It stands for the prerequisite, or the file which must be compiled in order to make a given object.

Note that, because `make` has some default rules defined in its configuration file, a single-file C program can be compiled very easily by typing

```
make filename.c
```

This is equivalent to

```
cc -c filename.c
cc -o filename filename.o
```

C.2 Perl

To summarize Perl, we need to know about the structure of a Perl program, the conditional constructs it has, its loops and its variables. In the latest versions of Perl (Perl 5), you can write object-oriented programs of great complexity. We shall not go into this depth, for the simple reason that Perl's strength is not as a general programming language but as a specialized language for text file handling. The syntax of Perl is in many ways like the C programming language, but there are important differences:

- Variables do not have *types*. They are interpreted in a context sensitive way. The operators which acts upon variables determine whether a variable is to be considered a string or as an integer, etc.
- Although there are no types, Perl defines *arrays* of different kinds. There are three different kinds of array, labelled by the symbols `$, @ and %.`
- Perl keeps a number of standard variables with special names, e.g. `$_ @ARGV` and `%ENV`. Special attention should be paid to these. They are very important!
- The shell reverse apostrophe notation `'command'` can be used to execute Unix programs and get the result into a Perl variable.

Here is a simple 'structured hello world' program in Perl. Notice that subroutines are called using the & symbol. There is no special way of marking the main program—it is simply that part of the program which starts at line 1.

```
#!/local/bin/perl
#
# Comments
#

&Hello();
&World;

# end of main

sub Hello
     {
     print "Hello ";
     }

sub World
     {
     print "World\n";
     }
```

The parentheses on subroutines are optional, if there are no parameters passed. Notice that each line must end in a semi-colon.

Scalar variables

In Perl, variables do not have to be declared before they are used. Whenever you use a new symbol, Perl automatically adds the symbol to its symbol table and initializes the variable to the empty string.

It is important to understand that there is no practical difference between zero and the empty string in Perl—except in the way that you, the user, choose to use it. Perl makes no distinction between strings and integers or any other types of data—except when it wants to interpret them. For instance, to compare two variables as strings is not the same as comparing them as integers, even if the string contains a textual representation of an integer. Take a look at the following program:

```
#!/local/bin/perl
#
# Nothing!
#

print "Nothing == $nothing\n";

print "Nothing is zero!\n" if ($nothing == 0);

if ($nothing eq "")
   {
   print STDERR "Nothing is really nothing!\n";
   }

$nothing = 0;

print "Nothing is now $nothing\n";
```

The output from this program is:

```
Nothing ==
Nothing is zero!
Nothing is really nothing!
Nothing is now 0
```

There are several important things to note here. First, we never declare the variable 'nothing'. When we try to write its value, Perl creates the name and associates a NULL value to it, i.e. the empty string. There is no error. Perl knows it is a variable because of the $ symbol in front of it. All *scalar* variables are identified by using the dollar symbol.

Next, we compare the value of `$nothing` to the integer '0' using the integer comparison symbol ==, and then we compare it to the empty string using the string comparison symbol `eq`. Both tests are true! That means that the empty string is interpreted as having a numerical value of zero. In fact *any string* which does not form a valid integer number has a numerical value of zero.

Finally, we can set `$nothing` explicitly to a valid integer string zero, which would now pass the first test, but fail the second.

As extra spice, this program also demonstrates two different ways of writing the `if` command in Perl.

The default scalar variable.

The special variable `$_` is used for many purposes in Perl. It is used as a buffer to contain the result of the last operation, the last line read in from a file, etc. It is so general that many functions which act on scalar variables work by default on `$_` if no other argument is specified. For example,

```
print;
```

is the same as

```
print $_;
```

Array (vector) variables

The complement of scalar variables is arrays. An array in Perl is identified by the `@symbol` and, like scalar variables, is allocated and initialized dynamically:

```
@array[0] = "This little piggy went to market";
@array[2] = "This little piggy stayed at home";

print "@array[0] @array[1] @array[2]";
```

The index of an array is always understood to be a number, not a string, so if you use a non-numerical string to refer to an array element, you will always get the zeroth element, since a non-numerical string has an integer value of zero.

An important array which every program defines is

```
@ARGV
```

This is the argument vector array, and contains the commands line arguments by analogy with the C-shell variable `$argv[]`.

Given an array, we can find the last element by using the `$#` operator. For example,

```
$last_element = $ARGV[$#ARGV];
```

Notice that each element in an array is a scalar variable. The `$` cannot be interpreted directly as the number of elements in the array, as it can in the C-shell. You should experiment with the value of this quantity – it often necessary to add 1 or 2 to its value in order to get the behaviour one is used to in the C-shell.

Perl does not support multiple-dimension arrays directly, but it is possible to simulate them yourself. (See the Perl book.)

Special array commands

The `shift` command acts on arrays and returns and removes the first element of the array. Afterwards, all of the elements are shifted down one place. So one way to read the elements of an array in order is to repeatedly call `shift`:

```
$next_element=shift(@myarray);
```

Note that, if the array argument is omitted, then `shift` works on `@ARGV` by default.

Another useful function is `split`, which takes a string and turns it into an array of strings. `split` works by choosing a character (usually a space) to delimit the array elements, so a string containing a sentence separated by spaces would be turned into an array of words. The syntax is

```
@array = split;                          # works with spaces on $_
@array = split(pattern,string);          # Breaks on pattern
($v1,$v2...) = split(pattern,string);# Name array elements
```

In the first of these cases, it is assumed that the variable `$_` is to be split on whitespace characters. In the second case, we decide on what character the split is to take place and on what string the function is to act. For instance

```
@new_array = split(":","name:pwd:uid:gid:gcos:home:shell");
```

The result is a seven element array called `@new_array`, where `$new_array[0]` is `name`, etc.

In the final example, the left-hand side shows that we wish to capture elements of the array in a named set of scalar variables. If the number of variables on the left-hand side is fewer than the number of strings which are generated on the right-hand side, they are discarded. If the number on the left-hand side is greater, then the remainder variables are empty.

Associated arrays

One of the very nice features of Perl is the ability to use one string as an index to another string in an array. For example, we can make a short encyclopaedia of zoo animals by constructing an associative array in which the keys (or indices) of the array are the names of animals, and the contents of the array are the information about them.

```
$animals{"Penguin"} = "Suspicious animal, good with cheese
  crackers...\n";
$animals{"dog"} = "Plays stupid, but could be a cover...\n";

if (index eq "fish")
   {
   $animals{$index = "Often comes in square boxes. Very
     cold.\n";
   }
```

An entire associated array is written `%array`, while the elements are `$array{$key}`.

Perl provides a special associative array for every program called `%ENV`. This contains the *environment variables* defined in the parent shell which is running the Perl program. For example

```
print "Username = $ENV{"USER"}\n";

$ld = "LD_LIBRARY_PATH";
print "The link editor path is $ENV{$ld}\n";
```

To get the current path into an ordinary array, one could write,

```
@path_array= split(":",$ENV{"PATH"});
```

Array example program

Here is an example which prints out a list of files in a specified directory, in order of their Unix protection bits. The *least* protected file files come first.

```
#!/local/bin/perl
#
# Demonstration of arrays and associated arrays.
# Print out a list of files, sorted by protection,
# so that the least secure files come first.
#
# e.g.   arrays <list of words>
#        arrays *.C
#
##############################################################

print "You typed in ",$#ARGV+1," arguments to command\n";

if ($#ARGV < 1)
   {
   print "That's not enough to do anything with!\n ";
   }

while ($next_arg = shift(@ARGV))
   {
   if ( ! ( -f $next_arg || -d $next_arg))
      {
      print "No such file: $next_arg\n ";
      next;
      }

   ($dev,$ino,$mode,$nlink,$uid,$gid,$rdev,$size) =
    stat $next_arg);
   $octalmode = sprintf("%o",$mode & 0777);

   $assoc_array{$octalmode} .= $next_arg.
        " : size (".$size."), mode (".$octalmode.")\n";
   }

print "In order: LEAST secure first!\n\n";

foreach $i (reverse sort keys (%assoc_array))
   {
   print %$assoc_array{$i};
   }
```

Loops and conditionals

Here are some of the most commonly used decision-making constructions and loops in Perl. The following is not a comprehensive list—for that, you will have to look in the Perl bible: *Programming Perl*, by Larry Wall and Randal Schwartz. The basic pattern follows the C programming language quite closely. In the case of the `for` loop, Perl has both the C-like version, called `for` and a `foreach` command which is like the C-shell implementation.

```
if( expression )
  {
  block;
  }
else
  {
  block;
  }

command if (expression);

unless (expression)
  {
  block;
  }
else
  {
  block;
  }

while (expression)
  {
  block;
  }

do
  {
  block;
  }
while (expression);

for (initializer; expression; statement)
  {
  block;
  }

foreach variable(array)
  {
  block;
  }
```

In all cases, the `else` clauses may be omitted. Strangely, Perl does not have a `switch` statement, but the Perl book describes how to make one using the features provided.

The for loop

The for loop is exactly like that in C or C++ and is used to iterate over a numerical index, like this:

```
for ($i = 0; $i < 10; $i++)
  {
  print $i, "\n ";
  }
```

The foreach loop

The foreach loop is like its counterpart in the C shell. It is used for reading elements one by one from a regular array. For example,

```
for each $i (@array)
  {
  print $i, "\n ";
  }
```

Iterating over elements in arrays

One of the main uses for `for` type loops is to iterate over successive values in an array. This can be done in two ways, which show the essential difference between `for` and `foreach`.

If we want to fetch each value in an array in turn, without caring about numerical indices, it is simplest to use the `foreach` loop.

```
@array = split(" ","a b c d e f g");

foreach $var ( @array )
  {
  print $var, "\n ";
  }
```

This example prints each letter on a separate line. If, on the other hand, we are interested in the index, for the purposes of some calculation, then the `for` loop is preferable.

```
@array = split(" ","a b c d e f g");

for ($i = 0; $i <= $#array; $i++)
  {
  print $array[$i], "\n ";
  }
```

Notice that, unlike the for-loop idiom in C/C++, the limit is `$ i <= $# array`, i.e. 'less than or equal to' rather than 'less than'. This is because the `$ \#` operator does not return the number of elements in the array, but rather the last element.

Associated arrays are slightly different, since they do not use numerical keys. Instead they use a set of strings, like in a database, so that you can use one string to look up another. To iterate over the values in the array we need to get a list of these strings. The `keys` command is used for this:

```
$assoc{"mark"} = "cool";
$assoc{"GNU"} = "brave";
$assoc{"zebra"} = "stripy";

foreach $var ( keys %assoc )
  {
  print "$var , $assoc {$var} \n";
  }
```

The order of the keys is not defined in the above example, but you can choose to sort them alphabetically by writing

```
foreach $var ( sort keys %assoc )
```

instead.

Iterating over lines in a file

Since Perl is about file handling we are very interested in reading files. Unlike C and C++, perl likes to read files line by line. The angle brackets are used for this (see section C.2). Assuming that we have some file handle `<file>`, for instance `<STDIN>`, we can always read the file line by line with a while-loop like this:

```
while ($line = <file>)
  {
  print $line;
  }
```

Note that `$line` includes the end of line character on the end of each line. If you want to remove it, you should add a `chomp` command:

```
while ($line = <file>)
  {
  chomp $line;
  print "line = ($line)\n";
  }
```

Files in Perl

Opening files is straightforward in Perl. Files must be opened and closed using—wait for it—the commands `open` and `close`. You should be careful to close files after you have finished with them—especially if you are writing *to* a file. Files are buffered, and often large parts of a file are not actually written until the `close` command is received.

Three files are, of course, always open for every program, namely `STDIN`, `STDOUT` and `STDERR`.

Formally, to open a file, we must obtain a file descriptor or file handle. This is done using `open`:

```
open (file_descrip,"Filename");
```

The angular brackets `<..>` are used to read from the file. For example,

```
$line = <file_descrip>;
```

reads one line from the file associated with `file_descrip`.

Let's look at some examples of filing opening. Here is how we can implement Unix's `cut` and `paste` commands in Perl:

```
#!/local/bin/perl
#
# Cut in perl
#

 # Cut second column
```

```
while (<>)
  {
  @cut_array = split;
  print "$cut_array[1]\n";
  }
```

This is the simplest way to open a file. The empty file descriptor <> tells Perl to take the argument of the command as a file name and open that file for reading. This is really short for `while($_=<STDIN>)` with the standard input redirected to the named file.

The `paste program` can be written as follows:

```
#!/local/bin/perl
#
# Paste in perl
#
# Two files only, syntax : paste file 1file2
#

open (file1,"@ARGV[0]") || die "Can't open @ARGV[0]\n";
open (file2,"@ARGV[1]") || die "Can't open @ARGV[1]\n";

while (($line1 = <file1>) || ($line2 = <file2>))
  {
  chop $line1;
  chop $line2;
  print "$line1  $line2\n"; # tab character between
  }
```

Here we see more formally how to read from two separate files at the same time. Notice that by putting the read commands into the test-expression for the `while loop`, we are using the fact that `<..>` returns a non-zero (true) value unless we have reached the end of the file.

To write and append to files, we use the shell redirection symbols inside the `open` command:

```
open(fd,"> filename"); # open file for writing
open(fd,">> filename"); # open file for appending
```

We can also open a pipe from an arbitrary Unix command and receive the output of that command as our input:

```
open (fd,"/bin/ps aux | ");
```

A simple Perl program

Let us now write the simplest Perl program which illustrates the way in which Perl can save time. We shall write it in three different ways to show what the short cuts mean. Let us implement the `cat` command, which copies files to the standard output. The simplest way to write this is Perl is the following:

```
#!/local/bin/perl
```

```
while (<>)
  {
  print;
  }
```

Here we have made heavy use of the many default assumptions which Perl makes. The program is simple, but difficult to understand for novices. First, we use the default file handle <> which means, take one line of input from a default file. This object returns true as long as it has not reached the end of the file, so this loop continues to read lines until it reaches the end of file. The default file is standard input, unless this script is invoked with a command line argument, in which case the argument is treated as a file name and Perl attempts to open the argument-file name for reading. The `print` statement has no argument telling it what to print, but Perl takes this to mean: print the default variable `$_`.

