

DragonFly Handbook

The DragonFly Documentation Project

DragonFly Handbook

by The DragonFly Documentation Project

Published June 2004

Copyright © 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, 2003, 2004 The FreeBSD Documentation Project

Copyright © 2004, 2005, 2006 The DragonFly Documentation Project

Welcome to DragonFly! This handbook covers the installation and day to day use of the *DragonFly* operating system. This manual is a *work in progress* and is the work of many individuals. Many sections do not yet exist and some of those that do exist need to be updated. If you are interested in helping with this project, send email to the DragonFly Documentation project mailing list (<http://leaf.dragonflybsd.org/mailarchive/>). The latest version of this document is always available from the DragonFly web site (<http://www.dragonflybsd.org/>) or mirror sites, in a variety of formats.

Portions of this document originally documented use of the FreeBSD (<http://www.freebsd.org/>) operating system. While many functions should be similar on DragonFly, some differences should be expected. If you find instructions here that no longer apply to DragonFly, please contact the documentation mailing list at DragonFly Documentation project mailing list (<http://leaf.dragonflybsd.org/mailarchive/>).

Redistribution and use in source (SGML DocBook) and 'compiled' forms (SGML, HTML, PDF, PostScript, RTF and so forth) with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code (SGML DocBook) must retain the above copyright notice, this list of conditions and the following disclaimer as the first lines of this file unmodified.
2. Redistributions in compiled form (transformed to other DTDs, converted to PDF, PostScript, RTF and other formats) must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Important: THIS DOCUMENTATION IS PROVIDED BY THE DRAGONFLYBSD PROJECT "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE DRAGONFLYBSD PROJECT BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS DOCUMENTATION, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

DragonFlyBSD is a registered trademark of the DragonFlyBSD Project.

FreeBSD is a registered trademark of Wind River Systems, Inc. This is expected to change soon.

NetBSD is a registered trademark of The NetBSD Foundation, Inc.

3Com and HomeConnect are registered trademarks of 3Com Corporation.

3ware and Escalade are registered trademarks of 3ware Inc.

ARM is a registered trademark of ARM Limited.

Adaptec is a registered trademark of Adaptec, Inc.

Adobe, Acrobat, Acrobat Reader, and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Apple, FireWire, Mac, Macintosh, Mac OS, Quicktime, and TrueType are trademarks of Apple Computer, Inc., registered in the United States and other countries.

Corel and WordPerfect are trademarks or registered trademarks of Corel Corporation and/or its subsidiaries in Canada, the United States and/or other countries.

Sound Blaster is a trademark of Creative Technology Ltd. in the United States and/or other countries.

CVSup is a registered trademark of John D. Polstra.

Heidelberg, Helvetica, Palatino, and Times Roman are either registered trademarks or trademarks of Heidelberger Druckmaschinen AG in the U.S. and other countries.

IBM, AIX, EtherJet, Netfinity, OS/2, PowerPC, PS/2, S/390, and ThinkPad are trademarks of International Business Machines Corporation in the United States, other countries, or both.

IEEE, POSIX, and 802 are registered trademarks of Institute of Electrical and Electronics Engineers, Inc. in the United States.

Intel, Celeron, EtherExpress, i386, i486, Itanium, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intuit and Quicken are registered trademarks and/or registered service marks of Intuit Inc., or one of its subsidiaries, in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States.

LSI Logic, AcceleRAID, eXtremeRAID, MegaRAID and Mylex are trademarks or registered trademarks of LSI Logic Corp.

M-Systems and DiskOnChip are trademarks or registered trademarks of M-Systems Flash Disk Pioneers, Ltd.

Macromedia, Flash, and Shockwave are trademarks or registered trademarks of Macromedia, Inc. in the United States and/or other countries.

Microsoft, FrontPage, MS-DOS, Outlook, Windows, Windows Media, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Netscape and the Netscape Navigator are registered trademarks of Netscape Communications Corporation in the U.S. and other countries.

GateD and NextHop are registered and unregistered trademarks of NextHop in the U.S. and other countries.

Motif, OSF/1, and UNIX are registered trademarks and IT DialTone and The Open Group are trademarks of The Open Group in the United States and other countries.

Oracle is a registered trademark of Oracle Corporation.

pkgsrc is a registered trademark of The NetBSD Foundation, Inc.

PowerQuest and PartitionMagic are registered trademarks of PowerQuest Corporation in the United States and/or other countries.

RealNetworks, RealPlayer, and RealAudio are the registered trademarks of RealNetworks, Inc.

Red Hat, RPM, are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

SAP, R/3, and mySAP are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Sun, Sun Microsystems, Java, Java Virtual Machine, JavaServer Pages, JDK, JSP, JVM, Netra, Solaris, StarOffice, Sun Blade, Sun Enterprise, Sun Fire, SunOS, and Ultra are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Symantec and Ghost are registered trademarks of Symantec Corporation in the United States and other countries.

MATLAB is a registered trademark of The MathWorks, Inc.

SpeedTouch is a trademark of Thomson

U.S. Robotics and Sportster are registered trademarks of U.S. Robotics Corporation.

VMware is a trademark of VMware, Inc.

Waterloo Maple and Maple are trademarks or registered trademarks of Waterloo Maple Inc.

Mathematica is a registered trademark of Wolfram Research, Inc.

XFree86 is a trademark of The XFree86 Project, Inc.

Ogg Vorbis and Xiph.Org are trademarks of Xiph.Org.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this document, and the FreeBSD Project was aware of the trademark claim, the designations have been followed by the “™” or the “®” symbol.

Table of Contents

Preface	xi
I. Getting Started	xvi
1 Introduction	1
1.1 Synopsis	1
1.2 Welcome to DragonFly!.....	1
1.3 About the DragonFly Project.....	3
2 Installation from CD.....	5
2.1 CD Installation Overview	5
2.2 CD Installation - Making room.....	5
2.3 CD Installation - Disk setup	6
2.4 Installing to Disk from CD	8
2.5 CD Installation - Post-install cleanup.....	8
2.6 CD Installation - New system setup	10
3 UNIX Basics	12
3.1 Synopsis.....	12
3.2 Virtual Consoles and Terminals.....	12
3.3 Permissions	15
3.4 Directory Structure	18
3.5 Disk Organization.....	19
3.6 Mounting and Unmounting File Systems	24
3.7 Processes.....	27
3.8 Daemons, Signals, and Killing Processes.....	28
3.9 Shells.....	30
3.10 Text Editors.....	32
3.11 Devices and Device Nodes	32
3.12 Binary Formats	33
3.13 For More Information.....	34
4 Installing Applications using NetBSD's pkgsrc framework	37
4.1 Synopsis.....	37
4.2 Overview of Software Installation	37
4.3 Finding Your Application	39
4.4 Using the Binary Packages System	39
4.5 Using the pkgsrc® Source Tree.....	41
4.6 Post-installation Activities.....	47
4.7 Dealing with Broken Packages	48
5 The X Window System	49
5.1 Synopsis.....	49
5.2 Understanding X.....	49
5.3 Installing X11	52
5.4 X11 Configuration	52
5.5 Using Fonts in X11	55
5.6 The X Display Manager.....	59
5.7 Desktop Environments.....	62

II. System Administration	66
6 Configuration and Tuning.....	67
6.1 Synopsis.....	67
6.2 Initial Configuration.....	67
6.3 Core Configuration.....	68
6.4 Application Configuration.....	69
6.5 Starting Services.....	69
6.6 Configuring the cron Utility.....	71
6.7 Using rc under DragonFly.....	72
6.8 Setting Up Network Interface Cards.....	74
6.9 Virtual Hosts.....	78
6.10 Configuration Files.....	79
6.11 Tuning with sysctl.....	82
6.12 Tuning Disks.....	83
6.13 Tuning Kernel Limits.....	86
6.14 Adding Swap Space.....	88
6.15 Power and Resource Management.....	89
6.16 Using and Debugging DragonFly ACPI.....	91
7 The DragonFly Booting Process.....	97
7.1 Synopsis.....	97
7.2 The Booting Problem.....	97
7.3 The Boot Manager and Boot Stages.....	98
7.4 Kernel Interaction During Boot.....	102
7.5 Init: Process Control Initialization.....	102
7.6 Shutdown Sequence.....	103
8 Users and Basic Account Management.....	105
8.1 Synopsis.....	105
8.2 Introduction.....	105
8.3 The Superuser Account.....	106
8.4 System Accounts.....	107
8.5 User Accounts.....	107
8.6 Modifying Accounts.....	107
8.7 Limiting Users.....	112
8.8 Personalizing Users.....	114
8.9 Groups.....	114
9 Configuring the DragonFly Kernel.....	116
9.1 Synopsis.....	116
9.2 Why Build a Custom Kernel?.....	116
9.3 Building and Installing a Custom Kernel.....	117
9.4 The Configuration File.....	118
9.5 Making Device Nodes.....	129
9.6 If Something Goes Wrong.....	130
10 Security.....	132
10.1 Synopsis.....	132
10.2 Introduction.....	132
10.3 Securing DragonFly.....	134
10.4 DES, MD5, and Crypt.....	140
10.5 One-time Passwords.....	141

10.6	Kerberos5	146
10.7	Firewalls.....	153
10.8	OpenSSL.....	163
10.9	VPN over IPsec.....	163
10.10	OpenSSH	173
11	Printing	179
11.1	Synopsis.....	179
11.2	Introduction.....	179
11.3	Basic Setup	180
11.4	Advanced Printer Setup	192
11.5	Using Printers	217
11.6	Alternatives to the Standard Spooler	224
11.7	Troubleshooting.....	224
12	Storage.....	228
12.1	Synopsis.....	228
12.2	Device Names	228
12.3	Adding Disks	229
12.4	RAID.....	230
12.5	Creating and Using Optical Media (CDs)	234
12.6	Creating and Using Optical Media (DVDs)	239
12.7	Creating and Using Floppy Disks.....	243
12.8	Creating and Using Data Tapes	245
12.9	Backups to Floppies.....	247
12.10	Backup Basics.....	248
12.11	Network, Memory, and File-Backed File Systems	255
12.12	File System Quotas.....	257
13	The Vinum Volume Manager	260
13.1	Synopsis.....	260
13.2	Disks Are Too Small.....	260
13.3	Access Bottlenecks	260
13.4	Data Integrity	262
13.5	Vinum Objects	263
13.6	Some Examples	265
13.7	Object Naming.....	271
13.8	Configuring Vinum.....	274
13.9	Using Vinum for the Root Filesystem	275
14	Localization - I18N/L10N Usage and Setup	281
14.1	Synopsis.....	281
14.2	The Basics.....	281
14.3	Using Localization.....	282
14.4	Compiling I18N Programs.....	287
14.5	Localizing DragonFly to Specific Languages.....	287
15	Desktop Applications	291
15.1	Synopsis.....	291
15.2	Browsers	291
15.3	Productivity.....	294
15.4	Document Viewers.....	297
15.5	Finance.....	298

15.6 Summary.....	300
16 Multimedia	301
16.1 Synopsis.....	301
16.2 Setting Up the Sound Card	301
16.3 MP3 Audio.....	305
16.4 Video Playback	308
16.5 Setting Up TV Cards	316
17 Serial Communications	319
17.1 Synopsis.....	319
17.2 Introduction.....	319
17.3 Terminals	323
17.4 Dial-in Service.....	327
17.5 Dial-out Service.....	335
17.6 Setting Up the Serial Console.....	338
18 PPP and SLIP	346
18.1 Synopsis.....	346
18.2 Using User PPP.....	346
18.3 Using Kernel PPP	358
18.4 Troubleshooting PPP Connections	365
18.5 Using PPP over Ethernet (PPPoE).....	368
18.6 Using SLIP.....	370
19 Advanced Networking.....	380
19.1 Synopsis.....	380
19.2 Gateways and Routes.....	380
19.3 Wireless Networking	386
19.4 Bluetooth.....	391
19.5 Bridging.....	399
19.6 NFS	401
19.7 Diskless Operation.....	406
19.8 ISDN	413
19.9 NIS/YP.....	417
19.10 DHCP.....	432
19.11 DNS	436
19.12 NTP.....	447
19.13 Network Address Translation	450
19.14 The inetd “Super-Server”.....	453
19.15 Parallel Line IP (PLIP)	457
19.16 IPv6.....	459
20 Electronic Mail.....	463
20.1 Synopsis.....	463
20.2 Using Electronic Mail.....	463
20.3 sendmail Configuration.....	465
20.4 Changing Your Mail Transfer Agent	468
20.5 Troubleshooting.....	470
20.6 Advanced Topics.....	472
20.7 SMTP with UUCP	474
20.8 Setting up to send only	476
20.9 Using Mail with a Dialup Connection.....	477

20.10 SMTP Authentication	478
20.11 Mail User Agents	479
20.12 Using fetchmail.....	486
20.13 Using procmail.....	487
21 Updating DragonFly.....	489
21.1 Initial Setup.....	489
21.2 Configuration	489
21.3 Preparing to Update	489
21.4 Updating the System.....	489
22 Linux Binary Compatibility	491
22.1 Synopsis.....	491
22.2 Installation	491
22.3 Installing Mathematica®	494
22.4 Installing Maple™	496
22.5 Installing MATLAB®.....	498
22.6 Installing Oracle®	501
22.7 Installing SAP® R/3®.....	504
22.8 Advanced Topics.....	524
III. Appendices	527
A. Obtaining DragonFly.....	528
A.1 CDROM and DVD Publishers	528
A.2 FTP Sites.....	528
A.3 Using CVSup	529
A.4 CVS Tags	536
B. Bibliography	537
B.1 Books & Magazines Specific to BSD	537
B.2 Users' Guides.....	538
B.3 Administrators' Guides	538
B.4 Programmers' Guides	539
B.5 Operating System Internals.....	539
B.6 Security Reference	540
B.7 Hardware Reference.....	540
B.8 UNIX History.....	540
B.9 Magazines and Journals	541
C. Resources on the Internet	542
C.1 Mailing Lists	542
C.2 Usenet Newsgroups.....	544
C.3 World Wide Web Servers	545
D. PGP Keys.....	546
D.1 Developers	546
Colophon.....	547

List of Tables

3-1. Disk Device Codes	24
12-1. Physical Disk Naming Conventions	228
13-1. Vinum Plex Organizations	264
19-1. Wiring a Parallel Cable for Networking	457
19-2. Reserved IPv6 addresses	460

Preface

Intended Audience

The DragonFly newcomer will find that the first section of this book guides the user through the DragonFly installation process and gently introduces the concepts and conventions that underpin UNIX®. Working through this section requires little more than the desire to explore, and the ability to take on board new concepts as they are introduced.

Once you have travelled this far, the second, far larger, section of the Handbook is a comprehensive reference to all manner of topics of interest to DragonFly system administrators. Some of these chapters may recommend that you do some prior reading, and this is noted in the synopsis at the beginning of each chapter.

For a list of additional sources of information, please see [Appendix B](#).

Organization of This Book

This book is split into three logically distinct sections. The first section, *Getting Started*, covers the installation and basic usage of DragonFly. It is expected that the reader will follow these chapters in sequence, possibly skipping chapters covering familiar topics. The second section, *System Administration*, covers a broad collection of subjects that are of interest to more advanced DragonFly users. Each section begins with a succinct synopsis that describes what the chapter covers and what the reader is expected to already know. This is meant to allow the casual reader to skip around to find chapters of interest. The third section contains appendices of reference information.

Chapter 1, Introduction

Introduces DragonFly to a new user. It describes the history of the DragonFly Project, its goals and development model.

Chapter 2, Installation

Walks a user through the entire installation process. Some advanced installation topics, such as installing through a serial console, are also covered.

Chapter 3, UNIX Basics

Covers the basic commands and functionality of the DragonFly operating system. If you are familiar with Linux or another flavor of UNIX then you can probably skip this chapter.

Chapter 4, Installing Applications

Covers the installation of third-party software using NetBSD®'s Packages Collection pkgsrc®.

Chapter 5, The X Window System

Describes the X Window System in general and using **XFree86™** and **X.org** on DragonFly in particular. Also describes common desktop environments such as **KDE** and **GNOME**.

Chapter 6, Configuration and Tuning

Describes the parameters available for system administrators to tune a DragonFly system for optimum performance. Also describes the various configuration files used in DragonFly and where to find them.

Chapter 7, Booting Process

Describes the DragonFly boot process and explains how to control this process with configuration options.

Chapter 8, Users and Basic Account Management

Describes the creation and manipulation of user accounts. Also discusses resource limitations that can be set on users and other account management tasks.

Chapter 9, Configuring the DragonFly Kernel

Explains why you might need to configure a new kernel and provides detailed instructions for configuring, building, and installing a custom kernel.

Chapter 10, Security

Describes many different tools available to help keep your DragonFly system secure, including Kerberos, IPsec, OpenSSH, and network firewalls.

Chapter 11, Printing

Describes managing printers on DragonFly, including information about banner pages, printer accounting, and initial setup.

Chapter 12, Storage

Describes how to manage storage media and filesystems with DragonFly. This includes physical disks, RAID arrays, optical and tape media, memory-backed disks, and network filesystems.

Chapter 13, Vinum

Describes how to use Vinum, a logical volume manager which provides device-independent logical disks, and software RAID-0, RAID-1 and RAID-5.

Chapter 14, Localization

Describes how to use DragonFly in languages other than English. Covers both system and application level localization.

Chapter 15, Desktop Applications

Lists some common desktop applications, such as web browsers and productivity suites, and describes how to install them on DragonFly.

Chapter 16, Multimedia

Shows how to set up sound and video playback support for your system. Also describes some sample audio and video applications.

Chapter 17, Serial Communications

Explains how to connect terminals and modems to your DragonFly system for both dial in and dial out connections.

Chapter 18, PPP and SLIP

Describes how to use PPP, SLIP, or PPP over Ethernet to connect to remote systems with DragonFly.

Chapter 19, Advanced Networking

Describes many networking topics, including sharing an Internet connection with other computers on your LAN, using network filesystems, sharing account information via NIS, setting up a name server, and much more.

Chapter 20, Electronic Mail

Explains the different components of an email server and dives into simple configuration topics for the most popular mail server software: **sendmail**.

Section 21.1, Updating DragonFly

Describes the development paths of DragonFly, and how to stay up-to-date.

Chapter 22, Linux Binary Compatibility

Describes the Linux compatibility features of DragonFly. Also provides detailed installation instructions for many popular Linux applications such as **Oracle®**, **SAP® R/3®**, and **Mathematica®**.

Appendix A, Obtaining DragonFly

Lists different sources for obtaining DragonFly media on CDROM or DVD as well as different sites on the Internet that allow you to download and install DragonFly.

Appendix B, Bibliography

This book touches on many different subjects that may leave you hungry for a more detailed explanation. The bibliography lists many excellent books that are referenced in the text.

Appendix C, Resources on the Internet

Describes the many forums available for DragonFly users to post questions and engage in technical conversations about DragonFly.

Appendix D, PGP Keys

Lists the PGP fingerprints of several DragonFly Developers.

Conventions used in this book

To provide a consistent and easy to read text, several conventions are followed throughout the book.

Typographic Conventions

Italic

An *italic* font is used for filenames, URLs, emphasized text, and the first usage of technical terms.

Monospace

A monospaced font is used for error messages, commands, environment variables, names of ports, hostnames, user names, group names, device names, variables, and code fragments.

Bold

A **bold** font is used for applications, commands, and keys.

User Input

Keys are shown in **bold** to stand out from other text. Key combinations that are meant to be typed simultaneously are shown with ‘+’ between the keys, such as:

Ctrl+Alt+Del

Meaning the user should type the **Ctrl**, **Alt**, and **Del** keys at the same time.

Keys that are meant to be typed in sequence will be separated with commas, for example:

Ctrl+X, Ctrl+S

Would mean that the user is expected to type the **Ctrl** and **X** keys simultaneously and then to type the **Ctrl** and **S** keys simultaneously.

Examples

Examples starting with # indicate a command that must be invoked as the superuser in DragonFly. You can login as `root` to type the command, or login as your normal account and use `su(1)` to gain superuser privileges.

```
# dd if=kern.flp of=/dev/fd0
```

Examples starting with % indicate a command that should be invoked from a normal user account. Unless otherwise noted, C-shell syntax is used for setting environment variables and other shell commands.

```
% top
```

Examples starting with E:\> indicate a MS-DOS® command. Unless otherwise noted, these commands may be executed from a “Command Prompt” window in a modern Microsoft® Windows® environment.

```
E:\> tools\fdimage floppies\kern.flp A:
```

Acknowledgments

The book you are holding represents the efforts of many hundreds of people around the world. Whether they sent in fixes for typos, or submitted complete chapters, all the contributions have been useful.

The DragonFly Handbook was originally built from an edition of the FreeBSD Handbook. The FreeBSD Handbook was created by the collective hard work of hundreds of people, and the DragonFly Documentation Team appreciates all their labor. Included here is a list of all individually identified people and corporations that contributed resources to this handbook.

Eric Anderson, Satoshi Asami, Bojan Bistrovic, Neil Blakey-Milner, Andrew Boothman, Harti Brandt, Jim Brown, BSDi, Andrey A. Chernov, Peter Childs, Munish Chopra, Joe Marcus Clarke, Nik Clayton, Mark Dapoz, Matt Dillon, Jean-François Dockès, Alex Dupre, Josef El-Rayes, Udo Erdelhoff, Marc Fonvieille, Dirk Frömberg, Robert Getschmann, James Gorham, Lucky Green, Coranth Gryphon, Jake Hamby, Brian N. Handy, Guy Helmer, Al Hoang, Tillman Hodgson, Jordan Hubbard, Robert Huff, Tom Hukins, Christophe Juniet, Poul-Henning Kamp, Aaron Kaplan, Martin Karlsson, Sean Kelly, Seth Kingsley, Holger Kipp, Nate Lawson, Chern Lee, Greg Lehey, John Lind, Ross Lippert, Bill Lloyd, Pav Lucistnik, Julio Merino, Mike Meyer, Hellmuth Michaelis, Jim Mock,

Marcel Moolenaar, Moses Moore, Bill Moran, Rich Murphey, Mark Murray, Alex Nash, Gregory Neil Shapiro, David O'Brien, Eric Ogren, Gary Palmer, Hiten M. Pandya, Bill Paul, Dan Pelleg, Steve Peterson, John Polstra, Andy Polyakov, Randy Pratt, Jeremy C. Reed, Tom Rhodes, Trev Roydhouse, Peter Schultz, Piero Serini, Christopher Shumway, Marc Silver, Mike Smith, Brian Somers, Gennady B. Sorokopud, Wylie Stilwell, Murray Stokely, Greg Sutter, Bill Swingle, Valentino Vaschetto, Robert Watson, Wind River Systems (<http://www.windriver.com>), Michael C. Wu, and Kazutaka YOKOTA.

I. Getting Started

This part of the DragonFly Handbook is for users and administrators who are new to DragonFly. These chapters:

- Introduce you to DragonFly.
- Guide you through the installation process.
- Teach you UNIX basics and fundamentals.
- Show you how to install the wealth of third party applications available for DragonFly.
- Introduce you to X, the UNIX windowing system, and detail how to configure a desktop environment that makes you more productive.

We have tried to keep the number of forward references in the text to a minimum so that you can read this section of the Handbook from front to back with the minimum page flipping required.

Chapter 1 Introduction

Restructured, reorganized, and parts rewritten by Jim Mock.

1.1 Synopsis

Thank you for your interest in DragonFly! The following chapter covers various aspects of the DragonFly Project, such as its history, goals, development model, and so on.

After reading this chapter, you will know:

- How DragonFly relates to other computer operating systems.
- The history of the DragonFly Project.
- The goals of the DragonFly Project.
- The basics of the DragonFly open-source development model.
- And of course: where the name “DragonFly” comes from.

1.2 Welcome to DragonFly!

DragonFly is a 4.4BSD-Lite based operating system for Intel (x86). A port to AMD64 is in progress. You can also read about the history of DragonFly, or the current release.

1.2.1 What Can DragonFly Do?

DragonFly has many noteworthy features. Some of these are:

- *Preemptive multitasking* with dynamic priority adjustment to ensure smooth and fair sharing of the computer between applications and users, even under the heaviest of loads.
- *Multi-user facilities* which allow many people to use a DragonFly system simultaneously for a variety of things. This means, for example, that system peripherals such as printers and tape drives are properly shared between all users on the system or the network and that individual resource limits can be placed on users or groups of users, protecting critical system resources from over-use.
- Strong *TCP/IP networking* with support for industry standards such as SLIP, PPP, NFS, DHCP, and NIS. This means that your DragonFly machine can interoperate easily with other systems as well as act as an enterprise server, providing vital functions such as NFS (remote file access) and email services or putting your organization on the Internet with WWW, FTP, routing and firewall (security) services.
- *Memory protection* ensures that applications (or users) cannot interfere with each other. One application crashing will not affect others in any way.
- DragonFly is a *32-bit* operating system.
- The industry standard *X Window System (X11R6)* provides a graphical user interface (GUI) for the cost of a common VGA card and monitor and comes with full sources.
- *Binary compatibility* with many programs built for Linux, SCO, SVR4, BSDI and NetBSD.

- Thousands of *ready-to-run* applications are available from the *pkgsrc packages* collection. Why search the net when you can find it all right here?
- Thousands of additional and *easy-to-port* applications are available on the Internet. DragonFly is source code compatible with most popular commercial UNIX systems and thus most applications require few, if any, changes to compile.
- Demand paged *virtual memory* and “merged VM/buffer cache” design efficiently satisfies applications with large appetites for memory while still maintaining interactive response to other users.
- *SMP* support for machines with multiple CPUs.
- A full complement of *C*, *C++*, *Fortran*, and *Perl* development tools. Many additional languages for advanced research and development are also available in the ports and packages collection.
- *Source code* for the entire system means you have the greatest degree of control over your environment. Why be locked into a proprietary solution at the mercy of your vendor when you can have a truly open system?
- Extensive *online documentation*.
- *And many more!*

DragonFly is based on the 4.4BSD-Lite release from Computer Systems Research Group (CSRG) at the University of California at Berkeley, along with later development of FreeBSD by the FreeBSD Project. It carries on the distinguished tradition of BSD systems development. In addition to the fine work provided by CSRG, the DragonFly Project has put in many thousands of hours in fine tuning the system for maximum performance and reliability in real-life load situations. As many of the commercial giants struggle to field PC operating systems with such features, performance and reliability, DragonFly can offer them *now!*

The applications to which DragonFly can be put are truly limited only by your own imagination. From software development to factory automation, inventory control to azimuth correction of remote satellite antennae; if it can be done with a commercial UNIX product then it is more than likely that you can do it with DragonFly too! DragonFly also benefits significantly from literally thousands of high quality applications developed by research centers and universities around the world, often available at little to no cost. Commercial applications are also available and appearing in greater numbers every day.

Because the source code for DragonFly itself is generally available, the system can also be customized to an almost unheard of degree for special applications or projects, and in ways not generally possible with operating systems from most major commercial vendors. Here is just a sampling of some of the applications in which people are currently using DragonFly:

- *Internet Services*: The robust TCP/IP networking built into DragonFly makes it an ideal platform for a variety of Internet services such as:
 - FTP servers
 - World Wide Web servers (standard or secure [SSL])
 - Firewalls and NAT (“IP masquerading”) gateways
 - Electronic Mail servers
 - USENET News or Bulletin Board Systems
 - And more...

With DragonFly, you can easily start out small with an inexpensive 386 class PC and upgrade all the way up to a quad-processor Xeon with RAID storage as your enterprise grows.

- *Education:* Are you a student of computer science or a related engineering field? There is no better way of learning about operating systems, computer architecture and networking than the hands on, under the hood experience that DragonFly can provide. A number of freely available CAD, mathematical and graphic design packages also make it highly useful to those whose primary interest in a computer is to get *other* work done!
- *Research:* With source code for the entire system available, DragonFly is an excellent platform for research in operating systems as well as other branches of computer science. DragonFly's freely available nature also makes it possible for remote groups to collaborate on ideas or shared development without having to worry about special licensing agreements or limitations on what may be discussed in open forums.
- *Networking:* Need a new router? A name server (DNS)? A firewall to keep people out of your internal network? DragonFly can easily turn that unused older PC sitting in the corner into an advanced router with sophisticated packet-filtering capabilities.
- *X Window workstation:* DragonFly is a fine choice for an inexpensive X terminal solution, either using the freely available XFree86 or X.org servers or one of the excellent commercial servers provided by Xi Graphics (<http://www.xig.com>). Unlike an X terminal, DragonFly allows many applications to be run locally if desired, thus relieving the burden on a central server. DragonFly can even boot "diskless", making individual workstations even cheaper and easier to administer.
- *Software Development:* The basic DragonFly system comes with a full complement of development tools including the renowned GNU C/C++ compiler and debugger.

DragonFly is available via anonymous FTP or CVS. Please see Appendix A for more information about obtaining DragonFly.

1.3 About the DragonFly Project

The following section provides some background information on the project, including a brief history, project goals, and the development model of the project.

1.3.1 A Brief History of DragonFly

Matthew Dillon, one of the developers for FreeBSD, was growing increasingly frustrated with the FreeBSD Project's direction for release 5. The FreeBSD 5 release had been delayed multiple times, and had performance problems compared to earlier releases of FreeBSD.

DragonFly was announced in June of 2003. The code base was taken from the 4.8 release of FreeBSD, which offered better performance and more complete features.

Development has proceeded at a very quick rate since then, with Matt Dillon and a small group of developers fixing longstanding BSD bugs and modernizing the new DragonFly system.

1.3.2 DragonFly Project Goals

DragonFly is an effort to maintain the traditional BSD format — lean, stable code — along with modern features such as lightweight threads, a workable packaging system, and a revised VFS. Underpinning all this work is efficient support for multiple processors, something rare among open source systems. Because DragonFly is built on an existing very stable code base, it is possible to make these radical changes as part of an incremental process.

1.3.3 The DragonFly Development Model

Written by Justin Sherrill.

DragonFly is developed by many people around the world. There is no qualification process; anyone may submit his or her code, documentation, or designs, for use in the Project. Here is a general description of the Project's organizational structure.

The CVS repository

Source for DragonFly is kept in CVS (<http://www.cvshome.org/>) (Concurrent Versions System), which is available with each DragonFly install. The primary CVS repository (<http://www.dragonflybsd.org/cvsweb>) resides on a machine in California, USA. Documentation on obtaining the DragonFly source is available elsewhere in this book.

Commit access

The best way of getting changes made to the DragonFly source is to mail the submit (<http://leaf.dragonflybsd.org/mailarchive/>) mailing list. Including desired source code changes (unified diff format is best) is the most useful format. A certain number of developers have access to commit changes to the DragonFly source, and can do so after review on that list.

The DragonFly development model is loose; changes to the code are generally peer-reviewed and added when any objections have been corrected. There is no formal entry/rejection process, though final say on all code submissions goes to Matt Dillon, as originator of this project.

1.3.4 The Current DragonFly Release

DragonFly is a freely available, full source 4.4BSD-Lite based release for Intel i386™, i486™, Pentium®, Pentium Pro, Celeron®, Pentium II, Pentium III, Pentium 4 (or compatible), and Xeon™ based computer systems. It is based primarily on FreeBSD 4.8, and includes enhancements from U.C. Berkeley's CSRG group, NetBSD, OpenBSD, 386BSD, and the Free Software Foundation.

A number of additional documents which you may find very helpful in the process of installing and using DragonFly may now also be found in the `/usr/share/doc` directory on any machine.

The most up-to-date documentation can be found at <http://www.dragonflybsd.org/> (`../../../../`).

1.3.5 "DragonFly" Origin

Matthew Dillon happened to take a picture of a dragonfly in his garden while trying to come up with a name for this new branch of BSD. Taking this as inspiration, "DragonFly" became the new name.

Chapter 2 Installation from CD

Written by Markus Schatzl and Justin Sherrill.

2.1 CD Installation Overview

This document describes the installation of DragonFly BSD on a plain i386 machine. This process uses a bootable DragonFly CD, usually referred to as a 'live CD'.

This CD is available at one of the current mirrors, which distribute the images by various protocols. The authoritative list can be found at the DragonFly website (<http://www.dragonflybsd.org/main/download.cgi>).

This document may be superseded by the /README file located on the live CD, which may reflect changes made after this document was last updated. Check that README for any last-minute changes and for an abbreviated version of this installation process.

The DragonFly development team is working on an automatic installation tool, which simplifies the partitioning and installation processes. Until this tool is in place, the manual process here is required. Some experience with BSD-style tools is recommended.

Caution: While this guide covers installing to a computer with an existing non-DragonFly operating system, take no chances! Back up any data on your disk drives that you want to save.

When installing to an old machine, it may not be possible to boot from a CD. Use a bootmanager on a floppy in those cases, such as Smart Bootmanager (<http://btmgr.sourceforge.net/>).

Caution: Always be sure of the target disk for any command. Unless otherwise specified, each command here assumes the first disk in the IDE chain is the target. (ad0) Adjust commands as needed.

2.2 CD Installation - Making room

2.2.1 DragonFly as the only operating system

If DragonFly is to be the only operating system on the target computer, preparing the disk is a short and simple process. Boot with the live CD, and log in as root to reach a command prompt.

First, the master boot record (MBR) must be cleared of any old information. This command clears all old data off your disk by writing zeros (if=/dev/zero) onto the system's master ata drive (of=/dev/ad0).

```
# dd if=/dev/zero of=/dev/ad0 bs=32k count=16
```

The now-empty disk must be formatted.

Important: This will destroy any existing data on a disk. Do this only if you plan to dedicate this disk to DragonFly.

```
# fdisk -I ad0
# fdisk -B ad0
```

2.2.2 Multiple operating systems on one hard disk

This example assumes that the target computer for installation has at least one operating system installed that needs to survive the installation process. A new partition for DragonFly needs to be created from the existing partition(s) that otherwise fill the disk. There must be unused space within the existing partition in order to resize it.

Important: The new partition is created from empty space in an existing partition. For example, an 18 gigabyte disk that has 17 gigabytes of existing data in the existing partition will only have 1 gigabyte available for the new partition.

Partition resizing needs to be accomplished with a third-party tool. Commercial programs such as Partition Magic (<http://www.symantec.com/partitionmagic/>) can accomplish these tasks. Free tools exist that can be adapted to this task, such as 'GNU parted', found on the Knoppix CD (<http://www.knopper.net/knoppix-mirrors/index-en.html>), or PAUD (<http://paud.sourceforge.net>).

Create a new partition of at least 5-6 gigabytes. It is possible to install within a smaller amount of disk space, but this will create problems not covered by this document. The newly created partition does not need to be formatted; the rest of the installation process treats that new partition as a new disk.

2.2.3 Multiple operating systems, multiple hard disks

Installing DragonFly to a separate disk removes the need for partition resizing, and is generally safer when trying to preserve an existing operating system installation.

This type of installation is very similar to installing DragonFly as the only operating system. The only difference is the disk named in each command.

2.3 CD Installation - Disk setup

2.3.1 Disk formatting

The slice layout on the newly formatted disk or partition needs to be set up, using this command.

```
# fdisk -u
```

If there are multiple operating systems on the disk, pick the correct partition judging by what partitions were created earlier with a resizing tool.

2.3.2 Boot block installation

The 'ad0' here refers to the first disk on the first IDE bus of a computer. Increment the number if the target disk is farther down the chain. For example, the master disk on the second IDE controller would be 'ad2'.

```
# boot0cfg -B ad0
# boot0cfg -v ad0
```

-s **SLICE**, where SLICE is a number, controls which slice on disk is used by boot0cfg to start from. By default, this number is 1, and will only need modification if a different slice contains DragonFly.

Use -o **packet** as an option to boot0cfg if the DragonFly partition is located beyond cylinder 1023 on the disk. This location problem usually only happens when another operating system is taking up more than the first 8 gigabytes of disk space. This problem cannot happen if DragonFly is installed to a dedicated disk

2.3.3 Disklabel

If DragonFly is installed anywhere but the first partition of the disk, the device entry for that partition will have to be created. Otherwise, the device entry is automatically created. Refer to this different partition instead of the 'ad0s1a' used in later examples.

```
# cd /dev; ./MAKEDEV ad0s2
```

The partition needs to be created on the DragonFly disk.

```
# disklabel -B -r -w ad0s1 auto
```

Using /etc/disklabel.ad0s1 as an example, issue the following command to edit the disklabel for the just-created partition.

```
# disklabel -e ad0s1
```

Partition	Size	Mountpoint
ad0s2a	256m	/
ad0s2b	1024m	swap
ad0s2c	leave alone	This represents the whole slice.
ad0s2d	256m	/var
ad0s2e	256m	/tmp !
ad0s2f	8192m	/usr - This should be at least 4096m
ad0s2g	*	/home - This holds 'everything else'

2.3.4 Partition Format

newfs will format each individual partition.

```
# newfs /dev/ad0s1a
# newfs -U /dev/ad0s1d
# newfs -U /dev/ad0s1e
# newfs -U /dev/ad0s1f
# newfs -U /dev/ad0s1g
```

Note: The -U option is not used for the root partition, since / is usually relatively small. Softupdates can cause it to run out of space while under a lot of disk activity, such as a buildworld.

Note: The command listing skips directly from ad0s1a to ad0s1d. This is because /dev/ad0s1b is used as swap and does not require formatting; ad0s1c refers to the entire disk and should not be formatted.

2.4 Installing to Disk from CD

Since the Live CD contains all needed data to create a running DragonFly system, the simplest installation possible is to copy the Live CD data to the newly formatted disk/partition created in previous steps.

These commands mount the newly created disk space and create the appropriate directories on it.

```
# mount /dev/ad0s1a /mnt
# mkdir /mnt/var
# mkdir /mnt/tmp
# mkdir /mnt/usr
# mkdir /mnt/home
# mount /dev/ad0s1d /mnt/var
# mount /dev/ad0s1e /mnt/tmp
# mount /dev/ad0s1f /mnt/usr
# mount /dev/ad0s1g /mnt/home
```

cpdup duplicates data from one volume to another. These commands copy data from the Live CD to the newly created directories on the mounted disk. Each step can take some time, depending on disk speed.

```
# cpdup / /mnt
# cpdup /var /mnt/var
# cpdup /etc /mnt/etc
# cpdup /dev /mnt/dev
# cpdup /usr /mnt/usr
```

Note: Nothing is copied to the /tmp directory that was created in the previous step. This is not an error, since /tmp is intended only for temporary storage.

2.5 CD Installation - Post-install cleanup

/tmp and /var/tmp are both often used as temporary directories. Since use is not consistent from application to application, it is worthwhile to create /tmp as a link to /var/tmp so space is not wasted in duplication.

```
# chmod 1777 /mnt/tmp
# rm -fr /mnt/var/tmp
# ln -s /tmp /mnt/var/tmp
```

Note: /tmp will not work until the computer is rebooted.

The file /etc/fstab describes the disk partition layout. However, the version copied to the target disk only reflects the Live CD layout. The installed /mnt/fstab.example can be used as a starting point for creating a new /etc/fstab.

```
# vi /mnt/etc/fstab.example
# mv /mnt/etc/fstab.example /mnt/etc/fstab
```

A corrupted disklabel will render a disk useless. While this is thankfully very rare, having a backup of the new install's disklabel may stave off disaster at some point in the future. This is optional. (Adjust the slice name to reflect the actual installation.)

```
# disklabel ad0s1 > /mnt/etc/disklabel.backup
```

Note: Nothing is copied to the /tmp directory that was created in the previous step. This is not an error, since /tmp is intended only for temporary storage.

Remove some unnecessary files copied over from the Live CD.

```
# rm /mnt/boot/loader.conf
# rm /mnt/boot.catalog
# rm -r /mnt/rr_moved
```

The system can now be rebooted. Be sure to remove the Live CD from the CDROM drive so that the computer can boot from the newly-installed disk.

```
# reboot
```

Note: Use the reboot command so that the disk can be unmounted cleanly. Hitting the power or reset buttons, while it won't hurt the Live CD, can leave the mounted disk in a inconsistent state.

If the system refuses to boot, there are several options to try:

- Old bootblocks can interfere with the initialization-process. To avoid this, zero-out the MBR. "of" should be changed to the correct disk entry if ad0 is not the targeted installation disk.

```
# dd if=/dev/zero of=/dev/ad0 bs=32 count=16
```

- It is possible that the DragonFly slice is beyond cylinder 1023 on the hard disk, and can't be detected. Packet mode can fix this problem.

```
# boot0cfg -o packet ad0
```

- If you can select CHS or LBA mode in your BIOS, try changing the mode to LBA.

After a successful boot from the newly installed hard drive, the timezone should be set. Use the command `tzsetup` to set the appropriate time zone.

```
# tzsetup
```

2.6 CD Installation - New system setup

Once the new DragonFly system is booting from disk, there are a number of steps that may be useful before working further. The file `/etc/rc.conf` controls a number of options for booting the system.

2.6.1 Setting up rc.conf

Depending on location, a different keyboard map may be needed. This is only necessary for computers outside of North America.

```
# kbdmap
```

Pick the appropriate keyboard map and remember the name. Place this name in `/etc/rc.conf`. For example:

```
keymap="german.iso.kbd"
```

The file `/etc/rc.conf` matches the one on the Live CD. Since it is now on an installed system are no longer running in a read-only environment, some changes should be made. Changes to this file will take effect after the next boot of the machine.

These lines can be removed.

```
syslogd_enable="NO"
```

```
xntpd_enable="NO"
```

```
cron_enable="NO"
```

For a system which uses USB, this line will need to be modified to a value of "YES":

```
usbd_enable="NO"
```

inetd controls various small servers like telnet or ftp. By default, all servers are off, and must be individually uncommented in `/etc/inetd.conf` to start them. This is optional.

```
inetd_enable="YES"           # Run the network daemon dispatcher (YES/NO).
inetd_program="/usr/sbin/inetd" # path to inetd, if you want a different one.
inetd_flags="-wW"           # Optional flags to inetd
```

2.6.2 Network Setup

For acquiring an IP address through DHCP, place this entry in `/etc/rc.conf`, using the appropriate card name. (ep0 is used as an example here.)

```
ifconfig_ep0="DHCP"
```

For a fixed IP, `/etc/rc.conf` requires a few more lines of data. (Again, ep0 is used as an example here.) Supply the correct local values for IP, netmask, and default router. The hostname should reflect what is entered in DNS for this computer.

```
ifconfig_ep0="inet 123.234.345.456 netmask 255.255.255.0"
hostname="myhostname"
defaultrouter="654.543.432.321"
```

Chapter 3 UNIX Basics

Rewritten by Chris Shumway.

3.1 Synopsis

The following chapter will cover the basic commands and functionality of the DragonFly operating system. Much of this material is relevant for any UNIX-like operating system. Feel free to skim over this chapter if you are familiar with the material. If you are new to DragonFly, then you will definitely want to read through this chapter carefully.

After reading this chapter, you will know:

- How to use the “virtual consoles” of DragonFly.
- How UNIX file permissions work along with understanding file flags in DragonFly.
- The default DragonFly file system layout.
- The DragonFly disk organization.
- How to mount and unmount file systems.
- What processes, daemons, and signals are.
- What a shell is, and how to change your default login environment.
- How to use basic text editors.
- What devices and device nodes are.
- What binary format is used under DragonFly.
- How to read manual pages for more information.

3.2 Virtual Consoles and Terminals

DragonFly can be used in various ways. One of them is typing commands to a text terminal. A lot of the flexibility and power of a UNIX operating system is readily available at your hands when using DragonFly this way. This section describes what “terminals” and “consoles” are, and how you can use them in DragonFly.

3.2.1 The Console

If you have not configured DragonFly to automatically start a graphical environment during startup, the system will present you with a login prompt after it boots, right after the startup scripts finish running. You will see something similar to:

```
Additional ABI support:.  
Local package initialization:.  
Additional TCP options:.
```

```
Fri Sep 20 13:01:06 EEST 2002
```

```
DragonFlyBSD/i386 (pc3.example.org) (ttyv0)
```

```
login:
```

The messages might be a bit different on your system, but you will see something similar. The last two lines are what we are interested in right now. The second last line reads:

```
DragonFlyBSD/i386 (pc3.example.org) (ttyv0)
```

This line contains some bits of information about the system you have just booted. You are looking at a “DragonFlyBSD” console, running on an Intel or compatible processor of the x86 architecture¹. The name of this machine (every UNIX machine has a name) is `pc3.example.org`, and you are now looking at its system console—the `ttyv0` terminal.

Finally, the last line is always:

```
login:
```

This is the part where you are supposed to type in your “username” to log into DragonFly. The next section describes how you can do this.

3.2.2 Logging into DragonFly

DragonFly is a multiuser, multiprocessing system. This is the formal description that is usually given to a system that can be used by many different people, who simultaneously run a lot of programs on a single machine.

Every multiuser system needs some way to distinguish one “user” from the rest. In DragonFly (and all the UNIX-like operating systems), this is accomplished by requiring that every user must “log into” the system before being able to run programs. Every user has a unique name (the “username”) and a personal, secret key (the “password”). DragonFly will ask for these two before allowing a user to run any programs.

Right after DragonFly boots and finishes running its startup scripts², it will present you with a prompt and ask for a valid username:

```
login:
```

For the sake of this example, let us assume that your username is `john`. Type `john` at this prompt and press **Enter**. You should then be presented with a prompt to enter a “password”:

```
login: john
Password:
```

Type in `john`’s password now, and press **Enter**. The password is *not echoed!* You need not worry about this right now. Suffice it to say that it is done for security reasons.

If you have typed your password correctly, you should by now be logged into DragonFly and ready to try out all the available commands.

You should see the MOTD or message of the day followed by a command prompt (a `#`, `$`, or `%` character). This indicates you have successfully logged into DragonFly.

3.2.3 Multiple Consoles

Running UNIX commands in one console is fine, but DragonFly can run many programs at once. Having one console where commands can be typed would be a bit of a waste when an operating system like DragonFly can run dozens of programs at the same time. This is where “virtual consoles” can be very helpful.

DragonFly can be configured to present you with many different virtual consoles. You can switch from one of them to any other virtual console by pressing a couple of keys on your keyboard. Each console has its own different output channel, and DragonFly takes care of properly redirecting keyboard input and monitor output as you switch from one virtual console to the next.

Special key combinations have been reserved by DragonFly for switching consoles³. You can use **Alt-F1**, **Alt-F2**, through **Alt-F8** to switch to a different virtual console in DragonFly.

As you are switching from one console to the next, DragonFly takes care of saving and restoring the screen output. The result is an “illusion” of having multiple “virtual” screens and keyboards that you can use to type commands for DragonFly to run. The programs that you launch on one virtual console do not stop running when that console is not visible. They continue running when you have switched to a different virtual console.

3.2.4 The `/etc/ttys` File

The default configuration of DragonFly will start up with eight virtual consoles. This is not a hardwired setting though, and you can easily customize your installation to boot with more or fewer virtual consoles. The number and settings of the virtual consoles are configured in the `/etc/ttys` file.

You can use the `/etc/ttys` file to configure the virtual consoles of DragonFly. Each uncommented line in this file (lines that do not start with a # character) contains settings for a single terminal or virtual console. The default version of this file that ships with DragonFly configures nine virtual consoles, and enables eight of them. They are the lines that start with `ttv`:

```
# name  getty                               type  status  comments
#
ttv0    "/usr/libexec/getty Pc"               cons25  on  secure
# Virtual terminals
ttv1    "/usr/libexec/getty Pc"               cons25  on  secure
ttv2    "/usr/libexec/getty Pc"               cons25  on  secure
ttv3    "/usr/libexec/getty Pc"               cons25  on  secure
ttv4    "/usr/libexec/getty Pc"               cons25  on  secure
ttv5    "/usr/libexec/getty Pc"               cons25  on  secure
ttv6    "/usr/libexec/getty Pc"               cons25  on  secure
ttv7    "/usr/libexec/getty Pc"               cons25  on  secure
ttv8    "/usr/pkg/xorg/bin/xdm -nodaemon"     xterm   off  secure
```

For a detailed description of every column in this file and all the options you can use to set things up for the virtual consoles, consult the `ttys(5)` manual page.

3.2.5 Single User Mode Console

A detailed description of what “single user mode” is can be found in Section 7.5.2. It is worth noting that there is only one console when you are running DragonFly in single user mode. There are no virtual consoles available. The

settings of the single user mode console can also be found in the `/etc/ttys` file. Look for the line that starts with `console`:

```
# name  getty                type    status    comments
#
# If console is marked "insecure", then init will ask for the root password
# when going to single-user mode.
console none                unknown off secure
```

Note: As the comments above the `console` line indicate, you can edit this line and change `secure` to `insecure`. If you do that, when DragonFly boots into single user mode, it will still ask for the `root` password.

Be careful when changing this to `insecure`. If you ever forget the `root` password, booting into single user mode is a bit involved. It is still possible, but it might be a bit hard for someone who is not very comfortable with the DragonFly booting process and the programs involved.

3.3 Permissions

DragonFly, being a direct descendant of BSD UNIX, is based on several key UNIX concepts. The first and most pronounced is that DragonFly is a multi-user operating system. The system can handle several users all working simultaneously on completely unrelated tasks. The system is responsible for properly sharing and managing requests for hardware devices, peripherals, memory, and CPU time fairly to each user.

Because the system is capable of supporting multiple users, everything the system manages has a set of permissions governing who can read, write, and execute the resource. These permissions are stored as three octets broken into three pieces, one for the owner of the file, one for the group that the file belongs to, and one for everyone else. This numerical representation works like this:

Value	Permission	Directory Listing
0	No read, no write, no execute	---
1	No read, no write, execute	--x
2	No read, write, no execute	-w-
3	No read, write, execute	-wx
4	Read, no write, no execute	r--
5	Read, no write, execute	r-x
6	Read, write, no execute	rw-
7	Read, write, execute	rwx

You can use the `-l` command line argument to `ls(1)` to view a long directory listing that includes a column with information about a file's permissions for the owner, group, and everyone else. For example, a `ls -l` in an arbitrary directory may show:

```
% ls -l
total 530
-rw-r--r-- 1 root  wheel    512 Sep  5 12:31 myfile
```

```
-rw-r--r-- 1 root wheel 512 Sep 5 12:31 otherfile
-rw-r--r-- 1 root wheel 7680 Sep 5 12:31 email.txt
...
```

Here is how the first column of `ls -l` is broken up:

```
-rw-r--r--
```

The first (leftmost) character tells if this file is a regular file, a directory, a special character device, a socket, or any other special pseudo-file device. In this case, the `-` indicates a regular file. The next three characters, `rw-` in this example, give the permissions for the owner of the file. The next three characters, `r--`, give the permissions for the group that the file belongs to. The final three characters, `r--`, give the permissions for the rest of the world. A dash means that the permission is turned off. In the case of this file, the permissions are set so the owner can read and write to the file, the group can read the file, and the rest of the world can only read the file. According to the table above, the permissions for this file would be `644`, where each digit represents the three parts of the file's permission.

This is all well and good, but how does the system control permissions on devices? DragonFly actually treats most hardware devices as a file that programs can open, read, and write data to just like any other file. These special device files are stored on the `/dev` directory.

Directories are also treated as files. They have read, write, and execute permissions. The executable bit for a directory has a slightly different meaning than that of files. When a directory is marked executable, it means it can be traversed into, that is, it is possible to “`cd`” (change directory) into it. This also means that within the directory it is possible to access files whose names are known (subject, of course, to the permissions on the files themselves).

In particular, in order to perform a directory listing, read permission must be set on the directory, whilst to delete a file that one knows the name of, it is necessary to have write *and* execute permissions to the directory containing the file.

There are more permission bits, but they are primarily used in special circumstances such as `setuid` binaries and sticky directories. If you want more information on file permissions and how to set them, be sure to look at the `chmod(1)` manual page.

3.3.1 Symbolic Permissions

Contributed by Tom Rhodes.

Symbolic permissions, sometimes referred to as symbolic expressions, use characters in place of octal values to assign permissions to files or directories. Symbolic expressions use the syntax of (who) (action) (permissions), where the following values are available:

Option	Letter	Represents
(who)	<code>u</code>	User
(who)	<code>g</code>	Group owner
(who)	<code>o</code>	Other
(who)	<code>a</code>	All (“world”)
(action)	<code>+</code>	Adding permissions
(action)	<code>-</code>	Removing permissions
(action)	<code>=</code>	Explicitly set permissions
(permissions)	<code>r</code>	Read

Option	Letter	Represents
(permissions)	w	Write
(permissions)	x	Execute
(permissions)	t	Sticky bit
(permissions)	s	Set UID or GID

These values are used with the `chmod(1)` command just like before, but with letters. For an example, you could use the following command to block other users from accessing `FILE`:

```
% chmod go= FILE
```

A comma separated list can be provided when more than one set of changes to a file must be made. For example the following command will remove the groups and “world” write permission on `FILE`, then it adds the execute permissions for everyone:

```
% chmod go-w,a+x FILE
```

3.3.2 DragonFly File Flags

Contributed by Tom Rhodes.

In addition to file permissions discussed previously, DragonFly supports the use of “file flags.” These flags add an additional level of security and control over files, but not directories.

These file flags add an additional level of control over files, helping to ensure that in some cases not even the `root` can remove or alter files.

File flags are altered by using the `chflags(1)` utility, using a simple interface. For example, to enable the system undeletable flag on the file `file1`, issue the following command:

```
# chflags sunlink
file1
```

And to disable the system undeletable flag, simply issue the previous command with “no” in front of the `sunlink`. Observe:

```
# chflags nosunlink
file1
```

To view the flags of this file, use the `ls(1)` with the `-lo` flags:

```
# ls -lo file1
```

The output should look like the following:

```
-rw-r--r-- 1 trhodes trhodes sunlnk 0 Mar 1 05:54
file1
```

Several flags may only added or removed to files by the `root` user. In other cases, the file owner may set these flags. It is recommended an administrator read over the `chflags(1)` and `chflags(2)` manual pages for more information.

3.4 Directory Structure

The DragonFly directory hierarchy is fundamental to obtaining an overall understanding of the system. The most important concept to grasp is that of the root directory, “/”. This directory is the first one mounted at boot time and it contains the base system necessary to prepare the operating system for multi-user operation. The root directory also contains mount points for every other file system that you may want to mount.

A mount point is a directory where additional file systems can be grafted onto the root file system. This is further described in Section 3.5. Standard mount points include `/usr`, `/var`, `/tmp`, `/mnt`, and `/cdrom`. These directories are usually referenced to entries in the file `/etc/fstab`. `/etc/fstab` is a table of various file systems and mount points for reference by the system. Most of the file systems in `/etc/fstab` are mounted automatically at boot time from the script `rc(8)` unless they contain the `noauto` option. Details can be found in Section 3.6.1.

A complete description of the file system hierarchy is available in `hier(7)`. For now, a brief overview of the most common directories will suffice.

Directory	Description
/	Root directory of the file system.
/bin/	User utilities fundamental to both single-user and multi-user environments.
/boot/	Programs and configuration files used during operating system bootstrap.
/boot/defaults/	Default bootstrapping configuration files; see <code>loader.conf(5)</code> .
/dev/	Device nodes; see <code>intro(4)</code> .
/etc/	System configuration files and scripts.
/etc/defaults/	Default system configuration files; see <code>rc(8)</code> .
/etc/mail/	Configuration files for mail transport agents such as <code>sendmail(8)</code> .
/etc/namedb/	named configuration files; see <code>named(8)</code> .
/etc/periodic/	Scripts that are run daily, weekly, and monthly, via <code>cron(8)</code> ; see <code>periodic(8)</code> .
/etc/ppp/	ppp configuration files; see <code>ppp(8)</code> .
/mnt/	Empty directory commonly used by system administrators as a temporary mount point.
/proc/	Process file system; see <code>procfs(5)</code> , <code>mount_procfs(8)</code> .
/root/	Home directory for the <code>root</code> account.
/sbin/	System programs and administration utilities fundamental to both single-user and multi-user environments.
/tmp/	Temporary files. The contents of <code>/tmp</code> are usually NOT preserved across a system reboot. A memory-based file system is often mounted at <code>/tmp</code> . This can be automated with an entry in <code>/etc/fstab</code> ; see <code>mysfs(8)</code> .
/usr/	The majority of user utilities and applications.
/usr/bin/	Common utilities, programming tools, and applications.
/usr/include/	Standard C include files.

Directory	Description
<code>/usr/lib/</code>	Archive libraries.
<code>/usr/libdata/</code>	Miscellaneous utility data files.
<code>/usr/libexec/</code>	System daemons & system utilities (executed by other programs).
<code>/usr/local/</code>	Local executables, libraries, etc. Within <code>/usr/local</code> , the general layout sketched out by <code>hier(7)</code> for <code>/usr</code> should be used. An exceptions is the <code>man</code> directory, which is directly under <code>/usr/local</code> rather than under <code>/usr/local/share</code> .
<code>/usr/obj/</code>	Architecture-specific target tree produced by building the <code>/usr/src</code> tree.
<code>/usr/pkg</code>	Used as the default destination for the files installed via the <code>pkgsrc</code> tree or <code>pkgsrc</code> packages (optional). The configuration directory is tunable, but the default location is <code>/usr/pkg/etc</code> .
<code>/usr/pkg/xorg/</code>	X11R6 distribution executables, libraries, etc (optional).
<code>/usr/pkgsrc</code>	The <code>pkgsrc</code> tree for installing packages (optional).
<code>/usr/sbin/</code>	System daemons & system utilities (executed by users).
<code>/usr/share/</code>	Architecture-independent files.
<code>/usr/src/</code>	BSD and/or local source files.
<code>/var/</code>	Multi-purpose log, temporary, transient, and spool files. A memory-based file system is sometimes mounted at <code>/var</code> . This can be automated with an entry in <code>/etc/fstab</code> ; see <code>mfs(8)</code> .
<code>/var/log/</code>	Miscellaneous system log files.
<code>/var/mail/</code>	User mailbox files.
<code>/var/spool/</code>	Miscellaneous printer and mail system spooling directories.
<code>/var/tmp/</code>	Temporary files. The files are usually preserved across a system reboot, unless <code>/var</code> is a memory-based file system.
<code>/var/yp</code>	NIS maps.

3.5 Disk Organization

The smallest unit of organization that DragonFly uses to find files is the filename. Filenames are case-sensitive, which means that `readme.txt` and `README.TXT` are two separate files. DragonFly does not use the extension (`.txt`) of a file to determine whether the file is a program, or a document, or some other form of data.

Files are stored in directories. A directory may contain no files, or it may contain many hundreds of files. A directory can also contain other directories, allowing you to build up a hierarchy of directories within one another. This makes it much easier to organize your data.

Files and directories are referenced by giving the file or directory name, followed by a forward slash, /, followed by any other directory names that are necessary. If you have directory `foo`, which contains directory `bar`, which contains the file `readme.txt`, then the full name, or *path* to the file is `foo/bar/readme.txt`.

Directories and files are stored in a file system. Each file system contains exactly one directory at the very top level, called the *root directory* for that file system. This root directory can then contain other directories.


So far this is probably similar to any other operating system you may have used. There are a few differences; for example, MS-DOS uses `\` to separate file and directory names, while Mac OS® uses `:`.

DragonFly does not use drive letters, or other drive names in the path. You would not write `c:/foo/bar/readme.txt` on DragonFly.


Instead, one file system is designated the *root file system*. The root file system's root directory is referred to as `/`. Every other file system is then *mounted* under the root file system. No matter how many disks you have on your DragonFly system, every directory appears to be part of the same disk.

Suppose you have three file systems, called `A`, `B`, and `C`. Each file system has one root directory, which contains two other directories, called `A1`, `A2` (and likewise `B1`, `B2` and `C1`, `C2`).

Call `A` the root file system. If you used the `ls` command to view the contents of this directory you would see two subdirectories, `A1` and `A2`. The directory tree looks like this:




A file system must be mounted on to a directory in another file system. So now suppose that you mount file system `B` on to the directory `A1`. The root directory of `B` replaces `A1`, and the directories in `B` appear accordingly:




Any files that are in the B1 or B2 directories can be reached with the path `/A1/B1` or `/A1/B2` as necessary. Any files that were in `/A1` have been temporarily hidden. They will reappear if B is *unmounted* from A.

If B had been mounted on A2 then the diagram would look like this:




and the paths would be `/A2/B1` and `/A2/B2` respectively.

File systems can be mounted on top of one another. Continuing the last example, the C file system could be mounted on top of the B1 directory in the B file system, leading to this arrangement:



Or C could be mounted directly on to the A file system, under the A1 directory:



If you are familiar with MS-DOS, this is similar, although not identical, to the `join` command.

This is not normally something you need to concern yourself with. Typically you create file systems when installing DragonFly and decide where to mount them, and then never change them unless you add a new disk.

It is entirely possible to have one large root file system, and not need to create any others. There are some drawbacks to this approach, and one advantage.

Benefits of Multiple File Systems

- Different file systems can have different *mount options*. For example, with careful planning, the root file system can be mounted read-only, making it impossible for you to inadvertently delete or edit a critical file. Separating user-writable file systems, such as `/home`, from other file systems also allows them to be mounted *nosuid*; this option prevents the *suid/guid* bits on executables stored on the file system from taking effect, possibly improving security.
- DragonFly automatically optimizes the layout of files on a file system, depending on how the file system is being used. So a file system that contains many small files that are written frequently will have a different optimization to one that contains fewer, larger files. By having one big file system this optimization breaks down.
- DragonFly's file systems are very robust should you lose power. However, a power loss at a critical point could still damage the structure of the file system. By splitting your data over multiple file systems it is more likely that the system will still come up, making it easier for you to restore from backup as necessary.

Benefit of a Single File System

- File systems are a fixed size. If you create a file system when you install DragonFly and give it a specific size, you may later discover that you need to make the partition bigger. The `growfs(8)` command makes it possible to increase the size of a file system on the fly.

File systems are contained in partitions. This does not have the same meaning as the common usage of the term partition (for example, MS-DOS partition), because of DragonFly's UNIX heritage. Each partition is identified by a letter from `a` through to `h`. Each partition can contain only one file system, which means that file systems are often described by either their typical mount point in the file system hierarchy, or the letter of the partition they are contained in.

DragonFly also uses disk space for *swap space*. Swap space provides DragonFly with *virtual memory*. This allows your computer to behave as though it has much more memory than it actually does. When DragonFly runs out of memory it moves some of the data that is not currently being used to the swap space, and moves it back in (moving something else out) when it needs it.

Some partitions have certain conventions associated with them.

Partition	Convention
<code>a</code>	Normally contains the root file system
<code>b</code>	Normally contains swap space
<code>c</code>	Normally the same size as the enclosing slice. This allows utilities that need to work on the entire slice (for example, a bad block scanner) to work on the <code>c</code> partition. You would not normally create a file system on this partition.
<code>d</code>	Partition <code>d</code> used to have a special meaning associated with it, although that is now gone. To this day, some tools may operate oddly if told to work on partition <code>d</code> .

Each partition-that-contains-a-file-system is stored in what DragonFly calls a *slice*. Slice is DragonFly's term for what the common call partitions, and again, this is because of DragonFly's UNIX background. Slices are numbered, starting at 1, through to 4.

Slice numbers follow the device name, prefixed with an `s`, starting at 1. So "`da0s1`" is the first slice on the first SCSI drive. There can only be four physical slices on a disk, but you can have logical slices inside physical slices of the appropriate type. These extended slices are numbered starting at 5, so "`ad0s5`" is the first extended slice on the first IDE disk. These devices are used by file systems that expect to occupy a slice.

Slices, “dangerously dedicated” physical drives, and other drives contain *partitions*, which are represented as letters from a to h. This letter is appended to the device name, so “da0a” is the a partition on the first da drive, which is “dangerously dedicated”. “ad1s3e” is the fifth partition in the third slice of the second IDE disk drive.

Finally, each disk on the system is identified. A disk name starts with a code that indicates the type of disk, and then a number, indicating which disk it is. Unlike slices, disk numbering starts at 0. Common codes that you will see are listed in Table 3-1.

When referring to a partition DragonFly requires that you also name the slice and disk that contains the partition, and when referring to a slice you should also refer to the disk name. Do this by listing the disk name, *s*, the slice number, and then the partition letter. Examples are shown in Example 3-1.

Example 3-2 shows a conceptual model of the disk layout that should help make things clearer.

In order to install DragonFly you must first configure the disk slices, then create partitions within the slice you will use for DragonFly, and then create a file system (or swap space) in each partition, and decide where that file system will be mounted.

Table 3-1. Disk Device Codes

Code	Meaning
ad	ATAPI (IDE) disk
da	SCSI direct access disk
acd	ATAPI (IDE) CDROM
cd	SCSI CDROM
fd	Floppy disk

Example 3-1. Sample Disk, Slice, and Partition Names

Name	Meaning
ad0s1a	The first partition (a) on the first slice (s1) on the first IDE disk (ad0).
da1s2e	The fifth partition (e) on the second slice (s2) on the second SCSI disk (da1).

Example 3-2. Conceptual Model of a Disk

This diagram shows DragonFly’s view of the first IDE disk attached to the system. Assume that the disk is 4 GB in size, and contains two 2 GB slices (MS-DOS partitions). The first slice contains a MS-DOS disk, *c:*, and the second slice contains a DragonFly installation. This example DragonFly installation has three partitions, and a swap partition.

The three partitions will each hold a file system. Partition a will be used for the root file system, e for the */var* directory hierarchy, and f for the */usr* directory hierarchy.

3.6 Mounting and Unmounting File Systems

The file system is best visualized as a tree, rooted, as it were, at `/`. `/dev`, `/usr`, and the other directories in the root directory are branches, which may have their own branches, such as `/usr/local`, and so on.

There are various reasons to house some of these directories on separate file systems. `/var` contains the directories `log/`, `spool/`, and various types of temporary files, and as such, may get filled up. Filling up the root file system is not a good idea, so splitting `/var` from `/` is often favorable.

Another common reason to contain certain directory trees on other file systems is if they are to be housed on separate physical disks, or are separate virtual disks, such as Network File System mounts, or CDROM drives.

3.6.1 The `fstab` File

During the boot process, file systems listed in `/etc/fstab` are automatically mounted (unless they are listed with the `noauto` option).

The `/etc/fstab` file contains a list of lines of the following format:

```
device          /mount-point fstype      options      dumpfreq     passno
```

`device`

A device name (which should exist), as explained in Section 12.2.

`mount-point`

A directory (which should exist), on which to mount the file system.

`fstype`

The file system type to pass to `mount(8)`. The default DragonFly file system is `ufs`.

`options`

Either `rw` for read-write file systems, or `ro` for read-only file systems, followed by any other options that may be needed. A common option is `noauto` for file systems not normally mounted during the boot sequence. Other options are listed in the `mount(8)` manual page.

`dumpfreq`

This is used by `dump(8)` to determine which file systems require dumping. If the field is missing, a value of zero is assumed.

`passno`

This determines the order in which file systems should be checked. File systems that should be skipped should have their `passno` set to zero. The root file system (which needs to be checked before everything else) should have its `passno` set to one, and other file systems' `passno` should be set to values greater than one. If more than one file systems have the same `passno` then `fsck(8)` will attempt to check file systems in parallel if possible.

Consult the `fstab(5)` manual page for more information on the format of the `/etc/fstab` file and the options it contains.

3.6.2 The `mount` Command

The `mount(8)` command is what is ultimately used to mount file systems.

In its most basic form, you use:

```
# mount device mountpoint
```

There are plenty of options, as mentioned in the `mount(8)` manual page, but the most common are:

Mount Options

`-a`

Mount all the file systems listed in `/etc/fstab`. Except those marked as “noauto”, excluded by the `-t` flag, or those that are already mounted.

`-d`

Do everything except for the actual mount system call. This option is useful in conjunction with the `-v` flag to determine what `mount(8)` is actually trying to do.

`-f`

Force the mount of an unclean file system (dangerous), or forces the revocation of write access when downgrading a file system’s mount status from read-write to read-only.

`-r`

Mount the file system read-only. This is identical to using the `rdonly` argument to the `-o` option.

`-t fstype`

Mount the given file system as the given file system type, or mount only file systems of the given type, if given the `-a` option.

“ufs” is the default file system type.

`-u`

Update mount options on the file system.

`-v`

Be verbose.

`-w`

Mount the file system read-write.

The `-o` option takes a comma-separated list of the options, including the following:

`nodev`

Do not interpret special devices on the file system. This is a useful security option.

noexec

Do not allow execution of binaries on this file system. This is also a useful security option.

nosuid

Do not interpret `setuid` or `setgid` flags on the file system. This is also a useful security option.

3.6.3 The `umount` Command

The `umount(8)` command takes, as a parameter, one of a mountpoint, a device name, or the `-a` or `-A` option.

All forms take `-f` to force unmounting, and `-v` for verbosity. Be warned that `-f` is not generally a good idea. Forcibly unmounting file systems might crash the computer or damage data on the file system.

`-a` and `-A` are used to unmount all mounted file systems, possibly modified by the file system types listed after `-t`. `-A`, however, does not attempt to unmount the root file system.

3.7 Processes

DragonFly is a multi-tasking operating system. This means that it seems as though more than one program is running at once. Each program running at any one time is called a *process*. Every command you run will start at least one new process, and there are a number of system processes that run all the time, keeping the system functional.

Each process is uniquely identified by a number called a *process ID*, or *PID*, and, like files, each process also has one owner and group. The owner and group information is used to determine what files and devices the process can open, using the file permissions discussed earlier. Most processes also have a parent process. The parent process is the process that started them. For example, if you are typing commands to the shell then the shell is a process, and any commands you run are also processes. Each process you run in this way will have your shell as its parent process. The exception to this is a special process called `init(8)`. `init` is always the first process, so its PID is always 1. `init` is started automatically by the kernel when DragonFly starts.

Two commands are particularly useful to see the processes on the system, `ps(1)` and `top(1)`. The `ps` command is used to show a static list of the currently running processes, and can show their PID, how much memory they are using, the command line they were started with, and so on. The `top` command displays all the running processes, and updates the display every few seconds, so that you can interactively see what your computer is doing.

By default, `ps` only shows you the commands that are running and are owned by you. For example:

```
% ps
  PID  TT  STAT      TIME COMMAND
   298  p0  Ss      0:01.10 tcsh
   7078 p0  S       2:40.88 xemacs mdoc.xsl (xemacs-21.1.14)
 37393 p0  I       0:03.11 xemacs freebsd.dsl (xemacs-21.1.14)
 48630 p0  S       2:50.89 /usr/local/lib/netscape-linux/navigator-linux-4.77.bi
 48730 p0  IW      0:00.00 (dns helper) (navigator-linux-)
 72210 p0  R+      0:00.00 ps
   390  p1  Is      0:01.14 tcsh
   7059 p2  Is+    1:36.18 /usr/local/bin/mutt -y
   6688 p3  IWs    0:00.00 tcsh
  10735 p4  IWs    0:00.00 tcsh
```

```

20256 p5 IWs 0:00.00 tcsh
  262 v0 IWs 0:00.00 -tcsh (tcsh)
  270 v0 IW+ 0:00.00 /bin/sh /usr/X11R6/bin/startx -- -bpp 16
  280 v0 IW+ 0:00.00 xinit /home/nik/.xinitrc -- -bpp 16
  284 v0 IW 0:00.00 /bin/sh /home/nik/.xinitrc
  285 v0 S 0:38.45 /usr/X11R6/bin/sawfish

```

As you can see in this example, the output from `ps(1)` is organized into a number of columns. `PID` is the process ID discussed earlier. PIDs are assigned starting from 1, go up to 99999, and wrap around back to the beginning when you run out. The `TT` column shows the tty the program is running on, and can safely be ignored for the moment. `STAT` shows the program's state, and again, can be safely ignored. `TIME` is the amount of time the program has been running on the CPU—this is usually not the elapsed time since you started the program, as most programs spend a lot of time waiting for things to happen before they need to spend time on the CPU. Finally, `COMMAND` is the command line that was used to run the program.

`ps(1)` supports a number of different options to change the information that is displayed. One of the most useful sets is `auxww`. `a` displays information about all the running processes, not just your own. `u` displays the username of the process' owner, as well as memory usage. `x` displays information about daemon processes, and `ww` causes `ps(1)` to display the full command line, rather than truncating it once it gets too long to fit on the screen.

The output from `top(1)` is similar. A sample session looks like this:

```

% top
last pid: 72257; load averages:  0.13,  0.09,  0.03   up 0+13:38:33  22:39:10
47 processes:  1 running, 46 sleeping
CPU states: 12.6% user,  0.0% nice,  7.8% system,  0.0% interrupt, 79.7% idle
Mem: 36M Active, 5256K Inact, 13M Wired, 6312K Cache, 15M Buf, 408K Free
Swap: 256M Total, 38M Used, 217M Free, 15% Inuse

  PID USERNAME PRI NICE  SIZE  RES STATE   TIME  WCPU   CPU COMMAND
72257 nik        28  0  1960K 1044K RUN     0:00 14.86%  1.42% top
 7078 nik         2  0 15280K 10960K select  2:54  0.88%  0.88% xemacs-21.1.14
  281 nik         2  0 18636K  7112K select  5:36  0.73%  0.73% XF86_SVGA
  296 nik         2  0  3240K 1644K select  0:12  0.05%  0.05% xterm
48630 nik         2  0 29816K  9148K select  3:18  0.00%  0.00% navigator-linu
  175 root         2  0   924K  252K select  1:41  0.00%  0.00% syslogd
 7059 nik         2  0   7260K 4644K poll   1:38  0.00%  0.00% mutt
...

```

The output is split into two sections. The header (the first five lines) shows the PID of the last process to run, the system load averages (which are a measure of how busy the system is), the system uptime (time since the last reboot) and the current time. The other figures in the header relate to how many processes are running (47 in this case), how much memory and swap space has been taken up, and how much time the system is spending in different CPU states.

Below that are a series of columns containing similar information to the output from `ps(1)`. As before you can see the PID, the username, the amount of CPU time taken, and the command that was run. `top(1)` also defaults to showing you the amount of memory space taken by the process. This is split into two columns, one for total size, and one for resident size—total size is how much memory the application has needed, and the resident size is how much it is actually using at the moment. In this example you can see that **Netscape®** has required almost 30 MB of RAM, but is currently only using 9 MB.

`top(1)` automatically updates this display every two seconds; this can be changed with the `s` option.

3.8 Daemons, Signals, and Killing Processes

When you run an editor it is easy to control the editor, tell it to load files, and so on. You can do this because the editor provides facilities to do so, and because the editor is attached to a *terminal*. Some programs are not designed to be run with continuous user input, and so they disconnect from the terminal at the first opportunity. For example, a web server spends all day responding to web requests, it normally does not need any input from you. Programs that transport email from site to site are another example of this class of application.

We call these programs *daemons*. Daemons were characters in Greek mythology; neither good or evil, they were little attendant spirits that, by and large, did useful things for mankind. Much like the web servers and mail servers of today do useful things. This is why the mascot for a number of BSD-based operating systems has, for a long time, been a cheerful looking daemon with sneakers and a pitchfork.

There is a convention to name programs that normally run as daemons with a trailing “d”. **BIND** is the Berkeley Internet Name Daemon (and the actual program that executes is called `named`), the **Apache** web server program is called `httpd`, the line printer spooling daemon is `lpd` and so on. This is a convention, not a hard and fast rule; for example, the main mail daemon for the **Sendmail** application is called `sendmail`, and not `maild`, as you might imagine.

Sometimes you will need to communicate with a daemon process. These communications are called *signals*, and you can communicate with a daemon (or with any other running process) by sending it a signal. There are a number of different signals that you can send—some of them have a specific meaning, others are interpreted by the application, and the application’s documentation will tell you how that application interprets signals. You can only send a signal to a process that you own. If you send a signal to someone else’s process with `kill(1)` or `kill(2)` permission will be denied. The exception to this is the `root` user, who can send signals to everyone’s processes.

DragonFly will also send applications signals in some cases. If an application is badly written, and tries to access memory that it is not supposed to, DragonFly sends the process the *Segmentation Violation* signal (`SIGSEGV`). If an application has used the `alarm(3)` system call to be alerted after a period of time has elapsed then it will be sent the Alarm signal (`SIGALRM`), and so on.

Two signals can be used to stop a process, `SIGTERM` and `SIGKILL`. `SIGTERM` is the polite way to kill a process; the process can *catch* the signal, realize that you want it to shut down, close any log files it may have open, and generally finish whatever it is doing at the time before shutting down. In some cases a process may even ignore `SIGTERM` if it is in the middle of some task that can not be interrupted.

`SIGKILL` can not be ignored by a process. This is the “I do not care what you are doing, stop right now” signal. If you send `SIGKILL` to a process then DragonFly will stop that process there and then⁴.

The other signals you might want to use are `SIGHUP`, `SIGUSR1`, and `SIGUSR2`. These are general purpose signals, and different applications will do different things when they are sent.

Suppose that you have changed your web server’s configuration file—you would like to tell the web server to re-read its configuration. You could stop and restart `httpd`, but this would result in a brief outage period on your web server, which may be undesirable. Most daemons are written to respond to the `SIGHUP` signal by re-reading their configuration file. So instead of killing and restarting `httpd` you would send it the `SIGHUP` signal. Because there is no standard way to respond to these signals, different daemons will have different behavior, so be sure and read the documentation for the daemon in question.

Signals are sent using the `kill(1)` command, as this example shows.

Sending a Signal to a Process

This example shows how to send a signal to `inetd(8)`. The `inetd` configuration file is `/etc/inetd.conf`, and `inetd` will re-read this configuration file when it is sent `SIGHUP`.

1. Find the process ID of the process you want to send the signal to. Do this using `ps(1)` and `grep(1)`. The `grep(1)` command is used to search through output, looking for the string you specify. This command is run as a normal user, and `inetd(8)` is run as `root`, so the `ax` options must be given to `ps(1)`.

```
% ps -ax | grep inetd
 198  ??  IWs    0:00.00 inetd -wW
```

So the `inetd(8)` PID is 198. In some cases the `grep inetd` command might also occur in this output. This is because of the way `ps(1)` has to find the list of running processes.

2. Use `kill(1)` to send the signal. Because `inetd(8)` is being run by `root` you must use `su(1)` to become `root` first.

```
% su
Password:
# /bin/kill -s HUP 198
```

In common with most UNIX commands, `kill(1)` will not print any output if it is successful. If you send a signal to a process that you do not own then you will see `kill: PID: Operation not permitted`. If you mistype the PID you will either send the signal to the wrong process, which could be bad, or, if you are lucky, you will have sent the signal to a PID that is not currently in use, and you will see `kill: PID: No such process`.

Why Use `/bin/kill`? Many shells provide the `kill` command as a built in command; that is, the shell will send the signal directly, rather than running `/bin/kill`. This can be very useful, but different shells have a different syntax for specifying the name of the signal to send. Rather than try to learn all of them, it can be simpler just to use the `/bin/kill ...` command directly.

Sending other signals is very similar, just substitute `TERM` or `KILL` in the command line as necessary.

Important: Killing random process on the system can be a bad idea. In particular, `init(8)`, process ID 1, is very special. Running `/bin/kill -s KILL 1` is a quick way to shutdown your system. *Always* double check the arguments you run `kill(1)` with *before* you press **Return**.

3.9 Shells

In DragonFly, a lot of everyday work is done in a command line interface called a shell. A shell's main job is to take commands from the input channel and execute them. A lot of shells also have built in functions to help everyday tasks such as file management, file globbing, command line editing, command macros, and environment variables. DragonFly comes with a set of shells, such as `sh`, the Bourne Shell, and `tcsh`, the improved C-shell. Many other shells are available from `pkgsrc`, such as `zsh` and `bash`.

Which shell do you use? It is really a matter of taste. If you are a C programmer you might feel more comfortable with a C-like shell such as `tcsh`. If you have come from Linux or are new to a UNIX command line interface you

might try `bash`. The point is that each shell has unique properties that may or may not work with your preferred working environment, and that you have a choice of what shell to use.

One common feature in a shell is filename completion. Given the typing of the first few letters of a command or filename, you can usually have the shell automatically complete the rest of the command or filename by hitting the **Tab** key on the keyboard. Here is an example. Suppose you have two files called `foobar` and `foo.bar`. You want to delete `foo.bar`. So what you would type on the keyboard is: `rm fo[Tab].[Tab]`.

The shell would print out `rm foo[BEEP].bar`.

The [BEEP] is the console bell, which is the shell telling me it was unable to totally complete the filename because there is more than one match. Both `foobar` and `foo.bar` start with `fo`, but it was able to complete to `foo`. If you type in `.`, then hit **Tab** again, the shell would be able to fill in the rest of the filename for you.

Another feature of the shell is the use of environment variables. Environment variables are a variable key pair stored in the shell's environment space. This space can be read by any program invoked by the shell, and thus contains a lot of program configuration. Here is a list of common environment variables and what they mean:

Variable	Description
USER	Current logged in user's name.
PATH	Colon separated list of directories to search for binaries.
DISPLAY	Network name of the X11 display to connect to, if available.
SHELL	The current shell.
TERM	The name of the user's terminal. Used to determine the capabilities of the terminal.
TERMCAP	Database entry of the terminal escape codes to perform various terminal functions.
OSTYPE	Type of operating system. e.g., DragonFly.
MACHTYPE	The CPU architecture that the system is running on.
EDITOR	The user's preferred text editor.
PAGER	The user's preferred text pager.
MANPATH	Colon separated list of directories to search for manual pages.

Setting an environment variable differs somewhat from shell to shell. For example, in the C-Style shells such as `tcsh` and `csh`, you would use `setenv` to set environment variables. Under Bourne shells such as `sh` and `bash`, you would use `export` to set your current environment variables. For example, to set or modify the `EDITOR` environment variable, under `csh` or `tcsh` a command like this would set `EDITOR` to `/usr/pkg/bin/emacs`:

```
% setenv EDITOR /usr/pkg/bin/emacs
```

Under Bourne shells:

```
% export EDITOR="/usr/pkg/bin/emacs"
```

You can also make most shells expand the environment variable by placing a `$` character in front of it on the command line. For example, `echo $TERM` would print out whatever `$TERM` is set to, because the shell expands `$TERM` and passes it on to `echo`.

Shells treat a lot of special characters, called meta-characters as special representations of data. The most common one is the `*` character, which represents any number of characters in a filename. These special meta-characters can be used to do filename globbing. For example, typing in `echo *` is almost the same as typing in `ls` because the shell takes all the files that match `*` and puts them on the command line for `echo` to see.

To prevent the shell from interpreting these special characters, they can be escaped from the shell by putting a backslash (`\`) character in front of them. `echo $TERM` prints whatever your terminal is set to. `echo \$TERM` prints `$TERM` as is.

3.9.1 Changing Your Shell

The easiest way to change your shell is to use the `chsh` command. Running `chsh` will place you into the editor that is in your `EDITOR` environment variable; if it is not set, you will be placed in `vi`. Change the “Shell:” line accordingly.

You can also give `chsh` the `-s` option; this will set your shell for you, without requiring you to enter an editor. For example, if you wanted to change your shell to `bash`, the following should do the trick:

```
% chsh -s /usr/pkg/bin/bash
```

Note: The shell that you wish to use *must* be present in the `/etc/shells` file. If you have installed a shell from the `pkgsrc` tree, then this should have been done for you already. If you installed the shell by hand, you must do this.

For example, if you installed `bash` by hand and placed it into `/usr/local/bin`, you would want to:

```
# echo "/usr/local/bin/bash" >> /etc/shells
```

Then rerun `chsh`.

3.10 Text Editors

A lot of configuration in DragonFly is done by editing text files. Because of this, it would be a good idea to become familiar with a text editor. DragonFly comes with a few as part of the base system, and many more are available in the `pkgsrc` tree.

The easiest and simplest editor to learn is an editor called `ee`, which stands for easy editor. To start `ee`, one would type at the command line `ee filename` where `filename` is the name of the file to be edited. For example, to edit `/etc/rc.conf`, type in `ee /etc/rc.conf`. Once inside of `ee`, all of the commands for manipulating the editor’s functions are listed at the top of the display. The caret `^` character represents the **Ctrl** key on the keyboard, so `^e` expands to the key combination **Ctrl+e**. To leave `ee`, hit the **Esc** key, then choose leave editor. The editor will prompt you to save any changes if the file has been modified.

DragonFly also comes with more powerful text editors such as `vi` as part of the base system, while other editors, like `emacs` and `vim`, are part of the `pkgsrc` tree. These editors offer much more functionality and power at the expense of being a little more complicated to learn. However if you plan on doing a lot of text editing, learning a more powerful editor such as `vim` or `emacs` will save you much more time in the long run.

3.11 Devices and Device Nodes

A device is a term used mostly for hardware-related activities in a system, including disks, printers, graphics cards, and keyboards. When DragonFly boots, the majority of what DragonFly displays are devices being detected. You can look through the boot messages again by viewing `/var/run/dmesg.boot`.

For example, `acd0` is the first IDE CDROM drive, while `kbd0` represents the keyboard.

Most of these devices in a UNIX operating system must be accessed through special files called device nodes, which are located in the `/dev` directory.

3.11.1 Creating Device Nodes with MAKEDEV

When adding a new device to your system, or compiling in support for additional devices, you may need to create one or more device nodes for the new devices.

Device nodes are created using the MAKEDEV(8) script as shown below:

```
# cd /dev
# sh MAKEDEV ad1
```

This example would make the proper device nodes for the second IDE drive when installed.

3.12 Binary Formats

To understand why DragonFly uses the `elf(5)` format, you must first know a little about the three currently “dominant” executable formats for UNIX:

- `a.out(5)`

The oldest and “classic” UNIX object format. It uses a short and compact header with a magic number at the beginning that is often used to characterize the format (see `a.out(5)` for more details). It contains three loaded segments: `.text`, `.data`, and `.bss` plus a symbol table and a string table.

- COFF

The SVR3 object format. The header now comprises a section table, so you can have more than just `.text`, `.data`, and `.bss` sections.

- `elf(5)`

The successor to COFF, featuring multiple sections and 32-bit or 64-bit possible values. One major drawback: ELF was also designed with the assumption that there would be only one ABI per system architecture. That assumption is actually quite incorrect, and not even in the commercial SYSV world (which has at least three ABIs: SVR4, Solaris, SCO) does it hold true.

DragonFly tries to work around this problem somewhat by providing a utility for *branding* a known ELF executable with information about the ABI it is compliant with. See the manual page for `brandelf(1)` for more information. DragonFly runs ELF.

So, why are there so many different formats?

Back in the dim, dark past, there was simple hardware. This simple hardware supported a simple, small system. `a.out` was completely adequate for the job of representing binaries on this simple system (a PDP-11). As people ported UNIX from this simple system, they retained the `a.out` format because it was sufficient for the early ports of UNIX to architectures like the Motorola 68k, VAXen, etc.

Then some bright hardware engineer decided that if he could force software to do some sleazy tricks, then he would be able to shave a few gates off the design and allow his CPU core to run faster. While it was made to work with this new kind of hardware (known these days as RISC), `a.out` was ill-suited for this hardware, so many formats were developed to get to a better performance from this hardware than the limited, simple `a.out` format could offer. Things like COFF, ECOFF, and a few obscure others were invented and their limitations explored before things seemed to settle on ELF.

In addition, program sizes were getting huge and disks (and physical memory) were still relatively small so the concept of a shared library was born. The VM system also became more sophisticated. While each one of these advancements was done using the `a.out` format, its usefulness was stretched more and more with each new feature. In addition, people wanted to dynamically load things at run time, or to junk parts of their program after the init code had run to save in core memory and swap space. Languages became more sophisticated and people wanted code called before main automatically. Lots of hacks were done to the `a.out` format to allow all of these things to happen, and they basically worked for a time. In time, `a.out` was not up to handling all these problems without an ever increasing overhead in code and complexity. While ELF solved many of these problems, it would be painful to switch from the system that basically worked. So ELF had to wait until it was more painful to remain with `a.out` than it was to migrate to ELF.

ELF is more expressive than `a.out` and allows more extensibility in the base system. The ELF tools are better maintained, and offer cross compilation support, which is important to many people. ELF may be a little slower than `a.out`, but trying to measure it can be difficult. There are also numerous details that are different between the two in how they map pages, handle init code, etc. None of these are very important, but they are differences.

3.13 For More Information

3.13.1 Manual Pages

The most comprehensive documentation on DragonFly is in the form of manual pages. Nearly every program on the system comes with a short reference manual explaining the basic operation and various arguments. These manuals can be viewed with the `man` command. Use of the `man` command is simple:

```
% man command
```

`command` is the name of the command you wish to learn about. For example, to learn more about `ls` command type:

```
% man ls
```

The online manual is divided up into numbered sections:

1. User commands.
2. System calls and error numbers.
3. Functions in the C libraries.
4. Device drivers.

5. File formats.
6. Games and other diversions.
7. Miscellaneous information.
8. System maintenance and operation commands.
9. Kernel internals.