We can therefore write this more explicitly as follows:

```
#!/local/bin/perl
open (HANDLE,"$ARGV[1]");
while (<HANDLE>)
  {
  print $_;
  }
```

Here we have simply filled in the assumptions explicitly. The command <HANDLE> now reads a single line from the named file-handle into the default variable `$_`. To make this program more general, we can elimiate the defaults entirely:

```
#!/local/bin/perl
open (HANDLE,"$ARGV[1]");
while ($line=<HANDLE>)
  {
  print $line;
  }
```

== **and** eq

Be careful to distinguish between the comparison operator for integers == and the corresponding operator for strings `eq`. These do not work in each other's places so if you get the wrong comparison operator your program might not work and it is quite difficult to find the error.

chop **and** chomp

The command `chop` cuts off the last character of a string. This is useful for removing newline characters when reading files, etc. The syntax is

```
chop;          # chop $_;
chop scalar; # remove last character in $scalar
```

A slightly more refined version which only chops off whitespace and end of line characters is the chomp function.

Perl subroutines

Subroutines are indicated, as in the example above, by the ampersand & symbol. When parameters are passed to a Perl subroutine, they are handed over as an array called @_, which is analogous to the $_ variable. Here is a simple example:

```
#!/local/bin/perl

$a="silver";
$b="gold";

&PrintArgs($a,$b);

# end of main

sub PrintArgs

  {
  ($local_a,$local_b) = @_;

  print "$local_a, $local_b\n";
  }
```

die - exit on error

When a program has to quit and give a message, the die command is normally used. If called without an argument, Perl generates its own message including a line number at which the error occurred. To include your own message, you write

```
die "My message....";
```

If the string is terminated with a \n newline character, the line number of the error is not printed, otherwise Perl appends the line number to your string.

When opening files, it is common to see the syntax:

```
open (filehandle,"Filename") || die "Can't open...";
```

The logical OR symbol is used, because open returns true if all goes well, in which case the right-hand side is never evaluated. If open is false, then die is executed. You can decide for yourself whether or not you think this is good programming style—we mention it here because it is common practice.

The stat() idiom

The Unix library function stat() is used to find out information about a given file. This function is available both in C and in Perl. In Perl, it returns an array of values. Usually, we are interested in knowing the access permissions of a file. stat() is called using the syntax

```
@array = stat ("file name");
```

or alternatively, using a named array

```
(device,$inode,$mode) = stat("file name");
```

The value returned in the *mode* variable is a bit-pattern. The most useful way of treating these bit patterns is to use octal numbers to interpret their meaning.

To find out whether a file is readable or writable to a group of users, we use a programming idiom which is very common for dealing with bit patterns: first we define a mask which zeroes out all of the bits in the mode string except those which we are specifically interested in. This is done by defining a mask value in which the bits we want are set to 1 and all others are set to zero. Then we AND the mask with the mode string. If the result is different from zero then we know that all of the bits were also set in the mode string. As in C, the bitwise AND operator in Perl is called &.

For example, to test whether a file is writable to other users in the same group as the file, we would write the following:

```
$mask = 020;    # Leading 0 means octal number
($device,$inode,$mode) = stat("file");
if ($mode $mask)
  {
  print "File is writable by the group";
  }
```

Here the 2 in the second octal number means 'write', the fact that it is the second octal number from the right means that it refers to 'group'. Thus, the result of the if-test is only true if that particular bit is true. We shall see this idiom in action below.

Perl Example Programs

The passwd program and `crypt()` function

Here is a simple implementation of the Unix `passwd` program in Perl.

```
#!/local/bin/perl
#
# A perl version of the passwd program.
#
# Note - the real passwd program needs to be much more
# secure than this one. This is just to demonstrate the
# use of the crypt() function.
#
#############################################################
print "Changing passwd for $ENV{'USER'} on $ENV{'HOST'}\n";
system 'stty','-echo';
print "Old passwd: ";
$oldpwd = <STDIN>;
chop $oldpwd;
($name,$coded_pwd,$uid,$gid,$x,$y,$z,$gcos,$home,$shell)
                                    = getpwnam($ENV"USER"});
```

```
if (crypt($oldpwd,$coded_pwd) ne $coded_pwd)
   {
   print "\nPasswd incorrect\n";
   exit (1)
   }

$oldpwd = "";                          # Destroy the evidence!

print "\nNew passwd: ";

$newpwd = <STDIN>;

print "\nRepeat new passwd: ";

$rnewpwd = <STDIN>;

chop $newpwd;
chop $rnewpwd;

if (newpwd ne $rnewpwd)
   {
   print "\n Incorrectly typed. Password unchanged.\n";
   exit (1);
   }

\index {{\code rand()}}
$salt = rand();
$new_coded_pwd = crypt($newpwd,$salt);

print "\n\n$name:$new_coded_pwd:$uid:$gid:$gcos:$home:
  $shell\n ";
```

Example with `fork()`

The following example uses the `fork` function to start a daemon which goes into the background and watches the system to which process is using the greatest amount of CPU time each minute. A pipe is opened from the BSD `ps` command.

```
#!/local/bin/perl
#
# A fork() demo. This program will sit in the background and
# make a list of the process which uses the maximum CPU average
# at 1 minute intervals. On a quiet BSD like system this will
# normally be the swapper (long term scheduler).
#

$true = 1;
$logfile="perl.cpu.logfile";

print "Max CPU logfile, forking daemon...\n ";

if (fork())
   {
   exit(0);
   }
```

```
while ($true)
   {
   open (logfile,">> $logfile") || die "Can't open logfile\n";
   open (ps,"/bin/ps aux |") || die "Couldn't open a pipe from ps
    !!";

   $skip_first_line = <ps>;
   $max_process = <ps>;
   close(ps);

   print logfile $max_process;
   close(logfile);
   sleep 60;

   ($a,$b,$c,$d,$e,$f,$g,$size) = stat($logfile);

   if ($size > 500)
   {
   print STDERR "Log file getting big, better quit!\n ";
   exit(0);
   }
}
```

Pattern matching and extraction

Perl has regular expression operators for identifying patterns. The operator

```
/regular expression/
```

returns true of false, depending on whether the regular expression matches the contents of `$_`. For example

```
if (/perl/)
   {
   print "String contains perl as a substring";
   }
if (/(Sat|Sun)day/)
   }
   print "Weekend day....";
   }
```

The effect is rather like the `grep` command. To use this operator on other variables you would write:

```
$variable =~ /regexp/;
```

Regular expression can contain parenthetic sub-expressions, e.g.

```
if (/(Sat|Sun)day (..)th (.*)/)
{
$first = $1;
$second = $2;
$third = $3;
}
```

in which case Perl places the objects matched by such sub-expressions in the variables `$1`, `$2`, etc.

Searching and replacing text

The sed-like function for replacing all occurances of a string is easily implemented in Perl using

```
while (<input>)
  {
  s/$search/$replace/g;
  print output;
  }
```

This example replaces the string inside the default variable. To replace in a general variable we use the operator = ~ with syntax:

```
$variable =~ s/search/replace/;
```

Here is an example of some of this operator in use. The following is a program which searches and replaces a string in several files. This is useful program for making a change globally in a group of files. The program is called 'file-replace'.

```
#!/local/bin/perl
##############################################################
#
# Look through files for findstring and change to newstring
# in all files.

# Define a temporary file and check it doesn't exist

$outputfile = "/tmp/file";
unlink $outputfile;

#
# Check command line for list of files
#

if ($#ARGV < 0)
   {
   die "Syntax: file-replace [file list]\n";
   }

print "Enter the string you want to find (Don't use quotes):
  \n\n:";
$findstring=<STDIN>;
chop $findstring;

print "Enter the string you want to replace with (Don't use
quotes): \n\n:";
$replacestring=<STDIN>;
chop $replacestring;

#

print "\nFind: $findstring\n ";
print "Replace: $replacestring\n";
print "\nConfirm(y/n) ";
$y = <STDIN>;
chop $y;
```

```
if ( $y ne "y")
   {
   die "Aborted -- nothing done.\n";
   }
else

   {
   print "Use CTRL-C to interrupt...\n";
   }

#
# Now shift default array @ARGV to get arguments 1 by 1
#

while ($file = shift)
   {
   if ($file eq "file-replace")
      print "Findmark will not operate on itself!";
      next;
      }

#
# Save existing mode of file for later
#

($dev,$ino,$mode)=stat($file);

open (INPUT,$file) || warn "Couldn't open $file\n";
open (OUTPUT,"> $outputfile") || warn "Can't open tmp";

$notify = 1;

while (<INPUT>)
   {
   if (/$findstring/ && $notify)
      {
      print "Fixing $file...\n";
      $notify = 0;
      }
   s/$findstring/$replacestring/g;
   print OUTPUT;
   }

close (OUTPUT);

#
# If nothing went wrong (if outfile not empty)
# move temp file to original and reset the
# file mode saved above
#

if (! -z $outputfile)
   {
   rename ($outputfile,$file);
   chmod ($mode,$file);
   }
else
```

```
    {
    print "Warning: file empty!\n.";
    }
}
```

Similarly, we can search for lines containing a string. Here is the grep program written in Perl:

```
#!/local/bin/perl
#
# grep as a perl program
#

# Check arguments etc

while (<>)
  {
  print if (/$ARGV[1]/);
  }
```

The operator `/search-string/` returns true if the search string is a substring of the default variable `$_`. To search an arbitrary string, we write

```
.... if (teststring = /search-string/);
```

Here *teststring* is searched for occurrances of *search-string*, and the result is true if one is found.

In Perl you can use regular expressions to search for text patterns. Note, however, that, like all regular expression dialects, Perl has its own conventions. For example, the dollar sign does not mean 'match the end of line' in Perl; instead one uses the \n symbol. Here is an example program which illustrates the use of regular expressions in Perl:

```
#!/local/bin/perl
#
# Test regular expressions in perl
#
# NB - careful with $ * symbols etc. Use " quotes since
#      the shell interprets these!
#

open (FILE,"regex_test");

$regex = $ARGV[$#ARGV];

# Looking for $ARGV[$#ARGV] in file...

while (<FILE>)

  if (/$regex/)
    {
    print;
    }
```

This can be tested with the following patterns:

`.*`	prints every line (matches everything)
`.`	all lines except those containing only blanks (. doesn't match ws/white-space)

`[a-z]`	matches any line containing lowercase
`[^a-z]`	matches any line containing something which is not lowercase a–z
`[A-Za-z]`	matches any line containing letters of any kind
`[0-9]`	match any line containing numbers
`#.*`	line containing a hash symbol followed by anything
`^#.*`	line starting with hash symbol (first char)
`;\n`	match line ending in a semi-colon

Try running this program with the test data on the following file which is called `regex_test` in the example program:

```
# A line beginning with a hash symbol

JUST UPPERCASE LETTERS

just lowercase letters

Letters and numbers 123456

123456

A line ending with a semi-colon;

Line with a comment # COMMENT...
```

Generate WWW pages auto-magically

The following program scans through the password database and build a standardized html-page for each user it finds there. It fills in the name of the user in each case. Note the use of the << operator for extended input, already used in the context of the shell. This allows us to format a whole passage of text, inserting variables at strategic places, and avoid having to the `print` over many `lines`.

```
#!/local/bin/perl
#
# Build a default home page for each user in /etc/passwd
#
#

$true = 1;
$false = 0;

# First build an associated array of users and full names

setpwent();

while ($true)
  {
  ($name,$passwd,$uid,$gid,$quota,$comment,$fullname) =
   getpwent;
  $FullName{$name} = $fullname;
  print "$name - $FullName{$name}\n";
  last if ($name eq "");
  }

print "\n";
```

```
# Now make a unique file name for each page and open a file

foreach $user (sort keys(%FullName))
  {
  next if ($user eq "");

  print "Making page for $user\n";
  $outputfile = "$user.html";

  open (OUT,"> $outputfile") || die "Can't open
   $outputfile\n";

  &MakePage;

  close (OUT);
  }

############################################################

sub MakePage

{
print OUT <<ENDMARKER;

<HTML>
<BODY>
<HEAD><TITLE>$FullName{$user}'s Home Page</TITLE></HEAD>
<H1>$FullName{$user}'s Home Page</H1>

Hi welcome to my home page. In case you hadn't
got it yet my name is: $FullName{$user}...

I study at <a href=http://www.iu.hioslo.no>Oslo College</a>.

</BODY>
</HTML>

ENDMARKER
}
```

Summary

Perl is a superior alternative to the shell which has much of the power of C and is therefore ideal for simple and even more complex system programming tasks. A Perl program is more efficient than a shell script, since it avoids large overheads associated with forking new processes and setting up pipes. The resident memory image of a Perl program is often smaller than that of a shell script when all of the sub-programs of a shell script are taken into account. We have barely scratched the surface of Perl here. If you intend to be a system administrator for Unix or NT systems, you could do much worse than to read the Perl book and learn Perl inside out.

C.3 WWW and CGI Programming

CGI stands for the Common Gateway Interface. It is the name given to scripts which can be executed from within pages of the World Wide Web. Although it is possible to use any

language in CGI programs (hence the word 'common'), the usual choice is Perl, because of the ease with which Perl can handle text.

The CGI interface is pretty unintelligent, in order to be as general as possible, so we need to do a bit of work to make scripts work.

Permissions

The key thing about the WWW which often causes a lot of confusion is that the WWW service runs with a user ID of `nobody` or `www`. The purpose of this is to ensure that no web user has the right to read or write files unless they are opened very explicitly to the world by the user who owns them.