In some cases, the same topic may appear in more than one section of the online manual. For example, there is a `chmod` user command and a `chmod()` system call. In this case, you can tell the `man` command which one you want by specifying the section:

```
% man 1 chmod
```

This will display the manual page for the user command `chmod`. References to a particular section of the online manual are traditionally placed in parenthesis in written documentation, so `chmod(1)` refers to the `chmod` user command and `chmod(2)` refers to the system call.

This is fine if you know the name of the command and simply wish to know how to use it, but what if you cannot recall the command name? You can use `man` to search for keywords in the command descriptions by using the `-k` switch:

```
% man -k mail
```

With this command you will be presented with a list of commands that have the keyword “mail” in their descriptions. This is actually functionally equivalent to using the `apropos` command.

So, you are looking at all those fancy commands in `/usr/bin` but do not have the faintest idea what most of them actually do? Simply do:

```
% cd /usr/bin
% man -f *
```

or

```
% cd /usr/bin
% whatis *
```

which does the same thing.

3.13.2 GNU Info Files

DragonFly includes many applications and utilities produced by the Free Software Foundation (FSF). In addition to manual pages, these programs come with more extensive hypertext documents called `info` files which can be viewed with the `info` command or, if you installed **emacs**, the `info` mode of **emacs**.

To use the `info(1)` command, simply type:

```
% info
```

For a brief introduction, type `h`. For a quick command reference, type `?`.

Notes

1. This is what `i386` means. Note that even if you are not running DragonFly on an Intel 386 CPU, this is going to be `i386`. It is not the type of your processor, but the processor “architecture” that is shown here.
2. Startup scripts are programs that are run automatically by DragonFly when booting. Their main function is to set things up for everything else to run, and start any services that you have configured to run in the background doing useful things.
3. A fairly technical and accurate description of all the details of the DragonFly console and keyboard drivers can be found in the manual pages of `syscons(4)`, `atkbd(4)`, `vidcontrol(1)` and `kbdcontrol(1)`. We will not expand on the details here, but the interested reader can always consult the manual pages for a more detailed and thorough explanation of how things work.
4. Not quite true—there are a few things that can not be interrupted. For example, if the process is trying to read from a file that is on another computer on the network, and the other computer has gone away for some reason (been turned off, or the network has a fault), then the process is said to be “uninterruptible”. Eventually the process will time out, typically after two minutes. As soon as this time out occurs the process will be killed.

Chapter 4 Installing Applications using NetBSD's pkgsrc framework

4.1 Synopsis

DragonFly is bundled with a rich collection of system tools as part of the base system. However, there is only so much one can do before needing to install an additional third-party application to get real work done. DragonFly utilizes NetBSD's pkgsrc framework ([pkgsrc.org \(http://www.pkgsrc.org/\)](http://www.pkgsrc.org/)) for installing third party software on your system. This system may be used to install the newest version of your favorite applications from local media or straight off the network.

After reading this chapter, you will know:

- How to install third-party binary software packages from the pkgsrc collection.
- How to build third-party software from the pkgsrc collection.
- Where to find DragonFly-specific changes to packages.
- How to remove previously installed packages.
- How to override the default values that the pkgsrc collection uses.
- How to upgrade your packages.

4.2 Overview of Software Installation

If you have used a UNIX system before you will know that the typical procedure for installing third party software goes something like this:

1. Download the software, which might be distributed in source code format, or as a binary.
2. Unpack the software from its distribution format (typically a tarball compressed with `compress(1)`, `gzip(1)`, or `bzip2(1)`).
3. Locate the documentation (perhaps an `INSTALL` or `README` file, or some files in a `doc/` subdirectory) and read up on how to install the software.
4. If the software was distributed in source format, compile it. This may involve editing a `Makefile`, or running a `configure` script, and other work.
5. Test and install the software.

And that is only if everything goes well. If you are installing a software package that was not deliberately ported to DragonFly you may even have to go in and edit the code to make it work properly.

Should you want to, you can continue to install software the “traditional” way with DragonFly. However, DragonFly provides technology from NetBSD, which can save you a lot of effort: pkgsrc. At the time of writing, over 6,000 third party applications have been made available in this way.

For any given application, the DragonFly Binary package for that application is a single file which you must download. The package contains pre-compiled copies of all the commands for the application, as well as any configuration files or documentation. A downloaded package file can be manipulated with DragonFly package management commands, such as `pkg_add(1)`, `pkg_delete(1)`, `pkg_info(1)`, and so on. Installing a new application can be carried out with a single command.

In addition the pkgsrc collection supplies a collection of files designed to automate the process of compiling an application from source code.

Remember that there are a number of steps you would normally carry out if you compiled a program yourself (downloading, unpacking, patching, compiling, installing). The files that make up a pkgsrc source collection contain all the necessary information to allow the system to do this for you. You run a handful of simple commands and the source code for the application is automatically downloaded, extracted, patched, compiled, and installed for you.

In fact, the pkgsrc source subsystem can also be used to generate packages which can later be manipulated with `pkg_add` and the other package management commands that will be introduced shortly.

Pkgsrc understands *dependencies*. Suppose you want to install an application that depends on a specific library being installed. Both the application and the library have been made available through the pkgsrc collection. If you use the `pkg_add` command or the pkgsrc subsystem to add the application, both will notice that the library has not been installed, and automatically install the library first.

You might be wondering why pkgsrc bothers with both. Binary packages and the source tree both have their own strengths, and which one you use will depend on your own preference.

Binary Package Benefits

- A compressed package tarball is typically smaller than the compressed tarball containing the source code for the application.
- Packages do not require any additional compilation. For large applications, such as **Mozilla**, **KDE**, or **GNOME** this can be important, particularly if you are on a slow system.
- Packages do not require any understanding of the process involved in compiling software on DragonFly.

Pkgsrc source Benefits

- Binary packages are normally compiled with conservative options, because they have to run on the maximum number of systems. By installing from the source, you can tweak the compilation options to (for example) generate code that is specific to a Pentium IV or Athlon processor.
- Some applications have compile time options relating to what they can and cannot do. For example, **Apache** can be configured with a wide variety of different built-in options. By building from the source you do not have to accept the default options, and can set them yourself.

In some cases, multiple packages will exist for the same application to specify certain settings. For example, **vim** is available as a `vim` package and a `vim-gtk` package, depending on whether you have installed an X11 server. This sort of rough tweaking is possible with packages, but rapidly becomes impossible if an application has more than one or two different compile time options.

- The licensing conditions of some software distributions forbid binary distribution. They must be distributed as source code.
- Some people do not trust binary distributions. With source code, it is possible to check for any vulnerabilities built into the program before installing it to an otherwise secure system. Few people perform this much review, however.
- If you have local patches, you will need the source in order to apply them.

- Some people like having code around, so they can read it if they get bored, hack it, borrow from it (license permitting, of course), and so on.

To keep track of updated pkgsrc releases subscribe to the NetBSD pkgsrc users mailing list (<http://www.netbsd.org/MailingLists/pkgsrc-users>) and the NetBSD pkgsrc users mailing list (<http://www.netbsd.org/MailingLists/tech-pkgsrc>). It's also useful to watch the DragonFly User related mailing list (<http://leaf.dragonflybsd.org/mailarchive/>) as errors with pkgsrc on DragonFly should be reported there.

Warning: Before installing any application, you should check <http://www.pkgsrc.org/> for security issues related to your application.

You can also install `security/audit-packages` which will automatically check all installed applications for known vulnerabilities, a check will be also performed before any application build. Meanwhile, you can use the command `audit-packages -d` after you have installed some packages.

The remainder of this chapter will explain how to use the pkgsrc system to install and manage third party software on DragonFly.

4.3 Finding Your Application

Before you can install any applications you need to know what you want, and what the application is called.

DragonFly's list of available applications is growing all the time. Fortunately, there are a number of ways to find what you want:

- There is a pkgsrc related web site that maintains an up-to-date searchable list of all the available applications, at <http://pkgsrc.se>. The packages and the corresponding source tree are divided into categories, and you may either search for an application by name (if you know it), or see all the applications available in a category.

4.4 Using the Binary Packages System

Original FreeBSD documentation contributed by DragonFly BSD customizations contributed by Chern Lee and Adrian Nida.

4.4.1 Installing a Binary Package

You can use the `pkg_add(1)` utility to install a pkgsrc software package from a local file or from a server on the network.

Example 4-1. Downloading a Package Manually and Installing It Locally

```
# ftp -a packages.stura.uni-rostock.de
Connected to fsr.uni-rostock.de.
220 packages.stura.uni-rostock.de FTP server (Version 6.00LS) ready.
331 Guest login ok, send your email address as password.
230 Guest login ok, access restrictions apply.
Remote system type is UNIX.
```

Using binary mode to transfer files.

```
ftp> cd /pkgsrc-current/DragonFly/RELEASE/i386/All/
250 CWD command successful.
ftp> get Overkill-0.15.tgz
local: Overkill-0.15.tgz remote: Overkill-0.15.tgz
229 Entering Extended Passive Mode (|||61652|)
150 Opening BINARY mode data connection for 'Overkill-0.15.tgz' (174638 bytes).
100% |*****| 170 KB 159.37 KB/s 00:00 ETA
226 Transfer complete.
174638 bytes received in 00:01 (159.30 KB/s)
ftp> exit
221 Goodbye.
# pkg_add Overkill-0.15.tgz
```

Note: It should be noted that simply issuing:

```
# pkg_add ftp://packages.stura.uni-rostock.de/pkgsrc-current/DragonFly/RELEASE/i386/All/Overkill-0.15.tgz
```

will yield the same result as the above example.

Unlike the FreeBSD version, the Pkgsrc `pkg_add(1)` does not need to be passed the `-r` option. As can be seen from the second example, you just need to pass in the URL of the package. The utility will also always automatically fetch and install all dependencies.

The example above would download the correct package and add it without any further user intervention. If you want to specify an alternative DragonFly Packages Mirror, instead of the main distribution site, you have to set `PACKAGESITE` accordingly, to override the default settings. `pkg_add(1)` uses `fetch(3)` to download the files, which honors various environment variables, including `FTP_PASSIVE_MODE`, `FTP_PROXY`, and `FTP_PASSWORD`. You may need to set one or more of these if you are behind a firewall, or need to use an FTP/HTTP proxy. See `fetch(3)` for the complete list.

Binary package files are distributed in `.tgz` formats. You can find them at the default location `ftp://goBSD.com/packages/`, among other sites. The layout of the packages is similar to that of the `/usr/pkgsrc` tree. Each category has its own directory, and every package can be found within the `All` directory.

The directory structure of the binary package system matches the source tree layout; they work with each other to form the entire package system.

4.4.2 Managing Packages

`pkg_info(1)` is a utility that lists and describes the various packages installed.

```
# pkg_info
digest-20050731      Message digest wrapper utility
screen-4.0.2nb4     Multi-screen window manager
...
```

`pkg_version(1)` is a utility that summarizes the versions of all installed packages. It compares the package version to the current version found in the ports tree.

4.4.3 Deleting a Package

To remove a previously installed software package, use the `pkg_delete(1)` utility.

```
# pkg_delete xchat-1.7.1
```

4.4.4 Miscellaneous

All package information is stored within the `/var/db/pkg` directory. The installed file list and descriptions of each package can be found within subdirectories of this directory.

4.5 Using the pkgsrc® Source Tree

The following sections provide basic instructions on using the pkgsrc source tree to install or remove programs from your system.

4.5.1 Obtaining the pkgsrc Source Tree

Before you can install pkgsrc packages from source, you must first obtain the pkgsrc source tree—which is essentially a set of `Makefiles`, patches, and description files placed in `/usr/pkgsrc`.

The primary method to obtain and keep your pkgsrc collection up to date is by using **CVS**

CVS

This is a quick method for getting the pkgsrc collection using **CVS**.

1. Run `cvs`:

```
# cd /usr/  
# cvs -d anoncvs@anoncvs.us.netbsd.org:/cvsroot co pkgsrc
```

2. Running the following command later will download and apply all the recent changes to your source tree.

```
# cd /usr/pkgsrc  
# cvs up
```

4.5.2 Installing Packages from Source

The first thing that should be explained when it comes to the source tree is what is actually meant by a “skeleton”. In a nutshell, a source skeleton is a minimal set of files that tell your DragonFly system how to cleanly compile and install a program. Each source skeleton should include:

- A `Makefile`. The `Makefile` contains various statements that specify how the application should be compiled and where it should be installed on your system.
- A `distinfo` file. This file contains information about the files that must be downloaded to build the port and their checksums, to verify that files have not been corrupted during the download using `md5(1)`.

- A `files` directory. This directory contains the application specific files that are needed for the programs appropriate run-time configuration.

This directory may also contain other files used to build the port.

- A `patches` directory. This directory contains patches to make the program compile and install on your DragonFly system. Patches are basically small files that specify changes to particular files. They are in plain text format, and basically say "Remove line 10" or "Change line 26 to this ...". Patches are also known as "diffs" because they are generated by the `diff(1)` program.
- A `DESCR` file. This is a more detailed, often multiple-line, description of the program.
- A `PLIST` file. This is a list of all the files that will be installed by the port. It also tells the `pkgsrc` system what files to remove upon deinstallation.

Some `pkgsrc` source skeletons have other files, such as `MESSAGE`. The `pkgsrc` system uses these files to handle special situations. If you want more details on these files, and on `pkgsrc` in general, check out The `pkgsrc` guide (<http://www.netbsd.org/Documentation/pkgsrc/>), available at the NetBSD website (<http://www.netbsd.org/>).

Now that you have enough background information to know what the `pkgsrc` source tree is used for, you are ready to install your first compiled package. There are two ways this can be done, and each is explained below.

Before we get into that, however, you will need to choose an application to install. There are a few ways to do this, with the easiest method being the `pkgsrc` listing on Joerg Sonnenberger's web site (<ftp://packages.stura.uni-rostock.de/pkgsrc-current/DragonFly/RELEASE/i386/All/>). You can browse through the packages listed there.

Another way to find a particular source tree is by using the `pkgsrc` collection's built-in search mechanism. To use the search feature, you will need to be in the `/usr/pkgsrc` directory. Once in that directory, run `bmake search key="program-name"` where `program-name` is the name of the program you want to find. This searches packages names, comments, descriptions and dependencies and can be used to find packages which relate to a particular subject if you don't know the name of the program you are looking for. For example, if you were looking for `apache2`:

```
# cd /usr/pkgsrc
# bmake search key="apache2"
Extracting complete dependency database. This may take a while...
.....
100
.....
200
<Snip />
5800
.....
5900
.....
Flattening dependencies
Flattening build dependencies
Generating INDEX file
Indexed 5999 packages
<Snip />
Pkg:    apache-2.0.55nb7
Path:   www/apache2
Info:   Apache HTTP (Web) server, version 2
Maint:  tron@NetBSD.org
```

```
Index: www
B-deps: perl>=5.0 apr>=0.9.7.2.0.55nb2 expat>=2.0.0nb1 libtool-base>=1.5.22nb1 gmake>=3.78 gett
R-deps: perl>=5.0 apr>=0.9.7.2.0.55nb2 expat>=2.0.0nb1
Arch: any
```

The part of the output you want to pay particular attention to is the “Path:” line, since that tells you where to find the source tree for the requested application. The other information provided is not needed in order to install the package, so it will not be covered here.

The search string is case-insensitive. Searching for “APACHE” will yield the same results as searching for “apache”.

Note: It should be noted that “Extracting [the] complete dependency database” does indeed take a while.

Note: You must be logged in as `root` to install packages.

Now that you have found an application you would like to install, you are ready to do the actual installation. The source package includes instructions on how to build source code, but does not include the actual source code. You can get the source code from a CD-ROM or from the Internet. Source code is distributed in whatever manner the software author desires. Frequently this is a tarred and gzipped file, but it might be compressed with some other tool or even uncompressed. The program source code, whatever form it comes in, is called a “distfile”. You can get the distfile from a CD-ROM or from the Internet.

Warning: Before installing any application, you should be sure to have an up-to-date source tree and you should check <http://www.pkgsrc.org/> for security issues related to your port.

A security vulnerabilities check can be automatically done by **audit-packages** before any new application installation. This tool can be found in the pkgsrc collection (`security/audit-packages`). Consider running `auditpackages -d` before installing a new package, to fetch the current vulnerabilities database. A security audit and an update of the database will be performed during the daily security system check. For more informations read the `audit-packages` and `periodic(8)` manual pages.

Note: It should be noted that the current setup of DragonFly requires the use of `bmake` instead of `make`. This is because the current version of `make` on DragonFly does not support all the parameters that NetBSD's does.

Note: You can save an extra step by just running `bmake install` instead of `bmake` and `bmake install` as two separate steps.

Note: Some shells keep a cache of the commands that are available in the directories listed in the `PATH` environment variable, to speed up lookup operations for the executable file of these commands. If you are using one of these shells, you might have to use the `rehash` command after installing a package, before the newly installed commands can be used. This is true for both shells that are part of the base-system (such as `tcsh`) and shells that are available as packages (for instance, `shells/zsh`).

4.5.2.1 Installing Packages from the Internet

As with the last section, this section makes an assumption that you have a working Internet connection. If you do not, you will need to put a copy of the distfile into `/usr/pkgsrc/distfiles` manually.

Installing a package from the Internet is done exactly the same way as it would be if you already had the distfile. The only difference between the two is that the distfile is downloaded from the Internet on demand.

Here are the steps involved:

```
# cd /usr/pkgsrc/chat/ircII
# bmake install clean
=> ircii-20040820.tar.bz2 doesn't seem to exist on this system.
=> Attempting to fetch ircii-20040820.tar.bz2 from ftp://ircii.warped.com/pub/ircII/.
=> [559843 bytes]
Connected to ircii.warped.com.
220 bungi.sjc.warped.net FTP server (tnftpd 20040810) ready.
331 Guest login ok, type your name as password.
230-
    A SERVICE OF WARPED.COM - FOR MORE INFORMATION: http://www.warped.com

230-
    Please read the file README
        it was last modified on Mon Feb  9 18:43:17 2004 - 794 days ago
230 Guest login ok, access restrictions apply.
Remote system type is UNIX.
Using binary mode to transfer files.
200 Type set to I.
250 CWD command successful.
250 CWD command successful.
local: ircii-20040820.tar.bz2 remote: ircii-20040820.tar.bz2
229 Entering Extended Passive Mode (|||60090|)
150 Opening BINARY mode data connection for 'ircii-20040820.tar.bz2' (559843 bytes).
100% |*****| 550 KB 110.34 KB/s 00:00 ETA
226 Transfer complete.
559843 bytes received in 00:04 (110.34 KB/s)
221-
    Data traffic for this session was 559843 bytes in 1 file.
    Total traffic for this session was 560993 bytes in 1 transfer.
221 Thank you for using the FTP service on bungi.sjc.warped.net.
=> Checksum SHA1 OK for ircii-20040820.tar.bz2.
=> Checksum RMD160 OK for ircii-20040820.tar.bz2.
work -> /usr/obj/pkgsrc/chat/ircII/work
==> Extracting for ircII-20040820
=====
The supported build options for this package are:

socks4 socks5
```

You can select which build options to use by setting `PKG_DEFAULT_OPTIONS` or the following variable. Its current value is shown:

```
PKG_OPTIONS.ircII (not defined)
```

```
=====
=====
The following variables will affect the build process of this package,
ircII-20040820.  Their current value is shown below:

* USE_INET6 = YES

You may want to abort the process now with CTRL-C and change their value
before continuing.  Be sure to run '/usr/pkg/bin/bmake clean' after
the changes.
=====
==> Patching for ircII-20040820
==> Applying pkgsrc patches for ircII-20040820
==> Overriding tools for ircII-20040820
==> Creating toolchain wrappers for ircII-20040820
==> Configuring for ircII-20040820
...
[configure output snipped]
...
==> Building for ircII-20040820
...
[compilation output snipped]
...
==> Installing for ircII-20040820
...
[installation output snipped]
...
==> [Automatic manual page handling]
==> Registering installation for ircII-20040820
==> Cleaning for ircII-20040820
#
```

As you can see, the only difference are the lines that tell you where the system is fetching the package's distfile from.

The pkgsrc system uses `ftp(1)` to download the files, which honors various environment variables, including `FTP_PASSIVE_MODE`, `FTP_PROXY`, and `FTP_PASSWORD`. You may need to set one or more of these if you are behind a firewall, or need to use an FTP/HTTP proxy. See `ftp(1)` for the complete list.

For users which cannot be connected all the time, the `bmake fetch` option is provided. Just run this command at the top level directory (`/usr/pkgsrc`) and the required files will be downloaded for you. This command will also work in the lower level categories, for example: `/usr/pkgsrc/net`. Note that if a package depends on libraries or other packages this will *not* fetch the distfiles of those packages as well.

Note: You can build all the packages in a category or as a whole by running `bmake` in the top level directory, just like the aforementioned `bmake fetch` method. This is dangerous, however, as some applications cannot co-exist. In other cases, some packages can install two different files with the same filename.

In some rare cases, users may need to acquire the tarballs from a site other than the `MASTER_SITES` (the location where files are downloaded from). You can override the `MASTER_SORT`, `MASTER_SORT_REGEX` and `INET_COUNTRY` options either within the `/etc/mk.conf`.

Note: Some packages allow (or even require) you to provide build options which can enable/disable parts of the application which are unneeded, certain security options, and other customizations. A few which come to mind are `www/mozilla`, `security/gpgme`, and `mail/sylpheed-claws`. To find out what build options the application you are installing requires type:

```
# bmake show-options
```

To change the build process, either change the values of `PKG_DEFAULT_OPTIONS` or `PKG_OPTIONS.PackageName` in `/etc/mk.conf` or on the commandline as so:

```
# bmake PKG_OPTIONS.ircII="-ssl"
```

An option is enabled if listed. It is disabled if it is prefixed by a minus sign.

4.5.2.2 Dealing with imake

Some applications that use `imake` (a part of the X Window System) do not work well with `PREFIX`, and will insist on installing under `/usr/X11R6`. Similarly, some Perl ports ignore `PREFIX` and install in the Perl tree. Making these applications respect `PREFIX` is a difficult or impossible job.

4.5.3 Removing Installed Packages

Now that you know how to install packages, you are probably wondering how to remove them, just in case you install one and later on decide that you installed the wrong program. We will remove our previous example (which was `ircII` for those of you not paying attention). As with installing packages, the first thing you must do is change to the package directory, `/usr/pkgsrc/chat/ircII`. After you change directories, you are ready to uninstall `ircII`. This is done with the `bmake deinstall` command:

```
# cd /usr/pkgsrc/chat/ircII
# make deinstall
==> Deinstalling for ircII-20040820
```

That was easy enough. You have removed `ircII` from your system. If you would like to reinstall it, you can do so by running `bmake reinstall` from the `/usr/pkgsrc/chat/ircII` directory.

The `bmake deinstall` and `bmake reinstall` sequence does not work once you have run `bmake clean`. If you want to deinstall a package after cleaning, use `pkg_delete(1)` as discussed in the `Pkgsrc` section of the Handbook.

4.5.4 Packages and Disk Space

Using the `pkgsrc` collection can definitely eat up your disk space. For this reason you should always remember to clean up the work directories using the `bmake clean` option. This will remove the `work` directory after a package

has been built, and installed. You can also remove the tar files from the `distfiles` directory, and remove the installed package when their use has delimited.

4.5.5 Upgrading Packages

Note: Once you have updated your pkgsrc collection, before attempting a package upgrade, you should check the `/usr/pkgsrc/UPDATING` file. This file describes various issues and additional steps users may encounter and need to perform when updating a port.

Keeping your packages up to date can be a tedious job. For instance, to upgrade a package you would go to the package directory, build the package, deinstall the old package, install the new package, and then clean up after the build. Imagine doing that for five packages, tedious right? This was a large problem for system administrators to deal with, and now we have utilities which do this for us. For instance the `pkg_chk` utility will do everything for you!

`pkg_chk` requires a few steps in order to work correctly. They are listed here.

```
# pkg_chk -g # make initial list of installed packages
# pkg_chk -r # remove all packages that are not up to date and packages that depend on them
# pkg_chk -a # install all missing packages (use binary packages, this is the default)
# pkg_chk -as # install all missing packages (build from source)
```

4.6 Post-installation Activities

After installing a new application you will normally want to read any documentation it may have included, edit any configuration files that are required, ensure that the application starts at boot time (if it is a daemon), and so on.

The exact steps you need to take to configure each application will obviously be different. However, if you have just installed a new application and are wondering “What now?” these tips might help:

- Use `pkg_info(1)` to find out which files were installed, and where. For example, if you have just installed `FooPackage` version 1.0.0, then this command

```
# pkg_info -L foopackage-1.0.0 | less
```

will show all the files installed by the package. Pay special attention to files in `man/` directories, which will be manual pages, `etc/` directories, which will be configuration files, and `doc/`, which will be more comprehensive documentation.

If you are not sure which version of the application was just installed, a command like this

```
# pkg_info | grep -i foopackage
```

will find all the installed packages that have `foopackage` in the package name. Replace `foopackage` in your command line as necessary.

- Once you have identified where the application’s manual pages have been installed, review them using `man(1)`. Similarly, look over the sample configuration files, and any additional documentation that may have been provided.

- If the application has a web site, check it for additional documentation, frequently asked questions, and so forth. If you are not sure of the web site address it may be listed in the output from

```
# pkg_info foopackage-1.0.0
```

A `WWW:` line, if present, should provide a URL for the application's web site.

- Packages that should start at boot (such as Internet servers) will usually install a sample script in `/usr/pkg/etc/rc.d`. You should review this script for correctness and edit or rename it if needed. See [Starting Services](#) for more information.

4.7 Dealing with Broken Packages

If you come across a package that does not work for you, there are a few things you can do, including:

1. Fix it! The pkgsrc Guide (<http://www.netbsd.org/Documentation/pkgsrc/>) includes detailed information on the “pkgsrc” infrastructure so that you can fix the occasional broken package or even submit your own!
2. Gripe—*by email only!* Send email to the maintainer of the package first. Type `bmake maintainer` or read the `Makefile` to find the maintainer's email address. Remember to include the name and version of the port (send the `$NetBSD:` line from the `Makefile`) and the output leading up to the error when you email the maintainer. If you do not get a response from the maintainer, you can try users (<http://leaf.dragonflybsd.org/mailarchive/>).
3. Grab the package from an FTP site near you. The “master” package collection is on `packages.stura.uni-rostock.de` in the All directory (<ftp://packages.stura.uni-rostock.de/pkgsrc-current/DragonFly/RELEASE/i386/All/>). These are more likely to work than trying to compile from source and are a lot faster as well. Use the `pkg_add(1)` program to install the package on your system.

Chapter 5 The X Window System

Updated for X.Org's X11 server by Ken Tom and Marc Fonvieille. Updated for DragonFly by Víctor Balada Díaz.

5.1 Synopsis

DragonFly uses X11 to provide users with a powerful graphical user interface. X11 is an open-source implementation of the X Window System that includes both **X.org** and **XFree86**. DragonFly default official flavor is **X.org**, the X11 server developed by the X.Org Foundation.

This chapter will cover the installation and configuration of X11 with emphasis on **X.org**.

For more information on the video hardware that X11 supports, check either the X.org (<http://www.x.org/>) or XFree86 (<http://www.XFree86.org/>) web sites.

After reading this chapter, you will know:

- The various components of the X Window System, and how they interoperate.
- How to install and configure X11.
- How to install and use different window managers.
- How to use TrueType® fonts in X11.
- How to set up your system for graphical logins (**XDM**).

Before reading this chapter, you should:

- Know how to install additional third-party software (Chapter 4).

Note: This chapter covers the installation and the configuration of both **X.org** and **XFree86** X11 servers. For the most part, configuration files, commands and syntaxes are identical. In the case where there are differences, both **X.org** and **XFree86** syntaxes will be shown.

5.2 Understanding X

Using X for the first time can be somewhat of a shock to someone familiar with other graphical environments, such as Microsoft Windows or Mac OS.

While it is not necessary to understand all of the details of various X components and how they interact, some basic knowledge makes it possible to take advantage of X's strengths.

5.2.1 Why X?

X is not the first window system written for UNIX, but it is the most popular of them. X's original development team had worked on another window system prior to writing X. That system's name was "W" (for "Window"). X was just the next letter in the Roman alphabet.

X can be called “X”, “X Window System”, “X11”, and a number of other terms. You may find that using the term “X Windows” to describe X11 can be offensive to some people; for a bit more insight on this, see X(7).

5.2.2 The X Client/Server Model

X was designed from the beginning to be network-centric, and adopts a “client-server” model.

In the X model, the “X server” runs on the computer that has the keyboard, monitor, and mouse attached. The server’s responsibility includes tasks such as managing the display, handling input from the keyboard and mouse, and so on. Each X application (such as **XTerm**, or **Netscape**) is a “client”. A client sends messages to the server such as “Please draw a window at these coordinates”, and the server sends back messages such as “The user just clicked on the OK button”.

In a home or small office environment, the X server and the X clients commonly run on the same computer. However, it is perfectly possible to run the X server on a less powerful desktop computer, and run X applications (the clients) on, say, the powerful and expensive machine that serves the office. In this scenario the communication between the X client and server takes place over the network.

This confuses some people, because the X terminology is exactly backward to what they expect. They expect the “X server” to be the big powerful machine down the hall, and the “X client” to be the machine on their desk.

It is important to remember that the X server is the machine with the monitor and keyboard, and the X clients are the programs that display the windows.

There is nothing in the protocol that forces the client and server machines to be running the same operating system, or even to be running on the same type of computer. It is certainly possible to run an X server on Microsoft Windows or Apple’s Mac OS, and there are various free and commercial applications available that do exactly that.

DragonFly will use by default **X.org** server. **X.org** is available for free, under a license very similar to the DragonFly license.

5.2.3 The Window Manager

The X design philosophy is much like the UNIX design philosophy, “tools, not policy”. This means that X does not try to dictate how a task is to be accomplished. Instead, tools are provided to the user, and it is the user’s responsibility to decide how to use those tools.

This philosophy extends to X not dictating what windows should look like on screen, how to move them around with the mouse, what keystrokes should be used to move between windows (i.e., **Alt+Tab**, in the case of Microsoft Windows), what the title bars on each window should look like, whether or not they have close buttons on them, and so on.

Instead, X delegates this responsibility to an application called a “Window Manager”. There are dozens of window managers available for X: **AfterStep**, **Blackbox**, **ctwm**, **Enlightenment**, **fvwm**, **Sawfish**, **twm**, **Window Maker**, and more. Each of these window managers provides a different look and feel; some of them support “virtual desktops”; some of them allow customized keystrokes to manage the desktop; some have a “Start” button or similar device; some are “themeable”, allowing a complete change of look-and-feel by applying a new theme. These window managers, and many more, are available in the `x11-wm` category of the Ports Collection.

In addition, the **KDE** and **GNOME** desktop environments both have their own window managers which integrate with the desktop.

Each window manager also has a different configuration mechanism; some expect configuration file written by hand, others feature GUI tools for most of the configuration tasks; at least one (**Sawfish**) has a configuration file written in a dialect of the Lisp language.

Focus Policy: Another feature the window manager is responsible for is the mouse “focus policy”. Every windowing system needs some means of choosing a window to be actively receiving keystrokes, and should visibly indicate which window is active as well.

A familiar focus policy is called “click-to-focus”. This is the model utilized by Microsoft Windows, in which a window becomes active upon receiving a mouse click.

X does not support any particular focus policy. Instead, the window manager controls which window has the focus at any one time. Different window managers will support different focus methods. All of them support click to focus, and the majority of them support several others.

The most popular focus policies are:

focus-follows-mouse

The window that is under the mouse pointer is the window that has the focus. This may not necessarily be the window that is on top of all the other windows. The focus is changed by pointing at another window, there is no need to click in it as well.

sloppy-focus

This policy is a small extension to focus-follows-mouse. With focus-follows-mouse, if the mouse is moved over the root window (or background) then no window has the focus, and keystrokes are simply lost. With sloppy-focus, focus is only changed when the cursor enters a new window, and not when exiting the current window.

click-to-focus

The active window is selected by mouse click. The window may then be “raised”, and appear in front of all other windows. All keystrokes will now be directed to this window, even if the cursor is moved to another window.

Many window managers support other policies, as well as variations on these. Be sure to consult the documentation for the window manager itself.

5.2.4 Widgets

The X approach of providing tools and not policy extends to the widgets seen on screen in each application.

“Widget” is a term for all the items in the user interface that can be clicked or manipulated in some way; buttons, check boxes, radio buttons, icons, lists, and so on. Microsoft Windows calls these “controls”.

Microsoft Windows and Apple’s Mac OS both have a very rigid widget policy. Application developers are supposed to ensure that their applications share a common look and feel. With X, it was not considered sensible to mandate a particular graphical style, or set of widgets to adhere to.

As a result, do not expect X applications to have a common look and feel. There are several popular widget sets and variations, including the original Athena widget set from MIT, **Motif**® (on which the widget set in Microsoft Windows was modeled, all bevelled edges and three shades of grey), **OpenLook**, and others.

Most newer X applications today will use a modern-looking widget set, either Qt, used by **KDE**, or GTK+, used by the **GNOME** project. In this respect, there is some convergence in look-and-feel of the UNIX desktop, which certainly makes things easier for the novice user.

5.3 Installing X11

X.org or **XFree86** may be installed on DragonFly. DragonFly doesn't force a default implementation, but recommends **X.org**. **X.org** is the X server of the open source X Window System implementation released by the X.Org Foundation. **X.org** is based on the code of **XFree86 4.4RC2** and X11R6.6. The X.Org Foundation released X11R6.7 in April 2004 and X11R6.8.2 in February 2005, this latter is the version currently available in the DragonFly pkgsrc framework.

To build and install **X.org** from the Ports Collection:

```
# cd /usr/pkgsrc/meta-pkgs/xorg
# bmake install clean
```

Note: To build **X.org** in its entirety, be sure to have at least 4 GB of free space available.

To build and install **XFree86** from the pkgsrc framework:

```
# echo "X11_TYPE=XFree86" >> /etc/mk.conf

# cd /usr/pkgsrc/meta-pkgs/XFree86
# bmake install clean
```

Alternatively, X11 can be installed directly from packages. Binary packages to use with `pkg_add(1)` tool are also available for X11. If you have configured `PKG_PATH` the remote fetching feature of `pkg_add(1)` is used, the version number of the package is not required. `pkg_add(1)` will automatically fetch the latest version of the application.

So to fetch and install the package of **X.org**, simply type:

```
# pkg_add xorg
```

The **XFree86 4.X** package can be installed by typing:

```
# pkg_add XFree86
```

Note: The examples above will install the complete X11 distribution including the servers, clients, fonts etc. Separate packages and ports of X11 are also available.

The rest of this chapter will explain how to configure X11, and how to set up a productive desktop environment.

5.4 X11 Configuration

Contributed by Christopher Shumway.

5.4.1 Before Starting

Before configuration of X11 the following information about the target system is needed:

- Monitor specifications
- Video Adapter chipset
- Video Adapter memory

The specifications for the monitor are used by X11 to determine the resolution and refresh rate to run at. These specifications can usually be obtained from the documentation that came with the monitor or from the manufacturer's website. There are two ranges of numbers that are needed, the horizontal scan rate and the vertical synchronization rate.

The video adapter's chipset defines what driver module X11 uses to talk to the graphics hardware. With most chipsets, this can be automatically determined, but it is still useful to know in case the automatic detection does not work correctly.

Video memory on the graphic adapter determines the resolution and color depth which the system can run at. This is important to know so the user knows the limitations of the system.

5.4.2 Configuring X11

Configuration of X11 is a multi-step process. The first step is to build an initial configuration file. As the super user, simply run:

```
# Xorg -configure
```

In the case of **XFree86** type:

```
# XFree86 -configure
```

This will generate an X11 configuration skeleton file in the `/root` directory called `xorg.conf.new` (whether you `su(1)` or do a direct login affects the inherited supervisor `$HOME` directory variable). For **XFree86**, this configuration file is called `XF86Config.new`. The X11 program will attempt to probe the graphics hardware on the system and write a configuration file to load the proper drivers for the detected hardware on the target system.

The next step is to test the existing configuration to verify that **X.org** can work with the graphics hardware on the target system. To perform this task, type:

```
# Xorg -config xorg.conf.new
```

XFree86 users will type:

```
# XFree86 -xf86config XF86Config.new
```

If a black and grey grid and an X mouse cursor appear, the configuration was successful. To exit the test, just press **Ctrl+Alt+Backspace** simultaneously.

Note: If the mouse does not work, you will need to first configure it before proceeding.

Next, tune the `xorg.conf.new` (or `XF86Config.new` if you are running **XFree86**) configuration file to taste. Open the file in a text editor such as `emacs(1)` or `ee(1)`. First, add the frequencies for the target system's monitor. These are usually expressed as a horizontal and vertical synchronization rate. These values are added to the `xorg.conf.new` file under the "Monitor" section:

```
Section "Monitor"
    Identifier      "Monitor0"
    VendorName     "Monitor Vendor"
    ModelName      "Monitor Model"
    HorizSync      30-107
    VertRefresh    48-120
EndSection
```

The `HorizSync` and `VertRefresh` keywords may be missing in the configuration file. If they are, they need to be added, with the correct horizontal synchronization rate placed after the `HorizSync` keyword and the vertical synchronization rate after the `VertRefresh` keyword. In the example above the target monitor's rates were entered.

X allows DPMS (Energy Star) features to be used with capable monitors. The `xset(1)` program controls the time-outs and can force standby, suspend, or off modes. If you wish to enable DPMS features for your monitor, you must add the following line to the monitor section:

```
Option          "DPMS"
```

While the `xorg.conf.new` (or `XF86Config.new`) configuration file is still open in an editor, select the default resolution and color depth desired. This is defined in the "Screen" section:

```
Section "Screen"
    Identifier      "Screen0"
    Device         "Card0"
    Monitor        "Monitor0"
    DefaultDepth   24
    SubSection     "Display"
        Viewport   0 0
        Depth      24
        Modes      "1024x768"
    EndSubSection
EndSection
```

The `DefaultDepth` keyword describes the color depth to run at by default. This can be overridden with the `-depth` command line switch to `Xorg(1)` (or `XFree86(1)`). The `Modes` keyword describes the resolution to run at for the given color depth. Note that only VESA standard modes are supported as defined by the target system's graphics hardware. In the example above, the default color depth is twenty-four bits per pixel. At this color depth, the accepted resolution is 1024 by 768 pixels.

Finally, write the configuration file and test it using the test mode given above.

Note: One of the tools available to assist you during troubleshooting process are the X11 log files, which contain information on each device that the X11 server attaches to. **X.org** log file names are in the format of

`/var/log/Xorg.0.log` (**XFree86** log file names follow the format of `XFree86.0.log`). The exact name of the log can vary from `Xorg.0.log` to `Xorg.8.log` and so forth.

If all is well, the configuration file needs to be installed in a common location where `Xorg(1)` (or `XFree86(1)`) can find it. This is typically `/etc/X11/xorg.conf` or `/usr/pkg/xorg/lib/X11/xorg.conf` (for **XFree86** it is called `/etc/X11/XF86Config` or `/usr/pkg/XFree86/lib/X11/XF86Config`).

```
# cp xorg.conf.new /etc/X11/xorg.conf
```

For **XFree86**:

```
# cp XF86Config.new /etc/X11/XF86Config
```

The X11 configuration process is now complete. You can start **XFree86 4.X** or **X.org** with `startx(1)`. The X11 server may also be started with the use of `xdm(1)`.

Note: There is also a graphical configuration tool, `xorgcfg(1)` (`xf86cfg(1)` for **XFree86**), that comes with the X11 distribution. It allows you to interactively define your configuration by choosing the appropriate drivers and settings. This program can be invoked from the console, by typing the command `xorgcfg -textmode`. For more details, refer to the `xorgcfg(1)` and `xf86cfg(1)` manual pages.

Alternatively, there is also a tool called `xorgconfig(1)` (`xf86config(1)` for **XFree86**), this program is a console utility that is less user friendly, but it may work in situations where the other tools do not.

5.4.3 Advanced Configuration Topics

5.4.3.1 Configuration with Intel® i810 Graphics Chipsets

Configuration with Intel® i810 integrated chipsets requires the `agpgart` AGP programming interface for X11 to drive the card. See the `agp(4)` driver manual page for more information.

This will allow configuration of the hardware as any other graphics board. Note on systems without the `agp(4)` driver compiled in the kernel, trying to load the module with `kldload(8)` will not work. This driver has to be in the kernel at boot time through being compiled in or using `/boot/loader.conf`.

If you are using **XFree86 4.1.0** (or later) and messages about unresolved symbols like `fbPictureInit` appear, try adding the following line after `Driver "i810"` in the X11 configuration file:

```
Option "NoDDC"
```

5.5 Using Fonts in X11

Contributed by Murray Stokely.

5.5.1 Type1 Fonts

The default fonts that ship with X11 are less than ideal for typical desktop publishing applications. Large presentation fonts show up jagged and unprofessional looking, and small fonts in **Netscape** are almost completely unintelligible. However, there are several free, high quality Type1 (PostScript®) fonts available which can be readily used with X11. For instance, the Freefonts collection (`fonts/freefonts`) includes a lot of fonts, but most of them are intended for use in graphics software such as the **Gimp**, and are not complete enough to serve as screen fonts. In addition, X11 can be configured to use TrueType fonts with a minimum of effort. For more details on this, see the X(7) manual page or the section on TrueType fonts.

To install the Freefonts font collection from the pkgsrc framework, run the following commands:

```
# cd /usr/pkgsrc/fonts/freefonts
# bmake install clean
```

And likewise with the other collections. To have the X server detect these fonts, add an appropriate line to the X server configuration file in `/etc/X11/` (`xorg.conf` for **X.org** and `XF86Config` for **XFree86**), which reads:

```
FontPath "/usr/pkg/lib/X11/fonts/freefont/"
```

Alternatively, at the command line in the X session run:

```
% xset fp+ /usr/pkg/lib/X11/fonts/freefont/
% xset fp rehash
```

This will work but will be lost when the X session is closed, unless it is added to the startup file (`~/.xinitrc` for a normal `startx` session, or `~/.xsession` when logging in through a graphical login manager like **XDM**). A third way is to use the new `/usr/pkg/xorg/etc/fonts/local.conf` file: see the section on anti-aliasing.

5.5.2 TrueType® Fonts

Both **XFree86 4.X** and **X.org** have built in support for rendering TrueType fonts. There are two different modules that can enable this functionality. The freetype module is used in this example because it is more consistent with the other font rendering back-ends. To enable the freetype module just add the following line to the "Module" section of the `/etc/X11/xorg.conf` or `/etc/X11/XF86Config` file.

```
Load "freetype"
```

For **XFree86 3.3.X**, a separate TrueType font server is needed. **Xfstt** is commonly used for this purpose. To install **Xfstt**, simply install the port `x11/xfstt`.

Now make a directory for the TrueType fonts (for example, `/usr/pkg/xorg/lib/X11/fonts/TrueType`) and copy all of the TrueType fonts into this directory. Keep in mind that TrueType fonts cannot be directly taken from a Macintosh®; they must be in UNIX/MS-DOS/Windows format for use by X11. Once the files have been copied into this directory, use **ttmkfdir** to create a `fonts.dir` file, so that the X font renderer knows that these new files have been installed. `ttmkfdir` is available from the pkgsrc framework as `fonts/ttmkfdir2`.


```
# cd /usr/pkg/xorg/lib/X11/fonts/TrueType
# ttmkfdir > fonts.dir
```

Now add the TrueType directory to the font path. This is just the same as described above for Type1 fonts, that is, use

```
% xset fp+ /usr/pkg/xorg/lib/X11/fonts/TrueType
% xset fp rehash
```

or add a `FontPath` line to the `xorg.conf` (or `XF86Config`) file.

That's it. Now **Netscape**, **Gimp**, **StarOffice**[™], and all of the other X applications should now recognize the installed TrueType fonts. Extremely small fonts (as with text in a high resolution display on a web page) and extremely large fonts (within **StarOffice**) will look much better now.

5.5.3 Anti-Aliased Fonts

Updated by Joe Marcus Clarke.