For files to be readable on the WWW, they must have file mode `644` and they must lie in a directory which has mode `755`. For a CGI program to be executable, it must have permission `755`, and for such a program to write to a file in a user's directory, it must be possible for the file to be created (if necessary) and everyone must be able to write to it. That means that files which are written to by the WWW must have mode `666`, and must either exist already or lie in a directory with permission `777`[1].

Protocols

CGI script programs communicate with W3 browsers using a very simple protocol. It goes like this:

- A web page sends data to a script using the 'forms' interface. Those data are concatenated into *a single line*. The data in separate fields of a form are separated by & signs. New lines are replaced by the text `%0D%0A`, which is the DOS ASCII representation of a newline, and spaces are replaced by + symbols.
- A CGI script reads this single line of text on the standard input.
- The CGI script replies to the web browser. The first line of the reply *must* be a line which tells the browser what mime-type the data are sent in. Usually, a CGI script replies in HTML code, in which case the first line in the reply must be:

```
Content-type: text/html
```

This must be followed by a blank line.

HTML coding of forms

To start a CGI program from a web page we use a *form* which is a part of the HTML code enclosed with the parentheses

```
<FORM method="POST" ACTION="/cgi-script-alias/program.pl">
  ...
</FORM>
```

The method 'post' means that the data which get typed into this form will be piped into the CGI program via its standard input. The 'action' specifies which program you want to start.

[1] You could also set the sticky bit `1777` to prevent malicious users from deleting your file.

Note that you cannot simply use the absolute path of the file, for security reasons. You must use something called a 'script alias' to tell the web browser where to find the program. If you do not have a script alias defined for you personally, then you need to get one from your system administrator. By using a script alias, no one from outside your site can see where your files are located, only that you have a 'cgi-bin' area somewhere on your system.

Within these parentheses, you can arrange to collect different kinds of input. The simplest kind of input is just a button which starts the CGI program. This has the form

```
<INPUT TYPE="submit" VALUE="Start my program">
```

This code creates a button. When you click on it the program in your 'action' string gets started. More generally, you will want to create input boxes where you can type in data. To create a single line input field, you use the following syntax:

```
<INPUT NAME="variable-name" SIZE=40>
```

This creates a single line text field of width 40 characters. This is not the limit on the length of the string which can be typed into the field, only a limit on the amount which is visible at any time. It is for visual formatting only. The NAME field is used to identify the data in the CGI script. The string you enter here will be sent to the CGI script in the form `variable-name=value of input....` Another type of input is a text area. This is a larger box where one can type in text on several lines. The syntax is:

```
<TEXTAREA NAME="variable-name" ROW=50 COLS=50>
```

which means: create a text area of fifty rows by fifty columns with a prompt to the left of the box. Again, the size has only to do with the visual formatting, not to do with limits on the amount of text which can be entered.

As an example, let's create a WWW page with a complete form which can be used to make a guest book, or order form.

```
<HTML>
<HEAD>
<TITLE>Example form</TITLE>
<!-- Comment: Mark Burgess, 27-Jan-1997 -->
<LINK REV="made" HREF="mailto:mark@iu.hioslo.no">
</HEAD>
<BODY>
<CENTER><H1>Write in my guest book...</H1></CENTER>
<HR>

<CENTER><H2>Please leave a comment using the form below.
  </H2><P>
<FORM method="POST" ACTION="/cgi-bin-mark/comment.pl">

Your Name/e-mail: <INPUT NAME="variable1" SIZE=40> <BR><BR>

<P>
<TEXTAREA NAME="variable2" cols=50 rows=8></TEXTAREA>
<P>

<INPUT TYPE=submit VALUE="Add message to book">
<INPUT TYPE=reset VALUE="Clear message">
</FORM>
```

```
<P>
</BODY>
</HTML>
```

The reset button clears the form. When the submit button is pressed, the CGI program is activated.

Perl and the Web

Interpreting data from forms

To interpret and respond to the data in a form, we must write a program which satisfies the protocol above (see section 3.6.4). We use Perl as a script language. The simplest valid CGI script is the following:

```
#!/local/bin/perl
 #
 # Reply with proper protocol
 #
print "Content-type: text/html\n\n";
 #
 # Get the data from the form ...
 #
$input = <STDIN>;
 #
 # ... and echo them back
 #
print $input, "\n Done! \n";
```

Although rather banal, this script is a useful starting point for CGI programming, because it shows you just how the input arrives at the script from the HTML form. The data arrive all in a single, enormously long line, full of funny characters. The first job of any script is to decode this line.

Before looking at how to decode the data, we should make an important point about the protocol line. If a web browser does not get this 'Content-type' line from the CGI script it returns with an error:

```
500 Server Error

The server encountered an internal error or misconfiguration
and was unable to complete your request.

Please contact the server administrator, and inform them of the
time the error occurred, and anything you might have done that
may have caused the error.

Error: HTTPd: malformed header from script www/cgi-bin/
  comment.pl
```

Before finishing your CGI script, you will probably encounter this error several times. A common reason for getting the error is a syntax error in your script. If your program contains an error, the first thing a browser gets in return is not the 'Content-type' line, but an error message. The browser does not pass on this error message, it just prints the uninformative message above.

If you can get the above script to work, then you are ready to decode the data which are sent to the script. The first thing is to use Perl to split the long line into an array of lines, by splitting on &. We can also convert all of the + symbols back into spaces. The script now looks like this:

```
#!/local/bin/perl

#
# Reply with proper protocol
#
print "Content-type: text/html\n\n";

#
# Get the data from the form ...
#
$input = <STDIN>;

#
# ... and echo them back
#
print "$input\n\n\n";
$input =~ s/\+/ /g;

#
# Now split the lines and convert
#
@array = split('&',$input);
foreach $var ( @array )
   {
   print "$var\n";
   }
print "Done! \n";
```

We now have a series of elements in our array. The output from this script is something like this:

```
variable1=Mark+Burgess+variable2=%OD%OAI+just+called+to+say+ (wrap)
....%OD%OA...hey+pig%2C+nothing%27s+working+out+the+way+I+planned
variable1=Mark Burgess variable2=%OD%OAI just called to say (wrap)
....%OD%OA...hey pig%2Cnothing%27s working out the way I planned Done!
```

As you can see, all control characters are converted into the form %XX. We should now try to do something with these. Since we are usually not interested in keeping new lines, or any other control codes, we can simply null-out these with a line of the form

```
$input =~ s/%../g;
```

The regular expression `%..` matches anything beginning with a percent symbol followed by two characters. The resulting output is then free of these symbols. We can then separate the variable contents from their names by splitting the input. Here is the complete code:

```
#!/local/bin/perl

#
# Reply with proper protocol
#

print "Content-type: text/html\n \n";

#
# Get the data from the form ...
#

$input = <STDIN>;

#
# ... and echo them back
#

print "$input\n \n \n";

$input = s/%..//g;

$input =~s/\+ / /g;

@array = split('&',$input);

foreach $var ( @array )
  {
  print "$var<br>";
  }

print "<hr>\n";

($name,$variable1) = split("variable1=",$array[0]);
($name,$variable2) = split("variable2=",$array[1]);

print "<br>var1 = variable1<br>";
print "<br>var2 = variable2<br>";

print "<br>Done! ";
```

and the output

```
variable1=Mark+Burgess&variable2=%OD%OAI+just+called+to
  +say (wrap)+....%OD%OA...hey+pig%2C+nothing%27s+working
  +out+the+way+I  +planned
variable1=Mark Burgess
variable2=I just called to say .......hey pig nothings working
  (wrap)
 out the way I planned

var1 = Mark Burgess

var2 = I just called to say .......hey pig nothings working out
(wrap) the way I planned

Done!
```

C.4 PHP and the Web

The PHP 3 language makes the whole business of web programming rather simpler than plain Perl does. It hides the business of translating variables from forms into new variables in a CGI program, and it even allows you to embed active code into your HTML pages. PHP has special support for querying data in an SQL database like MySQL or Oracle. PHP documentation lives at `http://www.php.net`.

Embedded PHP

PHP code can be embedded inside HTML pages provided your WWW server is configured with PHP support compiled in. PHP active pages are usually called `filename.phtml`. PHP code lives inside a tag with the general form

```
<?php code... ?>
```

For example, we could use this to import one file into another and print out a table of numbers:

```
<html>
<body>

<?php

include "file.html"

for ($i = 0; $i < 10; $i++)
  {
  print "Counting $i<br>";
  }

?>

</body>
</html>
```

This makes it easy to generate WWW pages with a fixed visual layout:

```
<?php
#
# Standard layout
#

# Set $title, $comment and $contents

############################################################

print "<body>\n";
print "<img src=img/header.gif>";

print "<h1>"$title</h1>";
print "<em>$comment</em>";
print "<blockquote>\n";

include $contents;
```

```
print ("</blockquote>\n");
print ("</body>\n");
print ("</html>\n");
```

Variables are easily set by calling PHP code in the form of a CGI program from a form.

PHP and forms

PHP is particularly good at dealing with forms, as a CGI scripting language. Consider the following form:

```
<html>
<body>
<form action="/cgi-bin-script-alias/spititout.php"
 method="post">

  Name: <input type="text" name="personal[name]"><br>
  Email: <input type="text" name="personal[email]"><br>
  Preferred language:
  <select multiple name="language[]">
      <option value="English">English
      <option value="Norwegian">Norwegian
      <option value="Gobbledigook">Gobbledigook
  </select>

  <input type=image src="image.gif" name="sub">

</form>
</body>
</html>
```

This produces a page into which one types a name and e-mail address and chooses a language from a list of three possible choices. When the user clicks on a button marked by the file `image.gif` the form is posted. Here is a program which unravels the data sent to the CGI program:

```
#!/local/bin/php

<?php
#
# A CGI program which handles a form
# Variables a translated automatically
#

$title = "This page title";
$comment = "This pages talks about the following.....";

############################################################

echo "<body>";
echo "<h1>$title</h1>";
echo "<em>$comment</em>";
echo "<blockquote>\n";

###
```

```
echo "Your name is $personal[name]<br><br>";
echo "Your email is $personal[email]<br><br>";

echo "Language options: ";
echo "<table> ";

for ($i = 0; strlen($language[$i]) > 0; $i++)
  {
  echo "<tr><td bgcolor=#ff0000>Variable language[$i] =
    $language[$i]</td></tr>";
  }

if ($language[0] == "Norwegian")
  {
  echo "Hei alle sammen<p>";
  }
else
  {
  echo "Greetings everyone, this page will be in English<p>";
  }

echo "</table> ";

###

echo ("</blockquote>\n");
echo ("</body>\n");
echo ("</html>\n");
?>
```

C.5 Cfengine

System maintenance involves a lot of jobs which are repetitive and menial. There are half a dozen languages and tools for writing programs which will automatically check the state of your system and perform a limited amount of routine maintenance automatically. Cfengine is an environment for turning system policy into automated action. It is a very high level *language* (much higher level than shell or Perl) and a *robot* for interpreting your programs and implementing them. Cfengine is a general tool for structuring, organizing and maintaining information system on a network. Because it is general, it does not try to solve every little problem you might come across; instead it provides you with a framework for solving all problems in a consistent and organized way. Cfengine's strength is that it encourages organization and consistency of practice – also, that it may easily be combined with other languages.

Cfengine is about (i) defining the way you want all hosts on your network to be set up (configured), (ii) writing this in a single 'program' which is read by every host on the network, (iii) running this program on every host in order to check and possibly fix the setup of the host. Cfengine programs make it easy to specify general rules for large groups of host and special rules for exceptional hosts. Here is a summary of cfengine's capabilities:

- Check and configure the network interface on network hosts.
- Edit text files for the system or for all users.

- Make and maintain symbolic links, including multiple links from a single command.
- Check and set the permissions and ownership of files.
- Tidy (delete) junk files which clutter the system.
- Systematic, automated (static) mounting of NFS file systems.
- Checking for the presence or absence of important files and file systems.
- Controlled execution of user scripts and shell commands.
- Process management.

By automating these procedures, you will save a lot of time and irritation, and make yourself available to do more interesting work.

A cfengine program is probably not like other programming languages you are used to. It is more like a Makefile. Instead of using low-level logic, it uses high-level classes to make decisions. Actions to be carried out are not written in the order in which they are to be carried out, but listed in bulk. The order in which commands are executed is specified in a special list called the *action-sequence*. A cfengine program is a free-format text file, usually called `cfengine.conf` and consisting of declarations of the form

```
action-type:
    classes::
        list of actions
```

The action type tells cfengine what the commands which follow do. The action type can be from the following list:

```
binservers
broadcast
control
copy
defaultroute
directories
disable
editfiles
files
groups
homeservers
ignore
import
links
mailserver
miscmounts
mountables
processes
required
resolve
shellcommands
tidy
unmount
```

You may run cfengine scripts/programs as often as you like. Each time you run a script, the engine determines whether anything needs to be done – if nothing needs to be done, nothing is done! If you use it to monitor and configure your entire network from a central file-base, then the natural thing is to run cfengine daily with the help of `cron`.

Cfengine configurations can save you an enormous amount of time by freeing you from repetitive tasks. Finally, you run your system chauffeur driven with your own programmable dog. Totally excellent.

The simplest way to use cfengine

The simplest cfengine configuration you can have consists of a control section and a shellcommands section, in which you collect together scripts and programs which should run on different hosts or host-types. Cfengine allows you to collect them all together in one file and label them in such a way that the right programs will be run on the right machines.

```
control:
  domain = ( mydomain )
  actionsequence = ( shellcommands )
shellcommands:
  # All GNU/Linux machines
  linux::
          "/usr/bin/updatedb"
  # Just one host
  myhost::
          "/bin/echo Hi there"
```

While this script does not make use of cfengine's special features, it shows you how you can control many machines from a single file. Cfengine reads the same file on every host and picks out only the commands which apply.