Anti-aliasing has been available in X11 since **XFree86** 4.0.2. However, font configuration was cumbersome before the introduction of **XFree86** 4.3.0. Beginning with **XFree86** 4.3.0, all fonts in X11 that are found in `/usr/pkg/xorg/lib/X11/fonts/` and `~/.fonts/` are automatically made available for anti-aliasing to Xft-aware applications. Not all applications are Xft-aware, but many have received Xft support. Examples of Xft-aware applications include Qt 2.3 and higher (the toolkit for the **KDE** desktop), GTK+ 2.0 and higher (the toolkit for the **GNOME** desktop), and **Mozilla** 1.2 and higher.

In order to control which fonts are anti-aliased, or to configure anti-aliasing properties, create (or edit, if it already exists) the file `/usr/pkg/xorg/lib/etc/fonts/local.conf`. Several advanced features of the Xft font system can be tuned using this file; this section describes only some simple possibilities. For more details, please see `fonts-conf(5)`.

This file must be in XML format. Pay careful attention to case, and make sure all tags are properly closed. The file begins with the usual XML header followed by a DOCTYPE definition, and then the `<fontconfig>` tag:

```
<?xml version="1.0"?>
<!DOCTYPE fontconfig SYSTEM "fonts.dtd">
<fontconfig>
```

As previously stated, all fonts in `/usr/pkg/xorg/lib/X11/fonts/` as well as `~/.fonts/` are already made available to Xft-aware applications. If you wish to add another directory outside of these two directory trees, add a line similar to the following to `/usr/pkg/lib/etc/fonts/local.conf`:

```
<dir>/path/to/my/fonts</dir>
```

After adding new fonts, and especially new font directories, you should run the following command to rebuild the font caches:

```
# fc-cache -f
```

Anti-aliasing makes borders slightly fuzzy, which makes very small text more readable and removes “staircases” from large text, but can cause eyestrain if applied to normal text. To exclude font sizes smaller than 14 point from anti-aliasing, include these lines:

```

<match target="font">
  <test name="size" compare="less">
    <double>14</double>
  </test>
  <edit name="antialias" mode="assign">
    <bool>>false</bool>
  </edit>
</match>
<match target="font">
  <test name="pixelsize" compare="less" qual="any">
    <double>14</double>
  </test>
  <edit mode="assign" name="antialias">
    <bool>>false</bool>
  </edit>
</match>

```

Spacing for some monospaced fonts may also be inappropriate with anti-aliasing. This seems to be an issue with **KDE**, in particular. One possible fix for this is to force the spacing for such fonts to be 100. Add the following lines:

```

<match target="pattern" name="family">
  <test qual="any" name="family">
    <string>fixed</string>
  </test>
  <edit name="family" mode="assign">
    <string>mono</string>
  </edit>
</match>
<match target="pattern" name="family">
  <test qual="any" name="family">
    <string>console</string>
  </test>
  <edit name="family" mode="assign">
    <string>mono</string>
  </edit>
</match>

```

(this aliases the other common names for fixed fonts as "mono"), and then add:

```

<match target="pattern" name="family">
  <test qual="any" name="family">
    <string>mono</string>
  </test>
  <edit name="spacing" mode="assign">
    <int>100</int>
  </edit>
</match>

```

Certain fonts, such as Helvetica, may have a problem when anti-aliased. Usually this manifests itself as a font that seems cut in half vertically. At worst, it may cause applications such as **Mozilla** to crash. To avoid this, consider adding the following to `local.conf`:

```

<match target="pattern" name="family">

```

```

<test qual="any" name="family">
  <string>Helvetica</string>
</test>
<edit name="family" mode="assign">
  <string>sans-serif</string>
</edit>
</match>

```

Once you have finished editing `local.conf` make sure you end the file with the `</fontconfig>` tag. Not doing this will cause your changes to be ignored.

The default font set that comes with X11 is not very desirable when it comes to anti-aliasing. A much better set of default fonts can be found in the `fonts/vera-ttf` port. This port will install a `/usr/pkg/lib/etc/fonts/local.conf` file if one does not exist already. If the file does exist, the port will create a `/usr/pkg/lib/etc/fonts/local.conf-vera` file. Merge the contents of this file into `/usr/pkg/lib/etc/fonts/local.conf`, and the Bitstream fonts will automatically replace the default X11 Serif, Sans Serif, and Monospaced fonts.

Finally, users can add their own settings via their personal `.fonts.conf` files. To do this, each user should simply create a `~/.fonts.conf`. This file must also be in XML format.

One last point: with an LCD screen, sub-pixel sampling may be desired. This basically treats the (horizontally separated) red, green and blue components separately to improve the horizontal resolution; the results can be dramatic. To enable this, add the line somewhere in the `local.conf` file:

```

<match target="font">
  <test qual="all" name="rgba">
    <const>unknown</const>
  </test>
  <edit name="rgba" mode="assign">
    <const>rgb</const>
  </edit>
</match>

```

Note: Depending on the sort of display, `rgb` may need to be changed to `bgr`, `vrgb` or `vbgr`: experiment and see which works best.

Anti-aliasing should be enabled the next time the X server is started. However, programs must know how to take advantage of it. At present, the Qt toolkit does, so the entire **KDE** environment can use anti-aliased fonts. **GTK+** and **GNOME** can also be made to use anti-aliasing via the “Font” capplet (see Section 5.7.1.3 for details). By default, **Mozilla** 1.2 and greater will automatically use anti-aliasing. To disable this, rebuild **Mozilla** with the `-DWITHOUT_XFT` flag.

5.6 The X Display Manager

Contributed by Seth Kingsley.

5.6.1 Overview

The X Display Manager (**XDM**) is an optional part of the X Window System that is used for login session management. This is useful for several types of situations, including minimal “X Terminals”, desktops, and large network display servers. Since the X Window System is network and protocol independent, there are a wide variety of possible configurations for running X clients and servers on different machines connected by a network. **XDM** provides a graphical interface for choosing which display server to connect to, and entering authorization information such as a login and password combination.

Think of **XDM** as providing the same functionality to the user as the `getty(8)` utility (see Section 17.3.2 for details). That is, it performs system logins to the display being connected to and then runs a session manager on behalf of the user (usually an X window manager). **XDM** then waits for this program to exit, signaling that the user is done and should be logged out of the display. At this point, **XDM** can display the login and display chooser screens for the next user to login.

5.6.2 Using XDM

The **XDM** daemon program is located in `/usr/pkg/xorg/bin/xdm`. This program can be run at any time as `root` and it will start managing the X display on the local machine. If **XDM** is to be run every time the machine boots up, a convenient way to do this is by adding an entry to `/etc/ttys`. For more information about the format and usage of this file, see Section 17.3.2.1. There is a line in the default `/etc/ttys` file for running the **XDM** daemon on a virtual terminal:

```
tttyv8  "/usr/pkg/xorg/bin/xdm -nodaemon"  xterm  off secure
```

By default this entry is disabled; in order to enable it change field 5 from `off` to `on` and restart `init(8)` using the directions in Section 17.3.2.2. The first field, the name of the terminal this program will manage, is `tttyv8`. This means that **XDM** will start running on the 9th virtual terminal.

5.6.3 Configuring XDM

The **XDM** configuration directory is located in `/var/lib/xdm`. The sample configuration files are in `/usr/pkg/share/examples/xorg/xdm/`, in this directory there are several files used to change the behavior and appearance of **XDM**. Typically these files will be found:

File	Description
<code>Xaccess</code>	Client authorization ruleset.
<code>Xresources</code>	Default X resource values.
<code>Xservers</code>	List of remote and local displays to manage.
<code>Xsession</code>	Default session script for logins.
<code>Xsetup_*</code>	Script to launch applications before the login interface.
<code>xdm-config</code>	Global configuration for all displays running on this machine.

File	Description
<code>xm-errors</code>	Errors generated by the server program.
<code>xm-pid</code>	The process ID of the currently running XDM.

Also in this directory are a few scripts and programs used to set up the desktop when **XDM** is running. The purpose of each of these files will be briefly described. The exact syntax and usage of all of these files is described in `xm(1)`.

The default configuration is a simple rectangular login window with the hostname of the machine displayed at the top in a large font and “Login:” and “Password:” prompts below. This is a good starting point for changing the look and feel of **XDM** screens.

5.6.3.1 Xaccess

The protocol for connecting to **XDM** controlled displays is called the X Display Manager Connection Protocol (XDMCP). This file is a ruleset for controlling XDMCP connections from remote machines. It is ignored unless the `xm-config` is changed to listen for remote connections. By default, it does not allow any clients to connect.

5.6.3.2 Xresources

This is an application-defaults file for the display chooser and the login screens. This is where the appearance of the login program can be modified. The format is identical to the app-defaults file described in the X11 documentation.

5.6.3.3 Xservers

This is a list of the remote displays the chooser should provide as choices.

5.6.3.4 Xsession

This is the default session script for **XDM** to run after a user has logged in. Normally each user will have a customized session script in `~/ .xsession` that overrides this script.

5.6.3.5 Xsetup_*

These will be run automatically before displaying the chooser or login interfaces. There is a script for each display being used, named `xsetup_` followed by the local display number (for instance `xsetup_0`). Typically these scripts will run one or two programs in the background such as `xconsole`.

5.6.3.6 xm-config

This contains settings in the form of app-defaults that are applicable to every display that this installation manages.

5.6.3.7 xm-errors

This contains the output of the X servers that **XDM** is trying to run. If a display that **XDM** is trying to start hangs for some reason, this is a good place to look for error messages. These messages are also written to the user’s `~/ .xsession-errors` file on a per-session basis.

5.6.4 Running a Network Display Server

In order for other clients to connect to the display server, edit the access control rules, and enable the connection listener. By default these are set to conservative values. To make **XDM** listen for connections, first comment out a line in the `xdm-config` file:

```
! SECURITY: do not listen for XDMCP or Chooser requests
! Comment out this line if you want to manage X terminals with xdm
DisplayManager.requestPort:      0
```

and then restart **XDM**. Remember that comments in app-defaults files begin with a “!” character, not the usual “#”. More strict access controls may be desired. Look at the example entries in `Xaccess`, and refer to the `xdm(1)` manual page.

5.6.5 Replacements for XDM

Several replacements for the default **XDM** program exist. One of them, **kdm** (bundled with **KDE**) is described later in this chapter. The **kdm** display manager offers many visual improvements and cosmetic frills, as well as the functionality to allow users to choose their window manager of choice at login time.

5.7 Desktop Environments

Contributed by Valentino Vaschetto.

This section describes the different desktop environments available for X on FreeBSD. A “desktop environment” can mean anything ranging from a simple window manager to a complete suite of desktop applications, such as **KDE** or **GNOME**.

5.7.1 GNOME

5.7.1.1 About GNOME

GNOME is a user-friendly desktop environment that enables users to easily use and configure their computers. **GNOME** includes a panel (for starting applications and displaying status), a desktop (where data and applications can be placed), a set of standard desktop tools and applications, and a set of conventions that make it easy for applications to cooperate and be consistent with each other. Users of other operating systems or environments should feel right at home using the powerful graphics-driven environment that **GNOME** provides.

5.7.1.2 Installing GNOME

GNOME can be easily installed from a package or from the `pkgsrc` framework:

To install the **GNOME** package from the network, simply type:

```
# pkg_add gnome
```

To build **GNOME** from source, use the ports tree:

```
# cd /usr/pkgsrc/meta-pkgs/gnome
# bmake install clean
```

Once **GNOME** is installed, the X server must be told to start **GNOME** instead of a default window manager.

The easiest way to start **GNOME** is with **GDM**, the GNOME Display Manager. **GDM**, which is installed as a part of the **GNOME** desktop (but is disabled by default), can be enabled by adding `gdm_enable="YES"` to `/etc/rc.conf`. Once you have rebooted, **GNOME** will start automatically once you log in — no further configuration is necessary.

GNOME may also be started from the command-line by properly configuring a file named `.xinitrc`. If a custom `.xinitrc` is already in place, simply replace the line that starts the current window manager with one that starts `/usr/pkg/bin/gnome-session` instead. If nothing special has been done to the configuration file, then it is enough simply to type:

```
% echo "/usr/pkg/bin/gnome-session" > ~/.xinitrc
```

Next, type `startx`, and the **GNOME** desktop environment will be started.

Note: If an older display manager, like **XDM**, is being used, this will not work. Instead, create an executable `.xsession` file with the same command in it. To do this, edit the file and replace the existing window manager command with `/usr/pkg/bin/gnome-session`:

```
% echo "#!/bin/sh" > ~/.xsession
% echo "/usr/pkg/bin/gnome-session" >> ~/.xsession
% chmod +x ~/.xsession
```

Yet another option is to configure the display manager to allow choosing the window manager at login time; the section on KDE details explains how to do this for **kdm**, the display manager of **KDE**.

5.7.1.3 Anti-aliased Fonts with GNOME

X11 supports anti-aliasing via its “RENDER” extension. GTK+ 2.0 and greater (the toolkit used by **GNOME**) can make use of this functionality. Configuring anti-aliasing is described in Section 5.5.3. So, with up-to-date software, anti-aliasing is possible within the **GNOME** desktop. Just go to Applications → Desktop Preferences → Font, and select either Best shapes, Best contrast, or Subpixel smoothing (LCDs). For a GTK+ application that is not part of the **GNOME** desktop, set the environment variable `GDK_USE_XFT` to 1 before launching the program.

5.7.2 KDE

5.7.2.1 About KDE

KDE is an easy to use contemporary desktop environment. Some of the things that **KDE** brings to the user are:

- A beautiful contemporary desktop
- A desktop exhibiting complete network transparency

- An integrated help system allowing for convenient, consistent access to help on the use of the **KDE** desktop and its applications
- Consistent look and feel of all **KDE** applications
- Standardized menu and toolbars, keybindings, color-schemes, etc.
- Internationalization: **KDE** is available in more than 40 languages
- Centralized consisted dialog driven desktop configuration
- A great number of useful **KDE** applications

KDE comes with a web browser called **Konqueror**, which represents a solid competitor to other existing web browsers on UNIX systems. More information on **KDE** can be found on the KDE website (<http://www.kde.org/>).

5.7.2.2 Installing KDE

Just as with **GNOME** or any other desktop environment, the easiest way to install **KDE** is through the `pkgsrc` framework or from a package:

To install the **KDE** package from the network, simply type:

```
# pkg_add kde
```

`pkg_add(1)` will automatically fetch the latest version of the application.

To build **KDE** from source, use the ports tree:

```
# cd /usr/pkgsrc/meta-pkgs/kde3
# bmake install clean
```

After **KDE** has been installed, the X server must be told to launch this application instead of the default window manager. This is accomplished by editing the `.xinitrc` file:

```
% echo "exec startkde" > ~/.xinitrc
```

Now, whenever the X Window System is invoked with `startx`, **KDE** will be the desktop.

If a display manager such as **XDM** is being used, the configuration is slightly different. Edit the `.xsession` file instead. Instructions for **kdm** are described later in this chapter.

5.7.3 More Details on KDE

Now that **KDE** is installed on the system, most things can be discovered through the help pages, or just by pointing and clicking at various menus. Windows or Mac® users will feel quite at home.

The best reference for **KDE** is the on-line documentation. **KDE** comes with its own web browser, **Konqueror**, dozens of useful applications, and extensive documentation. The remainder of this section discusses the technical items that are difficult to learn by random exploration.

5.7.3.1 The KDE Display Manager

An administrator of a multi-user system may wish to have a graphical login screen to welcome users. XDM can be used, as described earlier. However, **KDE** includes an alternative, **kdm**, which is designed to look more attractive and include more login-time options. In particular, users can easily choose (via a menu) which desktop environment (**KDE**, **GNOME**, or something else) to run after logging on.

To enable **kdm**, the `tttyv8` entry in `/etc/ttys` has to be adapted. The line should look as follows:

```
tttyv8 "/usr/pkg/bin/kdm -nodaemon" xterm on secure
```

5.7.4 XFce

5.7.4.1 About XFce

XFce is a desktop environment based on the GTK+ toolkit used by **GNOME**, but is much more lightweight and meant for those who want a simple, efficient desktop which is nevertheless easy to use and configure. Visually, it looks very much like **CDE**, found on commercial UNIX systems. Some of **XFce**'s features are:

- A simple, easy-to-handle desktop
- Fully configurable via mouse, with drag and drop, etc
- Main panel similar to **CDE**, with menus, applets and applications launchers
- Integrated window manager, file manager, sound manager, **GNOME** compliance module, and other things
- Themeable (since it uses GTK+)
- Fast, light and efficient: ideal for older/slower machines or machines with memory limitations

More information on **XFce** can be found on the XFce website (<http://www.xfce.org/>).

5.7.4.2 Installing XFce

A binary package for **XFce** exists (at the time of writing). To install, simply type:

```
# pkg_add -r xfce4
```

Alternatively, to build from source, use the pkgsrc framework:

```
# cd /usr/pkgsrc/meta-pkgs/xfce4
# make install clean
```

Now, tell the X server to launch **XFce** the next time X is started. Simply type this:

```
% echo "/usr/pkgsrc/bin/startxfce4" > ~/.xinitrc
```

The next time X is started, **XFce** will be the desktop. As before, if a display manager like **XDM** is being used, create an `.xsession`, as described in the section on **GNOME**, but with the `/usr/pkg/bin/startxfce4` command; or, configure the display manager to allow choosing a desktop at login time, as explained in the section on **kdm**.

II. System Administration

The remaining chapters of the DragonFly Handbook cover all aspects of DragonFly system administration. Each chapter starts by describing what you will learn as a result of reading the chapter, and also details what you are expected to know before tackling the material.

These chapters are designed to be read when you need the information. You do not have to read them in any particular order, nor do you need to read all of them before you can begin using DragonFly.

Chapter 6 Configuration and Tuning

Written by Chern Lee. Based on a tutorial written by Mike Smith. Also based on tuning(7) written by Matt Dillon.

6.1 Synopsis

One of the important aspects of DragonFly is system configuration. Correct system configuration will help prevent headaches during future upgrades. This chapter will explain much of the DragonFly configuration process, including some of the parameters which can be set to tune a DragonFly system.

After reading this chapter, you will know:

- How to efficiently work with file systems and swap partitions.
- The basics of `rc.conf` configuration and `rc.d` startup systems.
- How to configure and test a network card.
- How to configure virtual hosts on your network devices.
- How to use the various configuration files in `/etc`.
- How to tune DragonFly using `sysctl` variables.
- How to tune disk performance and modify kernel limitations.

Before reading this chapter, you should:

- Understand UNIX and DragonFly basics (Chapter 3).
- Be familiar with the basics of kernel configuration/compilation (Chapter 9).

6.2 Initial Configuration

6.2.1 Partition Layout

6.2.1.1 Base Partitions

When laying out file systems with `disklabel(8)` remember that hard drives transfer data faster from the outer tracks to the inner. Thus smaller and heavier-accessed file systems should be closer to the outside of the drive, while larger partitions like `/usr` should be placed toward the inner. It is a good idea to create partitions in a similar order to: `root`, `swap`, `/var`, `/usr`.

The size of `/var` reflects the intended machine usage. `/var` is used to hold mailboxes, log files, and printer spools. Mailboxes and log files can grow to unexpected sizes depending on how many users exist and how long log files are kept. Most users would never require a gigabyte, but remember that `/var/tmp` must be large enough to contain packages.

The `/usr` partition holds much of the files required to support the system, the `pkgsrc` collection (recommended) and the source code (optional). At least 2 gigabytes would be recommended for this partition.

When selecting partition sizes, keep the space requirements in mind. Running out of space in one partition while barely using another can be a hassle.

6.2.1.2 Swap Partition

As a rule of thumb, the swap partition should be about double the size of system memory (RAM). For example, if the machine has 128 megabytes of memory, the swap file should be 256 megabytes. Systems with less memory may perform better with more swap. Less than 256 megabytes of swap is not recommended and memory expansion should be considered. The kernel's VM paging algorithms are tuned to perform best when the swap partition is at least two times the size of main memory. Configuring too little swap can lead to inefficiencies in the VM page scanning code and might create issues later if more memory is added.

On larger systems with multiple SCSI disks (or multiple IDE disks operating on different controllers), it is recommend that a swap is configured on each drive (up to four drives). The swap partitions should be approximately the same size. The kernel can handle arbitrary sizes but internal data structures scale to 4 times the largest swap partition. Keeping the swap partitions near the same size will allow the kernel to optimally stripe swap space across disks. Large swap sizes are fine, even if swap is not used much. It might be easier to recover from a runaway program before being forced to reboot.

6.2.1.3 Why Partition?

Several users think a single large partition will be fine, but there are several reasons why this is a bad idea. First, each partition has different operational characteristics and separating them allows the file system to tune accordingly. For example, the root and `/usr` partitions are read-mostly, without much writing. While a lot of reading and writing could occur in `/var` and `/var/tmp`.

By properly partitioning a system, fragmentation introduced in the smaller write heavy partitions will not bleed over into the mostly-read partitions. Keeping the write-loaded partitions closer to the disk's edge, will increase I/O performance in the partitions where it occurs the most. Now while I/O performance in the larger partitions may be needed, shifting them more toward the edge of the disk will not lead to a significant performance improvement over moving `/var` to the edge. Finally, there are safety concerns. A smaller, neater root partition which is mostly read-only has a greater chance of surviving a bad crash.

6.3 Core Configuration

The principal location for system configuration information is within `/etc/rc.conf`. This file contains a wide range of configuration information, principally used at system startup to configure the system. Its name directly implies this; it is configuration information for the `rc*` files.

An administrator should make entries in the `rc.conf` file to override the default settings from `/etc/defaults/rc.conf`. The defaults file should not be copied verbatim to `/etc` - it contains default values, not examples. All system-specific changes should be made in the `rc.conf` file itself.

A number of strategies may be applied in clustered applications to separate site-wide configuration from system-specific configuration in order to keep administration overhead down. The recommended approach is to place

site-wide configuration into another file, such as `/etc/rc.conf.site`, and then include this file into `/etc/rc.conf`, which will contain only system-specific information.

As `rc.conf` is read by `sh(1)` it is trivial to achieve this. For example:

- `rc.conf`:

```
. rc.conf.site
hostname="node15.example.com"
network_interfaces="fxp0 lo0"
ifconfig_fxp0="inet 10.1.1.1"
```

- `rc.conf.site`:

```
defaultrouter="10.1.1.254"
saver="daemon"
blanktime="100"
```

The `rc.conf.site` file can then be distributed to every system using `rsync` or a similar program, while the `rc.conf` file remains unique.

Upgrading the system using `make world` will not overwrite the `rc.conf` file, so system configuration information will not be lost.

6.4 Application Configuration

Typically, installed applications have their own configuration files, with their own syntax, etc. It is important that these files be kept separate from the base system, so that they may be easily located and managed by the package management tools.

Typically, these files are installed in `/usr/local/etc`. In the case where an application has a large number of configuration files, a subdirectory will be created to hold them.

Normally, when a port or package is installed, sample configuration files are also installed. These are usually identified with a `.default` suffix. If there are no existing configuration files for the application, they will be created by copying the `.default` files.

For example, consider the contents of the directory `/usr/local/etc/apache`:

```
-rw-r--r--  1 root  wheel   2184 May 20  1998 access.conf
-rw-r--r--  1 root  wheel   2184 May 20  1998 access.conf.default
-rw-r--r--  1 root  wheel   9555 May 20  1998 httpd.conf
-rw-r--r--  1 root  wheel   9555 May 20  1998 httpd.conf.default
-rw-r--r--  1 root  wheel  12205 May 20  1998 magic
-rw-r--r--  1 root  wheel  12205 May 20  1998 magic.default
-rw-r--r--  1 root  wheel   2700 May 20  1998 mime.types
-rw-r--r--  1 root  wheel   2700 May 20  1998 mime.types.default
-rw-r--r--  1 root  wheel   7980 May 20  1998 srm.conf
-rw-r--r--  1 root  wheel   7933 May 20  1998 srm.conf.default
```

The file sizes show that only the `srm.conf` file has been changed. A later update of the **Apache** port would not overwrite this changed file.

6.5 Starting Services

It is common for a system to host a number of services. These may be started in several different fashions, each having different advantages.

Software installed from a port or the packages collection will often place a script in `/usr/local/etc/rc.d` which is invoked at system startup with a `start` argument, and at system shutdown with a `stop` argument. This is the recommended way for starting system-wide services that are to be run as `root`, or that expect to be started as `root`. These scripts are registered as part of the installation of the package, and will be removed when the package is removed.

A generic startup script in `/usr/local/etc/rc.d` looks like:

```
#!/bin/sh
echo -n ' FooBar'

case "$1" in
start)
    /usr/local/bin/foobar
    ;;
stop)
    kill -9 `cat /var/run/foobar.pid`
    ;;
*)
    echo "Usage: `basename $0` {start|stop}" >&2
    exit 64
    ;;
esac

exit 0
```

The startup scripts of DragonFly will look in `/usr/local/etc/rc.d` for scripts that have an `.sh` extension and are executable by `root`. Those scripts that are found are called with an option `start` at startup, and `stop` at shutdown to allow them to carry out their purpose. So if you wanted the above sample script to be picked up and run at the proper time during system startup, you should save it to a file called `FooBar.sh` in `/usr/local/etc/rc.d` and make sure it is executable. You can make a shell script executable with `chmod(1)` as shown below:

```
# chmod 755 FooBar.sh
```

Some services expect to be invoked by `inetd(8)` when a connection is received on a suitable port. This is common for mail reader servers (POP and IMAP, etc.). These services are enabled by editing the file `/etc/inetd.conf`. See `inetd(8)` for details on editing this file.

Some additional system services may not be covered by the toggles in `/etc/rc.conf`. These are traditionally enabled by placing the command(s) to invoke them in `/etc/rc.local` (which does not exist by default). Note that `rc.local` is generally regarded as the location of last resort; if there is a better place to start a service, do it there.

Note: Do *not* place any commands in `/etc/rc.conf`. To start daemons, or run any commands at boot time, place a script in `/usr/local/etc/rc.d` instead.

It is also possible to use the cron(8) daemon to start system services. This approach has a number of advantages, not least being that because cron(8) runs these processes as the owner of the `crontab`, services may be started and maintained by non-root users.

This takes advantage of a feature of cron(8): the time specification may be replaced by `@reboot`, which will cause the job to be run when cron(8) is started shortly after system boot.

6.6 Configuring the cron Utility

Contributed by Tom Rhodes.

One of the most useful utilities in DragonFly is cron(8). The cron utility runs in the background and constantly checks the `/etc/crontab` file. The cron utility also checks the `/var/cron/tabs` directory, in search of new crontab files. These crontab files store information about specific functions which cron is supposed to perform at certain times.

The cron utility uses two different types of configuration files, the system crontab and user crontabs. The only difference between these two formats is the sixth field. In the system crontab, the sixth field is the name of a user for the command to run as. This gives the system crontab the ability to run commands as any user. In a user crontab, the sixth field is the command to run, and all commands run as the user who created the crontab; this is an important security feature.

Note: User crontabs allow individual users to schedule tasks without the need for root privileges. Commands in a user's crontab run with the permissions of the user who owns the crontab.

The root user can have a user crontab just like any other user. This one is different from `/etc/crontab` (the system crontab). Because of the system crontab, there's usually no need to create a user crontab for root.

Let us take a look at the `/etc/crontab` file (the system crontab):

```
# /etc/crontab - root's crontab for DragonFly
#
# ❶
#
SHELL=/bin/sh
PATH=/etc:/bin:/sbin:/usr/bin:/usr/sbin ❷
HOME=/var/log
#
#
#minute hour mday month wday who command ❸
#
#
*/5 * * * * root /usr/libexec/atrun ❹
```

- ❶ Like most DragonFly configuration files, the # character represents a comment. A comment can be placed in the file as a reminder of what and why a desired action is performed. Comments cannot be on the same line as a command or else they will be interpreted as part of the command; they must be on a new line. Blank lines are ignored.

- ② First, the environment must be defined. The equals (=) character is used to define any environment settings, as with this example where it is used for the `SHELL`, `PATH`, and `HOME` options. If the shell line is omitted, `cron` will use the default, which is `sh`. If the `PATH` variable is omitted, no default will be used and file locations will need to be absolute. If `HOME` is omitted, `cron` will use the invoking users home directory.
- ③ This line defines a total of seven fields. Listed here are the values `minute`, `hour`, `mday`, `month`, `wday`, `who`, and `command`. These are almost all self explanatory. `minute` is the time in minutes the command will be run. `hour` is similar to the `minute` option, just in hours. `mday` stands for day of the month. `month` is similar to `hour` and `minute`, as it designates the month. The `wday` option stands for day of the week. All these fields must be numeric values, and follow the twenty-four hour clock. The `who` field is special, and only exists in the `/etc/crontab` file. This field specifies which user the command should be run as. When a user installs his or her `crontab` file, they will not have this option. Finally, the `command` option is listed. This is the last field, so naturally it should designate the command to be executed.
- ④ This last line will define the values discussed above. Notice here we have a `*/5` listing, followed by several more `*` characters. These `*` characters mean “first-last”, and can be interpreted as *every* time. So, judging by this line, it is apparent that the `atrun` command is to be invoked by `root` every five minutes regardless of what day or month it is. For more information on the `atrun` command, see the `atrun(8)` manual page.

Commands can have any number of flags passed to them; however, commands which extend to multiple lines need to be broken with the backslash “\” continuation character.

This is the basic set up for every `crontab` file, although there is one thing different about this one. Field number six, where we specified the username, only exists in the system `/etc/crontab` file. This field should be omitted for individual user `crontab` files.

6.6.1 Installing a Crontab

Important: You must not use the procedure described here to edit/install the system `crontab`. Simply use your favorite editor: the `cron` utility will notice that the file has changed and immediately begin using the updated version. If you use `crontab` to load the `/etc/crontab` file you may get an error like `root: not found` because of the system `crontab`'s additional user field.

To install a freshly written user `crontab`, first use your favorite editor to create a file in the proper format, and then use the `crontab` utility. The most common usage is:

```
% crontab crontab-file
```

In this example, `crontab-file` is the filename of a `crontab` that was previously created.

There is also an option to list installed `crontab` files: just pass the `-l` option to `crontab` and look over the output.

For users who wish to begin their own `crontab` file from scratch, without the use of a template, the `crontab -e` option is available. This will invoke the selected editor with an empty file. When the file is saved, it will be automatically installed by the `crontab` command.

If you later want to remove your user `crontab` completely, use `crontab` with the `-r` option.

6.7 Using rc under DragonFly

Contributed by Tom Rhodes.

DragonFly uses the NetBSD `rc.d` system for system initialization. Users should notice the files listed in the `/etc/rc.d` directory. Many of these files are for basic services which can be controlled with the `start`, `stop`, and `restart` options. For instance, `sshd(8)` can be restarted with the following command:

```
# /etc/rc.d/sshd restart
```

This procedure is similar for other services. Of course, services are usually started automatically as specified in `rc.conf(5)`. For example, enabling the Network Address Translation daemon at startup is as simple as adding the following line to `/etc/rc.conf`:

```
natd_enable="YES"
```

If a `natd_enable="NO"` line is already present, then simply change the `NO` to `YES`. The `rc` scripts will automatically load any other dependent services during the next reboot, as described below.

Since the `rc.d` system is primarily intended to start/stop services at system startup/shutdown time, the standard `start`, `stop` and `restart` options will only perform their action if the appropriate `/etc/rc.conf` variables are set. For instance the above `sshd restart` command will only work if `sshd_enable` is set to `YES` in `/etc/rc.conf`. To `start`, `stop` or `restart` a service regardless of the settings in `/etc/rc.conf`, the commands should be prefixed with “force”. For instance to restart `sshd` regardless of the current `/etc/rc.conf` setting, execute the following command:

```
# /etc/rc.d/sshd forcerestart
```

It is easy to check if a service is enabled in `/etc/rc.conf` by running the appropriate `rc.d` script with the option `rcvar`. Thus, an administrator can check that `sshd` is in fact enabled in `/etc/rc.conf` by running:

```
# /etc/rc.d/sshd rcvar
# sshd
$sshd_enable=YES
```

Note: The second line (`# sshd`) is the output from the `rc.d` script, not a `root` prompt.

To determine if a service is running, a `status` option is available. For instance to verify that `sshd` is actually started:

```
# /etc/rc.d/sshd status
sshd is running as pid 433.
```

It is also possible to `reload` a service. This will attempt to send a signal to an individual service, forcing the service to reload its configuration files. In most cases this means sending the service a `SIGHUP` signal.

The `rcNG` structure is used both for network services and system initialization. Some services are run only at boot; and the `RCNG` system is what triggers them.

Many system services depend on other services to function properly. For example, `NIS` and other `RPC`-based services may fail to start until after the `rpcbind` (`portmapper`) service has started. To resolve this issue, information about dependencies and other meta-data is included in the comments at the top of each startup script. The `rcorder(8)` program is then used to parse these comments during system initialization to determine the order in which system

services should be invoked to satisfy the dependencies. The following words may be included at the top of each startup file:

- **PROVIDE:** Specifies the services this file provides.
- **REQUIRE:** Lists services which are required for this service. This file will run *after* the specified services.
- **BEFORE:** Lists services which depend on this service. This file will run *before* the specified services.
- **KEYWORD:** When `rcorder(8)` uses the `-k` option, then only the `rc.d` files matching this keyword are used. ¹ For example, when using `-k shutdown`, only the `rc.d` scripts defining the `shutdown` keyword are used.

With the `-s` option, `rcorder(8)` will skip any `rc.d` script defining the corresponding keyword to skip. For example, scripts defining the `nostart` keyword are skipped at boot time.

By using this method, an administrator can easily control system services without the hassle of “runlevels” like some other UNIX operating systems.

Additional information about the DragonFly `rc.d` system can be found in the `rc(8)`, `rc.conf(5)`, and `rc.subr(8)` manual pages.

6.8 Setting Up Network Interface Cards

Contributed by Marc Fonvieille.

Nowadays we can not think about a computer without thinking about a network connection. Adding and configuring a network card is a common task for any DragonFly administrator.

6.8.1 Locating the Correct Driver

Before you begin, you should know the model of the card you have, the chip it uses, and whether it is a PCI or ISA card. DragonFly supports a wide variety of both PCI and ISA cards. Check the Hardware Compatibility List for your release to see if your card is supported.

Once you are sure your card is supported, you need to determine the proper driver for the card. The file `/usr/src/sys/i386/conf/LINT` will give you the list of network interfaces drivers with some information about the supported chipsets/cards. If you have doubts about which driver is the correct one, read the manual page of the driver. The manual page will give you more information about the supported hardware and even the possible problems that could occur.

If you own a common card, most of the time you will not have to look very hard for a driver. Drivers for common network cards are present in the `GENERIC` kernel, so your card should show up during boot, like so:

```
dc0: <82c169 PNIC 10/100BaseTX> port 0xa000-0xa0ff mem 0xd3800000-0xd38000ff irq 15 at device 11.0 on pci0
dc0: Ethernet address: 00:a0:cc:da:da:da
miibus0: <MII bus> on dc0
ukphy0: <Generic IEEE 802.3u media interface> on miibus0
ukphy0: 10baseT, 10baseT-FDX, 100baseTX, 100baseTX-FDX, auto
dc1: <82c169 PNIC 10/100BaseTX> port 0x9800-0x98ff mem 0xd3000000-0xd30000ff irq 11 at device 12.0 on pci0
dc1: Ethernet address: 00:a0:cc:da:da:db
miibus1: <MII bus> on dc1
```

```
ukphy1: <Generic IEEE 802.3u media interface> on miibus1
ukphy1: 10baseT, 10baseT-FDX, 100baseTX, 100baseTX-FDX, auto
```

In this example, we see that two cards using the dc(4) driver are present on the system.

To use your network card, you will need to load the proper driver. This may be accomplished in one of two ways. The easiest way is to simply load a kernel module for your network card with `kldload(8)`. A module is not available for all network card drivers (ISA cards and cards using the ed(4) driver, for example). Alternatively, you may statically compile the support for your card into your kernel. Check `/usr/src/sys/i386/conf/LINT` and the manual page of the driver to know what to add in your kernel configuration file. For more information about recompiling your kernel, please see Chapter 9. If your card was detected at boot by your kernel (`GENERIC`) you do not have to build a new kernel.

6.8.2 Configuring the Network Card

Once the right driver is loaded for the network card, the card needs to be configured. As with many other things, the network card may have been configured at installation time.

To display the configuration for the network interfaces on your system, enter the following command:

```
% ifconfig
dc0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    inet 192.168.1.3 netmask 0xffffffff broadcast 192.168.1.255
    ether 00:a0:cc:da:da:da
    media: Ethernet autoselect (100baseTX <full-duplex>)
    status: active
dc1: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    inet 10.0.0.1 netmask 0xffffffff broadcast 10.0.0.255
    ether 00:a0:cc:da:da:db
    media: Ethernet 10baseT/UTP
    status: no carrier
lp0: flags=8810<POINTOPOINT,SIMPLEX,MULTICAST> mtu 1500
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet 127.0.0.1 netmask 0xff000000
tun0: flags=8010<POINTOPOINT,MULTICAST> mtu 1500
```

Note: Note that entries concerning IPv6 (`inet6` etc.) were omitted in this example.

In this example, the following devices were displayed:

- `dc0`: The first Ethernet interface
- `dc1`: The second Ethernet interface
- `lp0`: The parallel port interface
- `lo0`: The loopback device
- `tun0`: The tunnel device used by **ppp**

DragonFly uses the driver name followed by the order in which one the card is detected at the kernel boot to name the network card, starting the count at zero. For example, `sis2` would be the third network card on the system using the `sis(4)` driver.

In this example, the `dc0` device is up and running. The key indicators are:

1. UP means that the card is configured and ready.
2. The card has an Internet (`inet`) address (in this case `192.168.1.3`).
3. It has a valid subnet mask (`netmask; 0xffffffff00` is the same as `255.255.255.0`).
4. It has a valid broadcast address (in this case, `192.168.1.255`).
5. The MAC address of the card (`ether`) is `00:a0:cc:da:da:da`
6. The physical media selection is on autoselection mode (`media: Ethernet autoselect (10baseTX <full-duplex>)`). We see that `dc1` was configured to run with `10baseT/UTP` media. For more information on available media types for a driver, please refer to its manual page.
7. The status of the link (`status`) is `active`, i.e. the carrier is detected. For `dc1`, we see `status: no carrier`. This is normal when an Ethernet cable is not plugged into the card.

If the `ifconfig(8)` output had shown something similar to:

```
dc0: flags=8843<BROADCAST,SIMPLEX,MULTICAST> mtu 1500
     ether 00:a0:cc:da:da:da
```

it would indicate the card has not been configured.

To configure your card, you need `root` privileges. The network card configuration can be done from the command line with `ifconfig(8)` as `root`.

```
# ifconfig dc0 inet 192.168.1.3 netmask 255.255.255.0
```

Manually configuring the care has the disadvantage that you would have to do it after each reboot of the system. The file `/etc/rc.conf` is where to add the network card's configuration.

Open `/etc/rc.conf` in your favorite editor. You need to add a line for each network card present on the system, for example in our case, we added these lines:

```
ifconfig_dc0="inet 192.168.1.3 netmask 255.255.255.0"
ifconfig_dc1="inet 10.0.0.1 netmask 255.255.255.0 media 10baseT/UTP"
```

You have to replace `dc0`, `dc1`, and so on, with the correct device for your cards, and the addresses with the proper ones. You should read the card driver and `ifconfig(8)` manual pages for more details about the allowed options and also `rc.conf(5)` manual page for more information on the syntax of `/etc/rc.conf`.

If you configured the network during installation, some lines about the network card(s) may be already present. Double check `/etc/rc.conf` before adding any lines.

You will also have to edit the file `/etc/hosts` to add the names and the IP addresses of various machines of the LAN, if they are not already there. For more information please refer to `hosts(5)` and to `/usr/share/examples/etc/hosts`.

6.8.3 Testing and Troubleshooting

Once you have made the necessary changes in `/etc/rc.conf`, you should reboot your system. This will allow the change(s) to the interface(s) to be applied, and verify that the system restarts without any configuration errors.

Once the system has been rebooted, you should test the network interfaces.

6.8.3.1 Testing the Ethernet Card

To verify that an Ethernet card is configured correctly, you have to try two things. First, ping the interface itself, and then ping another machine on the LAN.

First test the local interface:

```
% ping -c5 192.168.1.3
PING 192.168.1.3 (192.168.1.3): 56 data bytes
64 bytes from 192.168.1.3: icmp_seq=0 ttl=64 time=0.082 ms
64 bytes from 192.168.1.3: icmp_seq=1 ttl=64 time=0.074 ms
64 bytes from 192.168.1.3: icmp_seq=2 ttl=64 time=0.076 ms
64 bytes from 192.168.1.3: icmp_seq=3 ttl=64 time=0.108 ms
64 bytes from 192.168.1.3: icmp_seq=4 ttl=64 time=0.076 ms

--- 192.168.1.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.074/0.083/0.108/0.013 ms
```

Now we have to ping another machine on the LAN:

```
% ping -c5 192.168.1.2
PING 192.168.1.2 (192.168.1.2): 56 data bytes
64 bytes from 192.168.1.2: icmp_seq=0 ttl=64 time=0.726 ms
64 bytes from 192.168.1.2: icmp_seq=1 ttl=64 time=0.766 ms
64 bytes from 192.168.1.2: icmp_seq=2 ttl=64 time=0.700 ms
64 bytes from 192.168.1.2: icmp_seq=3 ttl=64 time=0.747 ms
64 bytes from 192.168.1.2: icmp_seq=4 ttl=64 time=0.704 ms

--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.700/0.729/0.766/0.025 ms
```

You could also use the machine name instead of `192.168.1.2` if you have set up the `/etc/hosts` file.

6.8.3.2 Troubleshooting

Troubleshooting hardware and software configurations is always a pain, and a pain which can be alleviated by checking the simple things first. Is your network cable plugged in? Have you properly configured the network services? Did you configure the firewall correctly? Is the card you are using supported by DragonFly? Always check the hardware notes before sending off a bug report. Update your version of DragonFly to the latest PREVIEW version. Check the mailing list archives, or perhaps search the Internet.

If the card works, yet performance is poor, it would be worthwhile to read over the tuning(7) manual page. You can also check the network configuration as incorrect network settings can cause slow connections.

Some users experience one or two “device timeouts”, which is normal for some cards. If they continue, or are bothersome, you may wish to be sure the device is not conflicting with another device. Double check the cable connections. Perhaps you may just need to get another card.

At times, users see a few `watchdog timeout` errors. The first thing to do here is to check your network cable. Many cards require a PCI slot which supports Bus Mastering. On some old motherboards, only one PCI slot allows it (usually slot 0). Check the network card and the motherboard documentation to determine if that may be the problem.

No route to host messages occur if the system is unable to route a packet to the destination host. This can happen if no default route is specified, or if a cable is unplugged. Check the output of `netstat -rn` and make sure there is a valid route to the host you are trying to reach. If there is not, read on to Chapter 19.

`ping: sendto: Permission denied` error messages are often caused by a misconfigured firewall. If `ipfw` is enabled in the kernel but no rules have been defined, then the default policy is to deny all traffic, even ping requests! Read on to Section 10.7 for more information.

Sometimes performance of the card is poor, or below average. In these cases it is best to set the media selection mode from `autoselect` to the correct media selection. While this usually works for most hardware, it may not resolve this issue for everyone. Again, check all the network settings, and read over the tuning(7) manual page.

6.9 Virtual Hosts

A very common use of DragonFly is virtual site hosting, where one server appears to the network as many servers. This is achieved by assigning multiple network addresses to a single interface.

A given network interface has one “real” address, and may have any number of “alias” addresses. These aliases are normally added by placing alias entries in `/etc/rc.conf`.

An alias entry for the interface `fxp0` looks like:

```
ifconfig_fxp0_alias0="inet xxx.xxx.xxx.xxx netmask xxx.xxx.xxx.xxx"
```

Note that alias entries must start with `alias0` and proceed upwards in order, (for example, `_alias1`, `_alias2`, and so on). The configuration process will stop at the first missing number.

The calculation of alias netmasks is important, but fortunately quite simple. For a given interface, there must be one address which correctly represents the network’s netmask. Any other addresses which fall within this network must have a netmask of all 1s (expressed as either `255.255.255.255` or `0xffffffff`).