A simple file for one host

Although cfengine is designed to organize all hosts on a network, you can also use it just on a single standalone host. In this case you don't need to know about classifying commands.

Let's write a simple file for checking the setup of your system. Here are some key points:

- Every cfengine must have a `control:` section with an `actionsequence` list, which tells it what to do, and in which order.
- You need to declare basic information about the way your system is set up. Try to keep this simple.

```
#!/usr/local/gnu/bin/cfengine -f
#
# Simple cfengine configuration file
#
```

```
control:

  actionsequence = ( checktimezone netconfig resolve files
   shellcommands )

  domain       = ( domain.country )
  netmask      = ( 255.255.255.0 )
  timezone     = ( MET )

#############################################################

broadcast:

  ones

defaultroute:

  cadeler30-gw

#############################################################

resolve:

  #
  # Add these nameservers to the /etc/resolv.conf file
  #

  128.39.89.10   # nexus
  158.36.85.10   # samson.hioslo.no
  129.241.1.99

#############################################################

files:

  /etc/passwd mode=644 owner=root action=fixall

#############################################################

shellcommands:

  Wednesday||Sunday::

    "/usr/local/bin/DoBackupScript"
```

A file for multiple hosts

If you want to have just a single file which describes all the hosts on your network, then you need to tell cfengine which commands are intended for which hosts. Having to mention every host explicitly would be a tedious business. Usually, though, we are trying to make hosts on a network basically the same as one another, so we can make generic rules which cover many hosts at a time. Nonetheless, there will still be a few obvious differences which need to be accounted for.

For example, the Solaris operating system is quite different from the GNU/Linux operating system, so some rules will apply to all hosts which run Solaris, whereas others will only apply to GNU/Linux. Cfengine uses classes like `solaris::` and `linux::` to label commands which apply only to these systems.

We might also want to make other differences, based not on operating system differences but on groups of hosts belonging to certain people, or with a special significance. We can therefore create classes using groups of hosts.

Classes

The idea of classes is central to the operation of cfengine. Saying that cfengine is 'class oriented' means that it doesn't make decisions using `if...then...else` constructions the way other languages do, but only carries out an action if the host running the program is in the same class as the action itself. To understand what this means, imagine sorting through a list of all the hosts at your site. Imagine also that you are looking for the *class* of hosts which belong to the computing department, which run GNU/Linux operating system and which have yellow spots! To figure out whether a particular host satisfies all of these criteria you first delete all of the hosts which are not GNU/Linux, then you delete all of the remaining ones which don't belong to the computing department, then you delete all the remaining ones which don't have yellow spots. If you are on the remaining list, then you are in the class of all computer-science-Linux-yellow-spotted hosts and you can carry out the action.

Cfengine works in this way, narrowing things down by asking if a host is in several classes at the same time. Although some information (like the kind of operating system you are running) can be obtained directly, clearly, to make this work we need to have lists of which hosts belong to the computer department and which ones have yellow spots.

So how does this work in a cfengine program? A program or configuration script consists of a set of declarations for what we refer to as *actions* which are to be carried out only for certain classes of host. Any host can execute a particular program, but only certain action are extracted—namely those which refer to that particular host. This happens automatically because cfengine builds up a list of the classes to which it belongs as it goes along, so it avoids having to make many decisions over and over again.

By defining classes which classify the hosts on your network in some easy to understand way, you can make a single action apply to many hosts in one go, i.e. just the hosts you need. You can make generic rules for specific type of operating system, you can group together clusters of workstations according to who will be using them, and you can paint yellow spots on them – whatever works for you.

A *cfengine action* looks like this:

```
action-type:
  compound-class::
    declaration
```

A single class can be one of several things:

- The name of an operating system architecture, e.g. `ultrix`, `sun4`, etc. This is referred to henceforth as a *hard class*.
- The (unqualified) name of a particular host. If your system returns a fully qualified domain name for your host, cfengine truncates it so as to unqualify the name.
- The name of a user-defined group of hosts.
- A day of the week (in the form `Monday Tuesday Wednesday..`).

- An hour of the day (in the form Hr00, Hr01 ... Hr23).
- Minutes in the hour (in the form Min00, Min17 ... Min45).
- A five minute interval in the hour (in the form Min00_05, Min05_10 ... Min55_00).
- A day of the month (in the form Day1 ... Day31).
- A month (in the form January, February, ... December).
- A year (in the form Yr1997, Yr2001).
- An arbitrary user-defined string.

A compound class is a sequence of simple classes connected by dots or 'pipe' symbols (vertical bars). For example:

```
myclass.sun4.Monday::
sun4|ultrix|osf::
```

A compound class evaluates to 'true' if all of the individual classes are separately true, thus in the above example the actions which follow `compound_class::` are only carried out if the host concerned is in `myclass`, is of type `sun4` and the day is Monday! In the second example, the host parsing the file must be either of type `sun4` *or* `ultrix` *or* `osf`. In other words, compound classes support two operators: AND and OR, written `.` and `|`, respectively. Cfengine doesn't care how many of these operators you use (since it skips over blank class names), so you could write either

```
solaris|irix::
```

or

```
solaris||irix::
```

depending on your taste. On the other hand, the order in which cfengine evaluates AND and OR operations *does* matter, and the rule is that AND takes priority over OR, so that `.` binds classes together tightly and all AND operations are evaluated before ORing the final results together. This is the usual behaviour in programming languages. You can use round parentheses in cfengine classes to override these preferences.

Cfengine allows you to define switch on and off dummy classes so that you can use them to select certain subsets of action. In particular, note that by defining your own classes, using them to make compound rules of this type, and then switching them on and off, you can also switch on and off the corresponding actions in a controlled way. The command line options `-D` and `-N` can be used for this purpose.

A logical NOT operator has been added to allow you to exclude certain specific hosts in a more flexible way. The logical NOT operator is (as in C and C++) `!`. For instance, the following example would allow all hosts except for `myhost`:

```
action:
!myhost::
    command
```

and similarly, so allow all hosts in a user-defined group `mygroup`, *except* for `myhost`, you would write

```
action:
mygroup.!myhost::
    command
```

which reads 'mygroup AND NOT myhost'. The NOT operator can also be combined with OR. For instance

```
class1|!class2
```

would select hosts which were either in class 1, or those which were not in class 2.

Finally, there is a number of reserved classes. The following are hard classes for various operating system architectures. They do not need to be defined because each host knows what operating system it is running. Thus, the appropriate one of these will always be defined on each host. Similarly, the day of the week is clearly not open to definition, unless you are running cfengine from outer space. The reserved classes are:

```
ultrix, sun4, sun3, hpux, hpux10, aix, solaris, osf, irix4,
  irix, irix64 freebsd, netbsd, openbsd, bsd4_3, newsos,
  solarisx86, aos, nextstep, bsdos, linux, debian, ray,
  unix_sv, GnU
```

If these classes are not sufficient to distinguish the hosts on your network, cfengine provides more specific classes which contain the name and release of the operating system. To find out what these look like for your systems you can run cfengine in 'parse-only-verbose' mode:

```
cfengine -p -v
```

and these will be displayed. For example, Solaris 2.4 systems generate the additional classes `sunos_5_4` and `sunos_sun4m`, `sunos_sun4m_5_4`.

Cfengine uses both the unqualified and fully host names as classes. Some sites and operating systems use fully qualified names for their hosts, i.e. `uname -n` returns to full domain qualified host name. This spoils the class matching algorithms for cfengine, so cfengine automatically truncates names which contain a dot '.' at the first '.' it encounters. If your host names contain dots (which do not refer to a domain name), then cfengine will be confused. The moral is: don't have dots in your host names! *NOTE: in order to ensure that the fully qualified name of the host becomes a class you must define the domain variable.* The dots in this string will be replaced by underscores.

In summary, the operator ordering in cfengine classes is as follows:

- () Parentheses override everything.
- ! The NOT operator binds tightest.
- . The AND operator binds more tightly than OR.
- | OR is the weakest operator.

We may now label actions by these classes to restrict their scope:

```
editfiles:
  solaris::
    /etc/motd
```

```
PrependIfNoSuchLine "Plan 9 was a better movie and OS!"

Rivals::
  /etc/motd
  AppendIfNoSuchLine "Your rpc.spray is so last month"
```

Actions or commands which work under a class operator like `solaris::` are only executed on hosts which belong to the given class. This is the way one makes decisions in cfengine: by class assignment rather than by `if..then..else` clauses.

Appendix D

Glossary

- *Atomic operation:* a basic, primitive operation which cannot be subdivided into smaller pieces, e.g. reading a block from a file.
- *Binaries:* files of compiled software in executable form. A compiler takes program sources and turns them into binaries.
- *BIND:* Berkeley Internet Name Domain. The library part of DNS, the routines which perform name service lookups.
- *Binary server:* a file server which makes available executable binaries for a given type of platform. A binary server is operating system specific, since software compiled on one type of system cannot be used on another. (See also Home server.)
- *Booting:* bootstrapping a machine. This comes from the expression 'to lift yourself by your bootstraps', which is supposed to reflect the way computers are able to start running from scratch, when they are powered up.
- *C/MOS:* complementary Metal Oxide Semiconductor. p-n back-to-back transistor technology, low dissipation.
- *Consolidated:* grouping resources in one place. A centralized mainframe type of solution for concentrating computing power in one place. This kind of solution makes sense for heavy calculations, performed in engineering of computer graphics.
- *Context switching:* time-sharing between processes. When the kernel switches between processes quickly in order to give the illusion of concurrency or multi-tasking.
- *Cracker:* a system intruder. Someone who cracks the system. A trespasser.
- *DAC:* Discretionary Access Control, i.e. optional rather than forced. (See MAC.)
- *Dataless client:* a client which has a disk and its own root partition, but which shares the `/usr` file tree using the NFS from a server.
- *Diskless client:* a client which has no disk at all, but which shares the its root and `/usr` file trees using the NFS from a server.
- *Distributed:* a decentralized solution, in which many workstations spread the computing power evenly throughout the network.
- *DNS:* the Domain Name Service, which converts internet names into IP addresses, and vice versa.

- *Domains:* a domain is a logical group of hosts. This word is used with several different meanings in connection with different software systems. The most common meaning is connected with DNS, the Domain Name Service. Here a domain refers to an Internet suffix, like `.domain.country`, or `.nasa.gov`. Internet domains denote organizations. Domain is also used in NT to refer to a group of hosts sharing the attributes of a common file server. Try not to confuse Domain Name Server (DNS) server with NT Domain server.
- *Enterprise:* a small business network environment. Enterprise management is a popular concept today because NT has been aimed at this market. Enterprise management typically involves running a web server, a database, a disk server and a group of workstations, and common resources like printers and so on. Many magazines think of enterprise management as the network model, but when people talk about Enterprise Management they are really thinking of small businesses with fairly uniform systems.
- *FQHN:* fully Qualified Host Name. The name of a host which is a sum of its unqualified name and its domain name, e.g. `host.domain.country`, of which `host` is the unqualified name and `domain.country` is the domain name.
- *Free software:* this usually refers to software published under the GNU Public License, Artistic License or derivative of these. Free software is not about money, but about the freedom to use, modify and redistribute software without restrictions over and above what normal courtesy to the author demands. Free software must always include human readable source code.
- *GUI:* Graphical User Interface.
- *Heterogeneous:* non-uniform. In a network context, a heterogeneous network is one which is composed of hosts with many different operating systems.
- *Home Server:* a file server which makes available users' home directories. A home server need not be operating system specific, provided it uses an commonly supported protocol, e.g. NFS, Samba. (See also Binary server.)
- *Homogeneous:* uniform. In a network context, a homogeneous network is one in which all of the hosts have the same operating system.
- *IMAP:* Internet Message Access Protocol. A modern approach to distributed e-mail services.
- *Inhomogeneous:* the opposite of homogeneous. See also *heterogeneous.*
- *Internetworking protocol:* a protocol which can send messages across quite different physical networks, binding them together into a unified communications base.
- *Index node (inode):* Unix's method of indexing files on a disk partition.
- *IP address:* Internet address. Something like `128.39.89.10`.
- *Latency:* the time you wait before receiving a reply during a transaction.
- *Legacy system:* an old computer or software package which a site has come to rely on, but which is otherwise outdated.
- *LISA:* Large Installation System Administration. This refers to environments with many (hundreds or thousands of) computers. The environments typically consist of many

different kinds of system from multiple vendors. These systems are usually owned by huge companies, organizations like NASA or universities.

MAC: Mandatory Access Control. (See DAC.)

MAC address: Media Access Control address. (e.g. Ethernet address). This is the hardware address which is burned into the network interface.

Memory image: A copy of some software in the actual RAM of the system. Often used to refer to the *resident size* of a program, or the amount of memory actually consumed by a program as it runs.

MFT: Master File Table. NTFS's system of indexing files on a disk partition.

NAT: Network Address Translator. A device which translates concealed, private IP addresses into public IP addresses. It can be used to increase the number of internal IP addresses possesed by an organization.

Open source: a software 'trademark' for software whose source files are made available to users. This is similar to the idea of free software, but it does not necessarily license users the abilitity to use and distribute the software with complete freedom. See http://www.OpenSource.com

Open systems: is a concept promoted originally by Sun Microsystems for Unix. It is about software systems being compatible through the use of freely available standards. Competitors are not prevented from knowing how to implement and include a technology in their products or from selling it under license.

PC: an Intel based personal computer, used by a single user.

Phreaker: phone phreaker. This is the name telephone network crackers used for themselves, before the so-called Hacker Crackdown of 1990.

PID: Process Identity Number.

Proprietary systems: is the opposite of open systems. These systems are secret and the details of their operation is not disclosed to competitors.

RAID: Redundant Array of Inexpensive Disks. A disk array with automatic redundancy and error correction. Can tolerate a failure of one disk in the array without loss of data.

SCSI: Small Computer Systems Interface. Used mainly for disks on multiuser systems and musical instruments.