For example, consider the case where the `fxp0` interface is connected to two networks, the `10.1.1.0` network with a netmask of `255.255.255.0` and the `202.0.75.16` network with a netmask of `255.255.255.240`. We want the system to appear at `10.1.1.1` through `10.1.1.5` and at `202.0.75.17` through `202.0.75.20`. As noted above, only the first address in a given network range (in this case, `10.1.1.1` and `202.0.75.17`) should have a real netmask; all the rest (`10.1.1.2` through `10.1.1.5` and `202.0.75.18` through `202.0.75.20`) must be configured with a netmask of `255.255.255.255`.

The following entries configure the adapter correctly for this arrangement:

```
ifconfig_fxp0="inet 10.1.1.1 netmask 255.255.255.0"
ifconfig_fxp0_alias0="inet 10.1.1.2 netmask 255.255.255.255"
ifconfig_fxp0_alias1="inet 10.1.1.3 netmask 255.255.255.255"
ifconfig_fxp0_alias2="inet 10.1.1.4 netmask 255.255.255.255"
```

```
ifconfig_fxp0_alias3="inet 10.1.1.5 netmask 255.255.255.255"
ifconfig_fxp0_alias4="inet 202.0.75.17 netmask 255.255.255.240"
ifconfig_fxp0_alias5="inet 202.0.75.18 netmask 255.255.255.255"
ifconfig_fxp0_alias6="inet 202.0.75.19 netmask 255.255.255.255"
ifconfig_fxp0_alias7="inet 202.0.75.20 netmask 255.255.255.255"
```

6.10 Configuration Files

6.10.1 /etc Layout

There are a number of directories in which configuration information is kept. These include:

<code>/etc</code>	Generic system configuration information; data here is system-specific.
<code>/etc/defaults</code>	Default versions of system configuration files.
<code>/etc/mail</code>	Extra sendmail(8) configuration, other MTA configuration files.
<code>/etc/ppp</code>	Configuration for both user- and kernel-ppp programs.
<code>/etc/namedb</code>	Default location for named(8) data. Normally named.conf and zone files are stored here.
<code>/usr/local/etc</code>	Configuration files for installed applications. May contain per-application subdirectories.
<code>/usr/local/etc/rc.d</code>	Start/stop scripts for installed applications.
<code>/var/db</code>	Automatically generated system-specific database files, such as the package database, the locate database, and so on

6.10.2 Hostnames

6.10.2.1 /etc/resolv.conf

`/etc/resolv.conf` dictates how DragonFly's resolver accesses the Internet Domain Name System (DNS).

The most common entries to `resolv.conf` are:

<code>nameserver</code>	The IP address of a name server the resolver should query. The servers are queried in the order listed with a maximum of three.
<code>search</code>	Search list for hostname lookup. This is normally determined by the domain of the local hostname.
<code>domain</code>	The local domain name.

A typical `resolv.conf`:

```
search example.com
nameserver 147.11.1.11
nameserver 147.11.100.30
```

Note: Only one of the `search` and `domain` options should be used.

If you are using DHCP, `dhclient(8)` usually rewrites `resolv.conf` with information received from the DHCP server.

6.10.2.2 `/etc/hosts`

`/etc/hosts` is a simple text database reminiscent of the old Internet. It works in conjunction with DNS and NIS providing name to IP address mappings. Local computers connected via a LAN can be placed in here for simplistic naming purposes instead of setting up a `named(8)` server. Additionally, `/etc/hosts` can be used to provide a local record of Internet names, reducing the need to query externally for commonly accessed names.

```
#
#
# Host Database
# This file should contain the addresses and aliases
# for local hosts that share this file.
# In the presence of the domain name service or NIS, this file may
# not be consulted at all; see /etc/nsswitch.conf for the resolution order.
#
#
::1                localhost localhost.my.domain myname.my.domain
127.0.0.1          localhost localhost.my.domain myname.my.domain

#
# Imaginary network.
#10.0.0.2           myname.my.domain myname
#10.0.0.3           myfriend.my.domain myfriend
#
# According to RFC 1918, you can use the following IP networks for
# private nets which will never be connected to the Internet:
#
#       10.0.0.0           -   10.255.255.255
#       172.16.0.0        -   172.31.255.255
#       192.168.0.0       -   192.168.255.255
#
# In case you want to be able to connect to the Internet, you need
# real official assigned numbers. PLEASE PLEASE PLEASE do not try
# to invent your own network numbers but instead get one from your
# network provider (if any) or from the Internet Registry (ftp to
# rs.internic.net, directory '/templates').
#
```

`/etc/hosts` takes on the simple format of:

```
[Internet address] [official hostname] [alias1] [alias2] ...
```

For example:

```
10.0.0.1 myRealHostname.example.com myRealHostname foobar1 foobar2
```

Consult `hosts(5)` for more information.

6.10.3 Log File Configuration

6.10.3.1 syslog.conf

`syslog.conf` is the configuration file for the `syslogd(8)` program. It indicates which types of `syslog` messages are logged to particular log files.

```
#
#
#     Spaces ARE valid field separators in this file. However,
#     other *nix-like systems still insist on using tabs as field
#     separators. If you are sharing this file between systems, you
#     may want to use only tabs as field separators here.
#     Consult the syslog.conf(5) manual page.
*.err;kern.debug;auth.notice;mail.crit      /dev/console
*.notice;kern.debug;lpr.info;mail.crit;news.err /var/log/messages
security.*                                  /var/log/security
mail.info                                   /var/log/maillog
lpr.info                                    /var/log/lpd-errs
cron.*                                      /var/log/cron
*.err                                       root
*.notice;news.err                          root
*.alert                                     root
*.emerg                                     *
# uncomment this to log all writes to /dev/console to /var/log/console.log
#console.info                               /var/log/console.log
# uncomment this to enable logging of all log messages to /var/log/all.log
#*. *                                       /var/log/all.log
# uncomment this to enable logging to a remote log host named loghost
#*. *                                       @loghost
# uncomment these if you're running inn
# news.crit                                  /var/log/news/news.crit
# news.err                                  /var/log/news/news.err
# news.notice                               /var/log/news/news.notice
!startslip
*. *                                       /var/log/slip.log
!ppp
*. *                                       /var/log/ppp.log
```

Consult the `syslog.conf(5)` manual page for more information.

6.10.3.2 newsyslog.conf

`newsyslog.conf` is the configuration file for `newsyslog(8)`, a program that is normally scheduled to run by `cron(8)`. `newsyslog(8)` determines when log files require archiving or rearranging. `logfile` is moved to `logfile.0`, `logfile.0` is moved to `logfile.1`, and so on. Alternatively, the log files may be archived in `gzip(1)` format causing them to be named: `logfile.0.gz`, `logfile.1.gz`, and so on.

`newsyslog.conf` indicates which log files are to be managed, how many are to be kept, and when they are to be touched. Log files can be rearranged and/or archived when they have either reached a certain size, or at a certain periodic time/date.

```
# configuration file for newsyslog
#
#
# filename          [owner:group]    mode count size when [ZB] [/pid_file] [sig_num]
/var/log/cron              600 3    100 *    Z
/var/log/amd.log           644 7    100 *    Z
/var/log/kerberos.log      644 7    100 *    Z
/var/log/lpd-errs          644 7    100 *    Z
/var/log/maillog           644 7    *    @T00 Z
/var/log/sendmail.st       644 10   *    168 B
/var/log/messages          644 5    100 *    Z
/var/log/all.log           600 7    *    @T00 Z
/var/log/slip.log          600 3    100 *    Z
/var/log/ppp.log           600 3    100 *    Z
/var/log/security          600 10   100 *    Z
/var/log/wtmp              644 3    *    @01T05 B
/var/log/daily.log         640 7    *    @T00 Z
/var/log/weekly.log        640 5    1    $W6D0 Z
/var/log/monthly.log       640 12   *    $M1D0 Z
/var/log/console.log       640 5    100 *    Z
```

Consult the newsyslog(8) manual page for more information.

6.10.4 sysctl.conf

sysctl.conf looks much like rc.conf. Values are set in a variable=value form. The specified values are set after the system goes into multi-user mode. Not all variables are settable in this mode.

A sample sysctl.conf turning off logging of fatal signal exits and letting Linux programs know they are really running under DragonFly:

```
kern.logsigexit=0          # Do not log fatal signal exits (e.g. sig 11)
compat.linux.osname=DragonFly
compat.linux.osrelease=4.3-STABLE
```

6.11 Tuning with sysctl

sysctl(8) is an interface that allows you to make changes to a running DragonFly system. This includes many advanced options of the TCP/IP stack and virtual memory system that can dramatically improve performance for an experienced system administrator. Over five hundred system variables can be read and set using sysctl(8).

At its core, sysctl(8) serves two functions: to read and to modify system settings.

To view all readable variables:

```
% sysctl -a
```

To read a particular variable, for example, kern.maxproc:

```
% sysctl kern.maxproc
```

```
kern.maxproc: 1044
```

To set a particular variable, use the intuitive *variable=value* syntax:

```
# sysctl kern.maxfiles=5000
kern.maxfiles: 2088 -> 5000
```

Settings of sysctl variables are usually either strings, numbers, or booleans (a boolean being 1 for yes or a 0 for no).

If you want to set automatically some variables each time the machine boots, add them to the `/etc/sysctl.conf` file. For more information see the `sysctl.conf(5)` manual page and the Section 6.10.4.

6.11.1 sysctl(8) Read-only

Contributed by Tom Rhodes.

In some cases it may be desirable to modify read-only sysctl(8) values. While this is not recommended, it is also sometimes unavoidable.

For instance on some laptop models the `cardbus(4)` device will not probe memory ranges, and fail with errors which look similar to:

```
cbb0: Could not map register memory
device_probe_and_attach: cbb0 attach returned 12
```

Cases like the one above usually require the modification of some default sysctl(8) settings which are set read only. To overcome these situations a user can put sysctl(8) “OIDs” in their local `/boot/loader.conf`. Default settings are located in the `/boot/defaults/loader.conf` file.

Fixing the problem mentioned above would require a user to set `hw.pci.allow_unsupported_io_range=1` in the aforementioned file. Now `cardbus(4)` will work properly.

6.12 Tuning Disks

6.12.1 Sysctl Variables

6.12.1.1 `vfs.vmiodirenable`

The `vfs.vmiodirenable` sysctl variable may be set to either 0 (off) or 1 (on); it is 1 by default. This variable controls how directories are cached by the system. Most directories are small, using just a single fragment (typically 1 K) in the file system and less (typically 512 bytes) in the buffer cache. With this variable turned off (to 0), the buffer cache will only cache a fixed number of directories even if you have a huge amount of memory. When turned on (to 1), this sysctl allows the buffer cache to use the VM Page Cache to cache the directories, making all the memory available for caching directories. However, the minimum in-core memory used to cache a directory is the physical page size (typically 4 K) rather than 512 bytes. We recommend keeping this option on if you are running any services which manipulate large numbers of files. Such services can include web caches, large mail systems, and news systems. Keeping this option on will generally not reduce performance even with the wasted memory but you should experiment to find out.

6.12.1.2 `vfs.write_behind`

The `vfs.write_behind` sysctl variable defaults to 1 (on). This tells the file system to issue media writes as full clusters are collected, which typically occurs when writing large sequential files. The idea is to avoid saturating the buffer cache with dirty buffers when it would not benefit I/O performance. However, this may stall processes and under certain circumstances you may wish to turn it off.

6.12.1.3 `vfs.hirunningspace`

The `vfs.hirunningspace` sysctl variable determines how much outstanding write I/O may be queued to disk controllers system-wide at any given instance. The default is usually sufficient but on machines with lots of disks you may want to bump it up to four or five *megabytes*. Note that setting too high a value (exceeding the buffer cache's write threshold) can lead to extremely bad clustering performance. Do not set this value arbitrarily high! Higher write values may add latency to reads occurring at the same time.

There are various other buffer-cache and VM page cache related sysctls. We do not recommend modifying these values. The VM system does an extremely good job of automatically tuning itself.

6.12.1.4 `vm.swap_idle_enabled`

The `vm.swap_idle_enabled` sysctl variable is useful in large multi-user systems where you have lots of users entering and leaving the system and lots of idle processes. Such systems tend to generate a great deal of continuous pressure on free memory reserves. Turning this feature on and tweaking the swapout hysteresis (in idle seconds) via `vm.swap_idle_threshold1` and `vm.swap_idle_threshold2` allows you to depress the priority of memory pages associated with idle processes more quickly than the normal pageout algorithm. This gives a helping hand to the pageout daemon. Do not turn this option on unless you need it, because the tradeoff you are making is essentially pre-page memory sooner rather than later; thus eating more swap and disk bandwidth. In a small system this option will have a determinable effect but in a large system that is already doing moderate paging this option allows the VM system to stage whole processes into and out of memory easily.

6.12.1.5 `hw.ata.wc`

IDE drives lie about when a write completes. With IDE write caching turned on, IDE hard drives not only write data to disk out of order, but will sometimes delay writing some blocks indefinitely when under heavy disk loads. A crash or power failure may cause serious file system corruption. Turning off write caching will remove the danger of this data loss, but will also cause disk operations to proceed *very slowly*. Change this only if prepared to suffer with the disk slowdown.

Changing this variable must be done from the boot loader at boot time. Attempting to do it after the kernel boots will have no effect.

For more information, please see `ata(4)` manual page.

6.12.2 Soft Updates

The `tunefs(8)` program can be used to fine-tune a file system. This program has many different options, but for now we are only concerned with toggling Soft Updates on and off, which is done by:

```
# tune2fs -n enable /filesystem
# tune2fs -n disable /filesystem
```

A filesystem cannot be modified with `tune2fs(8)` while it is mounted. A good time to enable Soft Updates is before any partitions have been mounted, in single-user mode.

Note: It is possible to enable Soft Updates at filesystem creation time, through use of the `-U` option to `newfs(8)`.

Soft Updates drastically improves meta-data performance, mainly file creation and deletion, through the use of a memory cache. We recommend to use Soft Updates on all of your file systems. There are two downsides to Soft Updates that you should be aware of: First, Soft Updates guarantees filesystem consistency in the case of a crash but could very easily be several seconds (even a minute!) behind updating the physical disk. If your system crashes you may lose more work than otherwise. Secondly, Soft Updates delays the freeing of filesystem blocks. If you have a filesystem (such as the root filesystem) which is almost full, performing a major update, such as `make installworld`, can cause the filesystem to run out of space and the update to fail.

6.12.2.1 More Details about Soft Updates

There are two traditional approaches to writing a file systems meta-data back to disk. (Meta-data updates are updates to non-content data like inodes or directories.)

Historically, the default behavior was to write out meta-data updates synchronously. If a directory had been changed, the system waited until the change was actually written to disk. The file data buffers (file contents) were passed through the buffer cache and backed up to disk later on asynchronously. The advantage of this implementation is that it operates safely. If there is a failure during an update, the meta-data are always in a consistent state. A file is either created completely or not at all. If the data blocks of a file did not find their way out of the buffer cache onto the disk by the time of the crash, `fsck(8)` is able to recognize this and repair the filesystem by setting the file length to 0. Additionally, the implementation is clear and simple. The disadvantage is that meta-data changes are slow. An `rm -r`, for instance, touches all the files in a directory sequentially, but each directory change (deletion of a file) will be written synchronously to the disk. This includes updates to the directory itself, to the inode table, and possibly to indirect blocks allocated by the file. Similar considerations apply for unrolling large hierarchies (`tar -x`).

The second case is asynchronous meta-data updates. This is the default for Linux/ext2fs and `mount -o async` for *BSD ufs. All meta-data updates are simply being passed through the buffer cache too, that is, they will be intermixed with the updates of the file content data. The advantage of this implementation is there is no need to wait until each meta-data update has been written to disk, so all operations which cause huge amounts of meta-data updates work much faster than in the synchronous case. Also, the implementation is still clear and simple, so there is a low risk for bugs creeping into the code. The disadvantage is that there is no guarantee at all for a consistent state of the filesystem. If there is a failure during an operation that updated large amounts of meta-data (like a power failure, or someone pressing the reset button), the filesystem will be left in an unpredictable state. There is no opportunity to examine the state of the filesystem when the system comes up again; the data blocks of a file could already have been written to the disk while the updates of the inode table or the associated directory were not. It is actually impossible to implement a `fsck` which is able to clean up the resulting chaos (because the necessary information is not available on the disk). If the filesystem has been damaged beyond repair, the only choice is to use `newfs(8)` on it and restore it from backup.

The usual solution for this problem was to implement *dirty region logging*, which is also referred to as *journaling*, although that term is not used consistently and is occasionally applied to other forms of transaction logging as well. Meta-data updates are still written synchronously, but only into a small region of the disk. Later on they will be

moved to their proper location. Because the logging area is a small, contiguous region on the disk, there are no long distances for the disk heads to move, even during heavy operations, so these operations are quicker than synchronous updates. Additionally the complexity of the implementation is fairly limited, so the risk of bugs being present is low. A disadvantage is that all meta-data are written twice (once into the logging region and once to the proper location) so for normal work, a performance “pessimization” might result. On the other hand, in case of a crash, all pending meta-data operations can be quickly either rolled-back or completed from the logging area after the system comes up again, resulting in a fast filesystem startup.

Kirk McKusick, the developer of Berkeley FFS, solved this problem with Soft Updates: all pending meta-data updates are kept in memory and written out to disk in a sorted sequence (“ordered meta-data updates”). This has the effect that, in case of heavy meta-data operations, later updates to an item “catch” the earlier ones if the earlier ones are still in memory and have not already been written to disk. So all operations on, say, a directory are generally performed in memory before the update is written to disk (the data blocks are sorted according to their position so that they will not be on the disk ahead of their meta-data). If the system crashes, this causes an implicit “log rewind”: all operations which did not find their way to the disk appear as if they had never happened. A consistent filesystem state is maintained that appears to be the one of 30 to 60 seconds earlier. The algorithm used guarantees that all resources in use are marked as such in their appropriate bitmaps: blocks and inodes. After a crash, the only resource allocation error that occurs is that resources are marked as “used” which are actually “free”. `fsck(8)` recognizes this situation, and frees the resources that are no longer used. It is safe to ignore the dirty state of the filesystem after a crash by forcibly mounting it with `mount -f`. In order to free resources that may be unused, `fsck(8)` needs to be run at a later time.

The advantage is that meta-data operations are nearly as fast as asynchronous updates (i.e. faster than with *logging*, which has to write the meta-data twice). The disadvantages are the complexity of the code (implying a higher risk for bugs in an area that is highly sensitive regarding loss of user data), and a higher memory consumption. Additionally there are some idiosyncrasies one has to get used to. After a crash, the state of the filesystem appears to be somewhat “older”. In situations where the standard synchronous approach would have caused some zero-length files to remain after the `fsck`, these files do not exist at all with a Soft Updates filesystem because neither the meta-data nor the file contents have ever been written to disk. Disk space is not released until the updates have been written to disk, which may take place some time after running `rm`. This may cause problems when installing large amounts of data on a filesystem that does not have enough free space to hold all the files twice.

6.13 Tuning Kernel Limits

6.13.1 File/Process Limits

6.13.1.1 `kern.maxfiles`

`kern.maxfiles` can be raised or lowered based upon your system requirements. This variable indicates the maximum number of file descriptors on your system. When the file descriptor table is full, `file: table is full` will show up repeatedly in the system message buffer, which can be viewed with the `dmesg` command.

Each open file, socket, or fifo uses one file descriptor. A large-scale production server may easily require many thousands of file descriptors, depending on the kind and number of services running concurrently.

`kern.maxfile`'s default value is dictated by the `MAXUSERS` option in your kernel configuration file.

`kern.maxfiles` grows proportionally to the value of `MAXUSERS`. When compiling a custom kernel, it is a good idea

to set this kernel configuration option according to the uses of your system. From this number, the kernel is given most of its pre-defined limits. Even though a production machine may not actually have 256 users connected at once, the resources needed may be similar to a high-scale web server.

Note: Setting `MAXUSERS` to 0 in your kernel configuration file will choose a reasonable default value based on the amount of RAM present in your system. It is set to 0 in the default `GENERIC` kernel.

6.13.1.2 `kern.ipc.somaxconn`

The `kern.ipc.somaxconn` sysctl variable limits the size of the listen queue for accepting new TCP connections. The default value of 128 is typically too low for robust handling of new connections in a heavily loaded web server environment. For such environments, it is recommended to increase this value to 1024 or higher. The service daemon may itself limit the listen queue size (e.g. `sendmail(8)`, or **Apache**) but will often have a directive in its configuration file to adjust the queue size. Large listen queues also do a better job of avoiding Denial of Service (DoS) attacks.

6.13.2 Network Limits

The `NMBCLUSTERS` kernel configuration option dictates the amount of network Mbufs available to the system. A heavily-trafficked server with a low number of Mbufs will hinder `DragonFly`'s ability. Each cluster represents approximately 2 K of memory, so a value of 1024 represents 2 megabytes of kernel memory reserved for network buffers. A simple calculation can be done to figure out how many are needed. If you have a web server which maxes out at 1000 simultaneous connections, and each connection eats a 16 K receive and 16 K send buffer, you need approximately 32 MB worth of network buffers to cover the web server. A good rule of thumb is to multiply by 2, so $2 \times 32 \text{ MB} / 2 \text{ KB} = 64 \text{ MB} / 2 \text{ kB} = 32768$. We recommend values between 4096 and 32768 for machines with greater amounts of memory. Under no circumstances should you specify an arbitrarily high value for this parameter as it could lead to a boot time crash. The `-m` option to `netstat(1)` may be used to observe network cluster use. `kern.ipc.nmbclusters` loader tunable should be used to tune this at boot time.

For busy servers that make extensive use of the `sendfile(2)` system call, it may be necessary to increase the number of `sendfile(2)` buffers via the `NSFBUFS` kernel configuration option or by setting its value in `/boot/loader.conf` (see `loader(8)` for details). A common indicator that this parameter needs to be adjusted is when processes are seen in the `sfbufa` state. The sysctl variable `kern.ipc.nsfbufs` is a read-only glimpse at the kernel configured variable. This parameter nominally scales with `kern.maxusers`, however it may be necessary to tune accordingly.

Important: Even though a socket has been marked as non-blocking, calling `sendfile(2)` on the non-blocking socket may result in the `sendfile(2)` call blocking until enough `struct sf_buf`'s are made available.

6.13.2.1 `net.inet.ip.portrange.*`

The `net.inet.ip.portrange.*` sysctl variables control the port number ranges automatically bound to TCP and UDP sockets. There are three ranges: a low range, a default range, and a high range. Most network programs use the default range which is controlled by the `net.inet.ip.portrange.first` and `net.inet.ip.portrange.last`, which default to 1024 and 5000, respectively. Bound port ranges are used for outgoing connections, and it is possible to run the system out of ports under certain circumstances. This most

commonly occurs when you are running a heavily loaded web proxy. The port range is not an issue when running servers which handle mainly incoming connections, such as a normal web server, or has a limited number of outgoing connections, such as a mail relay. For situations where you may run yourself out of ports, it is recommended to increase `net.inet.ip.portrange.last` modestly. A value of 10000, 20000 or 30000 may be reasonable. You should also consider firewall effects when changing the port range. Some firewalls may block large ranges of ports (usually low-numbered ports) and expect systems to use higher ranges of ports for outgoing connections — for this reason it is recommended that `net.inet.ip.portrange.first` be lowered.

6.13.2.2 TCP Bandwidth Delay Product

The TCP Bandwidth Delay Product Limiting is similar to TCP/Vegas in NetBSD. It can be enabled by setting `net.inet.tcp.inflight_enable` sysctl variable to 1. The system will attempt to calculate the bandwidth delay product for each connection and limit the amount of data queued to the network to just the amount required to maintain optimum throughput.

This feature is useful if you are serving data over modems, Gigabit Ethernet, or even high speed WAN links (or any other link with a high bandwidth delay product), especially if you are also using window scaling or have configured a large send window. If you enable this option, you should also be sure to set `net.inet.tcp.inflight_debug` to 0 (disable debugging), and for production use setting `net.inet.tcp.inflight_min` to at least 6144 may be beneficial. However, note that setting high minimums may effectively disable bandwidth limiting depending on the link. The limiting feature reduces the amount of data built up in intermediate route and switch packet queues as well as reduces the amount of data built up in the local host's interface queue. With fewer packets queued up, interactive connections, especially over slow modems, will also be able to operate with lower *Round Trip Times*. However, note that this feature only effects data transmission (uploading / server side). It has no effect on data reception (downloading).

Adjusting `net.inet.tcp.inflight_stab` is *not* recommended. This parameter defaults to 20, representing 2 maximal packets added to the bandwidth delay product window calculation. The additional window is required to stabilize the algorithm and improve responsiveness to changing conditions, but it can also result in higher ping times over slow links (though still much lower than you would get without the inflight algorithm). In such cases, you may wish to try reducing this parameter to 15, 10, or 5; and may also have to reduce `net.inet.tcp.inflight_min` (for example, to 3500) to get the desired effect. Reducing these parameters should be done as a last resort only.

6.14 Adding Swap Space

No matter how well you plan, sometimes a system does not run as you expect. If you find you need more swap space, it is simple enough to add. You have three ways to increase swap space: adding a new hard drive, enabling swap over NFS, and creating a swap file on an existing partition.

6.14.1 Swap on a New Hard Drive

The best way to add swap, of course, is to use this as an excuse to add another hard drive. You can always use another hard drive, after all. If you can do this, go reread the discussion about swap space in Section 6.2 for some suggestions on how to best arrange your swap.

6.14.2 Swapping over NFS

Swapping over NFS is only recommended if you do not have a local hard disk to swap to. Even though DragonFly has an excellent NFS implementation, NFS swapping will be limited by the available network bandwidth and puts an additional burden on the NFS server.

6.14.3 Swapfiles

You can create a file of a specified size to use as a swap file. In our example here we will use a 64MB file called `/usr/swap0`. You can use any name you want, of course.

Example 6-1. Creating a Swapfile

1. Be certain that your kernel configuration includes the vnode driver. It is *not* in recent versions of GENERIC.

```
pseudo-device  vn 1    #Vnode driver (turns a file into a device)
```
2. Create a vn-device:

```
# cd /dev
# sh MAKEDEV vn0
```
3. Create a swapfile (`/usr/swap0`):

```
# dd if=/dev/zero of=/usr/swap0 bs=1024k count=64
```
4. Set proper permissions on (`/usr/swap0`):

```
# chmod 0600 /usr/swap0
```
5. Enable the swap file in `/etc/rc.conf`:

```
swapfile="/usr/swap0"    # Set to name of swapfile if aux swapfile desired.
```
6. Reboot the machine or to enable the swap file immediately, type:

```
# vnconfig -e /dev/vn0b /usr/swap0 swap
```

6.15 Power and Resource Management

Written by Hiten Pandya and Tom Rhodes.

It is very important to utilize hardware resources in an efficient manner. Before ACPI was introduced, it was very difficult and inflexible for operating systems to manage the power usage and thermal properties of a system. The hardware was controlled by some sort of BIOS embedded interface, such as *Plug and Play BIOS (PNPBIOS)*, or *Advanced Power Management (APM)* and so on. Power and Resource Management is one of the key components of a modern operating system. For example, you may want an operating system to monitor system limits (and possibly alert you) in case your system temperature increased unexpectedly.

In this section, we will provide comprehensive information about ACPI. References will be provided for further reading at the end. Please be aware that ACPI is available on DragonFly systems as a default kernel module.

6.15.1 What Is ACPI?

Advanced Configuration and Power Interface (ACPI) is a standard written by an alliance of vendors to provide a standard interface for hardware resources and power management (hence the name). It is a key element in *Operating System-directed configuration and Power Management*, i.e.: it provides more control and flexibility to the operating system (OS). Modern systems “stretched” the limits of the current Plug and Play interfaces (such as APM), prior to the introduction of ACPI. ACPI is the direct successor to APM (Advanced Power Management).

6.15.2 Shortcomings of Advanced Power Management (APM)

The *Advanced Power Management (APM)* facility control’s the power usage of a system based on its activity. The APM BIOS is supplied by the (system) vendor and it is specific to the hardware platform. An APM driver in the OS mediates access to the *APM Software Interface*, which allows management of power levels.

There are four major problems in APM. Firstly, power management is done by the (vendor-specific) BIOS, and the OS does not have any knowledge of it. One example of this, is when the user sets idle-time values for a hard drive in the APM BIOS, that when exceeded, it (BIOS) would spin down the hard drive, without the consent of the OS. Secondly, the APM logic is embedded in the BIOS, and it operates outside the scope of the OS. This means users can only fix problems in their APM BIOS by flashing a new one into the ROM; which, is a very dangerous procedure, and if it fails, it could leave the system in an unrecoverable state. Thirdly, APM is a vendor-specific technology, which, means that there is a lot or parity (duplication of efforts) and bugs found in one vendor’s BIOS, may not be solved in others. Last but not the least, the APM BIOS did not have enough room to implement a sophisticated power policy, or one that can adapt very well to the purpose of the machine.

Plug and Play BIOS (PNPBIOS) was unreliable in many situations. PNPBIOS is 16-bit technology, so the OS has to use 16-bit emulation in order to “interface” with PNPBIOS methods.

The DragonFly APM driver is documented in the `apm(4)` manual page.

6.15.3 Configuring ACPI

The `acpi.ko` driver is loaded by default at start up by the `loader(8)` and should *not* be compiled into the kernel. The reasoning behind this is that modules are easier to work with, say if switching to another `acpi.ko` without doing a kernel rebuild. This has the advantage of making testing easier. Another reason is that starting ACPI after a system has been brought up is not too useful, and in some cases can be fatal. In doubt, just disable ACPI all together. This driver should not and can not be unloaded because the system bus uses it for various hardware interactions. ACPI can be disabled with the `acpicnf(8)` utility. In fact most of the interaction with ACPI can be done via `acpicnf(8)`. Basically this means, if anything about ACPI is in the `dmesg(8)` output, then most likely it is already running.

Note: ACPI and APM cannot coexist and should be used separately. The last one to load will terminate if the driver notices the other running.

In the simplest form, ACPI can be used to put the system into a sleep mode with `acpicnf(8)`, the `-s` flag, and a 1-5 option. Most users will only need 1. Option 5 will do a soft-off which is the same action as:

```
# halt -p
```

The other options are available. Check out the `acpicnf(8)` manual page for more information.

6.16 Using and Debugging DragonFly ACPI

Written by Nate Lawson. With contributions from Peter Schultz and Tom Rhodes.

ACPI is a fundamentally new way of discovering devices, managing power usage, and providing standardized access to various hardware previously managed by the BIOS. Progress is being made toward ACPI working on all systems, but bugs in some motherboards' *ACPI Machine Language* (AML) bytecode, incompleteness in DragonFly's kernel subsystems, and bugs in the Intel ACPI-CA interpreter continue to appear.

This document is intended to help you assist the DragonFly ACPI maintainers in identifying the root cause of problems you observe and debugging and developing a solution. Thanks for reading this and we hope we can solve your system's problems.

6.16.1 Submitting Debugging Information

Note: Before submitting a problem, be sure you are running the latest BIOS version and, if available, embedded controller firmware version.

For those of you that want to submit a problem right away, please send the following information to bugs (<http://leaf.dragonflybsd.org/mailarchive/>)

- Description of the buggy behavior, including system type and model and anything that causes the bug to appear. Also, please note as accurately as possible when the bug began occurring if it is new for you.
- The dmesg output after “boot -v”, including any error messages generated by you exercising the bug.
- dmesg output from “boot -v” with ACPI disabled, if disabling it helps fix the problem.
- Output from “sysctl hw.acpi”. This is also a good way of figuring out what features your system offers.
- URL where your *ACPI Source Language* (ASL) can be found. Do *not* send the ASL directly to the list as it can be very large. Generate a copy of your ASL by running this command:

```
# acpidump -t -d > name-system.asl
```

(Substitute your login name for *name* and manufacturer/model for *system*. Example: `njl-FooCo6000.asl`)

6.16.2 Background

ACPI is present in all modern computers that conform to the ia32 (x86), ia64 (Itanium), and amd64 (AMD) architectures. The full standard has many features including CPU performance management, power planes control, thermal zones, various battery systems, embedded controllers, and bus enumeration. Most systems implement less than the full standard. For instance, a desktop system usually only implements the bus enumeration parts while a laptop might have cooling and battery management support as well. Laptops also have suspend and resume, with their own associated complexity.

An ACPI-compliant system has various components. The BIOS and chipset vendors provide various fixed tables (e.g., FADT) in memory that specify things like the APIC map (used for SMP), config registers, and simple configuration values. Additionally, a table of bytecode (the *Differentiated System Description Table* DSDT) is provided that specifies a tree-like name space of devices and methods.

The ACPI driver must parse the fixed tables, implement an interpreter for the bytecode, and modify device drivers and the kernel to accept information from the ACPI subsystem. For DragonFly, Intel has provided an interpreter (ACPI-CA) that is shared with Linux and NetBSD. The path to the ACPI-CA source code is `src/sys/contrib/dev/acpica-unix-YYYYMMDD`, where YYYYMMDD is the release date of the ACPI-CA source code. The glue code that allows ACPI-CA to work on DragonFly is in `src/sys/dev/acpica5/Osd`. Finally, drivers that implement various ACPI devices are found in `src/sys/dev/acpica5`, and architecture-dependent code resides in `/sys/arch/acpica5`.

6.16.3 Common Problems

For ACPI to work correctly, all the parts have to work correctly. Here are some common problems, in order of frequency of appearance, and some possible workarounds or fixes.

6.16.3.1 Suspend/Resume

ACPI has three suspend to RAM (STR) states, S1-S3, and one suspend to disk state (STD), called S4. S5 is “soft off” and is the normal state your system is in when plugged in but not powered up. S4 can actually be implemented two separate ways. S4BIOS is a BIOS-assisted suspend to disk. S4OS is implemented entirely by the operating system.

Start by checking `sysctl hw.acpi` for the suspend-related items. Here are the results for my Thinkpad:

```
hw.acpi.supported_sleep_state: S3 S4 S5
hw.acpi.s4bios: 0
```

This means that I can use `acpicconf -s` to test S3, S4OS, and S5. If `s4bios` was one (1), I would have S4BIOS support instead of S4 OS.

When testing suspend/resume, start with S1, if supported. This state is most likely to work since it doesn't require much driver support. No one has implemented S2 but if you have it, it's similar to S1. The next thing to try is S3. This is the deepest STR state and requires a lot of driver support to properly reinitialize your hardware. If you have problems resuming, feel free to email the bugs (<http://leaf.dragonflybsd.org/mailarchive/>) list but do not expect the problem to be resolved since there are a lot of drivers/hardware that need more testing and work.

To help isolate the problem, remove as many drivers from your kernel as possible. If it works, you can narrow down which driver is the problem by loading drivers until it fails again. Typically binary drivers like `nvidia.ko`, **X11** display drivers, and USB will have the most problems while Ethernet interfaces usually work fine. If you can load/unload the drivers ok, you can automate this by putting the appropriate commands in `/etc/rc.suspend` and `/etc/rc.resume`. There is a commented-out example for unloading and loading a driver. Try setting `hw.acpi.reset_video` to zero (0) if your display is messed up after resume. Try setting longer or shorter values for `hw.acpi.sleep_delay` to see if that helps.

Another thing to try is load a recent Linux distribution with ACPI support and test their suspend/resume support on the same hardware. If it works on Linux, it's likely a DragonFly driver problem and narrowing down which driver causes the problems will help us fix the problem. Note that the ACPI maintainers do not usually maintain other drivers (e.g sound, ATA, etc.) so any work done on tracking down a driver problem should probably eventually be posted to the bugs (<http://leaf.dragonflybsd.org/mailarchive/>) list and mailed to the driver maintainer. If you are feeling adventurous, go ahead and start putting some debugging `printf(3)`s in a problematic driver to track down where in its resume function it hangs.

Finally, try disabling ACPI and enabling APM instead. If suspend/resume works with APM, you may be better off sticking with APM, especially on older hardware (pre-2000). It took vendors a while to get ACPI support correct and older hardware is more likely to have BIOS problems with ACPI.

6.16.3.2 System Hangs (temporary or permanent)

Most system hangs are a result of lost interrupts or an interrupt storm. Chipsets have a lot of problems based on how the BIOS configures interrupts before boot, correctness of the APIC (MADT) table, and routing of the *System Control Interrupt* (SCI).

Interrupt storms can be distinguished from lost interrupts by checking the output of `vmstat -i` and looking at the line that has `acpi0`. If the counter is increasing at more than a couple per second, you have an interrupt storm. If the system appears hung, try breaking to DDB (**CTRL+ALT+ESC** on console) and type `show interrupts`.

Your best hope when dealing with interrupt problems is to try disabling APIC support with `hint.apic.0.disabled="1"` in `loader.conf`.

6.16.3.3 Panics

Panics are relatively rare for ACPI and are the top priority to be fixed. The first step is to isolate the steps to reproduce the panic (if possible) and get a backtrace. Follow the advice for enabling `options DDB` and setting up a serial console (see Section 17.6.5.3) or setting up a `dump(8)` partition. You can get a backtrace in DDB with `tr`. If you have to handwrite the backtrace, be sure to at least get the lowest five (5) and top five (5) lines in the trace.

Then, try to isolate the problem by booting with ACPI disabled. If that works, you can isolate the ACPI subsystem by using various values of `debug.acpi.disable`. See the `acpi(4)` manual page for some examples.

6.16.3.4 System Powers Up After Suspend or Shutdown

First, try setting `hw.acpi.disable_on_poweroff="0"` in `loader.conf(5)`. This keeps ACPI from disabling various events during the shutdown process. Some systems need this value set to "1" (the default) for the same reason. This usually fixes the problem of a system powering up spontaneously after a suspend or poweroff.

6.16.3.5 Other Problems

If you have other problems with ACPI (working with a docking station, devices not detected, etc.), please email a description to the mailing list as well; however, some of these issues may be related to unfinished parts of the ACPI subsystem so they might take a while to be implemented. Please be patient and prepared to test patches we may send you.

6.16.4 ASL, `acpidump`, and IASL

The most common problem is the BIOS vendors providing incorrect (or outright buggy!) bytecode. This is usually manifested by kernel console messages like this:

```
ACPI-1287: *** Error: Method execution failed [\\_SB_.PCI0.LPC0.FIGD._STA] \\
(Node 0xc3f6d160), AE_NOT_FOUND
```

Often, you can resolve these problems by updating your BIOS to the latest revision. Most console messages are harmless but if you have other problems like battery status not working, they're a good place to start looking for problems in the AML. The bytecode, known as AML, is compiled from a source language called ASL. The AML is found in the table known as the DSDT. To get a copy of your ASL, use `acpidump(8)`. You should use both the `-t` (show contents of the fixed tables) and `-d` (disassemble AML to ASL) options. See the [Submitting Debugging Information](#) section for an example syntax.

The simplest first check you can do is to recompile your ASL to check for errors. Warnings can usually be ignored but errors are bugs that will usually prevent ACPI from working correctly. To recompile your ASL, issue the following command:

```
# iasl your.asl
```

6.16.5 Fixing Your ASL

In the long run, our goal is for almost everyone to have ACPI work without any user intervention. At this point, however, we are still developing workarounds for common mistakes made by the BIOS vendors. The Microsoft interpreter (`acpi.sys` and `acpiec.sys`) does not strictly check for adherence to the standard, and thus many BIOS vendors who only test ACPI under Windows never fix their ASL. We hope to continue to identify and document exactly what non-standard behavior is allowed by Microsoft's interpreter and replicate it so DragonFly can work without forcing users to fix the ASL. As a workaround and to help us identify behavior, you can fix the ASL manually. If this works for you, please send a `diff(1)` of the old and new ASL so we can possibly work around the buggy behavior in ACPI-CA and thus make your fix unnecessary.

Here is a list of common error messages, their cause, and how to fix them:

6.16.5.1 _OS dependencies

Some AML assumes the world consists of various Windows versions. You can tell DragonFly to claim it is any OS to see if this fixes problems you may have. An easy way to override this is to set `hw.acpi.osname="Windows 2001"` in `/boot/loader.conf` or other similar strings you find in the ASL.

6.16.5.2 Missing Return statements

Some methods do not explicitly return a value as the standard requires. While ACPI-CA does not handle this, DragonFly has a workaround that allows it to return the value implicitly. You can also add explicit Return statements where required if you know what value should be returned. To force `iasl` to compile the ASL, use the `-f` flag.

6.16.5.3 Overriding the Default AML

After you customize `your.asl`, you will want to compile it, run:

```
# iasl your.asl
```

You can add the `-f` flag to force creation of the AML, even if there are errors during compilation. Remember that some errors (e.g., missing Return statements) are automatically worked around by the interpreter.

`DSDT.aml` is the default output filename for `iasl`. You can load this instead of your BIOS's buggy copy (which is still present in flash memory) by editing `/boot/loader.conf` as follows:

```
acpi_dsdt_load="YES"
acpi_dsdt_name="/boot/DSDT.aml"
```

Be sure to copy your DSDT.aml to the /boot directory.

6.16.6 Getting Debugging Output From ACPI

The ACPI driver has a very flexible debugging facility. It allows you to specify a set of subsystems as well as the level of verbosity. The subsystems you wish to debug are specified as “layers” and are broken down into ACPI-CA components (ACPI_ALL_COMPONENTS) and ACPI hardware support (ACPI_ALL_DRIVERS). The verbosity of debugging output is specified as the “level” and ranges from ACPI_LV_ERROR (just report errors) to ACPI_LV_VERBOSE (everything). The “level” is a bitmask so multiple options can be set at once, separated by spaces. In practice, you will want to use a serial console to log the output if it is so long it flushes the console message buffer.

Debugging output is not enabled by default. To enable it, add options `ACPI_DEBUG` to your kernel config if ACPI is compiled into the kernel. You can add `ACPI_DEBUG=1` to your `/etc/make.conf` to enable it globally. If it is a module, you can recompile just your `acpi.ko` module as follows:

```
# cd /sys/dev/acpica5
&& make clean &&
make ACPI_DEBUG=1
```

Install `acpi.ko` in `/boot/kernel` and add your desired level and layer to `loader.conf`. This example enables debug messages for all ACPI-CA components and all ACPI hardware drivers (CPU, LID, etc.) It will only output error messages, the least verbose level.

```
debug.acpi.layer="ACPI_ALL_COMPONENTS ACPI_ALL_DRIVERS"
debug.acpi.level="ACPI_LV_ERROR"
```

If the information you want is triggered by a specific event (say, a suspend and then resume), you can leave out changes to `loader.conf` and instead use `sysctl` to specify the layer and level after booting and preparing your system for the specific event. The `sysctls` are named the same as the tunables in `loader.conf`.

6.16.7 References

More information about ACPI may be found in the following locations:

- The FreeBSD ACPI mailing list (<http://lists.FreeBSD.org/mailman/listinfo/freebsd-acpi>) (This is FreeBSD-specific; posting DragonFly questions here may not generate much of an answer.)
- The ACPI Mailing List Archives (FreeBSD) <http://lists.freebsd.org/pipermail/freebsd-acpi/>
- The old ACPI Mailing List Archives (FreeBSD) <http://home.jp.FreeBSD.org/mail-list/acpi-jp/>
- The ACPI 2.0 Specification <http://acpi.info/spec.htm>
- DragonFly Manual pages: `acpidump(8)`, `acpicnf(8)`, `acpidb(8)`
- DSDT debugging resource (http://www.cpqlinux.com/acpi-howto.html#fix_broken_dsdt). (Uses Compaq as an example but generally useful.)

Notes

1. Previously this was used to define *BSD dependent features.

Chapter 7 The DragonFly Booting Process

7.1 Synopsis

The process of starting a computer and loading the operating system is referred to as “the bootstrap process”, or simply “booting”. DragonFly’s boot process provides a great deal of flexibility in customizing what happens when you start the system, allowing you to select from different operating systems installed on the same computer, or even different versions of the same operating system or installed kernel.

This chapter details the configuration options you can set and how to customize the DragonFly boot process. This includes everything that happens until the DragonFly kernel has started, probed for devices, and started `init(8)`. If you are not quite sure when this happens, it occurs when the text color changes from bright white to grey.

After reading this chapter, you will know:

- What the components of the DragonFly bootstrap system are, and how they interact.
- The options you can give to the components in the DragonFly bootstrap to control the boot process.
- The basics of `device.hints(5)`.

x86 Only: This chapter only describes the boot process for DragonFly running on x86 systems.

7.2 The Booting Problem

Turning on a computer and starting the operating system poses an interesting dilemma. By definition, the computer does not know how to do anything until the operating system is started. This includes running programs from the disk. So if the computer can not run a program from the disk without the operating system, and the operating system programs are on the disk, how is the operating system started?

This problem parallels one in the book *The Adventures of Baron Munchausen*. A character had fallen part way down a manhole, and pulled himself out by grabbing his bootstraps, and lifting. In the early days of computing the term *bootstrap* was applied to the mechanism used to load the operating system, which has become shortened to “booting”.

On x86 hardware the Basic Input/Output System (BIOS) is responsible for loading the operating system. To do this, the BIOS looks on the hard disk for the Master Boot Record (MBR), which must be located on a specific place on the disk. The BIOS has enough knowledge to load and run the MBR, and assumes that the MBR can then carry out the rest of the tasks involved in loading the operating system possibly with the help of the BIOS.

The code within the MBR is usually referred to as a *boot manager*, especially when it interacts with the user. In this case the boot manager usually has more code in the first *track* of the disk or within some OS’s file system. (A boot manager is sometimes also called a *boot loader*, but FreeBSD uses that term for a later stage of booting.) Popular boot managers include **boot0** (a.k.a. **Boot Easy**, the standard DragonFly boot manager), **Grub**, **GAG**, and **LILO**. (Only **boot0** fits within the MBR.)

If you have only one operating system installed on your disks then a standard PC MBR will suffice. This MBR searches for the first bootable (a.k.a. active) slice on the disk, and then runs the code on that slice to load the

remainder of the operating system. The MBR installed by `fdisk(8)`, by default, is such an MBR. It is based on `/boot/mbr`.

If you have installed multiple operating systems on your disks then you can install a different boot manager, one that can display a list of different operating systems, and allows you to choose the one to boot from. Two of these are discussed in the next subsection.

The remainder of the DragonFly bootstrap system is divided into three stages. The first stage is run by the MBR, which knows just enough to get the computer into a specific state and run the second stage. The second stage can do a little bit more, before running the third stage. The third stage finishes the task of loading the operating system. The work is split into these three stages because the PC standards put limits on the size of the programs that can be run at stages one and two. Chaining the tasks together allows DragonFly to provide a more flexible loader.

The kernel is then started and it begins to probe for devices and initialize them for use. Once the kernel boot process is finished, the kernel passes control to the user process `init(8)`, which then makes sure the disks are in a usable state. `init(8)` then starts the user-level resource configuration which mounts file systems, sets up network cards to communicate on the network, and generally starts all the processes that usually are run on a DragonFly system at startup.

7.3 The Boot Manager and Boot Stages

7.3.1 The Boot Manager

The code in the MBR or boot manager is sometimes referred to as *stage zero* of the boot process. This subsection discusses two of the boot managers previously mentioned: **boot0** and **LILO**.

The boot0 Boot Manager: The MBR installed by FreeBSD's installer or `boot0cfg(8)`, by default, is based on `/boot/boot0`. (The **boot0** program is very simple, since the program in the MBR can only be 446 bytes long because of the slice table and `0x55AA` identifier at the end of the MBR.) If you have installed **boot0** and multiple operating systems on your hard disks, then you will see a display similar to this one at boot time:

Example 7-1. `boot0` Screenshot

```
F1 DOS
F2 FreeBSD
F3 Linux
F4 ??
F5 Drive 1

Default: F2
```

Other operating systems, in particular Windows, have been known to overwrite an existing MBR with their own. If this happens to you, or you want to replace your existing MBR with the DragonFly MBR then use the following command:

```
# fdisk -B -b /boot/boot0 device
```

where *device* is the device that you boot from, such as `ad0` for the first IDE disk, `ad2` for the first IDE disk on a second IDE controller, `da0` for the first SCSI disk, and so on. Or, if you want a custom configuration of the MBR, use `boot0cfg(8)`.

The LILO Boot Manager: To install this boot manager so it will also boot DragonFly, first start Linux and add the following to your existing `/etc/lilo.conf` configuration file:

```
other=/dev/hdXY
table=/dev/hdX
loader=/boot/chain.b
label=DragonFly
```

In the above, specify DragonFly's primary partition and drive using Linux specifiers, replacing *x* with the Linux drive letter and *y* with the Linux primary partition number. If you are using a SCSI drive, you will need to change `/dev/hd` to read something similar to `/dev/sd`. The `loader=/boot/chain.b` line can be omitted if you have both operating systems on the same drive. Now run `/sbin/lilo -v` to commit your new changes to the system; this should be verified by checking its screen messages.

7.3.2 Stage One, `/boot/boot1`, and Stage Two, `/boot/boot2`

Conceptually the first and second stages are part of the same program, on the same area of the disk. Because of space constraints they have been split into two, but you would always install them together. They are copied from the combined file `/boot/boot` by the installer or **disklabel** (see below).

They are located outside file systems, in the first track of the boot slice, starting with the first sector. This is where `boot0`, or any other boot manager, expects to find a program to run which will continue the boot process. The number of sectors used is easily determined from the size of `/boot/boot`.

They are found on the boot sector of the boot slice, which is where `boot0`, or any other program on the MBR expects to find the program to run to continue the boot process. The files in the `/boot` directory are copies of the real files, which are stored outside of the DragonFly file system.

`boot1` is very simple, since it can only be 512 bytes in size, and knows just enough about the DragonFly *disklabel*, which stores information about the slice, to find and execute `boot2`.

`boot2` is slightly more sophisticated, and understands the DragonFly file system enough to find files on it, and can provide a simple interface to choose the kernel or loader to run.

Since the loader is much more sophisticated, and provides a nice easy-to-use boot configuration, `boot2` usually runs it, but previously it was tasked to run the kernel directly.

Example 7-2. `boot2` Screenshot

```
>> DragonFly/i386 BOOT
Default: 0:ad(0,a)/boot/loader
boot:
```

If you ever need to replace the installed `boot1` and `boot2` use `disklabel(8)`:

```
# disklabel -B diskslice
```

where *diskslice* is the disk and slice you boot from, such as `ad0s1` for the first slice on the first IDE disk.

7.3.3 Stage Three, /boot/loader

The loader is the final stage of the three-stage bootstrap, and is located on the file system, usually as /boot/loader.

The loader is intended as a user-friendly method for configuration, using an easy-to-use built-in command set, backed up by a more powerful interpreter, with a more complex command set.

7.3.3.1 Loader Program Flow

During initialization, the loader will probe for a console and for disks, and figure out what disk it is booting from. It will set variables accordingly, and an interpreter is started where user commands can be passed from a script or interactively.

The loader will then read /boot/loader.rc, which by default reads in /boot/defaults/loader.conf which sets reasonable defaults for variables and reads /boot/loader.conf for local changes to those variables.

loader.rc then acts on these variables, loading whichever modules and kernel are selected.

Finally, by default, the loader issues a 10 second wait for key presses, and boots the kernel if it is not interrupted. If interrupted, the user is presented with a prompt which understands the easy-to-use command set, where the user may adjust variables, unload all modules, load modules, and then finally boot or reboot.

7.3.3.2 Loader Built-In Commands

These are the most commonly used loader commands. For a complete discussion of all available commands, please see loader(8).

`autoboot seconds`

Proceeds to boot the kernel if not interrupted within the time span given, in seconds. It displays a countdown, and the default time span is 10 seconds.

`boot [-options] [kernelname]`

Immediately proceeds to boot the kernel, with the given options, if any, and with the kernel name given, if it is.

`boot-conf`

Goes through the same automatic configuration of modules based on variables as what happens at boot. This only makes sense if you use `unload` first, and change some variables, most commonly `kernel`.

`help [topic]`

Shows help messages read from /boot/loader.help. If the topic given is `index`, then the list of available topics is given.

`include filename ...`

Processes the file with the given filename. The file is read in, and interpreted line by line. An error immediately stops the include command.

`load [-t type] filename`

Loads the kernel, kernel module, or file of the type given, with the filename given. Any arguments after filename are passed to the file.

`ls [-l] [path]`

Displays a listing of files in the given path, or the root directory, if the path is not specified. If `-l` is specified, file sizes will be shown too.

`lsdev [-v]`

Lists all of the devices from which it may be possible to load modules. If `-v` is specified, more details are printed.

`lsmod [-v]`

Displays loaded modules. If `-v` is specified, more details are shown.

`more filename`

Displays the files specified, with a pause at each `LINES` displayed.

`reboot`

Immediately reboots the system.

`set variable`

`set variable=value`

Sets the loader's environment variables.

`unload`

Removes all loaded modules.

7.3.3.3 Loader Examples

Here are some practical examples of loader usage:

- To simply boot your usual kernel, but in single-user mode:

```
boot -s
```

- To unload your usual kernel and modules, and then load just your old (or another) kernel:

```
unload
load kernel.old
```

You can use `kernel.GENERIC` to refer to the generic kernel that comes on the install disk, or `kernel.old` to refer to your previously installed kernel (when you have upgraded or configured your own kernel, for example).

Note: Use the following to load your usual modules with another kernel:

```
unload
set kernel="kernel.old"
boot-conf
```

- To load a kernel configuration script (an automated script which does the things you would normally do in the kernel boot-time configurator):

```
load -t userconfig_script /boot/kernel.conf
```

7.4 Kernel Interaction During Boot

Once the kernel is loaded by either loader (as usual) or boot2 (bypassing the loader), it examines its boot flags, if any, and adjusts its behavior as necessary.

7.4.1 Kernel Boot Flags

Here are the more common boot flags:

- a
during kernel initialization, ask for the device to mount as the root file system.
- C
boot from CDROM.
- c
run UserConfig, the boot-time kernel configurator
- s
boot into single-user mode
- v
be more verbose during kernel startup

Note: There are other boot flags; read `boot(8)` for more information on them.

7.5 Init: Process Control Initialization

Once the kernel has finished booting, it passes control to the user process `init(8)`, which is located at `/sbin/init`, or the program path specified in the `init_path` variable in `loader`.

7.5.1 Automatic Reboot Sequence

The automatic reboot sequence makes sure that the file systems available on the system are consistent. If they are not, and `fsck(8)` cannot fix the inconsistencies, `init(8)` drops the system into single-user mode for the system administrator to take care of the problems directly.

7.5.2 Single-User Mode

This mode can be reached through the automatic reboot sequence, or by the user booting with the `-s` option or setting the `boot_single` variable in `loader`.

It can also be reached by calling `shutdown(8)` without the `reboot (-r)` or `halt (-h)` options, from multi-user mode.

If the system `console` is set to `insecure` in `/etc/ttys`, then the system prompts for the `root` password before initiating single-user mode.

Example 7-3. An Insecure Console in `/etc/ttys`

```
# name  getty                                type    status      comments
#
# If console is marked "insecure", then init will ask for the root password
# when going to single-user mode.
console none                                unknown off insecure
```

Note: An `insecure` console means that you consider your physical security to the console to be insecure, and want to make sure only someone who knows the `root` password may use single-user mode, and it does not mean that you want to run your console insecurely. Thus, if you want security, choose `insecure`, not `secure`.

7.5.3 Multi-User Mode

If `init(8)` finds your file systems to be in order, or once the user has finished in single-user mode, the system enters multi-user mode, in which it starts the resource configuration of the system.

7.5.3.1 Resource Configuration (`rc`)

The resource configuration system reads in configuration defaults from `/etc/defaults/rc.conf`, and system-specific details from `/etc/rc.conf`, and then proceeds to mount the system file systems mentioned in `/etc/fstab`, start up networking services, start up miscellaneous system daemons, and finally runs the startup scripts of locally installed packages.

The `rc(8)` manual page is a good reference to the resource configuration system, as is examining the scripts themselves.

7.6 Shutdown Sequence

Upon controlled shutdown, via `shutdown(8)`, `init(8)` will attempt to run the script `/etc/rc.shutdown`, and then proceed to send all processes the `TERM` signal, and subsequently the `KILL` signal to any that do not terminate timely.

To power down a DragonFly machine on architectures and systems that support power management, simply use the command `shutdown -p now` to turn the power off immediately. To just reboot a DragonFly system, just use `shutdown -r now`. You need to be `root` or a member of `operator` group to run `shutdown(8)`. The `halt(8)` and

reboot(8) commands can also be used, please refer to their manual pages and to shutdown(8)'s one for more information.

Note: Power management requires acpi(4) support in the kernel or loaded as a module, or apm(4) support.

Chapter 8 Users and Basic Account Management

Contributed by Neil Blakey-Milner.

8.1 Synopsis

DragonFly allows multiple users to use the computer at the same time. Obviously, only one of those users can be sitting in front of the screen and keyboard at any one time ¹, but any number of users can log in through the network to get their work done. To use the system every user must have an account.

After reading this chapter, you will know:

- The differences between the various user accounts on a DragonFly system.
- How to add user accounts.
- How to remove user accounts.
- How to change account details, such as the user's full name, or preferred shell.
- How to set limits on a per-account basis, to control the resources such as memory and CPU time that accounts and groups of accounts are allowed to access.
- How to use groups to make account management easier.

Before reading this chapter, you should:

- Understand the basics of UNIX and DragonFly (Chapter 3).

8.2 Introduction

All access to the system is achieved via accounts, and all processes are run by users, so user and account management are of integral importance on DragonFly systems.

Every account on a DragonFly system has certain information associated with it to identify the account.

User name

The user name as it would be typed at the `login:` prompt. User names must be unique across the computer; you may not have two users with the same user name. There are a number of rules for creating valid user names, documented in `passwd(5)`; you would typically use user names that consist of eight or fewer all lower case characters.

Password

Each account has a password associated with it. The password may be blank, in which case no password will be required to access the system. This is normally a very bad idea; every account should have a password.

User ID (UID)

The UID is a number, traditionally from 0 to 65535², used to uniquely identify the user to the system. Internally, DragonFly uses the UID to identify users—any DragonFly commands that allow you to specify a user name will convert it to the UID before working with it. This means that you can have several accounts with different user names but the same UID. As far as DragonFly is concerned, these accounts are one user. It is unlikely you will ever need to do this.

Group ID (GID)

The GID is a number, traditionally from 0 to 65535², used to uniquely identify the primary group that the user belongs to. Groups are a mechanism for controlling access to resources based on a user's GID rather than their UID. This can significantly reduce the size of some configuration files. A user may also be in more than one group.

Login class

Login classes are an extension to the group mechanism that provide additional flexibility when tailoring the system to different users.

Password change time

By default DragonFly does not force users to change their passwords periodically. You can enforce this on a per-user basis, forcing some or all of your users to change their passwords after a certain amount of time has elapsed.

Account expiry time

By default DragonFly does not expire accounts. If you are creating accounts that you know have a limited lifespan, for example, in a school where you have accounts for the students, then you can specify when the account expires. After the expiry time has elapsed the account cannot be used to log in to the system, although the account's directories and files will remain.

User's full name

The user name uniquely identifies the account to DragonFly, but does not necessarily reflect the user's real name. This information can be associated with the account.

Home directory

The home directory is the full path to a directory on the system in which the user will start when logging on to the system. A common convention is to put all user home directories under `/home/username` or `/usr/home/username`. The user would store their personal files in their home directory, and any directories they may create in there.

User shell

The shell provides the default environment users use to interact with the system. There are many different kinds of shells, and experienced users will have their own preferences, which can be reflected in their account settings.

There are three main types of accounts: the Superuser, system users, and user accounts. The Superuser account, usually called `root`, is used to manage the system with no limitations on privileges. System users run services. Finally, user accounts are used by real people, who log on, read mail, and so forth.

8.3 The Superuser Account

The superuser account, usually called `root`, comes preconfigured to facilitate system administration, and should not be used for day-to-day tasks like sending and receiving mail, general exploration of the system, or programming.

This is because the superuser, unlike normal user accounts, can operate without limits, and misuse of the superuser account may result in spectacular disasters. User accounts are unable to destroy the system by mistake, so it is generally best to use normal user accounts whenever possible, unless you especially need the extra privilege.

You should always double and triple-check commands you issue as the superuser, since an extra space or missing character can mean irreparable data loss.

So, the first thing you should do after reading this chapter is to create an unprivileged user account for yourself for general usage if you have not already. This applies equally whether you are running a multi-user or single-user machine. Later in this chapter, we discuss how to create additional accounts, and how to change between the normal user and superuser.

8.4 System Accounts

System users are those used to run services such as DNS, mail, web servers, and so forth. The reason for this is security; if all services ran as the superuser, they could act without restriction.

Examples of system users are `daemon`, `operator`, `bind` (for the Domain Name Service), and `news`. Often sysadmins create `httpd` to run web servers they install.

`nobody` is the generic unprivileged system user. However, it is important to keep in mind that the more services that use `nobody`, the more files and processes that user will become associated with, and hence the more privileged that user becomes.

8.5 User Accounts

User accounts are the primary means of access for real people to the system, and these accounts insulate the user and the environment, preventing the users from damaging the system or other users, and allowing users to customize their environment without affecting others.

Every person accessing your system should have a unique user account. This allows you to find out who is doing what, prevent people from clobbering each others' settings or reading each others' mail, and so forth.

Each user can set up their own environment to accommodate their use of the system, by using alternate shells, editors, key bindings, and language.

8.6 Modifying Accounts

There are a variety of different commands available in the UNIX environment to manipulate user accounts. The most common commands are summarized below, followed by more detailed examples of their usage.

Command	Summary
<code>adduser(8)</code>	The recommended command-line application for adding new users.

Command	Summary
rmuser(8)	The recommended command-line application for removing users.
chpass(1)	A flexible tool to change user database information.
passwd(1)	The simple command-line tool to change user passwords.
pw(8)	A powerful and flexible tool to modify all aspects of user accounts.

8.6.1 adduser

adduser(8) is a simple program for adding new users. It creates entries in the system `passwd` and `group` files. It will also create a home directory for the new user, copy in the default configuration files (“dotfiles”) from `/usr/share/skel`, and can optionally mail the new user a welcome message.

To create the initial configuration file, use `adduser -s -config_create`.³ Next, we configure `adduser(8)` defaults, and create our first user account, since using `root` for normal usage is evil and nasty.

Example 8-1. Configuring `adduser` and adding a user

```
# adduser -v
Use option "--silent" if you don't want to see all warnings and questions.
Check /etc/shells
Check /etc/master.passwd
Check /etc/group
Enter your default shell: csh date no sh tcsh zsh [sh]: zsh
Your default shell is: zsh -> /usr/local/bin/zsh
Enter your default HOME partition: [/home]:
Copy dotfiles from: /usr/share/skel no [/usr/share/skel]:
Send message from file: /etc/adduser.message no
[/etc/adduser.message]: no
Do not send message
Use passwords (y/n) [y]: y

Write your changes to /etc/adduser.conf? (y/n) [n]: y

Ok, let's go.
Don't worry about mistakes. I will give you the chance later to correct any input.
Enter username [a-z0-9_-]: jru
Enter full name []: J. Random User
Enter shell csh date no sh tcsh zsh [zsh]:
Enter home directory (full path) [/home/jru]:
Uid [1001]:
Enter login class: default []:
Login group jru [jru]:
Login group is "jru". Invite jru into other groups: guest no
[no]: wheel
Enter password []:
Enter password again []:

Name: jru
Password: ****
```

```

Fullname: J. Random User
Uid:    1001
Gid:    1001 (jru)
Class:
Groups:  jru wheel
HOME:    /home/jru
Shell:   /usr/local/bin/zsh
OK? (y/n) [y]: y
Added user "jru"
Copy files from /usr/share/skel to /home/jru
Add another user? (y/n) [y]: n
Goodbye!
#

```

In summary, we changed the default shell to **zsh** (an additional shell found in `pkgsrc`), and turned off the sending of a welcome mail to added users. We then saved the configuration, created an account for `jru`, and made sure `jru` is in `wheel` group (so that she may assume the role of `root` with the `su(1)` command.)

Note: The password you type in is not echoed, nor are asterisks displayed. Make sure you do not mistype the password twice.

Note: Just use `adduser(8)` without arguments from now on, and you will not have to go through changing the defaults. If the program asks you to change the defaults, exit the program, and try the `-s` option.

8.6.2 `rmuser`

You can use `rmuser(8)` to completely remove a user from the system. `rmuser(8)` performs the following steps:

1. Removes the user's `crontab(1)` entry (if any).
2. Removes any `at(1)` jobs belonging to the user.
3. Kills all processes owned by the user.
4. Removes the user from the system's local password file.
5. Removes the user's home directory (if it is owned by the user).
6. Removes the incoming mail files belonging to the user from `/var/mail`.
7. Removes all files owned by the user from temporary file storage areas such as `/tmp`.
8. Finally, removes the username from all groups to which it belongs in `/etc/group`.

Note: If a group becomes empty and the group name is the same as the username, the group is removed; this complements the per-user unique groups created by `adduser(8)`.

`rmuser(8)` cannot be used to remove superuser accounts, since that is almost always an indication of massive destruction.

By default, an interactive mode is used, which attempts to make sure you know what you are doing.

Example 8-2. `rmuser` Interactive Account Removal

```
# rmuser jru
Matching password entry:
jru:*:1001:1001::0:0:J. Random User:/home/jru:/usr/local/bin/zsh
Is this the entry you wish to remove? y
Remove user's home directory (/home/jru)? y
Updating password file, updating databases, done.
Updating group file: trusted (removing group jru -- personal group is empty) done.
Removing user's incoming mail file /var/mail/jru: done.
Removing files belonging to jru from /tmp: done.
Removing files belonging to jru from /var/tmp: done.
Removing files belonging to jru from /var/tmp/vi.recover: done.
#
```

8.6.3 `chpass`

`chpass(1)` changes user database information such as passwords, shells, and personal information.

Only system administrators, as the superuser, may change other users' information and passwords with `chpass(1)`.

When passed no options, aside from an optional username, `chpass(1)` displays an editor containing user information.

When the user exists from the editor, the user database is updated with the new information.

Example 8-3. Interactive `chpass` by Superuser

```
#Changing user database information for jru.
Login: jru
Password: *
Uid [#]: 1001
Gid [# or name]: 1001
Change [month day year]:
Expire [month day year]:
Class:
Home directory: /home/jru
Shell: /usr/local/bin/zsh
Full Name: J. Random User
Office Location:
Office Phone:
Home Phone:
Other information:
```

The normal user can change only a small subset of this information, and only for themselves.

Example 8-4. Interactive `chpass` by Normal User

```
#Changing user database information for jru.
Shell: /usr/local/bin/zsh
Full Name: J. Random User
Office Location:
Office Phone:
Home Phone:
Other information:
```

Note: `chfn(1)` and `chsh(1)` are just links to `chpass(1)`, as are `ypchpass(1)`, `ypchfn(1)`, and `ypchsh(1)`. NIS support is automatic, so specifying the `yp` before the command is not necessary. If this is confusing to you, do not worry, NIS will be covered in Chapter 19.

8.6.4 `passwd`

`passwd(1)` is the usual way to change your own password as a user, or another user's password as the superuser.

Note: To prevent accidental or unauthorized changes, the original password must be entered before a new password can be set.

Example 8-5. Changing Your Password

```
% passwd
Changing local password for jru.
Old password:
New password:
Retype new password:
passwd: updating the database...
passwd: done
```

Example 8-6. Changing Another User's Password as the Superuser

```
# passwd jru
Changing local password for jru.
New password:
Retype new password:
passwd: updating the database...
passwd: done
```

Note: As with `chpass(1)`, `yppasswd(1)` is just a link to `passwd(1)`, so NIS works with either command.

8.6.5 pw

pw(8) is a command line utility to create, remove, modify, and display users and groups. It functions as a front end to the system user and group files. pw(8) has a very powerful set of command line options that make it suitable for use in shell scripts, but new users may find it more complicated than the other commands presented here.

8.7 Limiting Users

If you have users, the ability to limit their system use may have come to mind. DragonFly provides several ways an administrator can limit the amount of system resources an individual may use. These limits are divided into two sections: disk quotas, and other resource limits.

Disk quotas limit disk usage to users, and they provide a way to quickly check that usage without calculating it every time. Quotas are discussed in Section 12.12.

The other resource limits include ways to limit the amount of CPU, memory, and other resources a user may consume. These are defined using login classes and are discussed here.

Login classes are defined in `/etc/login.conf`. The precise semantics are beyond the scope of this section, but are described in detail in the `login.conf(5)` manual page. It is sufficient to say that each user is assigned to a login class (`default` by default), and that each login class has a set of login capabilities associated with it. A login capability is a `name=value` pair, where `name` is a well-known identifier and `value` is an arbitrary string processed accordingly depending on the name. Setting up login classes and capabilities is rather straight-forward and is also described in `login.conf(5)`.

Resource limits are different from plain vanilla login capabilities in two ways. First, for every limit, there is a soft (current) and hard limit. A soft limit may be adjusted by the user or application, but may be no higher than the hard limit. The latter may be lowered by the user, but never raised. Second, most resource limits apply per process to a specific user, not the user as a whole. Note, however, that these differences are mandated by the specific handling of the limits, not by the implementation of the login capability framework (i.e., they are not *really* a special case of login capabilities).

And so, without further ado, below are the most commonly used resource limits (the rest, along with all the other login capabilities, may be found in `login.conf(5)`).

`coredumpsize`

The limit on the size of a core file generated by a program is, for obvious reasons, subordinate to other limits on disk usage (e.g., `filesize`, or disk quotas). Nevertheless, it is often used as a less-severe method of controlling disk space consumption: since users do not generate core files themselves, and often do not delete them, setting this may save them from running out of disk space should a large program (e.g., **emacs**) crash.

`cputime`

This is the maximum amount of CPU time a user's process may consume. Offending processes will be killed by the kernel.

Note: This is a limit on CPU *time* consumed, not percentage of the CPU as displayed in some fields by `top(1)` and `ps(1)`. A limit on the latter is, at the time of this writing, not possible, and would be rather useless: legitimate use of a compiler, for instance, can easily use almost 100% of a CPU for some time.

`filesize`

This is the maximum size of a file the user may possess. Unlike disk quotas, this limit is enforced on individual files, not the set of all files a user owns.

`maxproc`

This is the maximum number of processes a user may be running. This includes foreground and background processes alike. For obvious reasons, this may not be larger than the system limit specified by the `kern.maxproc` `sysctl(8)`. Also note that setting this too small may hinder a user's productivity: it is often useful to be logged in multiple times or execute pipelines. Some tasks, such as compiling a large program, also spawn multiple processes (e.g., `make(1)`, `cc(1)`, and other intermediate preprocessors).

`memorylocked`

This is the maximum amount of memory a process may have requested to be locked into main memory (e.g., see `mlock(2)`). Some system-critical programs, such as `amd(8)`, lock into main memory such that in the event of being swapped out, they do not contribute to a system's trashing in time of trouble.

`memoryuse`

This is the maximum amount of memory a process may consume at any given time. It includes both core memory and swap usage. This is not a catch-all limit for restricting memory consumption, but it is a good start.

`openfiles`

This is the maximum amount of files a process may have open. In DragonFly, files are also used to represent sockets and IPC channels; thus, be careful not to set this too low. The system-wide limit for this is defined by the `kern.maxfiles` `sysctl(8)`.

`sbsize`

This is the limit on the amount of network memory, and thus mbufs, a user may consume. This originated as a response to an old DoS attack by creating a lot of sockets, but can be generally used to limit network communications.

`stacksize`

This is the maximum size a process' stack may grow to. This alone is not sufficient to limit the amount of memory a program may use; consequently, it should be used in conjunction with other limits.

There are a few other things to remember when setting resource limits. Following are some general tips, suggestions, and miscellaneous comments.

- Processes started at system startup by `/etc/rc` are assigned to the `daemon` login class.
- Although the `/etc/login.conf` that comes with the system is a good source of reasonable values for most limits, only you, the administrator, can know what is appropriate for your system. Setting a limit too high may open your system up to abuse, while setting it too low may put a strain on productivity.
- Users of the X Window System (X11) should probably be granted more resources than other users. X11 by itself takes a lot of resources, but it also encourages users to run more programs simultaneously.
- Remember that many limits apply to individual processes, not the user as a whole. For example, setting `openfiles` to 50 means that each process the user runs may open up to 50 files. Thus, the gross amount of files a user may open is the value of `openfiles` multiplied by the value of `maxproc`. This also applies to memory consumption.

For further information on resource limits and login classes and capabilities in general, please consult the relevant manual pages: `cap_mkdb(1)`, `getrlimit(2)`, `login.conf(5)`.

8.8 Personalizing Users

Localization is an environment set up by the system administrator or user to accommodate different languages, character sets, date and time standards, and so on. This is discussed in the localization chapter.

8.9 Groups

A group is simply a list of users. Groups are identified by their group name and GID (Group ID). In DragonFly (and most other UNIX like systems), the two factors the kernel uses to decide whether a process is allowed to do something is its user ID and list of groups it belongs to. Unlike a user ID, a process has a list of groups associated with it. You may hear some things refer to the “group ID” of a user or process; most of the time, this just means the first group in the list.

The group name to group ID map is in `/etc/group`. This is a plain text file with four colon-delimited fields. The first field is the group name, the second is the encrypted password, the third the group ID, and the fourth the comma-delimited list of members. It can safely be edited by hand (assuming, of course, that you do not make any syntax errors!). For a more complete description of the syntax, see the `group(5)` manual page.

If you do not want to edit `/etc/group` manually, you can use the `pw(8)` command to add and edit groups. For example, to add a group called `teamtwo` and then confirm that it exists you can use:

Example 8-7. Adding a Group Using `pw(8)`

```
# pw groupadd teamtwo
```

```
# pw groupshow teamtwo
teamtwo:*:1100:
```

The number 1100 above is the group ID of the group `teamtwo`. Right now, `teamtwo` has no members, and is thus rather useless. Let's change that by inviting `jru` to the `teamtwo` group.

Example 8-8. Adding Somebody to a Group Using `pw(8)`

```
# pw groupmod teamtwo -M jru
# pw groupshow teamtwo
teamtwo:*:1100:jru
```

The argument to the `-M` option is a comma-delimited list of users who are members of the group. From the preceding sections, we know that the password file also contains a group for each user. The latter (the user) is automatically added to the group list by the system; the user will not show up as a member when using the `groupshow` command to `pw(8)`, but will show up when the information is queried via `id(1)` or similar tool. In other words, `pw(8)` only manipulates the `/etc/group` file; it will never attempt to read additionally data from `/etc/passwd`.

Example 8-9. Using `id(1)` to Determine Group Membership

```
% id jru
uid=1001(jru) gid=1001(jru) groups=1001(jru), 1100(teamtwo)
```

As you can see, `jru` is a member of the groups `jru` and `teamtwo`.

For more information about `pw(8)`, see its manual page, and for more information on the format of `/etc/group`, consult the `group(5)` manual page.

Notes

1. Well, unless you hook up multiple terminals, but we will save that for Chapter 17.
2. It is possible to use UID/GIDs as large as 4294967295, but such IDs can cause serious problems with software that makes assumptions about the values of IDs.
3. The `-s` makes `adduser(8)` default to quiet. We use `-v` later when we want to change defaults.

Chapter 9 Configuring the DragonFly Kernel

Updated and restructured by Jim Mock. Originally contributed by Jake Hamby.

9.1 Synopsis

The kernel is the core of the DragonFly operating system. It is responsible for managing memory, enforcing security controls, networking, disk access, and much more. While more and more of DragonFly becomes dynamically configurable it is still occasionally necessary to reconfigure and recompile your kernel.

After reading this chapter, you will know:

- Why you might need to build a custom kernel.
- How to write a kernel configuration file, or alter an existing configuration file.
- How to use the kernel configuration file to create and build a new kernel.
- How to install the new kernel.
- How to create any entries in `/dev` that may be required.
- How to troubleshoot if things go wrong.

9.2 Why Build a Custom Kernel?

Traditionally, DragonFly has had what is called a “monolithic” kernel. This means that the kernel was one large program, supported a fixed list of devices, and if you wanted to change the kernel’s behavior then you had to compile a new kernel, and then reboot your computer with the new kernel.

Today, DragonFly is rapidly moving to a model where much of the kernel’s functionality is contained in modules which can be dynamically loaded and unloaded from the kernel as necessary. This allows the kernel to adapt to new hardware suddenly becoming available (such as PCMCIA cards in a laptop), or for new functionality to be brought into the kernel that was not necessary when the kernel was originally compiled. This is known as a modular kernel. Colloquially these are called KLDs.

Despite this, it is still necessary to carry out some static kernel configuration. In some cases this is because the functionality is so tied to the kernel that it can not be made dynamically loadable. In others it may simply be because no one has yet taken the time to write a dynamic loadable kernel module for that functionality yet.

Building a custom kernel is one of the most important rites of passage nearly every UNIX user must endure. This process, while time consuming, will provide many benefits to your DragonFly system. Unlike the `GENERIC` kernel, which must support a wide range of hardware, a custom kernel only contains support for *your* PC’s hardware. This has a number of benefits, such as:

- Faster boot time. Since the kernel will only probe the hardware you have on your system, the time it takes your system to boot will decrease dramatically.
- Less memory usage. A custom kernel often uses less memory than the `GENERIC` kernel, which is important because the kernel must always be present in real memory. For this reason, a custom kernel is especially useful on a system with a small amount of RAM.

- Additional hardware support. A custom kernel allows you to add in support for devices such as sound cards, which are not present in the `GENERIC` kernel.

9.3 Building and Installing a Custom Kernel

First, let us take a quick tour of the kernel build directory. All directories mentioned will be relative to the main `/usr/src/sys` directory, which is also accessible through `/sys`. There are a number of subdirectories here representing different parts of the kernel, but the most important, for our purposes, are `arch/conf`, where you will edit your custom kernel configuration, and `compile`, which is the staging area where your kernel will be built. `arch` represents either `i386` or `amd64`, at this point in development. Everything inside a particular architecture's directory deals with that architecture only; the rest of the code is common to all platforms to which DragonFly could potentially be ported. Notice the logical organization of the directory structure, with each supported device, file system, and option in its own subdirectory.

Note: If there is *not* a `/usr/src/sys` directory on your system, then the kernel source has not been installed. The easiest way to do this is via `cvsup`.

Next, move to the `arch/conf` directory and copy the `GENERIC` configuration file to the name you want to give your kernel. For example:

```
# cd /usr/src/sys/i386/conf
# cp GENERIC MYKERNEL
```

Traditionally, this name is in all capital letters and, if you are maintaining multiple DragonFly machines with different hardware, it is a good idea to name it after your machine's hostname. We will call it `MYKERNEL` for the purpose of this example.

Tip: Storing your kernel config file directly under `/usr/src` can be a bad idea. If you are experiencing problems it can be tempting to just delete `/usr/src` and start again. Five seconds after you do that you realize that you have deleted your custom kernel config file. Do not edit `GENERIC` directly, as it may get overwritten the next time you update your source tree, and your kernel modifications will be lost.

You might want to keep your kernel config file elsewhere, and then create a symbolic link to the file in the `i386` directory.

For example:

```
# cd /usr/src/sys/i386/conf
# mkdir /root/kernels
# cp GENERIC /root/kernels/MYKERNEL
# ln -s /root/kernels/MYKERNEL
```

Note: You must execute these and all of the following commands under the `root` account or you will get permission denied errors.

Now, edit `MYKERNEL` with your favorite text editor. If you are just starting out, the only editor available will probably be `vi`, which is too complex to explain here, but is covered well in many books in the bibliography. However, DragonFly does offer an easier editor called `ee` which, if you are a beginner, should be your editor of choice. Feel free to change the comment lines at the top to reflect your configuration or the changes you have made to differentiate it from `GENERIC`.

If you have built a kernel under SunOS™ or some other BSD operating system, much of this file will be very familiar to you. If you are coming from some other operating system such as DOS, on the other hand, the `GENERIC` configuration file might seem overwhelming to you, so follow the descriptions in the Configuration File section slowly and carefully.

Note: Be sure to always check the file `/usr/src/UPDATING`, before you perform any update steps, in the case you sync your source tree with the latest sources of the DragonFly project. In this file all important issues with updating DragonFly are typed out. `/usr/src/UPDATING` always fits your version of the DragonFly source, and is therefore more accurate for new information than the handbook.

Building a Kernel

1. Change to the `/usr/src` directory.

```
# cd /usr/src
```

2. Compile the kernel.

```
# make buildkernel KERNCONF=MYKERNEL
```

3. Install the new kernel.

```
# make installkernel KERNCONF=MYKERNEL
```

If you have *not* upgraded your source tree in any way since the last time you successfully completed a `buildworld-installworld` cycle (you have not run `CVSup`), then it is safe to use the `quickworld` and `quickkernel`, `buildworld`, `buildkernel` sequence.

The new kernel will be copied to the root directory as `/kernel` and the old kernel will be moved to `/kernel.old`. Now, shutdown the system and reboot to use your new kernel. In case something goes wrong, there are some **troubleshooting** instructions at the end of this chapter. Be sure to read the section which explains how to recover in case your new kernel does not boot.

Note: If you have added any new devices (such as sound cards), you may have to add some device nodes to your `/dev` directory before you can use them. For more information, take a look at Making Device Nodes section later on in this chapter.

9.4 The Configuration File

The general format of a configuration file is quite simple. Each line contains a keyword and one or more arguments. For simplicity, most lines only contain one argument. Anything following a `#` is considered a comment and ignored. The following sections describe each keyword, generally in the order they are listed in `GENERIC`, although some related keywords have been grouped together in a single section (such as Networking) even though they are actually

scattered throughout the `GENERIC` file. An exhaustive list of options and more detailed explanations of the device lines is present in the `LINT` configuration file, located in the same directory as `GENERIC`. If you are in doubt as to the purpose or necessity of a line, check first in `LINT`.

The following is an example `GENERIC` kernel configuration file with various additional comments where needed for clarity. This example should match your copy in `/usr/src/sys/i386/conf/GENERIC` fairly closely. For details of all the possible kernel options, see `/usr/src/sys/i386/conf/LINT`.

```
#
#
# GENERIC -- Generic kernel configuration file for DragonFly/i386
#
# Check the LINT configuration file in sys/i386/conf, for an
# exhaustive list of options.
#
# $DragonFly: src/sys/i386/conf/GENERIC,v 1.17 2004/06/25 05:09:38 hmp Exp $
```

The following are the mandatory keywords required in *every* kernel you build:

```
machine i386
```

This is the machine architecture. It must be either `i386`, or `amd64`.

```
cpu          I386_CPU
cpu          I486_CPU
cpu          I586_CPU
cpu          I686_CPU
```

The above option specifies the type of CPU you have in your system. You may have multiple instances of the CPU line (i.e., you are not sure whether you should use `I586_CPU` or `I686_CPU`), however, for a custom kernel, it is best to specify only the CPU you have. If you are unsure of your CPU type, you can check the `/var/run/dmesg.boot` file to view your boot up messages.

```
ident          GENERIC
```

This is the identification of the kernel. You should change this to whatever you named your kernel, i.e. `MYKERNEL` if you have followed the instructions of the previous examples. The value you put in the `ident` string will print when you boot up the kernel, so it is useful to give the new kernel a different name if you want to keep it separate from your usual kernel (i.e. you want to build an experimental kernel).

```
maxusers      n
```

The `maxusers` option sets the size of a number of important system tables. This number is supposed to be roughly equal to the number of simultaneous users you expect to have on your machine.

(Recommended) The system will auto-tune this setting for you if you explicitly set it to `0`¹. If you want to manage it yourself you will want to set `maxusers` to at least 4, especially if you are using the X Window System or compiling software. The reason is that the most important table set by `maxusers` is the maximum number of processes, which is set to $20 + 16 * \text{maxusers}$, so if you set `maxusers` to 1, then you can only have 36 simultaneous processes, including the 18 or so that the system starts up at boot time, and the 15 or so you will probably create when you start the X Window System. Even a simple task like reading a manual page will start up nine processes to filter, decompress, and view it. Setting `maxusers` to 64 will allow you to have up to 1044 simultaneous processes, which should be enough for nearly all uses. If, however, you see the dreaded `proc table full` error when trying to start

another program, or are running a server with a large number of simultaneous users, you can always increase the number and rebuild.

Note: `maxusers` does *not* limit the number of users which can log into your machine. It simply sets various table sizes to reasonable values considering the maximum number of users you will likely have on your system and how many processes each of them will be running. One keyword which *does* limit the number of simultaneous *remote logins and X terminal windows* is `pseudo-device pty 16`.

```
# Floating point support - do not disable.
device      npx0      at nexus? port IO_NPX irq 13
```

`npx0` is the interface to the floating point math unit in DragonFly, which is either the hardware co-processor or the software math emulator. This is *not* optional.

```
# Pseudo devices - the number indicates how many units to allocate.
pseudo-device loop          # Network loopback
```

This is the generic loopback device for TCP/IP. If you telnet or FTP to `localhost` (a.k.a., `127.0.0.1`) it will come back at you through this device. This is *mandatory*.

Everything that follows is more or less optional. See the notes underneath or next to each option for more information.

```
#makeoptions      DEBUG=-g          #Build kernel with gdb(1) debug symbols
```

The normal build process of the DragonFly does not include debugging information when building the kernel and strips most symbols after the resulting kernel is linked, to save some space at the install location. If you are going to do tests of kernels in the `DEVELOPMENT` branch or develop changes of your own for the DragonFly kernel, you might want to uncomment this line. It will enable the use of the `-g` option which enables debugging information when passed to `gcc(1)`.

```
options          MATH_EMULATE      #Support for x87 emulation
```

This line allows the kernel to simulate a math co-processor if your computer does not have one (386 or 486SX). If you have a 486DX, or a 386 or 486SX (with a separate 387 or 487 chip), or higher (Pentium, Pentium II, etc.), you can comment this line out.

Note: The normal math co-processor emulation routines that come with DragonFly are *not* very accurate. If you do not have a math co-processor, and you need the best accuracy, it is recommended that you change this option to `GPL_MATH_EMULATE` to use the GNU math support, which is not included by default for licensing reasons.

```
options          INET              #InterNETworking
```

Networking support. Leave this in, even if you do not plan to be connected to a network. Most programs require at least loopback networking (i.e., making network connections within your PC), so this is essentially mandatory.

```
options          INET6             #IPv6 communications protocols
```

This enables the IPv6 communication protocols.


```
options      FFS          #Berkeley Fast Filesystem
options      FFS_ROOT     #FFS usable as root device [keep this!]
```

This is the basic hard drive Filesystem. Leave it in if you boot from the hard disk.

```
options      UFS_DIRHASH #Improve performance on big directories
```

This option includes functionality to speed up disk operations on large directories, at the expense of using additional memory. You would normally keep this for a large server, or interactive workstation, and remove it if you are using DragonFly on a smaller system where memory is at a premium and disk access speed is less important, such as a firewall.

```
options      SOFTUPDATES #Enable FFS Soft Updates support
```

This option enables Soft Updates in the kernel, this will help speed up write access on the disks. Even when this functionality is provided by the kernel, it must be turned on for specific disks. Review the output from `mount(8)` to see if Soft Updates is enabled for your system disks. If you do not see the `soft-updates` option then you will need to activate it using the `tunefs(8)` (for existing filesystems) or `newfs(8)` (for new filesystems) commands.

```
options      MFS          #Memory Filesystem
options      MD_ROOT     #MD is a potential root device
```

This is the memory-mapped filesystem. This is basically a RAM disk for fast storage of temporary files, useful if you have a lot of swap space that you want to take advantage of. A perfect place to mount an MFS partition is on the `/tmp` directory, since many programs store temporary data here. To mount an MFS RAM disk on `/tmp`, add the following line to `/etc/fstab`:

```
/dev/ad1s2b /tmp mfs rw 0 0
```

Now you simply need to either reboot, or run the command `mount /tmp`.

```
options      NFS          #Network Filesystem
options      NFS_ROOT     #NFS usable as root device, NFS required
```

The network Filesystem. Unless you plan to mount partitions from a UNIX file server over TCP/IP, you can comment these out.

```
options      MSDOSFS     #MSDOS Filesystem
```

The MS-DOS Filesystem. Unless you plan to mount a DOS formatted hard drive partition at boot time, you can safely comment this out. It will be automatically loaded the first time you mount a DOS partition, as described above. Also, the excellent **mtools** software (in `pkgsrc`) allows you to access DOS floppies without having to mount and unmount them (and does not require `MSDOSFS` at all).

```
options      CD9660      #ISO 9660 Filesystem
options      CD9660_ROOT #CD-ROM usable as root, CD9660 required
```

The ISO 9660 Filesystem for CDRoms. Comment it out if you do not have a CDRom drive or only mount data CDs occasionally (since it will be dynamically loaded the first time you mount a data CD). Audio CDs do not need this Filesystem.

```
options      PROCFS      #Process filesystem
```

The process filesystem. This is a “pretend” filesystem mounted on `/proc` which allows programs like `ps(1)` to give you more information on what processes are running.

```
options          COMPAT_43      #Compatible with BSD 4.3 [KEEP THIS!]
```

Compatibility with 4.3BSD. Leave this in; some programs will act strangely if you comment this out.

```
options          SCSI_DELAY=15000  #Delay (in ms) before probing SCSI
```

This causes the kernel to pause for 15 seconds before probing each SCSI device in your system. If you only have IDE hard drives, you can ignore this, otherwise you will probably want to lower this number, perhaps to five seconds (5000 ms), to speed up booting. Of course, if you do this, and DragonFly has trouble recognizing your SCSI devices, you will have to raise it back up.

```
options          UCONSOLE          #Allow users to grab the console
```

Allow users to grab the console, which is useful for X users. For example, you can create a console **xterm** by typing `xterm -C`, which will display any `write(1)`, `talk(1)`, and any other messages you receive, as well as any console messages sent by the kernel.

```
options          USERCONFIG        #boot -c editor
```

This option allows you to boot the configuration editor from the boot menu.

```
options          VISUAL_USERCONFIG  #visual boot -c editor
```

This option allows you to boot the visual configuration editor from the boot menu.

```
options          KTRACE            #ktrace(1) support
```

This enables kernel process tracing, which is useful in debugging.

```
options          SYSVSHM           #SYSV-style shared memory
```

This option provides for System V shared memory. The most common use of this is the XSHM extension in X, which many graphics-intensive programs will automatically take advantage of for extra speed. If you use X, you will definitely want to include this.

```
options          SYSVSEM           #SYSV-style semaphores
```

Support for System V semaphores. Less commonly used but only adds a few hundred bytes to the kernel.

```
options          SYSVMSG           #SYSV-style message queues
```

Support for System V messages. Again, only adds a few hundred bytes to the kernel.