Server: a process (a daemon) which implemenents a particular service. Services can be local to one host, or netwide.

Server-host: the host on which a server process runs. This is often abbrieviated simply to 'server', causing much confusion.

SID: security Identity number (NT).

SIMM: Memory chip arrays.

Spoofing: impersonation, faking, posing as a false identity.

SSL: Secure Socket Layer. A security wrapper which makes used of public-private key encryption to create a Virtual Private Network (VPN) link between two hosts. The SSL, developed by Netscape, has become the standard for secure communication.

- *Striping:* a way of spreading data over several disk controllers to increase throughput. Striping can be dangerous on disk failure, since files are stored over several disks, meaning that if one disk fails, all data are lost.
- *Superuser:* the root or Administrator or privileged user account.
- *SVR4:* System 5 release 4 Unix. AT&T's code release.
- *TTL:* Time To Live or Transistor-Transistor Logic.
- *UID:* User Identity Number (Unix).
- *Unqualified name:* See FQHN.
- *URL:* Uniform Resource Locator. A network 'file name' including the name of the host on which the resource resides and the network service (port number) which provides it.
- *Vendor:* a company which sells hardware or software. A *seller.*
- *Workstation:* a desktop computer which might be used by several users. Workstations can be based on for example SPARC (Sun Microsystems) or Alpha (Digital/Compaq) chip sets.
- *X11:* the Unix windows system.

Appendix E

Recommended Reading

1 *Unix System Administration Handbook, second edition*, E. Nemeth, G. Synder, S. Seebass and T.R. Hein, Prentice Hall.
2 *Essential System Administration, Second Edition*, Æ. Frisch, O'Reilly & Assoc.
3 *Windows NT: User Administration*, A.J. Meggitt and T.D. Ritchey, O'Reilly & Assoc.
4 *Network Administration Survival Guide*, S. Plumley, J. Wiley & Sons.
5 *Mastering Netware 5*, J.E. Gaskin, Sybex Network Press.
6 *Computer Networks, A Systems Approach*, L.L. Peterson and B.S. Davie, Morgan Kaufman.
7 *TCP/IP Network Administration*, Craig Hunt, O'Reilly & Assoc.
8 *DNS and BIND*, Paul Albitz and Cricket Liu, O'Reilly & Assoc.
9 *The Hacker Crackdown*, B. Sterling. Bantam Book.
10 *Computer Security*, D. Gollmann, J. Wiley & Sons.
11 *Practical Unix Security*, Simson Garfinkel and Gene Spafford, O'Reilly & Assoc.
12 *Building Internet Firewalls*, D.B. Chapman and E. D. Zwicky, O'Reilly & Assoc.
13 *Security reference*, `http://www.rootshell.com`

Bibliography

[1] RFC 1244. Site security handbook.

[2] J. Abbate. User account administration at project athena. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 28, 1987.

[3] J. Abbey. The group administration shell and the gash network computing environment. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 191, 1994.

[4] System administration and network security organization. http://www.sans.org.

[5] P. Albitz and C. Liu. *DNS and BIND.* O'Reilley & Assoc., California, 1992.

[6] D. Alter. Electronic mail gone wild. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 24, 1987.

[7] E. Anderson and D. Patterson. Extensible, scalable monitoring for clusters of computers. *Proceedings of the 11th Systems Administration conference (LISA)*, vol. 9, 1997.

[8] P. Anderson. Managing program binaries in a heterogeneous unix network. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 1, 1991.

[9] P. Anderson. Effective use of personal workstation disks in an NFS network. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 1, 1992.

[10] P. Anderson. Towards a high level machine configuration system. *Proceedings of the 8th Systems Administration conference (LISA)*, 1994.

[11] J. Apisdort, K. Claffy, K. Thompson and R. Wilder. Oc3mon: Flexible, affordable, high performance statistics collection. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 97, 1996.

[12] B. Archer. Towards a *POSIX* standard for software administration. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 67, 1993.

[13] B. Arnold. If you've seen one Unix, you've seen them all. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 11, 1991.

[14] B. Arnold. Accountworks: users create accounts on SQ1, Notes, NT and Unix. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 49, 1998.

[15] E. Arnold and C. Ruff. Configuration control and management. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 195, 1991.

[16] SAGE/Usenix association. *http://www.usenix.org.*

[17] Usenix Association. A guide to developing computing security documents, http://www.usenix.-org.

[18] ATM. Asychronous Transfer Mode. *http://www.atmforum.com.*

[19] AT&T. Virtual network computing. *http://www.uk.research.att.com/vnc.*

[20] M.R. Barber. Increased server availability and flexibility through failover capability. *Proceedings of the 11th systems administration Conference (LISA)*, vol. 89, 1997.

[21] J. Becker-Berlin. Software synchronization at the federal judicial center. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 12, 1987.
[22] B. Beecher. Dealing with lame delegations. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 127, 1992.
[23] S.M. Bellovin. Security problems in the TCP/IP protocol suite. *Computer Communications Review*, 19:2:32–48, http://www.research.att.com/~smb/papers/ipext.pdf, 1989.
[24] S.M. Bellovin. Using the domain name system for system break-ins. *Proceedings of the 5th USENIX Security Symposium*, vol. 199, 1995.
[25] BIND. *http://www.isc.org*.
[26] M. Bishop. Sharing accounts. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 36, 1987.
[27] D.R. Brownbridge and L.F. Marshall. The Newcastle connection or Unixes of the world unite. *Software—Practice and Experience*, 12:1147, 1982.
[28] P. Bumbulis, D. Cowan, E. Giguère and T. Stepien. Integrating Unix within a microcomputer oriented development environment. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 29, 1991.
[29] M. Burgess. *http://www.iu.hioslo.no/~mark/lectures*.
[30] M. Burgess. Cfengine www site. *http://www.iu.hioslo.no/cfengine*.
[31] M. Burgess. Operating systems. *http://www.iu.hioslo.no/~mark/lectures/OSindex.html*.
[32] M. Burgess. A site configuration engine. *Computing Systems*, 8:309, 1995.
[33] M. Burgess. Automated system administration with feedback regulation. *Software—Practice and Experience*, 28:1519, 1998.
[34] M. Burgess. Cfengine as a component of computer immune- systems. *Proceedings of the Norwegian conference on Informatics*, 1998.
[35] M. Burgess. Computer immunology. *Proceedings of the 12th systems administration conference (LISA)*, page 283, 1998.
[36] M. Burgess, H. Haugerud and S. Straumsness. Measuring host normality. *Software—Practice and Experience* (submitted), 1999.
[37] M. Burgess and R. Ralston. Distributed resource administration using cfengine. *Software—Practice and Experience*, 27:1083, 1997.
[38] M. Burgess and D. Skipitaris. Adaptive locks for frequently scheduled tasks with unpredictable runtimes. *Proceedings of the 11th systems administration conference (LISA)*, page 113, 1997.
[39] Mark Burgess. Managing OS security with cfengine. *;login:*, 1999.
[40] Mark Burgess. Talk at the CERN HEPIX meeting, France. October 1994.
[41] Caldera. *COAS project*. http://www.caldera.com.
[42] S. Carter. Standards and guidelines for Unix workstation installations. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 51, 1988.
[43] R. Chahley. Next generation planning tool. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 19, 1987.
[44] D.B. Chapman and E.D.Zwicky. *Building Internet Firewalls*. O'Reilley & Assoc., California, 1995.
[45] T. Christiansen. Op: a flexible tool for restricted superuser access. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 89, 1989.
[46] T. Kovacs, C.J. Yashinovitz and J. Kalucki. An optical disk backup/restore system. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 123, 1989.
[47] NTP client software. Clock synchronization software. *http://www.eecis.udel.edu/ntp/software.html*.
[48] P.R. Cohen. *Empirical Methods for Artificial Intelligence*. MIT Press, Cambridge, MA, 1995.
[49] W. Colyer and W. Wong. Depot: a tool for managing software environments. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 151, 1992.
[50] The computer incident center. *http://www.ciac.11nl.gov*.

[51] W.C. Connelly. Unix login administration at bellcore. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 13, 1988.

[52] Virtual Private Network Consortium. *http://www.vpnc.org*.

[53] M.A. Cooper. Spm: system for password management. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 149, 1995.

[54] P. Coq and S. Jean. Sysview: a user-friendly environment for administration of distributed unix systems. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 143, 1992.

[55] B. Corbridge, R. Henig and C. Slater. Packet filtering in an ip router. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 227, 1991.

[56] P. Cottrell. Password management at the University of Maryland. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 32, 1987.

[57] A.L. Couch. Visualizing huge tracefiles with xscal. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 51, 1996.

[58] N.H. Cuccia. The design and implementation of a mailhub electronic mail environment. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 37, 1991.

[59] D.A. Curry, S.D. Kimery, K.C. De La Croix and J.R. Schwab. Acmaint: an account creation and maintenance system for distributed Unix systems. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 1, 1990.

[60] M.S. Cyganik. System administration in the Andrew File System. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 67, 1988.

[61] G.E. da Silveria. A configuration distribution system for heterogeneous networks. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 109, 1998.

[62] M. Dagenais, S. Boucher, R. Gérin-Lajoie, P. Laplante and P. Mailhot. Lude: a distributed software library. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 25, 1993.

[63] T. Darmohray. A sendmail.cf scheme for a large network. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 45, 1991.

[64] Sleepcat Berkeley db project. *http://www.sleepycat.com*.

[65] A. de Leon. From thinnet to 10base-t from sys admin to network manager. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 229, 1995.

[66] L. de Leon, M. Rodriquez and B. Thompson. Our users have root! *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 17, 1993.

[67] S. DeSimone and C. Lombardi. Sysctl: A distributed control package. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 131, 1993.

[68] P.D'haeseleer. An immunological approach to change detection: Theoretical results. *9th IEEE Computer Security Foundations Workshop*, 1996.

[69] P. D'haeseleer, S. Forrest and P. Helman. *ACM Transactions on Information System Security*, submitted 1997.

[70] J. Dunham. A guide to large database tuning. *Performance Computing*, 35: May 1999.

[71] D. Eadline. Extreme linux performance tuning. *Proceedings of the second workshop on extreme Linux. http://www.extremelinux.org*.

[72] T. Eirich. Beam: a tool for flexible software update. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 75, 1994.

[73] R. Elling and M. Long. User-setup: a system for custom configuration of user environments, or helping users help themselves. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 215, 1992.

[74] R. Emmaus, T.V. Erlandsen and G.J. Kristiansen. *Network log analysis*. Oslo College dissertation, Oslo, 1998.

[75] M.J. Ranum *et al.* Implementing a generalized tool for network monitoring. *Proceedings of the 11th systems administration conference (LISA)*, page 1, 1997.
[76] Ethernet. *http://www.gigabit-ethernet.org*.
[77] Ethics. *http://www.acm.org/constitution/code.html*.
[78] Ethics. *http://www4.ncsu.edu/unity/users/j/jherkert/ethics.html*.
[79] R. Evard. Managing the ever growing to-do list. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 111, 1994.
[80] R. Evard. Tenwen: the re-engineering of a computing environment. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 37, 1994.
[81] R. Evard. An analysis of unix system configuration. *Proceedings of the 11th Systems Administration conference (LISA)*, page 179, 1997.
[82] R. Evard and R. Leslie. Soft: a software environment abstraction mechanism. *Proceedings of the eight systems administration conference LISA, (SAGE/USENIX)*, page 65, 1994.
[83] Host factory software system. *URL: http://www.wv.com*.
[84] M.K. Fenlon. A case study of network management. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 2, 1987.
[85] J. Finke. Automation of site configuration management. *Proceedings of the 11th systems administration conference (LISA)*, page 155, 1997.
[86] R. Finkel and B. Sturgill. Tools for system administration in a heterogeneous environment. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 15, 1989.
[87] TIS firewall toolkit. *http://www.tis.com*.
[88] M. Fisk. Automating the administration of heterogeneous lans. *Proceedings of the 10th systems administration conference (LISA)*, 1996.
[89] M. Fletcher. Doit: a network software management tool. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 189, 1992.
[90] M. Fletcher. An lpd for the 90s. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 73, 1996.
[91] S. Forrest, S. Hofmeyr and A. Somayaji. *Communications of the ACM*, 40: 88, 1997.
[92] S. Forrest, S. A. Hofmeyr, A. Somayaji and T.A. Longstaff. *Proceedings of 1996 IEEE Symposium on Computer Security and Privacy*, 1996.
[93] S. Forrest, A.S. Perelson, L. Allen and R. Cherukuri. *Proceedings of the 1994 IEEE symposium on research in security and privacy*, Los Alamitos, CA, IEEE Computer Society Press, 1994.
[94] S. Forrest, A. Somayaji and D. Ackley. *Proceedings of the sixth workshop on hot topics in operating systems*, Computer Society Press, Los Alamitos, CA:67–72, 1997.
[95] Æ. Frisch. *Essential System Administration, Second Edition*. O' Reilly, 1995.
[96] Æ. Frisch. *Essential Windows NT System Administration*. O'Reilly, 1998.
[97] J.L. Furlani. Modules: providing a flexible user environment. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 141, 1991.
[98] S. Garfinkel and G. Spafford. *Practical UNIX Security (2nd Edition)*. O'Reilley & Assoc., California, 1998.
[99] J.E. Gaskin. *Mastering Netware 5*. Sybex, Network Press, Alameda, 1999.
[100] L. Girardin and D. Brodbeck. A visual approach for monitoring logs. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 299, 1998.
[101] X. Gittler, W.P. Moore and J. Rambhaskar. Morgan Stanley's Aurora system. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 47, 1995.
[102] N. Goldenfeld and N.P. Kadanoff. Lessons from complexity. *Science*, 284:87, 1999.
[103] D. Gollmann. *Computer Security*. J. Wiley & Sons, Chichester, 1999.
[104] M. Gomberg, R. Evard and C. Stacey. A comparison of large- scale software installation methods on nt and unix. *Proceedings of the large installation system administration of Windows NT conference (SAGE/USENIX)*, page 37, 1998.