Note: The `ipcs(1)` command will list any processes using each of these System V facilities.

```
options          P1003_1B          #Posix P1003_1B real-time extensions
```

```
options          _KPOSIX_PRIORITY_SCHEDULING
```

Real-time extensions added in the 1993 POSIX®. Certain applications in the ports collection use these (such as **StarOffice**).

```
options ICMP_BANDLIM #Rate limit bad replies
```

This option enables ICMP error response bandwidth limiting. You typically want this option as it will help protect the machine from denial of service packet attacks.

```
# To make an SMP kernel, the next two are needed
#options SMP # Symmetric MultiProcessor Kernel
#options APIC_IO # Symmetric (APIC) I/O
```

The above are both required for SMP support.

```
device isa
```

All PCs supported by DragonFly have one of these. Do not remove, even if you have no ISA slots. If you have an IBM PS/2 (Micro Channel Architecture), DragonFly provides some limited support at this time. For more information about the MCA support, see `/usr/src/sys/i386/conf/LINT`.

```
device eisa
```

Include this if you have an EISA motherboard. This enables auto-detection and configuration support for all devices on the EISA bus.

```
device pci
```

Include this if you have a PCI motherboard. This enables auto-detection of PCI cards and gatewaying from the PCI to ISA bus.

```
device agp
```

Include this if you have an AGP card in the system. This will enable support for AGP, and AGP GART for boards which have these features.

```
# Floppy drives
device fdc0 at isa? port IO_FD1 irq 6 drq 2
device fd0 at fdc0 drive 0
device fd1 at fdc0 drive 1
```

This is the floppy drive controller. `fd0` is the A: floppy drive, and `fd1` is the B: drive.

```
device ata
```

This driver supports all ATA and ATAPI devices. You only need one `device ata` line for the kernel to detect all PCI ATA/ATAPI devices on modern machines.

```
device atadisk # ATA disk drives
```

This is needed along with `device ata` for ATA disk drives.

```
device atapid # ATAPI CDROM drives
```

This is needed along with `device ata` for ATAPI CDROM drives.

```
device      atapifd          # ATAPI floppy drives
```

This is needed along with device `ata` for ATAPI floppy drives.

```
device      atapist         # ATAPI tape drives
```

This is needed along with device `ata` for ATAPI tape drives.

```
options     ATA_STATIC_ID   #Static device numbering
```

This makes the controller number static (like the old driver) or else the device numbers are dynamically allocated.

```
# ATA and ATAPI devices
```

```
device      ata0           at isa? port IO_WD1 irq 14
```

```
device      ata1           at isa? port IO_WD2 irq 15
```

Use the above for older, non-PCI systems.

```
# SCSI Controllers
```

```
device      ahb            # EISA AHA1742 family
```

```
device      ahc            # AHA2940 and onboard AIC7xxx devices
```

```
device      amd            # AMD 53C974 (Teckram DC-390(T))
```

```
device      dpt            # DPT Smartcache - See LINT for options!
```

```
device      isp            # Qlogic family
```

```
device      ncr            # NCR/Symbios Logic
```

```
device      sym            # NCR/Symbios Logic (newer chipsets)
```

```
device      adv0           at isa?
```

```
device      adw
```

```
device      bt0           at isa?
```

```
device      aha0           at isa?
```

```
device      aic0           at isa?
```

SCSI controllers. Comment out any you do not have in your system. If you have an IDE only system, you can remove these altogether.

```
# SCSI peripherals
```

```
device      scbus         # SCSI bus (required)
```

```
device      da            # Direct Access (disks)
```

```
device      sa            # Sequential Access (tape etc)
```

```
device      cd            # CD
```

```
device      pass          # Passthrough device (direct SCSI
access)
```

SCSI peripherals. Again, comment out any you do not have, or if you have only IDE hardware, you can remove them completely.

Note: The USB `umass(4)` driver (and a few other drivers) use the SCSI subsystem even though they are not real SCSI devices. Therefore make sure not to remove SCSI support, if any such drivers are included in the kernel configuration.

```
# RAID controllers
```

```
device      ida          # Compaq Smart RAID
device      amr          # AMI MegaRAID
device      mlx          # Mylex DAC960 family
```

Supported RAID controllers. If you do not have any of these, you can comment them out or remove them.

```
# atkbd0 controls both the keyboard and the PS/2 mouse
device      atkbd0      at isa? port IO_KBD
```

The keyboard controller (`atkbd0`) provides I/O services for the AT keyboard and PS/2 style pointing devices. This controller is required by the keyboard driver (`atkbd`) and the PS/2 pointing device driver (`psm`).

```
device      atkbd0      at atkbd? irq 1
```

The `atkbd` driver, together with `atkbd0` controller, provides access to the AT 84 keyboard or the AT enhanced keyboard which is connected to the AT keyboard controller.

```
device      psm0        at atkbd? irq 12
```

Use this device if your mouse plugs into the PS/2 mouse port.

```
device      vga0        at isa?
```

The video card driver.

```
# splash screen/screen saver
pseudo-device splash
```

Splash screen at start up! Screen savers require this too.

```
# syscons is the default console driver, resembling an SCO console
device      sc0         at isa?
```

`sc0` is the default console driver, which resembles a SCO console. Since most full-screen programs access the console through a terminal database library like `termcap`, it should not matter whether you use this or `vt0`, the VT220 compatible console driver. When you log in, set your `TERM` variable to `scoansi` if full-screen programs have trouble running under this console.

```
# Enable this and PCVT_FREEBSD for pcvt vt220 compatible console driver
#device      vt0         at isa?
#options     XSERVER          # support for X server on a vt console
#options     FAT_CURSOR       # start with block cursor
# If you have a ThinkPAD, uncomment this along with the rest of the PCVT lines
#options     PCVT_SCANSET=2   # IBM keyboards are non-std
```

This is a VT220-compatible console driver, backward compatible to VT100/102. It works well on some laptops which have hardware incompatibilities with `sc0`. Also set your `TERM` variable to `vt100` or `vt220` when you log in. This driver might also prove useful when connecting to a large number of different machines over the network, where `termcap` or `terminfo` entries for the `sc0` device are often not available — `vt100` should be available on virtually any platform.

```
# Power management support (see LINT for more options)
device      apm0        at nexus? disable flags 0x20 # Advanced Power Management
```

Advanced Power Management support. Useful for laptops.

```
# PCCARD (PCMCIA) support
device      card
device      pcic0    at isa? irq 10 port 0x3e0 iomem 0xd0000
device      pcic1    at isa? irq 11 port 0x3e2 iomem 0xd4000 disable
```

PCMCIA support. You want this if you are using a laptop.

```
# Serial (COM) ports
device      sio0      at isa? port IO_COM1 flags 0x10 irq 4
device      sio1      at isa? port IO_COM2 irq 3
device      sio2      at isa? disable port IO_COM3 irq 5
device      sio3      at isa? disable port IO_COM4 irq 9
```

These are the four serial ports referred to as COM1 through COM4 in the MS-DOS/Windows world.

Note: If you have an internal modem on COM4 and a serial port at COM2, you will have to change the IRQ of the modem to 2 (for obscure technical reasons, IRQ2 = IRQ 9) in order to access it from DragonFly. If you have a multiport serial card, check the manual page for `sio(4)` for more information on the proper values for these lines. Some video cards (notably those based on S3 chips) use IO addresses in the form of `0x*2e8`, and since many cheap serial cards do not fully decode the 16-bit IO address space, they clash with these cards making the COM4 port practically unavailable.

Each serial port is required to have a unique IRQ (unless you are using one of the multiport cards where shared interrupts are supported), so the default IRQs for COM3 and COM4 cannot be used.

```
# Parallel port
device      ppc0      at isa? irq 7
```

This is the ISA-bus parallel port interface.

```
device      ppbus     # Parallel port bus (required)
```

Provides support for the parallel port bus.

```
device      lpt       # Printer
```

Support for parallel port printers.

Note: All three of the above are required to enable parallel printer support.

```
device      plip      # TCP/IP over parallel
```

This is the driver for the parallel network interface.

```
device      ppi       # Parallel port interface device
```

The general-purpose I/O (“geek port”) + IEEE1284 I/O.

```
#device     vpo       # Requires scbus and da
```

This is for an Iomega Zip drive. It requires `scbus` and `da` support. Best performance is achieved with ports in EPP 1.9 mode.

```
# PCI Ethernet NICs.
device      de          # DEC/Intel DC21x4x ("Tulip")
device      fxp         # Intel EtherExpress PRO/100B (82557, 82558)
device      tx          # SMC 9432TX (83c170 "EPIC")
device      vx          # 3Com 3c590, 3c595 ("Vortex")
device      wx          # Intel Gigabit Ethernet Card ("Wiseman")
```

Various PCI network card drivers. Comment out or remove any of these not present in your system.

```
# PCI Ethernet NICs that use the common MII bus controller code.
device      miibus      # MII bus support
```

MII bus support is required for some PCI 10/100 Ethernet NICs, namely those which use MII-compliant transceivers or implement transceiver control interfaces that operate like an MII. Adding `device miibus` to the kernel config pulls in support for the generic `miibus` API and all of the PHY drivers, including a generic one for PHYs that are not specifically handled by an individual driver.

```
device      dc          # DEC/Intel 21143 and various workalikes
device      rl          # RealTek 8129/8139
device      sf          # Adaptec AIC-6915 ("Starfire")
device      sis         # Silicon Integrated Systems SiS 900/SiS 7016
device      ste         # Sundance ST201 (D-Link DFE-550TX)
device      tl          # Texas Instruments ThunderLAN
device      vr          # VIA Rhine, Rhine II
device      wb          # Winbond W89C840F
device      xl          # 3Com 3c90x ("Boomerang", "Cyclone")
```

Drivers that use the MII bus controller code.

```
# ISA Ethernet NICs.
device      ed0         at isa? port 0x280 irq 10 iomem 0xd8000
device      ex
device      ep
# WaveLAN/IEEE 802.11 wireless NICs. Note: the WaveLAN/IEEE really
# exists only as a PCMCIA device, so there is no ISA attachment needed
# and resources will always be dynamically assigned by the pccard code.
device      wi
# Aironet 4500/4800 802.11 wireless NICs. Note: the declaration below will
# work for PCMCIA and PCI cards, as well as ISA cards set to ISA PnP
# mode (the factory default). If you set the switches on your ISA
# card for a manually chosen I/O address and IRQ, you must specify
# those parameters here.
device      an
# The probe order of these is presently determined by i386/isa/isa_compat.c.
device      ie0         at isa? port 0x300 irq 10 iomem 0xd0000
device      fe0         at isa? port 0x300
device      le0         at isa? port 0x300 irq 5 iomem 0xd0000
device      lnc0        at isa? port 0x280 irq 10 drq 0
device      cs0         at isa? port 0x300
device      sn0         at isa? port 0x300 irq 10
```

```
# requires PCCARD (PCMCIA) support to be activated
#device      xe0      at isa?
```

ISA Ethernet drivers. See `/usr/src/sys/i386/conf/LINT` for which cards are supported by which driver.

```
pseudo-device  ether          # Ethernet support
```

`ether` is only needed if you have an Ethernet card. It includes generic Ethernet protocol code.

```
pseudo-device  sl          1      # Kernel SLIP
```

`sl` is for SLIP support. This has been almost entirely supplanted by PPP, which is easier to set up, better suited for modem-to-modem connection, and more powerful. The *number* after `sl` specifies how many simultaneous SLIP sessions to support.

```
pseudo-device  ppp          1      # Kernel PPP
```

This is for kernel PPP support for dial-up connections. There is also a version of PPP implemented as a userland application that uses `tun` and offers more flexibility and features such as demand dialing. The *number* after `ppp` specifies how many simultaneous PPP connections to support. .

```
device  tun          # Packet tunnel.
```

This is used by the userland PPP software. A *number* after `tun` specifies the number of simultaneous PPP sessions to support. See the PPP section of this book for more information.

```
pseudo-device  pty          # Pseudo-ttys (telnet etc)
```

This is a “pseudo-terminal” or simulated login port. It is used by incoming `telnet` and `rlogin` sessions, **xterm**, and some other applications such as **Emacs**. The *number* after `pty` indicates the number of `ptys` to create. If you need more than the default of 16 simultaneous **xterm** windows and/or remote logins, be sure to increase this number accordingly, up to a maximum of 256.

```
pseudo-device  md          # Memory “disks”
```

Memory disk pseudo-devices.

```
pseudo-device  gif          # IPv6 and IPv4 tunneling
```

This implements IPv6 over IPv4 tunneling, IPv4 over IPv6 tunneling, IPv4 over IPv4 tunneling, and IPv6 over IPv6 tunneling.

```
pseudo-device  faith       # IPv6-to-IPv4 relaying (translation)
```

This pseudo-device captures packets that are sent to it and diverts them to the IPv4/IPv6 translation daemon.

```
# The 'bpf' device enables the Berkeley Packet Filter.
# Be aware of the administrative consequences of enabling this!
pseudo-device  bpf          # Berkeley packet filter
```

This is the Berkeley Packet Filter. This pseudo-device allows network interfaces to be placed in promiscuous mode, capturing every packet on a broadcast network (e.g., an Ethernet). These packets can be captured to disk and or examined with the `tcpdump(1)` program.

Note: The `bpf(4)` device is also used by `dhclient(8)` to obtain the IP address of the default router (gateway) and so on. If you use DHCP, leave this uncommented.

```
# USB support
#device      uhci          # UHCI PCI->USB interface
#device      ohci          # OHCI PCI->USB interface
#device      usb           # USB Bus (required)
#device      ugen          # Generic
#device      uhid          # "Human Interface Devices"
#device      ukbd          # Keyboard
#device      ulpt          # Printer
#device      umass         # Disks/Mass storage - Requires scbus and da
#device      ums           # Mouse
# USB Ethernet, requires mii
#device      aue           # ADMtek USB ethernet
#device      cue           # CATC USB ethernet
#device      kue           # Kawasaki LSI USB ethernet
```

Support for various USB devices.

For more information and additional devices supported by DragonFly, see `/usr/src/sys/i386/conf/LINT`.

9.5 Making Device Nodes

Almost every device in the kernel has a corresponding “node” entry in the `/dev` directory. These nodes look like regular files, but are actually special entries into the kernel which programs use to access the device. The shell script `/dev/MAKEDEV`, which is executed when you first install the operating system, creates nearly all of the device nodes supported. However, it does not create *all* of them, so when you add support for a new device, it pays to make sure that the appropriate entries are in this directory, and if not, add them. Here is a simple example:

Suppose you add the IDE CD-ROM support to the kernel. The line to add is:

```
device acd0
```

This means that you should look for some entries that start with `acd0` in the `/dev` directory, possibly followed by a letter, such as `c`, or preceded by the letter `r`, which means a “raw” device. It turns out that those files are not there, so you must change to the `/dev` directory and type:

```
# sh MAKEDEV acd0
```

When this script finishes, you will find that there are now `acd0c` and `racd0c` entries in `/dev` so you know that it executed correctly.

For sound cards, the following command creates the appropriate entries:

```
# sh MAKEDEV snd0
```

Note: When creating device nodes for devices such as sound cards, if other people have access to your machine, it may be desirable to protect the devices from outside access by adding them to the `/etc/fstab` file. See `fstab(5)` for more information.

Follow this simple procedure for any other non-GENERIC devices which do not have entries.

Note: All SCSI controllers use the same set of `/dev` entries, so you do not need to create these. Also, network cards and SLIP/PPP pseudo-devices do not have entries in `/dev` at all, so you do not have to worry about these either.

9.6 If Something Goes Wrong

There are five categories of trouble that can occur when building a custom kernel. They are:

`config` fails:

If the `config(8)` command fails when you give it your kernel description, you have probably made a simple error somewhere. Fortunately, `config(8)` will print the line number that it had trouble with, so you can quickly skip to it with `vi`. For example, if you see:

```
config: line 17: syntax error
```

You can skip to the problem in `vi` by typing `17G` in command mode. Make sure the keyword is typed correctly, by comparing it to the `GENERIC` kernel or another reference.

`make` fails:

If the `make` command fails, it usually signals an error in your kernel description, but not severe enough for `config(8)` to catch it. Again, look over your configuration, and if you still cannot resolve the problem, send mail to the DragonFly Bugs mailing list (<http://leaf.dragonflybsd.org/mailarchive/>) with your kernel configuration, and it should be diagnosed very quickly.

Installing the new kernel fails:

If the kernel compiled fine, but failed to install (the `make install` or `make installkernel` command failed), the first thing to check is if your system is running at `securelevel 1` or higher (see `init(8)`). The kernel installation tries to remove the immutable flag from your kernel and set the immutable flag on the new one. Since `securelevel 1` or higher prevents unsetting the immutable flag for any files on the system, the kernel installation needs to be performed at `securelevel 0` or lower.

The kernel does not boot:

If your new kernel does not boot, or fails to recognize your devices, do not panic! Fortunately, DragonFly has an excellent mechanism for recovering from incompatible kernels. Simply choose the kernel you want to boot from at the DragonFly boot loader. You can access this when the system counts down from 10. Hit any key except for the **Enter** key, type `unload` and then type `boot kernel.old`, or the filename of any other kernel that will boot properly. When reconfiguring a kernel, it is always a good idea to keep a kernel that is known to work on hand.

After booting with a good kernel you can check over your configuration file and try to build it again. One helpful resource is the `/var/log/messages` file which records, among other things, all of the kernel messages from every successful boot. Also, the `dmesg(8)` command will print the kernel messages from the current boot.

Note: If you are having trouble building a kernel, make sure to keep a `GENERIC`, or some other kernel that is known to work on hand as a different name that will not get erased on the next build. You cannot rely on `kernel.old` because when installing a new kernel, `kernel.old` is overwritten with the last installed kernel which may be non-functional. Also, as soon as possible, move the working kernel to the proper `kernel` location or commands such as `ps(1)` will not work properly. The proper command to “unlock” the kernel file that `make` installs (in order to move another kernel back permanently) is:

```
# chflags noschg /kernel
```

If you find you cannot do this, you are probably running at a `securelevel(8)` greater than zero. Edit `kern_securelevel` in `/etc/rc.conf` and set it to `-1`, then reboot. You can change it back to its previous setting when you are happy with your new kernel.

And, if you want to “lock” your new kernel into place, or any file for that matter, so that it cannot be moved or tampered with:

```
# chflags schg /kernel
```

The kernel works, but `ps(1)` does not work any more:

If you have installed a different version of the kernel from the one that the system utilities have been built with, many system-status commands like `ps(1)` and `vmstat(8)` will not work any more. You must recompile the `libkvm` library as well as these utilities. This is one reason it is not normally a good idea to use a different version of the kernel from the rest of the operating system.

Notes

1. The auto-tuning algorithm sets `maxuser` equal to the amount of memory in the system, with a minimum of 32, and a maximum of 384.

Chapter 10 Security

Much of this chapter has been taken from the security(7) manual page by Matthew Dillon.

10.1 Synopsis

This chapter will provide a basic introduction to system security concepts, some general good rules of thumb, and some advanced topics under DragonFly. A lot of the topics covered here can be applied to system and Internet security in general as well. The Internet is no longer a “friendly” place in which everyone wants to be your kind neighbor. Securing your system is imperative to protect your data, intellectual property, time, and much more from the hands of hackers and the like.

DragonFly provides an array of utilities and mechanisms to ensure the integrity and security of your system and network.

After reading this chapter, you will know:

- Basic system security concepts, in respect to DragonFly.
- About the various crypt mechanisms available in DragonFly, such as DES and MD5.
- How to set up one-time password authentication.
- How to set up **KerberosIV**.
- How to set up **Kerberos5**.
- How to create firewalls using IPFW.
- How to configure IPsec and create a VPN between DragonFly/Windows machines.
- How to configure and use **OpenSSH**, DragonFly’s SSH implementation.

Before reading this chapter, you should:

- Understand basic DragonFly and Internet concepts.

10.2 Introduction

Security is a function that begins and ends with the system administrator. While all BSD UNIX multi-user systems have some inherent security, the job of building and maintaining additional security mechanisms to keep those users “honest” is probably one of the single largest undertakings of the sysadmin. Machines are only as secure as you make them, and security concerns are ever competing with the human necessity for convenience. UNIX systems, in general, are capable of running a huge number of simultaneous processes and many of these processes operate as servers — meaning that external entities can connect and talk to them. As yesterday’s mini-computers and mainframes become today’s desktops, and as computers become networked and internetworked, security becomes an even bigger issue.

Security is best implemented through a layered “onion” approach. In a nutshell, what you want to do is to create as many layers of security as are convenient and then carefully monitor the system for intrusions. You do not want to overbuild your security or you will interfere with the detection side, and detection is one of the single most important aspects of any security mechanism. For example, it makes little sense to set the `schg` flags (see `chflags(1)`) on every

system binary because while this may temporarily protect the binaries, it prevents an attacker who has broken in from making an easily detectable change that may result in your security mechanisms not detecting the attacker at all.

System security also pertains to dealing with various forms of attack, including attacks that attempt to crash, or otherwise make a system unusable, but do not attempt to compromise the `root` account (“break root”). Security concerns can be split up into several categories:

1. Denial of service attacks.
2. User account compromises.
3. Root compromise through accessible servers.
4. Root compromise via user accounts.
5. Backdoor creation.

A denial of service attack is an action that deprives the machine of needed resources. Typically, DoS attacks are brute-force mechanisms that attempt to crash or otherwise make a machine unusable by overwhelming its servers or network stack. Some DoS attacks try to take advantage of bugs in the networking stack to crash a machine with a single packet. The latter can only be fixed by applying a bug fix to the kernel. Attacks on servers can often be fixed by properly specifying options to limit the load the servers incur on the system under adverse conditions. Brute-force network attacks are harder to deal with. A spoofed-packet attack, for example, is nearly impossible to stop, short of cutting your system off from the Internet. It may not be able to take your machine down, but it can saturate your Internet connection.

A user account compromise is even more common than a DoS attack. Many sysadmins still run standard **telnetd**, **rlogind**, **rshd**, and **ftpd** servers on their machines. These servers, by default, do not operate over encrypted connections. The result is that if you have any moderate-sized user base, one or more of your users logging into your system from a remote location (which is the most common and convenient way to login to a system) will have his or her password sniffed. The attentive system admin will analyze his remote access logs looking for suspicious source addresses even for successful logins.

One must always assume that once an attacker has access to a user account, the attacker can break `root`. However, the reality is that in a well secured and maintained system, access to a user account does not necessarily give the attacker access to `root`. The distinction is important because without access to `root` the attacker cannot generally hide his tracks and may, at best, be able to do nothing more than mess with the user’s files, or crash the machine. User account compromises are very common because users tend not to take the precautions that sysadmins take.

System administrators must keep in mind that there are potentially many ways to break `root` on a machine. The attacker may know the `root` password, the attacker may find a bug in a root-run server and be able to break `root` over a network connection to that server, or the attacker may know of a bug in a `suid-root` program that allows the attacker to break `root` once he has broken into a user’s account. If an attacker has found a way to break `root` on a machine, the attacker may not have a need to install a backdoor. Many of the `root` holes found and closed to date involve a considerable amount of work by the attacker to cleanup after himself, so most attackers install backdoors. A backdoor provides the attacker with a way to easily regain `root` access to the system, but it also gives the smart system administrator a convenient way to detect the intrusion. Making it impossible for an attacker to install a backdoor may actually be detrimental to your security, because it will not close off the hole the attacker found to break in the first place.

Security remedies should always be implemented with a multi-layered “onion peel” approach and can be categorized as follows:

1. Securing `root` and staff accounts.

2. Securing `root` — root-run servers and `suid/sgid` binaries.
3. Securing user accounts.
4. Securing the password file.
5. Securing the kernel core, raw devices, and filesystems.
6. Quick detection of inappropriate changes made to the system.
7. Paranoia.

The next section of this chapter will cover the above bullet items in greater depth.

10.3 Securing DragonFly

Command vs. Protocol: Throughout this document, we will use **bold** text to refer to a command or application. This is used for instances such as `ssh`, since it is a protocol as well as command.

The sections that follow will cover the methods of securing your DragonFly system that were mentioned in the last section of this chapter.

10.3.1 Securing the `root` Account and Staff Accounts

First off, do not bother securing staff accounts if you have not secured the `root` account. Most systems have a password assigned to the `root` account. The first thing you do is assume that the password is *always* compromised. This does not mean that you should remove the password. The password is almost always necessary for console access to the machine. What it does mean is that you should not make it possible to use the password outside of the console or possibly even with the `su(1)` command. For example, make sure that your `pty`'s are specified as being insecure in the `/etc/ttys` file so that direct `root` logins via `telnet` or `rlogin` are disallowed. If using other login services such as **sshd**, make sure that direct `root` logins are disabled there as well. You can do this by editing your `/etc/ssh/sshd_config` file, and making sure that `PermitRootLogin` is set to `NO`. Consider every access method — services such as FTP often fall through the cracks. Direct `root` logins should only be allowed via the system console.

Of course, as a sysadmin you have to be able to get to `root`, so we open up a few holes. But we make sure these holes require additional password verification to operate. One way to make `root` accessible is to add appropriate staff accounts to the `wheel` group (in `/etc/group`). The staff members placed in the `wheel` group are allowed to `su` to `root`. You should never give staff members native `wheel` access by putting them in the `wheel` group in their password entry. Staff accounts should be placed in a `staff` group, and then added to the `wheel` group via the `/etc/group` file. Only those staff members who actually need to have `root` access should be placed in the `wheel` group. It is also possible, when using an authentication method such as Kerberos, to use Kerberos' `.k5login` file in the `root` account to allow a `ksu(1)` to `root` without having to place anyone at all in the `wheel` group. This may be the better solution since the `wheel` mechanism still allows an intruder to break `root` if the intruder has gotten hold of your password file and can break into a staff account. While having the `wheel` mechanism is better than having nothing at all, it is not necessarily the safest option.

An indirect way to secure staff accounts, and ultimately `root` access is to use an alternative login access method and do what is known as “starring” out the encrypted password for the staff accounts. Using the `vipw(8)` command, one

can replace each instance of an encrypted password with a single “*” character. This command will update the `/etc/master.passwd` file and user/password database to disable password-authenticated logins.

A staff account entry such as:

```
foobar:R9DT/Fa1/LV9U:1000:1000::0:0:Foo Bar:/home/foobar:/usr/local/bin/tcsh
```

Should be changed to this:

```
foobar*:1000:1000::0:0:Foo Bar:/home/foobar:/usr/local/bin/tcsh
```

This change will prevent normal logins from occurring, since the encrypted password will never match “*”. With this done, staff members must use another mechanism to authenticate themselves such as `kerberos(1)` or `ssh(1)` using a public/private key pair. When using something like Kerberos, one generally must secure the machines which run the Kerberos servers and your desktop workstation. When using a public/private key pair with `ssh`, one must generally secure the machine used to login *from* (typically one’s workstation). An additional layer of protection can be added to the key pair by password protecting the key pair when creating it with `ssh-keygen(1)`. Being able to “star” out the passwords for staff accounts also guarantees that staff members can only login through secure access methods that you have set up. This forces all staff members to use secure, encrypted connections for all of their sessions, which closes an important hole used by many intruders: sniffing the network from an unrelated, less secure machine.

The more indirect security mechanisms also assume that you are logging in from a more restrictive server to a less restrictive server. For example, if your main box is running all sorts of servers, your workstation should not be running any. In order for your workstation to be reasonably secure you should run as few servers as possible, up to and including no servers at all, and you should run a password-protected screen blanker. Of course, given physical access to a workstation an attacker can break any sort of security you put on it. This is definitely a problem that you should consider, but you should also consider the fact that the vast majority of break-ins occur remotely, over a network, from people who do not have physical access to your workstation or servers.

Using something like Kerberos also gives you the ability to disable or change the password for a staff account in one place, and have it immediately affect all the machines on which the staff member may have an account. If a staff member’s account gets compromised, the ability to instantly change his password on all machines should not be underrated. With discrete passwords, changing a password on N machines can be a mess. You can also impose re-passwording restrictions with Kerberos: not only can a Kerberos ticket be made to timeout after a while, but the Kerberos system can require that the user choose a new password after a certain period of time (say, once a month).

10.3.2 Securing Root-run Servers and SUID/SGID Binaries

The prudent sysadmin only runs the servers he needs to, no more, no less. Be aware that third party servers are often the most bug-prone. For example, running an old version of **imapd** or **popper** is like giving a universal `root` ticket out to the entire world. Never run a server that you have not checked out carefully. Many servers do not need to be run as `root`. For example, the **ntalk**, **comsat**, and **finger** daemons can be run in special user *sandboxes*. A sandbox is not perfect, unless you go through a large amount of trouble, but the onion approach to security still stands: If someone is able to break in through a server running in a sandbox, they still have to break out of the sandbox. The more layers the attacker must break through, the lower the likelihood of his success. Root holes have historically been found in virtually every server ever run as `root`, including basic system servers. If you are running a machine through which people only login via **sshd** and never login via **telnetd** or **rshd** or **rlogind**, then turn off those services!

DragonFly now defaults to running **ntalkd**, **comsat**, and **finger** in a sandbox. Another program which may be a candidate for running in a sandbox is `named(8)`. `/etc/defaults/rc.conf` includes the arguments necessary to run **named** in a sandbox in a commented-out form. Depending on whether you are installing a new system or

upgrading an existing system, the special user accounts used by these sandboxes may not be installed. The prudent sysadmin would research and implement sandboxes for servers whenever possible.

There are a number of other servers that typically do not run in sandboxes: **sendmail**, **popper**, **imapd**, **ftpd**, and others. There are alternatives to some of these, but installing them may require more work than you are willing to perform (the convenience factor strikes again). You may have to run these servers as `root` and rely on other mechanisms to detect break-ins that might occur through them.

The other big potential `root` holes in a system are the `suid-root` and `sgid` binaries installed on the system. Most of these binaries, such as **rlogin**, reside in `/bin`, `/sbin`, `/usr/bin`, or `/usr/sbin`. While nothing is 100% safe, the system-default `suid` and `sgid` binaries can be considered reasonably safe. Still, `root` holes are occasionally found in these binaries. A `root` hole was found in `xlib` in 1998 that made **xterm** (which is typically `suid`) vulnerable. It is better to be safe than sorry and the prudent sysadmin will restrict `suid` binaries, that only staff should run, to a special group that only staff can access, and get rid of (`chmod 000`) any `suid` binaries that nobody uses. A server with no display generally does not need an **xterm** binary. `Sgid` binaries can be almost as dangerous. If an intruder can break an `sgid-kmem` binary, the intruder might be able to read `/dev/kmem` and thus read the encrypted password file, potentially compromising any passworded account. Alternatively an intruder who breaks `group kmem` can monitor keystrokes sent through `pty`'s, including `pty`'s used by users who login through secure methods. An intruder that breaks the `tty` group can write to almost any user's `tty`. If a user is running a terminal program or emulator with a keyboard-simulation feature, the intruder can potentially generate a data stream that causes the user's terminal to echo a command, which is then run as that user.

10.3.3 Securing User Accounts

User accounts are usually the most difficult to secure. While you can impose Draconian access restrictions on your staff and “star” out their passwords, you may not be able to do so with any general user accounts you might have. If you do have sufficient control, then you may win out and be able to secure the user accounts properly. If not, you simply have to be more vigilant in your monitoring of those accounts. Use of `ssh` and Kerberos for user accounts is more problematic, due to the extra administration and technical support required, but still a very good solution compared to a crypted password file.

10.3.4 Securing the Password File

The only sure fire way is to * out as many passwords as you can and use `ssh` or Kerberos for access to those accounts. Even though the encrypted password file (`/etc/spwd.db`) can only be read by `root`, it may be possible for an intruder to obtain read access to that file even if the attacker cannot obtain `root-write` access.

Your security scripts should always check for and report changes to the password file (see the Checking file integrity section below).

10.3.5 Securing the Kernel Core, Raw Devices, and Filesystems

If an attacker breaks `root` he can do just about anything, but there are certain conveniences. For example, most modern kernels have a packet sniffing device driver built in. Under DragonFly it is called the `bpf` device. An intruder will commonly attempt to run a packet sniffer on a compromised machine. You do not need to give the intruder the capability and most systems do not have the need for the `bpf` device compiled in.

But even if you turn off the `bpf` device, you still have `/dev/mem` and `/dev/kmem` to worry about. For that matter, the intruder can still write to raw disk devices. Also, there is another kernel feature called the module loader, `kldload(8)`. An enterprising intruder can use a KLD module to install his own `bpf` device, or other sniffing device, on a running kernel. To avoid these problems you have to run the kernel at a higher secure level, at least `securelevel 1`. The `securelevel` can be set with a `sysctl` on the `kern.securelevel` variable. Once you have set the `securelevel` to 1, write access to raw devices will be denied and special `chflags` flags, such as `schg`, will be enforced. You must also ensure that the `schg` flag is set on critical startup binaries, directories, and script files — everything that gets run up to the point where the `securelevel` is set. This might be overdoing it, and upgrading the system is much more difficult when you operate at a higher secure level. You may compromise and run the system at a higher secure level but not set the `schg` flag for every system file and directory under the sun. Another possibility is to simply mount `/` and `/usr` read-only. It should be noted that being too Draconian in what you attempt to protect may prevent the all-important detection of an intrusion.

10.3.6 Checking File Integrity: Binaries, Configuration Files, Etc.

When it comes right down to it, you can only protect your core system configuration and control files so much before the convenience factor rears its ugly head. For example, using `chflags` to set the `schg` bit on most of the files in `/` and `/usr` is probably counterproductive, because while it may protect the files, it also closes a detection window. The last layer of your security onion is perhaps the most important — detection. The rest of your security is pretty much useless (or, worse, presents you with a false sense of safety) if you cannot detect potential incursions. Half the job of the onion is to slow down the attacker, rather than stop him, in order to give the detection side of the equation a chance to catch him in the act.

The best way to detect an incursion is to look for modified, missing, or unexpected files. The best way to look for modified files is from another (often centralized) limited-access system. Writing your security scripts on the extra-secure limited-access system makes them mostly invisible to potential attackers, and this is important. In order to take maximum advantage you generally have to give the limited-access box significant access to the other machines in the business, usually either by doing a read-only NFS export of the other machines to the limited-access box, or by setting up `ssh` key-pairs to allow the limited-access box to `ssh` to the other machines. Except for its network traffic, NFS is the least visible method — allowing you to monitor the filesystems on each client box virtually undetected. If your limited-access server is connected to the client boxes through a switch, the NFS method is often the better choice. If your limited-access server is connected to the client boxes through a hub, or through several layers of routing, the NFS method may be too insecure (network-wise) and using `ssh` may be the better choice even with the audit-trail tracks that `ssh` lays.

Once you give a limited-access box, at least read access to the client systems it is supposed to monitor, you must write scripts to do the actual monitoring. Given an NFS mount, you can write scripts out of simple system utilities such as `find(1)` and `md5(1)`. It is best to physically `md5` the client-box files at least once a day, and to test control files such as those found in `/etc` and `/usr/local/etc` even more often. When mismatches are found, relative to the base `md5` information the limited-access machine knows is valid, it should scream at a `sysadmin` to go check it out. A good security script will also check for inappropriate `suid` binaries and for new or deleted files on system partitions such as `/` and `/usr`.

When using `ssh` rather than NFS, writing the security script is much more difficult. You essentially have to `scp` the scripts to the client box in order to run them, making them visible, and for safety you also need to `scp` the binaries (such as `find`) that those scripts use. The `ssh` client on the client box may already be compromised. All in all, using `ssh` may be necessary when running over insecure links, but it is also a lot harder to deal with.

A good security script will also check for changes to user and staff members access configuration files: `.rhosts`, `.shosts`, `.ssh/authorized_keys` and so forth... files that might fall outside the purview of the `MD5` check.

If you have a huge amount of user disk space, it may take too long to run through every file on those partitions. In this case, setting mount flags to disallow suid binaries and devices on those partitions is a good idea. The `nodev` and `nosuid` options (see `mount(8)`) are what you want to look into. You should probably scan them anyway, at least once a week, since the object of this layer is to detect a break-in whether or not the break-in is effective.

Process accounting (see `accton(8)`) is a relatively low-overhead feature of the operating system which might help as a post-break-in evaluation mechanism. It is especially useful in tracking down how an intruder has actually broken into a system, assuming the file is still intact after the break-in occurs.

Finally, security scripts should process the log files, and the logs themselves should be generated in as secure a manner as possible — remote syslog can be very useful. An intruder tries to cover his tracks, and log files are critical to the sysadmin trying to track down the time and method of the initial break-in. One way to keep a permanent record of the log files is to run the system console to a serial port and collect the information on a continuing basis through a secure machine monitoring the consoles.

10.3.7 Paranoia

A little paranoia never hurts. As a rule, a sysadmin can add any number of security features, as long as they do not affect convenience, and can add security features that *do* affect convenience with some added thought. Even more importantly, a security administrator should mix it up a bit — if you use recommendations such as those given by this document verbatim, you give away your methodologies to the prospective attacker who also has access to this document.

10.3.8 Denial of Service Attacks

This section covers Denial of Service attacks. A DoS attack is typically a packet attack. While there is not much you can do about modern spoofed packet attacks that saturate your network, you can generally limit the damage by ensuring that the attacks cannot take down your servers.

1. Limiting server forks.
2. Limiting springboard attacks (ICMP response attacks, ping broadcast, etc.).
3. Kernel Route Cache.

A common DoS attack is against a forking server that attempts to cause the server to eat processes, file descriptors, and memory, until the machine dies. **inetd** (see `inetd(8)`) has several options to limit this sort of attack. It should be noted that while it is possible to prevent a machine from going down, it is not generally possible to prevent a service from being disrupted by the attack. Read the **inetd** manual page carefully and pay specific attention to the `-c`, `-C`, and `-R` options. Note that spoofed-IP attacks will circumvent the `-C` option to **inetd**, so typically a combination of options must be used. Some standalone servers have self-fork-limitation parameters.

Sendmail has its `-OMaxDaemonChildren` option, which tends to work much better than trying to use sendmail's load limiting options due to the load lag. You should specify a `MaxDaemonChildren` parameter, when you start **sendmail**, high enough to handle your expected load, but not so high that the computer cannot handle that number of **sendmails** without falling on its face. It is also prudent to run sendmail in queued mode

(`-ODeliveryMode=queued`) and to run the daemon (`sendmail -bd`) separate from the queue-runs (`sendmail -q15m`). If you still want real-time delivery you can run the queue at a much lower interval, such as `-q1m`, but be sure to specify a reasonable `MaxDaemonChildren` option for *that* sendmail to prevent cascade failures.

Syslogd can be attacked directly and it is strongly recommended that you use the `-s` option whenever possible, and the `-a` option otherwise.

You should also be fairly careful with connect-back services such as **tcpwrapper**'s `reverse-identd`, which can be attacked directly. You generally do not want to use the `reverse-ident` feature of **tcpwrappers** for this reason.

It is a very good idea to protect internal services from external access by firewalling them off at your border routers. The idea here is to prevent saturation attacks from outside your LAN, not so much to protect internal services from network-based `root` compromise. Always configure an exclusive firewall, i.e., “firewall everything *except* ports A, B, C, D, and M-Z”. This way you can firewall off all of your low ports except for certain specific services such as **named** (if you are primary for a zone), **ntalkd**, **sendmail**, and other Internet-accessible services. If you try to configure the firewall the other way — as an inclusive or permissive firewall, there is a good chance that you will forget to “close” a couple of services, or that you will add a new internal service and forget to update the firewall. You can still open up the high-numbered port range on the firewall, to allow permissive-like operation, without compromising your low ports. Also take note that DragonFly allows you to control the range of port numbers used for dynamic binding, via the various `net.inet.ip.portrange` `sysctl`'s (`sysctl -a | fgrep portrange`), which can also ease the complexity of your firewall's configuration. For example, you might use a normal `first/last` range of 4000 to 5000, and a `hiport` range of 49152 to 65535, then block off everything under 4000 in your firewall (except for certain specific Internet-accessible ports, of course).

Another common DoS attack is called a springboard attack — to attack a server in a manner that causes the server to generate responses which overloads the server, the local network, or some other machine. The most common attack of this nature is the *ICMP ping broadcast attack*. The attacker spoofs ping packets sent to your LAN's broadcast address with the source IP address set to the actual machine they wish to attack. If your border routers are not configured to stomp on ping's to broadcast addresses, your LAN winds up generating sufficient responses to the spoofed source address to saturate the victim, especially when the attacker uses the same trick on several dozen broadcast addresses over several dozen different networks at once. Broadcast attacks of over a hundred and twenty megabits have been measured. A second common springboard attack is against the ICMP error reporting system. By constructing packets that generate ICMP error responses, an attacker can saturate a server's incoming network and cause the server to saturate its outgoing network with ICMP responses. This type of attack can also crash the server by running it out of `mbuf`'s, especially if the server cannot drain the ICMP responses it generates fast enough. The DragonFly kernel has a new kernel compile option called `ICMP_BANDLIM` which limits the effectiveness of these sorts of attacks. The last major class of springboard attacks is related to certain internal **inetd** services such as the `udp echo` service. An attacker simply spoofs a UDP packet with the source address being server A's echo port, and the destination address being server B's echo port, where server A and B are both on your LAN. The two servers then bounce this one packet back and forth between each other. The attacker can overload both servers and their LANs simply by injecting a few packets in this manner. Similar problems exist with the internal **chargen** port. A competent sysadmin will turn off all of these `inetd`-internal test services.

Spoofed packet attacks may also be used to overload the kernel route cache. Refer to the `net.inet.ip.rtxpire`, `rtminexpire`, and `rtmaxcache` `sysctl` parameters. A spoofed packet attack that uses a random source IP will cause the kernel to generate a temporary cached route in the route table, viewable with `netstat -rna | fgrep w3`. These routes typically timeout in 1600 seconds or so. If the kernel detects that the cached route table has gotten too big it will dynamically reduce the `rtxpire` but will never decrease it to less than `rtminexpire`. There are two problems:

1. The kernel does not react quickly enough when a lightly loaded server is suddenly attacked.
2. The `rtminexpire` is not low enough for the kernel to survive a sustained attack.

If your servers are connected to the Internet via a T3 or better, it may be prudent to manually override both

`rtexpire` and `rtminexpire` via `sysctl(8)`. Never set either parameter to zero (unless you want to crash the machine). Setting both parameters to two seconds should be sufficient to protect the route table from attack.

10.3.9 Access Issues with Kerberos and SSH

There are a few issues with both Kerberos and `ssh` that need to be addressed if you intend to use them. Kerberos V is an excellent authentication protocol, but there are bugs in the kerberized `telnet` and `rlogin` applications that make them unsuitable for dealing with binary streams. Also, by default Kerberos does not encrypt a session unless you use the `-x` option. `ssh` encrypts everything by default.