[105] W.H. Gray and A.K. Powers. Project accounting on a large-scale unix system. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 7, 1988.

[106] J. Greely. A flexible filesystem cleanup utility. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 105, 1991.

[107] J-C Grégoire. Delegation: uniformity in heterogeneous distributed administration. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 113, 1993.

[108] G.R. Grimmett and D.R. Stirzaker. *Probability and Random Processes.* Oxford Scientific Publications, Oxford, 1982.

[109] M. Grubb. How to get there from here: scaling the enterprise-wide mail infrastructure. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 131, 1996.

[110] B. Hagemark. Site: a language and system for configuring many computers as one computer site. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 1, 1989.

[111] P. Hall. Resource duplication for 100 percent uptime. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 43, 1987.

[112] S. Hambridge and J.C. Sedayao. Horses and barn doors: evolution of corporate guidelines for internet usage. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 9, 1993.

[113] S.E. Hansen and E.T. Atkins. Automated system monitoring and notification with swatch. *Proceedings of the 7th Systems Administration conference (LISA)*, 1993.

[114] D.R. Hardy and H.M. Morreale. Buzzerd: automated system monitoring with notification in a network environment. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 203, 1992.

[115] R. Harker. Selectively rejecting spam using sendmail. *Proceedings of the eleventh systems administration conference LISA, (SAGE/USENIX)*, page 205, 1997.

[116] K. Harkness. A centralized multi-system problem tracking system. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 40, 1987.

[117] M. Harlander. Central system administration in a heterogeneous unix environmentl genuadmin. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 1, 1994.

[118] J.A. Harris. The design and implementation of a network account management system. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 33, 1996.

[119] H.E. Harrison. Maintaining a consistent software environment. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 16, 1987.

[120] H.E. Harrison. A flexible backup system for large disk farms, or what to do with 20 gigabytes. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 33, 1988.

[121] H.E. Harrison. So many workstations, so little time. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 79, 1992.

[122] H.E. Harrison, M.C. Mitchell and M.E. Shaddock. Pong: a flexible network services monitoring system. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 167, 1994.

[123] S. Hecht. The Andrew backup system. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 35, 1988.

[124] E. Heilman. Priv: an exercise in administrative expansion. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 38, 1987.

[125] J. Hietaniemi. ipasswd: Proactive password security. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 105, 1992.

[126] N. Hillary. Implementing a consistent system over many hosts. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 69, 1989.

[127] S.A. Hofmeyr and S. Forrest. Immunizing computer networks: Getting all the machines in your network to fight the hacker disease. *1999 IEEE Symposium on Security and Privacy*, 9–12 May 1999.

[128] S. A. Hofmeyr, S. Forrest and P. D'haeseleer. An immunological approach to distributed network intrusion detection. *Paper presented at RAID'98 – First International Workshop on the Recent Advances in Intrusion Detection*, Louvain-la-Neuve, Belgium, September 1998.

[129] S. A. Hofmeyr, A. Somayaji and S. Forrest. Intrusion detection using sequences of system calls. *Journal of Computer Security* (in press).

[130] C. Hogan. Decentralising distributed systems administration. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 139, 1995.

[131] C.B. Hommel. System backup in a distributed responsibility environment. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 8, 1987.

[132] P. Hoogenboom and J. Lepreau. Computer system performance problem detection using time series models. *Proceedings of the USENIX technical conference, Summer 1993*, page 15, 1993.

[133] J.D. Howard. An analysis of security incidents on the internet. *http://www.cert.org/research/JHThesis/Start.html*, 1997.

[134] B. Howell and B. Satdeva. We have met the enemy. an informal survey of policy practices in the internetworked community. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 159, 1991.

[135] D. Hughes. Using visualization in system administration. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 59, 1996.

[136] B.H. Hunter. Password administration for multiple large scale systems. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 1, 1988.

[137] T. Hunter and S. Wanatabe. Guerilla system administration: scaling small group administration to a larger installed base. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 99, 1993.

[138] IANA. Internet Assigned Numbers Authority (port number delegation). *http://www.iana.org.*

[139] IEEE. *A standard classification for software anomalies.* IEEE Computer Society Press 1992.

[140] B. Jacob and N. Shoemaker. The Myer-Briggs type indicator: an interpersonal tool for system administrators. *Proceedings of the seventh systems administration conference LISA (supplement), (SAGE/USENIX)*, page 7, 1993.

[141] H. Jaffee. Restoring from multiple tape dumps. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 9, 1987.

[142] D. Joiret. Administration of a Unix machine network. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 1, 1987.

[143] G.M. Jones and S.M. Romig. Cloning customized hosts (or customizing cloned hosts). *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 233, 1991.

[144] V. Jones and D. Schrodel. Balancing security and convenience. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 5, 1987.

[145] W.H. Bent Jr. System administration as a user interface: an extended metaphor. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 209, 1993.

[146] H. Kaplan. Highly automated low personell system administration in a wall street environment. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 185, 1994.

[147] J.O. Kephart. A biologically inspired immune system for computers. *Proceedings of the fourth international workshop on the synthesis and simulation of living systems.* MIT Press, Cambridge MA, page 130, 1994.

[148] Linux kernel site. *http://www.linux.org.*

[149] Y.W. Kim. Electronic mail maintenance/distribution. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 27, 1987.

[150] R.W. Kint. Scrape: System configuration resource and process exception. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 217, 1991.

[151] R.W. Kint, C.V. Gale and A.B. Liwen. Administration of a dynamic heterogeneous network. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 59, 1989.

[152] K. Kistlitzin. Network monitoring by scripts. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 101, 1990.

[153] D. Koblas and P.M. Moriarty. Pits: a request management system. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 197, 1992.

[154] C. Koenigsberg. Release of replicated software in the vice file system. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 14, 1987.

[155] R. Kolstad. A next step in backup and restore technology. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 73, 1991.

[156] R. Kolstad. Tuning sendmail for large mailing lists. *Proceedings of the 11th systems administration conference (LISA)*, vol. 195, 1997.

[157] C. Kubicki. Customer satisfaction metrics and measurement. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 63, 1992.

[158] C. Kubicki. The System Administration Maturity Model: SAMM. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 213, 1993.

[159] D. Kuncicky and B.A. Wynn. *Educating and Training System Administrators: a survey.* SAGE, Short Topics in System Administration, 1998.

[160] The 10pht. *http://www.10pht.com.*

[161] E.C. Leeper. Login management for large installations. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 35, 1987.

[162] R. Lehman, G. Carpenter and N. Hien. Concurrent network management with a distributed management tool. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 235, 1992.

[163] D. Libes. Using *expect* to automate system administration tasks. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 107, 1990.

[164] libgd. Graphical library. *http://www.boutell.com/gd.*

[165] D. Lilly. Administration of network password files and NFS file access. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 3, 1988.

[166] T. Limoncelli, T. Reingold, R. Narayan and R. Loura. Creating a network for lucent bell labs south. *Proceedings of the 11th systems administration conference (LISA)*, vol. 123, 1997.

[167] L.K.C. Leighton. NT domains for Unix. *Proceedings of the large installation system administration of windows NT conference (SAGE/USENIX)*, page 85, 1998.

[168] S.W. Lodin. The corporate software bank. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 33, 1993.

[169] M. Loukides. *System Performance Tuning.* O'Reilley, California, 1990.

[170] K. Manheimer, B.A. Warsaw, S.N. Clark and W. Rowe. The depot: a framework for sharing software installation across organizational and Unix platform boundaries. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 37, 1990.

[171] P. Maniago, Consulting via mail at Andrew. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 22, 1987.

[172] C. Manning and T. Irvin. Upgrading 150 workstations in a single sitting. *Proceedings of the seventh systems administration conference LISA (supplement), (SAGE/USENIX)*, page 17, 1993.

[173] D. McNutt. Role based system administration or who, what, where, how. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 107, 1993.

[174] D. McNutt. Where did all the bytes go? *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 157, 1993.

[175] S. McRobert. From twisting country lanes to multilane Ethernet superhighways. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 221, 1995.

[176] J.T. Meek, E.S. Eichert and K. Takayama. Wide area network ecology. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 149, 1998.

[177] A.J. Meggitt and T.D. Ritchey. *Windows NT User Administration.* O'Reilly, 1997.

[178] E.S. Menter. Managing the mission critical environment. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 81, 1993.

[179] M. Metz and H. Kaye. Deejay: The dump jockey: a heterogeneous network backup system. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 115, 1992.

[180] Sun Microsystems. *Solstice system documentation. http://www.sun.com.*

[181] M.M. Midden. Academic Computing Services and Systems (ACSS). *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 30, 1987.

[182] J.E. Miller. Managing an ever-changing user database. *Proceedings of the seventh systems administration conference LISA (supplement), (SAGE/USENIX)*, page 1, 1993.

[183] T. Miller, C. Stirlen and E. Nemeth. Satool: A system administrator's cockpit, an implementation. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 119, 1993.

[184] K. Montgomery and D. Reynolds. Filesystem backups in a heterogeneous environment. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 95, 1989.

[185] R.T. Morris. A weakness in the 4.2 BSD Unix TCP/IP software. *Computer Science Technical Report*, 117:ftp://ftp.research.att.com/dist/internet_security/117.ps.Z.

[186] A. Mott. Link globally, act locally: a centrally maintained database of symlinks. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 127, 1991.

[187] E. Nemeth, G. Synder, S. Seebass and T.R. Hein. *Unix System Administration Hand-book, second edition.* Prentice Hall, 1995.

[188] E.W. Norwood. Transitioning users to a supported environment. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 45, 1988.

[189] T. Oetiker, Mrtg—the multi router traffic grapher. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 141, 1998.

[190] J. Finke. Monitoring usage of workstations with a relational database. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 149, 1994.

[191] J. Okamoto. Nightly: how to handle multiple scripts on multiple machines with one configuration file. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 171, 1992.

[192] Hewlett Packard, *Openview.*

[193] Palantir. The palantir was a project run by the University of Oslo Centre for Information Technology (USIT). Details can be obtained from *palantirusit.uio.no.* and *http://www.palantir.-uio.no.* I am informed that this project is now terminated.

[194] P.E. Pareseghian. A simple incremental file backup system. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 41, 1988.

[195] D.A. Patterson and J.L. Hennessy. *Computer Organization and Design.* Morgan-Kaufmann, San Franciso, CA, 1998.

[196] V. Paxson. Bro: A system for detecting network intruders in real time. *Proceedings of the 7th USENIX security symposium*, 1998.

[197] V. Paxson and S. Floyd. Wide area traffic: the failure of poisson modelling. *IEEE/ACM Transactions on Networking*, 3(3):226, 1995.

[198] P. D'haeseleer, S. Forrest, and P. Helman. An immunological approach to change detection: algorithms, analysis, and implications. *Proceedings of the 1996 IEEE symposium on computer security and privacy*, 1996.

[199] PHRACK. *http://www.phrack.com.*

[200] C. Pierce. The igor system administration tool. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 9, 1996.

[201] M. Poepping. Backup and restore for Unix system. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 10, 1987.

[202] Security policy. *gopher://gopher.eff.org/11/caf/policies.*

[203] Security policy. *http://musie.phlab.missouri.edu/policy/copies/tamucollection 1.html.*

[204] H. Pomeranz. Plod: keep track of what you are doing. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 183, 1993.

[205] P. Powell and J. Mason. Lprng – an enhanced print spooler system. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 13, 1995.

[206] W. Curtis Preston. Using gigabyte ethernet to backup six terabytes. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 87, 1998.

[207] Spam prevention. *http://www.pobox.com/gsutter/junkfilter/* for details.

[208] FreeS/WAN project. *http://www.xs4all.nl/freeswan.*

[209] Webmin project. *http://www.webmin.com.*

[210] D. Pukatzki and J. Schumann. Autoload: the network management system. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 97, 1992.

[211] Walters R. Tracking hardware configurations in a heterogenous network with syslogd. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 241, 1995.

[212] I. Reguero, D. Foster and I. Deloose. Large scale print spool service. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 229, 1998.

[213] P. Riddle, P. Danckeart and M. Metaferia. Agus: an automatic multiplatform account generation system. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 171, 1995.

[214] Token ring. *http://www.data.com/tutorials/tokenring.html.*

[215] M. Rodriquez, Software distribution in a network environment. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 20, 1987.

[216] S.M. Romig. Backup at Ohio State, take 2. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 137, 1990.

[217] M. Rosenstein and E. Peisach. Mkserv: worstation customization and privatization. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 89, 1992.

[218] J.P. Rouillard and R.B. Martin. Config: a mechanism for installing and tracking system configurations. *Proceedings of the 8th systems administration conference (LISA)*, 1994.

[219] J.P. Rouillard and R.B. Martin. Config: a mechanism for installing and tracking system configurations. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 9, 1994.

[220] J.P. Rouillard and R.B. Martin. Depot-lite: a mechanism for managing software. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 83, 1994.

[221] G. Rudorfer. Managing PC operating systems a revision control system. *Proceedings of the 11th systems administration conference (LISA)*, Vol. 79, 1997.

[222] C. Ruefenacht. Rust: Managing problem reports and to-do lists. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 81, 1996.

[223] N. Sammons. Multi-platform interrogation and reporting with rscan. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, 1995.

[224] B. Satdeva and P.M. Moriarty. Fdist: A domain based file distribution system for a heterogeneous environment. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 109, 1991.

[225] S.P. Schaefer and S.R. Vemulakonda. newu: Musti-host user setup. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 23, 1990.

[226] P. Schafer. Is centralized system administration the answer? *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 55, 1992.

[227] G.L. Schaps and P. Bishop. A practical approach to nfs reponse time monitoring. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 165, 1993.

[228] J. Scharf and P. Vixie. Sends: a tool for managing domain naming and electronic mail in a large organization. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 93, 1994.

[229] J. Schönwälder and H. Langendörfer. How to keep track of your network configuration. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 189, 1993.

[230] K. Schwartz, L. Cottrell and M Dart. Advenures in evolution of a high-bandwidth network for central servers. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 159, 1994.

[231] K.L. Schwartz. Optimal routing of ip packets to multi-homed hosts. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 9, 1992.

[232] J. Sellens. Software maintenance in a campus environment: the xhier approach. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 21, 1991.

[233] Sendmail. *http://www.sendmail.org.*

[234] M.E. Shaddock, M.C. Mitchell and H.E. Harrison. How to upgrade 1500 workstations on Saturday and till have time to mow the yard on Sunday. *Proceedings of the ninth systems administration conference LISA, (SAGE/USENIX)*, page 139, 1995.

[235] C.E. Shannon and W. Weaver. *The Mathematical Theory of Communication.* University of Illinois Press, Urbana, IL, 1949.

[236] J.M. Sharp. Request: a tool for training new sys admins and managing old ones. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 69, 1992.

[237] Root shell security site. *http://www.rootshell.com.*

[238] C. Shipley and C. Wang. Monitoring activity on a large unix network with Perl and syslogd. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 209, 1991.

[239] S. Shumway. Issues in on-line backup. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 81, 1991.

[240] T. Sigmon. Automatic software distribution. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 21, 1987.

[241] J. Da Silva and Ólafur Guðmundsson. The Amanda network backup manager. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 171, 1993.

[242] R. Silverberg. *Shadrach in the Furnace.* Gollancz, 1976.

[243] N. Simicich. Yabs. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 109, 1989.

[244] S. Simmons. Making a large network reliable. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 47, 1988.

[245] S. Simmons. Life without root. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 89, 1990.

[246] K.C. Smallwood, Guidelines and tools for software maintenance. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 47, 1990.

[247] J.M. Smith. Creating an environment for novice users. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 37, 1987.

[248] T. Smith. Excelan administration. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 4, 1987.

[249] A. Somayaji, S. Hofmeyr, and S. Forrest. Principles of a computer immune system. *New security paradigms workshop*, September 1997.

[250] B. Spence. Spy: a unix file system security monitor. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 75, 1989.

[251] H.L. Stern and B.L. Wong. NFS performance and network loading. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 33, 1992.
[252] R. Stevens. *http://www.kohala.com/rstevens/shimomura.95jan25.txt.*
[253] R. Stevens. *Advanced programming in the UNIX environment.* Addison-Wesley, Reading, MA, 1992.
[254] R. Stevens. *TCP/IP Illustrated Vols 1–3.* Addison Wesley, Reading, MA, 1994–1996.
[255] R.J. Stolfa. Simplifying system administration tasks: the UAMS approach. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 203, 1993.
[256] K. Stone. System cloning at HP-SDD. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 18, 1987.
[257] Tivoli systems/IBM. *Tivoli software products. http://www.tivoli.com.*
[258] L.W. Taylor and J.R. Hayes. An automated student account system. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 29, 1987.
[259] The Computer Emergency Response Team. *http://www.cert.org.*
[260] G.S. Thomas, J.O. Schroeder, M.E. Orcutt, D.C. Johnson, J.T. Simmelink and J.P. Moore. Unix host administration in a heterogeneous distributed computing environment. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 43, 1996.
[261] W.F. Tichy, RCS-a system for version control. *Software—Practice and Experience*, 15: 637, 1985.
[262] Tuning tips. *http://ps-ax.com.*
[263] S. Traugott and J. Huddleston. Bootstrapping an infrastructure. *Proceedings of the 12th USENIX/ LISA conference on system administration*, vol. 181, 1998.
[264] M. Urban, UDB: Rand's group and user database. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 11, 1990.
[265] D.L. Urner. Pinpointing system performance issues. *Proceedings of the 11th systems administration conference (LISA)*, vol. 141, 1997.
[266] P. van Epp and B. Baines. Dropping the mainframe without crushing the userss: mainframe to distributed unix in nine months. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 39, 1992.
[267] R.R. Vangala, M.J. Cripps and R.G. Varadarajan. Software distribution and management in a networked environment. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 163, 1992.
[268] A. Vasilatos. Automated dumping at project athena. *Proceedings of the first systems administration conference LISA, (SAGE/USENIX)*, page 7, 1987.
[269] Wietse Venema. TCP wrappers. *http://ciac.llnl.gov/ciac/ToolsUnixNetSec.html.*
[270] J.S. Vöckler, *http://www.rvs.uni- hannover.de/people/voeckler/tune/en/tune.html.*
[271] J. von Neumann. The general and logical theory of automata. *Reprinted in vol 5 of his Collected Works*, 1948.
[272] J. von Neumann, Probabilistic logics and the synthesis of reliable organisms from unreliable components. *Reprinted in vol 5 of his Collected Works*, 1952.
[273] W.A. Doster, Y-H. Leong, and S.J. Matteson. Uniqname overview. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 27, 1990.
[274] C. Warrender, S. Forrest, and B. Pearlmutter. Detecting intrusions using system calls: Alternative data models. *IEEE symposium on security and privacy*, submitted 9–12 May 1999.
[275] A. Watson and B. Nelson. Laddis: A multi-vendor and vendor- neutral spec nfs benchmark. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 17, 1992.
[276] L.Y. Weissler. Backup without tapes. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 191, 1991.
[277] E.T. Whittaker and G. Robinson. *Calculus of observations.* Blackie and Son Ltd., London, 1929.
[278] W. Willinger and V. Paxson. Where mathematics meets the Internet. *Notices of the Am. Math. Soc.*, 45(8):961, 1998.

[279] W. Willinger, V. Paxson and M.S. Taqqu. Self-similarity and heavy tails: structural modelling of network traffic. *A practical guide to heavy tails: statistical techniques and applications*, Birkhauses, Boston, pages 27–53, 1996.

[280] C.E. Wills, K. Cadwell and W. Marrs. Customizing in a Unix computing environment. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 43, 1993.

[281] I.S. Winkler and B. Dealy. Information security technology? Don't rely on it. A case study in social engineering. *Proceedings of the 5th USENIX security symposium*, page 1, 1995.

[282] W.C. Wong. Local disk depot: customizing the software environment. *Proceedings of the seventh systems administration conference LISA, (SAGE/USENIX)*, page 51, 1993.

[283] B. Woodard, Building an enterprise printing system. *Proceedings of the twelfth systems administration conference LISA, (SAGE/USENIX)*, page 219, 1998.

[284] H.Y. Yeom, J. Ha and I. Kim. IP multiplexing by transparent port-address translator. *Proceedings of the tenth systems administration conference LISA, (SAGE/USENIX)*, page 113, 1996.

[285] T. Ylonen. SSH – secure login connections over the internet. *Proceedings of the 6th USENIX Security Symposium*, vol. 37, 1996.

[286] M.V. Zelkowitz and D.R. Wallace. Experimental models for validating technology. *IEEE Computer*, page 23, 1998.

[287] E. D. Zwicky, Backup at Ohio state. *Proceedings of the second systems administration conference LISA, (SAGE/USENIX)*, page 43, 1988.

[288] E.D. Zwicky. Disk space management without quotas. *Proceedings of the third systems administration conference LISA, (SAGE/USENIX)*, page 41, 1989.

[289] E.D. Zwicky. Enhancing your apparent psychic abilities. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 171, 1991.

[290] E.D. Zwicky. Torture testing backup and archive programs: things you ought to know but probably would rather not. *Proceedings of the fifth systems administration conference LISA, (SAGE/USENIX)*, page 181, 1991.

[291] E.D. Zwicky. Typecast: beyond cloned hosts. *Proceedings of the sixth systems administration conference LISA, (SAGE/USENIX)*, page 73, 1992.

[292] E.D. Zwicky. Getting more work out of work tracking systems. *Proceedings of the eighth systems administration conference LISA, (SAGE/USENIX)*, page 105, 1994.

[293] E.D. Zwicky, S. Simmons and R. Dalton. Policy as a system administration tool. *Proceedings of the fourth systems administration conference LISA, (SAGE/USENIX)*, page 115, 1990.

Index

/var/spool/mail, 349
<> *filehandle in Perl*, 363
==, 357
/bin, 17
/devices, 17
/dev, 17
/etc/aliases, 226
/etc/checklist, 228
/etc/checklist, *HPUX*, 84
/etc/dfs/dfstab, 31, 227
/etc/ethers, 62, 93
/etc/exports, 31, 227
/etc/filesystems, AIX, 84
/etc/filesystems, 228
/etc/fstab, 84, 104, 228
/etc/group, 19
/etc/hosts.allow, 229, 276
/etc/hosts.deny, 276
/etc/inetd.conf. 183, 276
/etc/inittab, 79
/etc/named.boot, 189
/etc/named.conf, 189
/etc/nsswitch.conf, 92
/etc/printcap, 230
/etc/rc.local, 184
/etc/rc *files*, 184
/etc/resolv.conf, 65, 91
/etc/services, 183
/etc/system, 177
/etc/vfstab, 84, 228
/etc, 17
/export, 17
/home, 17
/sbin, 17
/sys, 17
/users, 17
/usr/bin, 17
/usr/etc/resolv.conf *on IRIX*, 91
/usr/local/gnu, 98
/usr/local/site, 98
/usr/local, 17, 96
/usr, 107
/var/adm, 18
/var/mail, 349
/var/spool, 18
/var, 18
-D option, 387
-N option, 387
cshrc, 114, 117
directory, 18
directory, 18
fvwm2rc, 118
.fvwm95rc, 118
.mwmrc, 118
.profile, 114, 117
.rhosts, 346
.xinitrc, 118
.xsession, 118
in make, 35
$ < *in make*, 355
$? *in make*, 355

a record, 195
access bits, 19
  octal form, 20
  text form, 20
access control-for services, 276
  lists, 23
access rights, 19
access to files, 19
ACEs in NT, 27
ACLs, 23

in NT, 27
network services, 186
actionsequence, 383
administrator account, 35
AFS, 29, 113
agents, 182
aliases-in mail, 226
DNS, 45
alive, checking a host, 350
analyzing security, 271
andrew filesystem, 29
application layer, 50
arch *program*, 64
argument vector in Perl, 356, 358
ARP/RARP, 62
arrays – associated in Perl, 359
(normal) in Perl, 358
and split, 359
in Perl, 356
associated arrays, iteration, 362
AT & T, 11
athena, 113
attacks, links, 102
auspex, 147
Average time before failure, 299

backdoors, 242
background process, NT, 33
backup, 241
schedule, 265
tools, 265
backups, 264
big endian, 55
binary server, 71
BIND, 349
version, 189
setting up, 91
binding socket service, 186
biod, 229
BIOS, 10, 37, 81
block, disk, 179
blocks, 83
boot-loader, 108
scripts, 79
booting-Unix, 79
NT, 81
BOOTP protocol, 47
bootstrapping an infrastructure, 139
bridge, 53
broadcast address,
BSD-4.3, 154
Unix, 11
byte order, 55

cache-file, DNS, 188
poisoning, 259
CACLS *command*, 27
cancel, 234
canonical name, 45, 195
canonical names, 188
Catman *command*, 348
Causality, 319
cfdisk,
Cfengine, 271, 382
conf, 383
checksums, 270
prevention, 162
specialized hosts, 38
chgrp command, 21
CGI protocol, 374
checking – the mode of installed software, 98
whether host is alive, 350
Checksums, 270
chgrp *command*, 21
chmod *command*, 20
chown *command*, 21
chomp *command in Perl*, 365
chop *command in Perl*, 365
close *command in Perl*, 363
Class A, B, C, D, E networks, 56
Classes, 386
compound, 386
defining and undefining, 387
clock synchronization, 153
cloning NT, 95
closed system, 315
CNAME, 195
collisions, 175
command-interpreter, 13
CALCS *see* CALCS *command*
catman *see* catman *command*
chgrp *see* chgrp *command*
chmod *see* chmod *command*
chown *see* chown *command*
cp *see* cp *command*
crontab *see* crontab *command*
df *see* df *command*
dump *see* dump *command*
du *see* du command
etherfind *see* etherfind *command*

exportfs *see* exportfs *command*
find *see* find *command*
fconfig *see* cf config *command*
iostat *see* iostat *command*
kill *see* kill *command*
ldconfig *see* ldconfig *command*
lod *see* lod *command*
locate *see* locate *command*
mkfile *see* mkfile *command*
netstat *see* netstat *command*
newfs *see* newfs *command*
nfsstat *see* nfsstat *command*
ping *see* ping *command*
ps *see* ps *command*
rdump *see* rdump *command*
renice *see* renice *command*
restore *see* restore *command*
rlogin *see* rlogin *command*
rm -i *see* rm -i *command*
route *see* route *command*
rsh *see* rsh *command*
snoop *see* snoop *command*
su -c *see* su -c *command*
swapon *see* swapon *command*
tar *see* tar *command*
Telnet *see* Telnet *command*
traceroute *see* traceroute *command*
ufsdump *see* ufsdump *command*
vmstat *see* vmstat *command*
what is *see* what is *command*
which *see* which *command*
whois *see* whois *command*
line arguments in Perl, 358
community-string, 275
strings, 145
components, handling, 36
compound classes, 386
computer immunology, 136
configure, 98
connection times, TCP, 177
contact with the outside world, 61
contention, 75
convergence, 136, 137, 314
corruption in file system, 170
cp *command* 347
crack, *passwords*, 273
cron, 154, 348, 384
controlling with cfengine, 158
Crontab *command*, 154, 349
crypt, 367
Cut as a Perl script, 363
Cygwin Unix compatibility for NT, 98

Daemon, 183
Daemons, 43
and services, 183
starting without privilege, 102
data links layer, 50
day of the week, 160
DCE, 29, 113
death to the users, 69
dfstab, 31
default name server, 66
printer, 230
route, 57, 350
defunct process, 33
delegation, 58
delta distribution, 323
demultiplexing, 143
denial of service attack, 241, 256
dependencies in make files, 353
dependency, 70
problems, 169
depot, 98
deterministic system, 316
devices, 104
df command, 346
DFS, 29
NT, 29
diagnostics, 164
die, 366
differences, hosts, 33
dig, 65
disk-backups, 264
doctor, 349
installing, 104
mirroring, 75, 263
partition names, 106
performance, 174
quotas, 125
repair, 349
statistics, 349
striping, 174
distributed computing environment, 29
distribution, measurements, 321
DNS, 63, 64, 187, 349
aliases, 45
cache file, 188
BIND setup, 91
mail records, 195

revoking is rights, 193
dnsquery, 349
domain, 64
listing hosts in, 67
name, 65
name, definition, 91
name service, 187
NT, 47
OS, 23
D option *see* -D option
do..while *in Perl*, 361
DOS, 10
DOS attack, 256
dots in hostnames, 388
down, checking a host, 350
downtime, 299
drive letter assignment 84
dump command, 347
du *command*, 347
dynamical systems, 331

eeprom, 37
encryption, 367
entropy, 317
entry points to OS code, 9
environment variables, 33
in Perl, 356, 358
error-law, 326
reporting, 164
in Perl, 366
etherfind *command*, 350
eq, 357
eq *and* == *in Perl*, 365
executable, making programs, 21
exiting on errors in Perl, 366
exporting-files, 230
file systems, Unix, 31
on GNU/Linux, 227
exports *command*, 227
external hosts do not seem to exist, 62

fail-over, 143
fault tolerance, policy, 162
fdisk, 83
feedback regulation, 136
file-access permission, 19
handles in Perl, 363
hierarchy, Unix, 17
protection bits, 19
sharing, Windows/Unix, 148
system table, 104
type problem in WWW, 213
Files-in Perl, 363
iterating over lines, 363
find *command*, 348
Finding – a mail server, 66
domain information, 349
the name server for other domains, 67
firewalls, 93, 281
for loops in Perl, 361
foreach loop, 362
for *loop in Perl*, 360
Un Perl, 360
Fork, 368
forking new processes, 368
format *program, Sun*, 349
formatting a file system, 349
forms in HTML, 375
fourier analysis, 330
FQHN, 187
fractal nature of network traffic, 302
fragment, of block, 179
fragmentation of IP, 258
free software foundation, 97
fsck *program*, 349
FTP, 98, 181
Fully qualified names, 388
fvwm *see window manager*
fvwm 2 *see window manager*
fvwm 95 *see window manager*

game theory, 331
gateway, 350
gaussian distribution, 327
glue record, DNS, 200
GNU software, 97
grouping time values, 160
groups, 19, 115
and time intervals, 160
in cfengine, 385
guest accounts, 117

handling components, 36
handshaking, 50
hangup signal, 348
hard links, 18
NT, 26
heavy-tailed distribution, 329
help desk, 124
Hewlett-Packard, 11

hierarchy, file, 17
HINFO, 195
hme fast ethernet interface, 175
home directory, 114
homogeneity, 138
host name – gets truncated, 388
  lookup, 65, 91
HTTP, 181
HTTPS, 182
hub, 53

IBM AS/400s, 10
IBM S/370, 10
IBM S/390, 10
IDE disks, 37
if config command 60, 350
if *in Perl*, 360
IMAP, 182
immune system, 135
immunity model, 136
immunology, 135
incremental backup, 266
index nodes, 17, 18
inetd master – daemon, 184
inheritance of environment, 34
inode corruption, 170
inodes, 17, 18
in rarpd, 93
inctd, 348
installboot, *SunOS*, 108
INSTALL, 98
installing a new disk, 104
interface configuration, 60, 350
internet domain, 64
interpretation of values in Perl, 357
interrupts, 10
intranet, 203, 209
iostat *command*, 350
IP address, 56, 65, 187
  lookup, 65
  setting, 60
  slash notation, 192
  v6, 56
ISO, 49
iterating over files, 363
iteration over arrays, 362

junkfilter, 224

Kerberos, 113
kernel, 18
  architecture, 106
  configuration, 107
  tuning, Solaris, 177
keys, 362
kill *command*, 348
  NT process, 33

labelling a disk, 83
lame delegation, DNS, 201
latency, 179
law of errors, 326
LDAP, 182
ldconfig *command* 348
ldd *command*, 347
License servers, 99
link attacks, 102
linux, 11
  exports, 227
little endian, 55
lmgrd, license server, 92
ln, 18
ln -s, 18
loadlin, 92
local variables in Perl, 366
locate *command*, 348
log rotation, 170
logical NOT, 387
login directory, 114
long file listing, 19
looking up name/domain information, 349
lookup hosts in a domain, 67
loopback-address, 57, 91
  network in DNS, 188
lost + found, 84
pc, 233
lp, 233
  default printer, 230
lpd, 233
lpq, 233
lprm, 233
lpr, 233
lpsched, 234
lpshut, 234
lpstat -a, 234
lpstat -o all, 234
lp, 233
ls -1, 19

mach *program*, 64

Macintosh, 10, 48, 146
magic numbers, 19
mailaddress of administrator, 68
  aliases, 226
  exchangers, 66
  queue, 347
  records in DNS, 175
  relaying, 217
  spool directory, 349
  finding the server, 66
mailbox system, 216
make, 98
make program for configuration, 141
management information base, 145
making programs executable, 21
master boot record, 81
mean downtime, 299
mean-time before failure, 299
  value, 325
memory leak, 172
MIB, 145
mime types in WWW, 375
mirroring of disks, 75, 266
mission critical systems, 235
mkfile *command*, 93
mkfs command, 84
modular kernel, 178
month, 159
mount -a, 228
mountd, 229
mount *command*, 228
mounting-file systems, 31, 104, 347
  problems, 229
MTS, 10
multicast address, 57
multi-port repeater, 53
multi user OS, 10
multi-user mode, 79
multitasking system, 10
mwm *window manager*, 118
MX, 195
  records, 196
MySQL, 102, 214

name service, 43
  lookups, 65
name server for other domains, 67
  list, 91
naming scheme for Internet, 64
NAT, 58
ncftp, 78
ndd-command, Solaris, 107, 177
  Kernel parameters, 107
Netmask, 58
  examples, 57
netstat -r-*and routing table*, 350
  command, 61
netstat *command*, 350
network address translator, 58
  appliance, 48
  byte order, 55
  information service, 65, 112
  interface, 36, 50
  interfaces, 350
  layer, 50
  numbers, 188
  transmission method, 51
networks, 55
Newcastle file system, 28
newfs, 84, 349
newsprint, 234
NFS, 28
  lient/server statistics, 350
  root access, 272
nfsd, 229
nfsiod, 229
nfsstat *command*, 350
nice, 348
NIS, 65, 91, 112
  plus, 91
nmap program, 278
  port scanner, 70
no contact with outside world, 61
non-repudiation, 34
N option *see* N option
normal – distribution, 326
  error law, 326
NOT operator, 387
Novell, 28, 47
  disk purge, 167
NS, 195
nslookup, 65, 349
NT, 47
  installation, 90
  ACL/ACEs, 27
  drive letter assignment, 84
  install, 95
NTP, 182
null client (mail), 197

one time passwords, 274–5
open *command in Perl*, 363
open system, 315
operating system, 9
operator ordering, 388
option *see* -D option, -N option
Oracle, 214
OS/2 boot manager, 92
OSI model, 49
overheads, performance, 174

paging, 85
PAM, 148
parallelism, 143
parameters in Perl functions, 366
pareto distribution, 329
partitions, 104
password cracking, NT, 273
  file, 367
  sniffing, 273
paste as a Perl script, 364
pattern matching in Perl, 369, 370
  replacement in Perl, 369
PCNFS, 146
Perl, 160, 271, 356
  strings and scalar, 357
  truncating strings, 365
  variables and types, 356
permissions – on files, 19
  on installed software, 99
persistent connections, 207
PHP, 204, 214
physical layer, 50
PID, 32
Ping-attacks, 256
  command, 256, 300
pluggable authentication modules, 148
police service, policy, 161
policy, 148
  formalizing, 144
  user support, 124
port, 186
  scanning, 63, 70
  sniffing, 277
portmapper, 229
Posix ACLs, 23
presentation layer, 50
prey – Predator models, 136
principle of uniformity, 138
print services, 229
  spool area, 230
print-queue listing, 233–4
  remove job, 233–4
  start, 233–4
  stop, 233–4
printer, 233
  choosing a default, 230
registration, 230–1
privileged users, 123
probability distributions, 326
probe-scsi, *Sun*, 104
process-ID, 32
  starvation, 164
procmail, 224
protection bits, 19
protocols, 50
proxy, 182
  firewall, 283
ps *command*, 348
PTR records, 198
pty's increasing number, 177
PwDump, NT, 273

q = any, nslookup, 66
q = mx, nslookup, 66
q = ns, nslookup, 67
queso program, 278
quotas, 125

race conditions, 102
RAID, 263
RARP, 62, 93
rdump *command*, 347
README, 98
rc *files*, 79
real time systems, 235
redundancy, 143, 239
registering a printer, 230–1
registry, NT, 95
regulation, feedback, 136
relaying, mail, 217
renice *command*, 348
repairing a damaged disk, 349
repeater, 53
resolver, setting up, 91
resources, competition, 138–139
restarting daemons, 348
restore *command*, 347
restricting privilege, 9, 12, 19, 27, 33, 339
reverse lookup, DNS, 188

rlogin *command*, 346
rm -i *command*, 263
root – account, 35
  partition, 105
rotation, logs, 170
route *command*, 350
router, 53
router/switch difference, 53
routing – information, 350
  table, 61, 350
RPG, 182
  service not registered error, 229
rpc.mountd, 229
rpc.nfsd, 229
rsh *command*, 346
run-cfengine *file*, 159
running jobs at specified times, 154

s-bit, 22
S-HTTP, 181
S/KEY, 274
Samba, 148
scalar variables in Perl, 357
scheduling priority, 349
scheduling service, NT, 155
script aliases in WWW, 376
scripts, 13
SCSI-disks, 37
  probe on SunOS, 104
searching and replacing in Perl (example), 369
sectors, 83
security – holes, 242
  analysis, 271
sendmail, 349
sed *as a perl script*, 370
sequence guessing, 258
serial number, DNS, 199
server message block, 148
service – configuration, 183
  packs, NT, 90
services, 43, 138
  daemons, 183
  starting without privilege, 101
session layer, 50
setgid bit, 19
setuid bit, 19
  programs, 242
  software, 101
shadow password files, 273
shannon entropy, 317
share, 227
shareall, 228
sharing filesystems, Unix, 30
shell, 13, 31
shift *and arrays*, 358
short cuts, NT, 26
shutdown, NT, 81
SIMM, 37
simple network management protocol, 145
single task system, 10
single user OS, 10
single-user mode, 79
site specific data, 99
slash notation, IP, 192
slowly running systems, 167
SMB protocol, 148
smurf attack, 258
SNMP, 145, 275
  security, 145
snoop *command*, 350
SOA, 195
socket connections, 350
sockets, 50
soft links, 18
spectrum of frequencies, 330
split *and arrays*, 359
split *command*, 359
SSH, 182
  command, 346
standard – deviation, 325
  standard error of the mean, 327
standard I/O in Perl, 363
standardization, 38
start up files for Unix, 79
startx, 118
starvation of process, 164
static kernel, 178
statistics, disks, 350
  NFS, 350
  virtual memory, 350
sticky bit, 22
strings in Perl, 356
stty *and switching off term echo*,
subnets, 58
subroutines in Perl, 366
su -c *command*, 102
suffix rules in Makefiles, 353
Sun Microsystems, 11
superuser, 35, 123
support, 124

SVR4, 393
swapon command, 349
swap-partition, 106
  space, 349
swapping, 85
  switching on, 349
switch/router difference, 53
switched networks, 53
sybase, 214
symbolic-ink attacks, 102
  links, 18
SYN flooding, 258
System 5/System V, 11
system – accounting, 125
  policy, 144, 147
  registry, 95
  type, 62, 64

tar *command*, 347
t-bit, 22
TCP-tuning, 177
  wrappers, 186, 276
  IP, 49
  IP security, privilege, 36
  IP spoofing, 257
tcpd, 276
teardrop, 258
Telnet *command*, 346
terminal echo and stty, 367
text form of access bits, 20
thrashing, 176
time-classes, 159
  executing jobs at specified, 153
  service, 153
timezone, 63
traceroute *command*, 350
traffic analysis, 329
transceiver, 36
transport layer, 50
tripwire, 270
  troubleshooting, 164
truncating strings in Perl, 365
trust relationship, 112
trusted ports, 36
TTL, 393
types in Perl, 356

ufsdump *command*, 347
uid, 113
Ultrix, 154
umask, 21
uname, 64
undeleting files, 263
uniformity, 38, 138, 140
unless in *Perl*, 361
up, checking a host, 350
updatedb *script*, 348
URL 71
usage patterns, understanding, 139
user, 9
  name, 12
  support, 124
  id, 113
UWIN Unix toolkit for NT, 98

virtual-machine model, 141
  memory statistics, 350
  Network Computing, 124
  private network, 279
vmstat *command*, 350
vmunix, 18
VNC, 124
VPN, 279

weather, 38
whatis *command*, 348
which *command*, 348
whois *command*, 349
while *in Perl*, 361
Window manager
  fvwm, 118
  fvwm 2, 118
  fvwm 95, 118
Windows, 10
workstation, NT, 13
WWW security, 279

xdm 118
xhost-access control, 278
  command 228
xntpd, 153

years, 159
Yellow Pages, 112
YP, see Yellow pages

zombie process, 33