`ssh` works quite well in every respect except that it forwards encryption keys by default. What this means is that if you have a secure workstation holding keys that give you access to the rest of the system, and you `ssh` to an insecure machine, your keys are usable. The actual keys themselves are not exposed, but `ssh` installs a forwarding port for the duration of your login, and if an attacker has broken `root` on the insecure machine he can utilize that port to use your keys to gain access to any other machine that your keys unlock.

We recommend that you use `ssh` in combination with Kerberos whenever possible for staff logins. `ssh` can be compiled with Kerberos support. This reduces your reliance on potentially exposable `ssh` keys while at the same time protecting passwords via Kerberos. `ssh` keys should only be used for automated tasks from secure machines (something that Kerberos is unsuited to do). We also recommend that you either turn off key-forwarding in the `ssh` configuration, or that you make use of the `from=IP/DOMAIN` option that `ssh` allows in its `authorized_keys` file to make the key only usable to entities logging in from specific machines.

10.4 DES, MD5, and Crypt

Parts rewritten and updated by Bill Swingle.

Every user on a UNIX system has a password associated with their account. It seems obvious that these passwords need to be known only to the user and the actual operating system. In order to keep these passwords secret, they are encrypted with what is known as a “one-way hash”, that is, they can only be easily encrypted but not decrypted. In other words, what we told you a moment ago was obvious is not even true: the operating system itself does not *really* know the password. It only knows the *encrypted* form of the password. The only way to get the “plain-text” password is by a brute force search of the space of possible passwords.

Unfortunately the only secure way to encrypt passwords when UNIX came into being was based on DES, the Data Encryption Standard. This was not such a problem for users resident in the US, but since the source code for DES could not be exported outside the US, DragonFly had to find a way to both comply with US law and retain compatibility with all the other UNIX variants that still used DES.

The solution was to divide up the encryption libraries so that US users could install the DES libraries and use DES but international users still had an encryption method that could be exported abroad. This is how DragonFly came to use MD5 as its default encryption method. MD5 is believed to be more secure than DES, so installing DES is offered primarily for compatibility reasons.

10.4.1 Recognizing Your Crypt Mechanism

`libcrypt.a` provides a configurable password authentication hash library. Currently the library supports DES, MD5 and Blowfish hash functions. By default DragonFly uses MD5 to encrypt passwords.

It is pretty easy to identify which encryption method DragonFly is set up to use. Examining the encrypted passwords in the `/etc/master.passwd` file is one way. Passwords encrypted with the MD5 hash are longer than those encrypted with the DES hash and also begin with the characters `1`. Passwords starting with `$2a$` are encrypted with the Blowfish hash function. DES password strings do not have any particular identifying characteristics, but they are shorter than MD5 passwords, and are coded in a 64-character alphabet which does not include the `$` character, so a relatively short string which does not begin with a dollar sign is very likely a DES password.

The password format used for new passwords is controlled by the `passwd_format` login capability in `/etc/login.conf`, which takes values of `des`, `md5` or `blf`. See the `login.conf(5)` manual page for more information about login capabilities.

10.5 One-time Passwords

S/Key is a one-time password scheme based on a one-way hash function. DragonFly uses the MD4 hash for compatibility but other systems have used MD5 and DES-MAC. S/Key is part of the FreeBSD base system, and is also used on a growing number of other operating systems. S/Key is a registered trademark of Bell Communications Research, Inc.

There are three different sorts of passwords which we will discuss below. The first is your usual UNIX style or Kerberos password; we will call this a “UNIX password”. The second sort is the one-time password which is generated by the S/Key `key` program or the OPIE `opiekey(1)` program and accepted by the `keyinit` or `opiepasswd(1)` programs and the login prompt; we will call this a “one-time password”. The final sort of password is the secret password which you give to the `key/opiekey` programs (and sometimes the `keyinit/opiepasswd` programs) which it uses to generate one-time passwords; we will call it a “secret password” or just unqualified “password”.

The secret password does not have anything to do with your UNIX password; they can be the same but this is not recommended. S/Key and OPIE secret passwords are not limited to eight characters like old UNIX passwords¹, they can be as long as you like. Passwords of six or seven word long phrases are fairly common. For the most part, the S/Key or OPIE system operates completely independently of the UNIX password system.

Besides the password, there are two other pieces of data that are important to S/Key and OPIE. One is what is known as the “seed” or “key”, consisting of two letters and five digits. The other is what is called the “iteration count”, a number between 1 and 100. S/Key creates the one-time password by concatenating the seed and the secret password, then applying the MD4/MD5 hash as many times as specified by the iteration count and turning the result into six short English words. These six English words are your one-time password. The authentication system (primarily PAM) keeps track of the last one-time password used, and the user is authenticated if the hash of the user-provided password is equal to the previous password. Because a one-way hash is used it is impossible to generate future one-time passwords if a successfully used password is captured; the iteration count is decremented after each successful login to keep the user and the login program in sync. When the iteration count gets down to 1, S/Key and OPIE must be reinitialized.

There are three programs involved in each system which we will discuss below. The `key` and `opiekey` programs accept an iteration count, a seed, and a secret password, and generate a one-time password or a consecutive list of one-time passwords. The `keyinit` and `opiepasswd` programs are used to initialize S/Key and OPIE respectively, and to change passwords, iteration counts, or seeds; they take either a secret passphrase, or an iteration count, seed, and one-time password. The `keyinfo` and `opieinfo` programs examine the relevant credentials files (`/etc/skeykeys` or `/etc/opiekeys`) and print out the invoking user’s current iteration count and seed.

There are four different sorts of operations we will cover. The first is using `keyinit` or `opiepasswd` over a secure connection to set up one-time-passwords for the first time, or to change your password or seed. The second operation is using `keyinit` or `opiepasswd` over an insecure connection, in conjunction with `key` or `opiekey` over a secure connection, to do the same. The third is using `key/opiekey` to log in over an insecure connection. The fourth is using `key` or `opiekey` to generate a number of keys which can be written down or printed out to carry with you when going to some location without secure connections to anywhere.

10.5.1 Secure Connection Initialization

To initialize S/Key for the first time, change your password, or change your seed while logged in over a secure connection (e.g., on the console of a machine or via `ssh`), use the `keyinit` command without any parameters while logged in as yourself:

```
% keyinit
Adding unfurl:
Reminder - Only use this method if you are directly connected.
If you are using telnet or rlogin exit with no password and use keyinit -s.
Enter secret password:
Again secret password:

ID unfurl s/key is 99 to17757
DEFY CLUB PRO NASH LACE SOFT
```

For OPIE, `opiepasswd` is used instead:

```
% opiepasswd -c
[grimreaper] ~ $ opiepasswd -f -c
Adding unfurl:
Only use this method from the console; NEVER from remote. If you are using
telnet, xterm, or a dial-in, type ^C now or exit with no password.
Then run opiepasswd without the -c parameter.
Using MD5 to compute responses.
Enter new secret pass phrase:
Again new secret pass phrase:
ID unfurl OTP key is 499 to4268
MOS MALL GOAT ARM AVID COED
```

At the `Enter new secret pass phrase:` or `Enter secret password:` prompts, you should enter a password or phrase. Remember, this is not the password that you will use to login with, this is used to generate your one-time login keys. The “ID” line gives the parameters of your particular instance: your login name, the iteration count, and seed. When logging in the system will remember these parameters and present them back to you so you do not have to remember them. The last line gives the particular one-time password which corresponds to those parameters and your secret password; if you were to re-login immediately, this one-time password is the one you would use.

10.5.2 Insecure Connection Initialization

To initialize or change your secret password over an insecure connection, you will need to already have a secure connection to some place where you can run `key` or `opiekey`; this might be in the form of a desk accessory on a Macintosh, or a shell prompt on a machine you trust. You will also need to make up an iteration count (100 is

probably a good value), and you may make up your own seed or use a randomly-generated one. Over on the insecure connection (to the machine you are initializing), use the `keyinit -s` command:

```
% keyinit -s
Updating unfurl:
Old key: to17758
Reminder you need the 6 English words from the key command.
Enter sequence count from 1 to 9999: 100
Enter new key [default to17759]:
s/key 100 to 17759
s/key access password:
s/key access password:CURE MIKE BANE HIM RACY GORE
```

For OPIE, you need to use `opiepasswd`:

```
% opiepasswd

Updating unfurl:
You need the response from an OTP generator.
Old secret pass phrase:
    otp-md5 498 to4268 ext
    Response: GAME GAG WELT OUT DOWN CHAT
New secret pass phrase:
    otp-md5 499 to4269
    Response: LINE PAP MILK NELL BUOY TROY
```

```
ID mark OTP key is 499 gr4269
LINE PAP MILK NELL BUOY TROY
```

To accept the default seed (which the `keyinit` program confusingly calls a key), press **Return**. Then before entering an access password, move over to your secure connection or S/Key desk accessory, and give it the same parameters:

```
% key 100 to17759
Reminder - Do not use this program while logged in via telnet or rlogin.
Enter secret password: <secret password>
CURE MIKE BANE HIM RACY GORE
```

Or for OPIE:

```
% opiekey 498 to4268
Using the MD5 algorithm to compute response.
Reminder: Don't use opiekey from telnet or dial-in sessions.
Enter secret pass phrase:
GAME GAG WELT OUT DOWN CHAT
```

Now switch back over to the insecure connection, and copy the one-time password generated over to the relevant program.

10.5.3 Generating a Single One-time Password

Once you have initialized S/Key, when you login you will be presented with a prompt like this:

```
% telnet example.com
Trying 10.0.0.1...
Connected to example.com
Escape character is '^]'.

DragonFly/i386 (example.com) (tty)

login: <username>
s/key 97 fw13894
Password:
```

Or for OPIE:

```
% telnet example.com
Trying 10.0.0.1...
Connected to example.com
Escape character is '^]'.

DragonFly/i386 (example.com) (tty)

login: <username>
otp-md5 498 gr4269 ext
Password:
```

As a side note, the S/Key and OPIE prompts have a useful feature (not shown here): if you press **Return** at the password prompt, the prompter will turn echo on, so you can see what you are typing. This can be extremely useful if you are attempting to type in a password by hand, such as from a printout.

At this point you need to generate your one-time password to answer this login prompt. This must be done on a trusted system that you can run `key` or `opiekey` on. (There are versions of these for DOS, Windows and Mac OS as well.) They need both the iteration count and the seed as command line options. You can cut-and-paste these right from the login prompt on the machine that you are logging in to.

On the trusted system:

```
% key 97 fw13894
Reminder - Do not use this program while logged in via telnet or rlogin.
Enter secret password:
WELD LIP ACTS ENDS ME HAAG
```

For OPIE:

```
% opiekey 498 to4268
Using the MD5 algorithm to compute response.
Reminder: Don't use opiekey from telnet or dial-in sessions.
Enter secret pass phrase:
GAME GAG WELT OUT DOWN CHAT
```

Now that you have your one-time password you can continue logging in:

```
login: <username>
```



```
s/key 97 fw13894
Password: <return to enable echo>
s/key 97 fw13894
Password [echo on]: WELD LIP ACTS ENDS ME HAAG
Last login: Tue Mar 21 11:56:41 from 10.0.0.2 ...
```

10.5.4 Generating Multiple One-time Passwords

Sometimes you have to go places where you do not have access to a trusted machine or secure connection. In this case, it is possible to use the `key` and `opiekey` commands to generate a number of one-time passwords beforehand to be printed out and taken with you. For example:

```
% key -n 5 30 zz99999
Reminder - Do not use this program while logged in via telnet or rlogin.
Enter secret password: <secret password>
26: SODA RUDE LEA LIND BUDD SILT
27: JILT SPY DUTY GLOW COWL ROT
28: THEM OW COLA RUNT BONG SCOT
29: COT MASH BARR BRIM NAN FLAG
30: CAN KNEE CAST NAME FOLK BILK
```

Or for OPIE:

```
% opiekey -n 5 30 zz99999
Using the MD5 algorithm to compute response.
Reminder: Don't use opiekey from telnet or dial-in sessions.
Enter secret pass phrase: <secret password>
26: JOAN BORE FOSS DES NAY QUIT
27: LATE BIAS SLAY FOLK MUCH TRIG
28: SALT TIN ANTI LOON NEAL USE
29: RIO ODIN GO BYE FURY TIC
30: GREW JIVE SAN GIRD BOIL PHI
```

The `-n 5` requests five keys in sequence, the `30` specifies what the last iteration number should be. Note that these are printed out in *reverse* order of eventual use. If you are really paranoid, you might want to write the results down by hand; otherwise you can cut-and-paste into `lpr`. Note that each line shows both the iteration count and the one-time password; you may still find it handy to scratch off passwords as you use them.

10.5.5 Restricting Use of UNIX® Passwords

S/Key can place restrictions on the use of UNIX passwords based on the host name, user name, terminal port, or IP address of a login session. These restrictions can be found in the configuration file `/etc/skey.access`. The `skey.access(5)` manual page has more information on the complete format of the file and also details some security cautions to be aware of before depending on this file for security.

If there is no `/etc/skey.access` file (this is the default), then all users will be allowed to use UNIX passwords. If the file exists, however, then all users will be required to use S/Key unless explicitly permitted to do otherwise by configuration statements in the `skey.access` file. In all cases, UNIX passwords are permitted on the console.

Here is a sample `skey.access` configuration file which illustrates the three most common sorts of configuration statements:

```
permit internet 192.168.0.0 255.255.0.0
permit user fnord
permit port ttyd0
```

The first line (`permit internet`) allows users whose IP source address (which is vulnerable to spoofing) matches the specified value and mask, to use UNIX passwords. This should not be considered a security mechanism, but rather, a means to remind authorized users that they are using an insecure network and need to use S/Key for authentication.

The second line (`permit user`) allows the specified username, in this case `fnord`, to use UNIX passwords at any time. Generally speaking, this should only be used for people who are either unable to use the `key` program, like those with dumb terminals, or those who are uneducable.

The third line (`permit port`) allows all users logging in on the specified terminal line to use UNIX passwords; this would be used for dial-ups.

Here is a sample `opieaccess` file:

```
permit 192.168.0.0 255.255.0.0
```

This line allows users whose IP source address (which is vulnerable to spoofing) matches the specified value and mask, to use UNIX passwords at any time.

If no rules in `opieaccess` are matched, the default is to deny non-OPIE logins.

10.6 Kerberos5

Contributed by Tillman Hodgson. Based on a contribution by Mark Murray.

The following information only applies to **Kerberos5**. Users who wish to use the **KerberosIV** package may install the `security/krb4` port.

Kerberos is a network add-on system/protocol that allows users to authenticate themselves through the services of a secure server. Services such as remote login, remote copy, secure inter-system file copying and other high-risk tasks are made considerably safer and more controllable.

Kerberos can be described as an identity-verifying proxy system. It can also be described as a trusted third-party authentication system. **Kerberos** provides only one function — the secure authentication of users on the network. It does not provide authorization functions (what users are allowed to do) or auditing functions (what those users did). After a client and server have used **Kerberos** to prove their identity, they can also encrypt all of their communications to assure privacy and data integrity as they go about their business.

Therefore it is highly recommended that **Kerberos** be used with other security methods which provide authorization and audit services.

The following instructions can be used as a guide on how to set up **Kerberos** as distributed for DragonFly. However, you should refer to the relevant manual pages for a complete description.

For purposes of demonstrating a **Kerberos** installation, the various namespaces will be handled as follows:

- The DNS domain (“zone”) will be example.org.
- The **Kerberos** realm will be EXAMPLE.ORG.

Note: Please use real domain names when setting up **Kerberos** even if you intend to run it internally. This avoids DNS problems and assures inter-operation with other **Kerberos** realms.

10.6.1 History

Kerberos was created by MIT as a solution to network security problems. The **Kerberos** protocol uses strong cryptography so that a client can prove its identity to a server (and vice versa) across an insecure network connection.

Kerberos is both the name of a network authentication protocol and an adjective to describe programs that implement the program (**Kerberos** telnet, for example). The current version of the protocol is version 5, described in RFC 1510.

Several free implementations of this protocol are available, covering a wide range of operating systems. The Massachusetts Institute of Technology (MIT), where **Kerberos** was originally developed, continues to develop their **Kerberos** package. It is commonly used in the US as a cryptography product, as such it has historically been affected by US export regulations. The MIT **Kerberos** is available as a port (`security/krb5`). Heimdal **Kerberos** is another version 5 implementation, and was explicitly developed outside of the US to avoid export regulations (and is thus often included in non-commercial UNIX variants). The Heimdal **Kerberos** distribution is available as a port (`security/heimdal`), and a minimal installation of it is included in the base DragonFly install.

In order to reach the widest audience, these instructions assume the use of the Heimdal distribution included in DragonFly.

10.6.2 Setting up a Heimdal KDC

The Key Distribution Center (KDC) is the centralized authentication service that **Kerberos** provides — it is the computer that issues **Kerberos** tickets. The KDC is considered “trusted” by all other computers in the **Kerberos** realm, and thus has heightened security concerns.

Note that while running the **Kerberos** server requires very few computing resources, a dedicated machine acting only as a KDC is recommended for security reasons.

To begin setting up a KDC, ensure that your `/etc/rc.conf` file contains the correct settings to act as a KDC (you may need to adjust paths to reflect your own system):

```
kerberos5_server_enable="YES"
kadmind5_server_enable="YES"
kerberos_stash="YES"
```

Next we will set up your **Kerberos** config file, `/etc/krb5.conf`:

```
[libdefaults]
    default_realm = EXAMPLE.ORG
[realms]
    EXAMPLE.ORG = {
        kdc = kerberos.example.org
    }
```

```
[domain_realm]
    .example.org = EXAMPLE.ORG
```

Note that this `/etc/krb5.conf` file implies that your KDC will have the fully-qualified hostname of `kerberos.example.org`. You will need to add a CNAME (alias) entry to your zone file to accomplish this if your KDC has a different hostname.

Note: For large networks with a properly configured BIND DNS server, the above example could be trimmed to:

```
[libdefaults]
    default_realm = EXAMPLE.ORG
```

With the following lines being appended to the `example.org` zonefile:

```
_kerberos._udp      IN  SRV      01 00 88 kerberos.example.org.
_kerberos._tcp      IN  SRV      01 00 88 kerberos.example.org.
_kpasswd._udp       IN  SRV      01 00 464 kerberos.example.org.
_kerberos-adm._tcp  IN  SRV      01 00 749 kerberos.example.org.
_kerberos           IN  TXT      EXAMPLE.ORG.
```

Next we will create the **Kerberos** database. This database contains the keys of all principals encrypted with a master password. You are not required to remember this password, it will be stored in a file (`/var/heimdal/m-key`). To create the master key, run `kstash` and enter a password.

Once the master key has been created, you can initialize the database using the `kadmin` program with the `-l` option (standing for “local”). This option instructs `kadmin` to modify the database files directly rather than going through the `kadmind` network service. This handles the chicken-and-egg problem of trying to connect to the database before it is created. Once you have the `kadmin` prompt, use the `init` command to create your realms initial database.

Lastly, while still in `kadmin`, create your first principal using the `add` command. Stick to the defaults options for the principal for now, you can always change them later with the `modify` command. Note that you can use the `?` command at any prompt to see the available options.

A sample database creation session is shown below:

```
# kstash
Master key: xxxxxxxx
Verifying password - Master key: xxxxxxxx

# kadmin -l
kadmin> init EXAMPLE.ORG
Realm max ticket life [unlimited]:
kadmin> add tillman
Max ticket life [unlimited]:
Max renewable life [unlimited]:
Attributes []:
Password: xxxxxxxx
Verifying password - Password: xxxxxxxx
```

Now it is time to start up the KDC services. Run `/etc/rc.d/kerberos start` and `/etc/rc.d/kadmind start` to bring up the services. Note that you won't have any kerberized daemons running at this point but you

should be able to confirm that the KDC is functioning by obtaining and listing a ticket for the principal (user) that you just created from the command-line of the KDC itself:

```
% k5init tillman
tillman@EXAMPLE.ORG's Password:

% k5list
Credentials cache: FILE:/tmp/krb5cc_500
Principal: tillman@EXAMPLE.ORG

    Issued            Expires            Principal
Aug 27 15:37:58    Aug 28 01:37:58    krbtgt/EXAMPLE.ORG@EXAMPLE.ORG
```

10.6.3 Kerberos enabling a server with Heimdal services

First, we need a copy of the **Kerberos** configuration file, `/etc/krb5.conf`. To do so, simply copy it over to the client computer from the KDC in a secure fashion (using network utilities, such as `scp(1)`, or physically via a floppy disk).

Next you need a `/etc/krb5.keytab` file. This is the major difference between a server providing **Kerberos** enabled daemons and a workstation — the server must have a `keytab` file. This file contains the servers host key, which allows it and the KDC to verify each others identity. It must be transmitted to the server in a secure fashion, as the security of the server can be broken if the key is made public. This explicitly means that transferring it via a clear text channel, such as FTP, is a very bad idea.

Typically, you transfer the `keytab` to the server using the `kadmin` program. This is handy because you also need to create the host principal (the KDC end of the `krb5.keytab`) using `kadmin`.

Note that you must have already obtained a ticket and that this ticket must be allowed to use the `kadmin` interface in the `kadmind.acl`. See the section titled “Remote administration” in the Heimdal info pages (`info heimdal`) for details on designing access control lists. If you do not want to enable remote `kadmin` access, you can simply securely connect to the KDC (via local console, `ssh(1)` or **Kerberos** `telnet(1)`) and perform administration locally using `kadmin -l`.

After installing the `/etc/krb5.conf` file, you can use `kadmin` from the **Kerberos** server. The `add --random-key` command will let you add the servers host principal, and the `ext` command will allow you to extract the servers host principal to its own `keytab`. For example:

```
# kadmin
kadmin> add --random-key host/myserver.example.org
Max ticket life [unlimited]:
Max renewable life [unlimited]:
Attributes []:
kadmin> ext host/myserver.example.org
kadmin> exit
```

Note that the `ext` command (short for “extract”) stores the extracted key in `/etc/krb5.keytab` by default.

If you do not have `kadmind` running on the KDC (possibly for security reasons) and thus do not have access to `kadmin` remotely, you can add the host principal (`host/myserver.EXAMPLE.ORG`) directly on the KDC and then extract it to a temporary file (to avoid over-writing the `/etc/krb5.keytab` on the KDC) using something like this:

```
# kadmin
kadmin> ext --keytab=/tmp/example.keytab host/myserver.example.org
kadmin> exit
```

You can then securely copy the keytab to the server computer (using `scp` or a floppy, for example). Be sure to specify a non-default keytab name to avoid over-writing the keytab on the KDC.

At this point your server can communicate with the KDC (due to its `krb5.conf` file) and it can prove its own identity (due to the `krb5.keytab` file). It is now ready for you to enable some **Kerberos** services. For this example we will enable the `telnet` service by putting a line like this into your `/etc/inetd.conf` and then restarting the `inetd(8)` service with `/etc/rc.d/inetd restart`:

```
telnet    stream  tcp      nowait  root    /usr/libexec/telnetd  telnetd  -a user
```

The critical bit is that the `-a` (for authentication) type is set to `user`. Consult the `telnetd(8)` manual page for more details.

10.6.4 Kerberos enabling a client with Heimdal

Setting up a client computer is almost trivially easy. As far as **Kerberos** configuration goes, you only need the **Kerberos** configuration file, located at `/etc/krb5.conf`. Simply securely copy it over to the client computer from the KDC.

Test your client computer by attempting to use `kinit`, `klist`, and `kdestroy` from the client to obtain, show, and then delete a ticket for the principal you created above. You should also be able to use **Kerberos** applications to connect to **Kerberos** enabled servers, though if that does not work and obtaining a ticket does the problem is likely with the server and not with the client or the KDC.

When testing an application like `telnet`, try using a packet sniffer (such as `tcpdump(1)`) to confirm that your password is not sent in the clear. Try using `telnet` with the `-x` option, which encrypts the entire data stream (similar to `ssh`).

The core **Kerberos** client applications (traditionally named `kinit`, `klist`, `kdestroy`, and `kpasswd`) are installed in the base DragonFly install. Note that DragonFly versions prior to 5.0 renamed them to `k5init`, `k5list`, `k5destroy`, `k5passwd`, and `k5stash` (though it is typically only used once).

Various non-core **Kerberos** client applications are also installed by default. This is where the “minimal” nature of the base Heimdal installation is felt: `telnet` is the only **Kerberos** enabled service.

The Heimdal port adds some of the missing client applications: **Kerberos** enabled versions of `ftp`, `rsh`, `rcp`, `rlogin`, and a few other less common programs. The MIT port also contains a full suite of **Kerberos** client applications.

10.6.5 User configuration files: `.k5login` and `.k5users`

Users within a realm typically have their **Kerberos** principal (such as `tillman@EXAMPLE.ORG`) mapped to a local user account (such as a local account named `tillman`). Client applications such as `telnet` usually do not require a user name or a principal.

Occasionally, however, you want to grant access to a local user account to someone who does not have a matching **Kerberos** principal. For example, `tillman@EXAMPLE.ORG` may need access to the local user account `webdevelopers`. Other principals may also need access to that local account.

The `.k5login` and `.k5users` files, placed in a user's home directory, can be used similar to a powerful combination of `.hosts` and `.rhosts`, solving this problem. For example, if a `.k5login` with the following contents:

```
tillman@example.org
jdoe@example.org
```

Were to be placed into the home directory of the local user `webdevelopers` then both principals listed would have access to that account without requiring a shared password.

Reading the manual pages for these commands is recommended. Note that the `ksu` manual page covers `.k5users`.

10.6.6 Kerberos Tips, Tricks, and Troubleshooting

- When using either the Heimdal or MIT **Kerberos** ports ensure that your `PATH` environment variable lists the **Kerberos** versions of the client applications before the system versions.
- Is your time in sync? Are you sure? If the time is not in sync (typically within five minutes) authentication will fail.
- MIT and Heimdal inter-operate nicely. Except for `kadmin`, the protocol for which is not standardized.
- If you change your hostname, you also need to change your `host/` principal and update your keytab. This also applies to special keytab entries like the `www/` principal used for Apache's `www/mod_auth_kerb`.
- All hosts in your realm must be resolvable (both forwards and reverse) in DNS (or `/etc/hosts` as a minimum). CNAMEs will work, but the A and PTR records must be correct and in place. The error message isn't very intuitive: `Kerberos5 refuses authentication because Read req failed: Key table entry not found`.
- Some operating systems that may be acting as clients to your KDC do not set the permissions for `ksu` to be `setuid root`. This means that `ksu` does not work, which is a good security idea but annoying. This is not a KDC error.
- With MIT **Kerberos**, if you want to allow a principal to have a ticket life longer than the default ten hours, you must use `modify_principal` in `kadmin` to change the `maxlife` of both the principal in question and the `krbtgt` principal. Then the principal can use the `-l` option with `kinit` to request a ticket with a longer lifetime.
-

Note: If you run a packet sniffer on your KDC to add in troubleshooting and then run `kinit` from a workstation, you will notice that your TGT is sent immediately upon running `kinit` — even before you type your password! The explanation is that the **Kerberos** server freely transmits a TGT (Ticket Granting Ticket) to any unauthorized request; however, every TGT is encrypted in a key derived from the user's password. Therefore, when a user types their password it is not being sent to the KDC, it is being used to decrypt the TGT that `kinit` already obtained. If the decryption process results in a valid ticket with a valid time stamp, the user has valid **Kerberos** credentials. These credentials include a session key for establishing secure communications with the **Kerberos** server in the future, as well as the actual ticket-granting ticket, which is actually encrypted with the **Kerberos** server's own key. This second layer of encryption is unknown to the user, but it is what allows the **Kerberos** server to verify the authenticity of each TGT.

- You have to keep the time in sync between all the computers in your realm. NTP is perfect for this. For more information on NTP, see Section 19.12.
- If you want to use long ticket lifetimes (a week, for example) and you are using **OpenSSH** to connect to the machine where your ticket is stored, make sure that **Kerberos** `TicketCleanup` is set to `no` in your `sshd_config` or else your tickets will be deleted when you log out.
- Remember that host principals can have a longer ticket lifetime as well. If your user principal has a lifetime of a week but the host you are connecting to has a lifetime of nine hours, you will have an expired host principal in your cache and the ticket cache will not work as expected.
- When setting up a `krb5.dict` file to prevent specific bad passwords from being used (the manual page for `kadmind` covers this briefly), remember that it only applies to principals that have a password policy assigned to them. The `krb5.dict` files format is simple: one string per line. Creating a symbolic link to `/usr/share/dict/words` might be useful.

10.6.7 Differences with the MIT port

The major difference between the MIT and Heimdal installs relates to the `kadmin` program which has a different (but equivalent) set of commands and uses a different protocol. This has a large implications if your KDC is MIT as you will not be able to use the Heimdal `kadmin` program to administer your KDC remotely (or vice versa, for that matter).

The client applications may also take slightly different command line options to accomplish the same tasks. Following the instructions on the MIT **Kerberos** web site (<http://web.mit.edu/Kerberos/www/>) is recommended. Be careful of path issues: the MIT port installs into `/usr/local/` by default, and the “normal” system applications may be run instead of MIT if your `PATH` environment variable lists the system directories first.

Note: With the MIT `security/krb5` port that is provided by DragonFly, be sure to read the `/usr/local/share/doc/krb5/README.FreeBSD` file installed by the port if you want to understand why logins via `telnetd` and `klogind` behave somewhat oddly. Most importantly, correcting the “incorrect permissions on cache file” behavior requires that the `login.krb5` binary be used for authentication so that it can properly change ownership for the forwarded credentials.

10.6.8 Mitigating limitations found in Kerberos

10.6.8.1 Kerberos is an all-or-nothing approach

Every service enabled on the network must be modified to work with **Kerberos** (or be otherwise secured against network attacks) or else the users credentials could be stolen and re-used. An example of this would be **Kerberos** enabling all remote shells (via `rsh` and `telnet`, for example) but not converting the POP3 mail server which sends passwords in plaintext.

10.6.8.2 Kerberos is intended for single-user workstations

In a multi-user environment, **Kerberos** is less secure. This is because it stores the tickets in the `/tmp` directory, which is readable by all users. If a user is sharing a computer with several other people simultaneously (i.e. multi-user), it is possible that the user's tickets can be stolen (copied) by another user.

This can be overcome with the `-c` filename command-line option or (preferably) the `KRB5CCNAME` environment variable, but this is rarely done. In principal, storing the ticket in the users home directory and using simple file permissions can mitigate this problem.

10.6.8.3 The KDC is a single point of failure

By design, the KDC must be as secure as the master password database is contained on it. The KDC should have absolutely no other services running on it and should be physically secured. The danger is high because **Kerberos** stores all passwords encrypted with the same key (the "master" key), which in turn is stored as a file on the KDC.

As a side note, a compromised master key is not quite as bad as one might normally fear. The master key is only used to encrypt the **Kerberos** database and as a seed for the random number generator. As long as access to your KDC is secure, an attacker cannot do much with the master key.

Additionally, if the KDC is unavailable (perhaps due to a denial of service attack or network problems) the network services are unusable as authentication can not be performed, a recipe for a denial-of-service attack. This can be alleviated with multiple KDCs (a single master and one or more slaves) and with careful implementation of secondary or fall-back authentication (PAM is excellent for this).

10.6.8.4 Kerberos Shortcomings

Kerberos allows users, hosts and services to authenticate between themselves. It does not have a mechanism to authenticate the KDC to the users, hosts or services. This means that a trojanned `kinit` (for example) could record all user names and passwords. Something like `security/tripwire` or other file system integrity checking tools can alleviate this.

10.6.9 Resources and further information

- The **Kerberos** FAQ (<http://www.faqs.org/faqs/Kerberos-faq/general/preamble.html>)
- Designing an Authentication System: a Dialogue in Four Scenes (<http://web.mit.edu/Kerberos/www/dialogue.html>)
- RFC 1510, The **Kerberos** Network Authentication Service (V5) (<http://www.ietf.org/rfc/rfc1510.txt?number=1510>)
- MIT **Kerberos** home page (<http://web.mit.edu/Kerberos/www/>)
- Heimdal **Kerberos** home page (<http://www.pdc.kth.se/heimdal/>)

10.7 Firewalls

Contributed by Gary Palmer and Alex Nash.

Firewalls are an area of increasing interest for people who are connected to the Internet, and are even finding applications on private networks to provide enhanced security. This section will hopefully explain what firewalls are, how to use them, and how to use the facilities provided in the DragonFly kernel to implement them.

Note: People often think that having a firewall between your internal network and the “Big Bad Internet” will solve all your security problems. It may help, but a poorly set up firewall system is more of a security risk than not having one at all. A firewall can add another layer of security to your systems, but it cannot stop a really determined cracker from penetrating your internal network. If you let internal security lapse because you believe your firewall to be impenetrable, you have just made the crackers job that much easier.

10.7.1 What Is a Firewall?

There are currently two distinct types of firewalls in common use on the Internet today. The first type is more properly called a *packet filtering router*. This type of firewall utilizes a multi-homed machine and a set of rules to determine whether to forward or block individual packets. A multi-homed machine is simply a device with multiple network interfaces. The second type, known as a *proxy server*, relies on daemons to provide authentication and to forward packets, possibly on a multi-homed machine which has kernel packet forwarding disabled.

Sometimes sites combine the two types of firewalls, so that only a certain machine (known as a *bastion host*) is allowed to send packets through a packet filtering router onto an internal network. Proxy services are run on the bastion host, which are generally more secure than normal authentication mechanisms.

DragonFly comes with a kernel packet filter (known as IPFW), which is what the rest of this section will concentrate on. Proxy servers can be built on DragonFly from third party software, but there is such a variety of proxy servers available that it would be impossible to cover them in this section.

10.7.1.1 Packet Filtering Routers

A router is a machine which forwards packets between two or more networks. A packet filtering router is programmed to compare each packet to a list of rules before deciding if it should be forwarded or not. Most modern IP routing software includes packet filtering functionality that defaults to forwarding all packets. To enable the filters, you need to define a set of rules.

To decide whether a packet should be passed on, the firewall looks through its set of rules for a rule which matches the contents of the packet’s headers. Once a match is found, the rule action is obeyed. The rule action could be to drop the packet, to forward the packet, or even to send an ICMP message back to the originator. Only the first match counts, as the rules are searched in order. Hence, the list of rules can be referred to as a “rule chain”.

The packet-matching criteria varies depending on the software used, but typically you can specify rules which depend on the source IP address of the packet, the destination IP address, the source port number, the destination port number (for protocols which support ports), or even the packet type (UDP, TCP, ICMP, etc).

10.7.1.2 Proxy Servers

Proxy servers are machines which have had the normal system daemons (**telnetd**, **ftpd**, etc) replaced with special servers. These servers are called *proxy servers*, as they normally only allow onward connections to be made. This

enables you to run (for example) a proxy **telnet** server on your firewall host, and people can **telnet** in to your firewall from the outside, go through some authentication mechanism, and then gain access to the internal network (alternatively, proxy servers can be used for signals coming from the internal network and heading out).

Proxy servers are normally more secure than normal servers, and often have a wider variety of authentication mechanisms available, including “one-shot” password systems so that even if someone manages to discover what password you used, they will not be able to use it to gain access to your systems as the password expires immediately after the first use. As they do not actually give users access to the host machine, it becomes a lot more difficult for someone to install backdoors around your security system.

Proxy servers often have ways of restricting access further, so that only certain hosts can gain access to the servers. Most will also allow the administrator to specify which users can talk to which destination machines. Again, what facilities are available depends largely on what proxy software you choose.

10.7.2 What Does IPFW Allow Me to Do?

IPFW, the software supplied with DragonFly, is a packet filtering and accounting system which resides in the kernel, and has a user-land control utility, `ipfw(8)`. Together, they allow you to define and query the rules used by the kernel in its routing decisions.

There are two related parts to IPFW. The firewall section performs packet filtering. There is also an IP accounting section which tracks usage of the router, based on rules similar to those used in the firewall section. This allows the administrator to monitor how much traffic the router is getting from a certain machine, or how much WWW traffic it is forwarding, for example.

As a result of the way that IPFW is designed, you can use IPFW on non-router machines to perform packet filtering on incoming and outgoing connections. This is a special case of the more general use of IPFW, and the same commands and techniques should be used in this situation.

10.7.3 Enabling IPFW on DragonFly

As the main part of the IPFW system lives in the kernel, you will need to add one or more options to your kernel configuration file, depending on what facilities you want, and recompile your kernel. See "Reconfiguring your Kernel" (Chapter 9) for more details on how to recompile your kernel.

Warning: IPFW defaults to a policy of `deny ip from any to any`. If you do not add other rules during startup to allow access, *you will lock yourself out* of the server upon rebooting into a firewall-enabled kernel. We suggest that you set `firewall_type=open` in your `/etc/rc.conf` file when first enabling this feature, then refining the firewall rules in `/etc/rc.firewall` after you have tested that the new kernel feature works properly. To be on the safe side, you may wish to consider performing the initial firewall configuration from the local console rather than via **ssh**. Another option is to build a kernel using both the `IPFWALL` and `IPFWALL_DEFAULT_TO_ACCEPT` options. This will change the default rule of IPFW to `allow ip from any to any` and avoid the possibility of a lockout.

There are currently four kernel configuration options relevant to IPFW:

```
options IPFWALL
```

Compiles into the kernel the code for packet filtering.

```
options IPFWALL_VERBOSE
```

Enables code to allow logging of packets through `syslogd(8)`. Without this option, even if you specify that packets should be logged in the filter rules, nothing will happen.

```
options IPFWALL_VERBOSE_LIMIT=10
```

Limits the number of packets logged through `syslogd(8)` on a per entry basis. You may wish to use this option in hostile environments in which you want to log firewall activity, but do not want to be open to a denial of service attack via `syslog` flooding.

When a chain entry reaches the packet limit specified, logging is turned off for that particular entry. To resume logging, you will need to reset the associated counter using the `ipfw(8)` utility:

```
# ipfw zero 4500
```

Where 4500 is the chain entry you wish to continue logging.

```
options IPFWALL_DEFAULT_TO_ACCEPT
```

This changes the default rule action from “deny” to “allow”. This avoids the possibility of locking yourself out if you happen to boot a kernel with `IPFWALL` support but have not configured your firewall yet. It is also very useful if you often use `ipfw(8)` as a filter for specific problems as they arise. Use with care though, as this opens up the firewall and changes the way it works.

10.7.4 Configuring IPFW

The configuration of the IPFW software is done through the `ipfw(8)` utility. The syntax for this command looks quite complicated, but it is relatively simple once you understand its structure.

There are currently four different command categories used by the utility: addition/deletion, listing, flushing, and clearing. Addition/deletion is used to build the rules that control how packets are accepted, rejected, and logged. Listing is used to examine the contents of your rule set (otherwise known as the chain) and packet counters (accounting). Flushing is used to remove all entries from the chain. Clearing is used to zero out one or more accounting entries.

10.7.4.1 Altering the IPFW Rules

The syntax for this form of the command is:

```
ipfw [-N] command [index] action [log] protocol addresses [options]
```

There is one valid flag when using this form of the command:

`-N`

Resolve addresses and service names in output.

The *command* given can be shortened to the shortest unique form. The valid *commands* are:

add

Add an entry to the firewall/accounting rule list

delete

Delete an entry from the firewall/accounting rule list

Previous versions of IPFW used separate firewall and accounting entries. The present version provides packet accounting with each firewall entry.

If an *index* value is supplied, it is used to place the entry at a specific point in the chain. Otherwise, the entry is placed at the end of the chain at an index 100 greater than the last chain entry (this does not include the default policy, rule 65535, deny).

The `log` option causes matching rules to be output to the system console if the kernel was compiled with `IPFIREWALL_VERBOSE`.

Valid *actions* are:

reject

Drop the packet, and send an ICMP host or port unreachable (as appropriate) packet to the source.

allow

Pass the packet on as normal. (aliases: `pass`, `permit`, and `accept`)

deny

Drop the packet. The source is not notified via an ICMP message (thus it appears that the packet never arrived at the destination).

count

Update packet counters but do not allow/deny the packet based on this rule. The search continues with the next chain entry.

Each *action* will be recognized by the shortest unambiguous prefix.

The *protocols* which can be specified are:

all

Matches any IP packet

icmp

Matches ICMP packets

tcp

Matches TCP packets

udp

Matches UDP packets

The *address* specification is:

from *address/mask* [*port*] to *address/mask* [*port*] [*via interface*]

You can only specify *port* in conjunction with *protocols* which support ports (UDP and TCP).

The *via* is optional and may specify the IP address or domain name of a local IP interface, or an interface name (e.g. *ed0*) to match only packets coming through this interface. Interface unit numbers can be specified with an optional wildcard. For example, *ppp** would match all kernel PPP interfaces.

The syntax used to specify an *address/mask* is:

address

or

address/mask-bits

or

address:mask-pattern

A valid hostname may be specified in place of the IP address. *mask-bits* is a decimal number representing how many bits in the address mask should be set. e.g. specifying *192.216.222.1/24* will create a mask which will allow any address in a class C subnet (in this case, *192.216.222*) to be matched. *mask-pattern* is an IP address which will be logically AND'ed with the address given. The keyword *any* may be used to specify "any IP address".

The port numbers to be blocked are specified as:

port [*,port* [*,port* [...]]]

to specify either a single port or a list of ports, or

port-port

to specify a range of ports. You may also combine a single range with a list, but the range must always be specified first.

The *options* available are:

frag

Matches if the packet is not the first fragment of the datagram.

in

Matches if the packet is on the way in.

out

Matches if the packet is on the way out.

ipoptions *spec*

Matches if the IP header contains the comma separated list of options specified in *spec*. The supported IP options are: *ssrr* (strict source route), *lssrr* (loose source route), *rr* (record packet route), and *ts* (time stamp). The absence of a particular option may be specified with a leading *!*.

established

Matches if the packet is part of an already established TCP connection (i.e. it has the RST or ACK bits set). You can optimize the performance of the firewall by placing *established* rules early in the chain.

setup

Matches if the packet is an attempt to establish a TCP connection (the SYN bit is set but the ACK bit is not).

tcpflags *flags*

Matches if the TCP header contains the comma separated list of *flags*. The supported flags are *fin*, *syn*, *rst*, *psh*, *ack*, and *urg*. The absence of a particular flag may be indicated by a leading *!*.

icmptypes *types*

Matches if the ICMP type is present in the list *types*. The list may be specified as any combination of ranges and/or individual types separated by commas. Commonly used ICMP types are: 0 echo reply (ping reply), 3 destination unreachable, 5 redirect, 8 echo request (ping request), and 11 time exceeded (used to indicate TTL expiration as with `traceroute(8)`).

10.7.4.2 Listing the IPFW Rules

The syntax for this form of the command is:

```
ipfw [-a] [-c] [-d] [-e] [-t] [-N] [-S] list
```

There are seven valid flags when using this form of the command:

-a

While listing, show counter values. This option is the only way to see accounting counters.

-c

List rules in compact form.

-d

Show dynamic rules in addition to static rules.

-e

If **-d** was specified, also show expired dynamic rules.

-t

Display the last match times for each chain entry. The time listing is incompatible with the input syntax used by the `ipfw(8)` utility.

-N

Attempt to resolve given addresses and service names.

-S

Show the set each rule belongs to. If this flag is not specified, disabled rules will not be listed.

10.7.4.3 Flushing the IPFW Rules

The syntax for flushing the chain is:

```
ipfw flush
```

This causes all entries in the firewall chain to be removed except the fixed default policy enforced by the kernel (index 65535). Use caution when flushing rules; the default deny policy will leave your system cut off from the network until allow entries are added to the chain.

10.7.4.4 Clearing the IPFW Packet Counters

The syntax for clearing one or more packet counters is:

```
ipfw zero [index]
```

When used without an *index* argument, all packet counters are cleared. If an *index* is supplied, the clearing operation only affects a specific chain entry.

10.7.5 Example Commands for ipfw

This command will deny all packets from the host `evil.crackers.org` to the telnet port of the host `nice.people.org`:

```
# ipfw add deny tcp from evil.crackers.org to nice.people.org 23
```

The next example denies and logs any TCP traffic from the entire `crackers.org` network (a class C) to the `nice.people.org` machine (any port).

```
# ipfw add deny log tcp from evil.crackers.org/24 to nice.people.org
```

If you do not want people sending X sessions to your internal network (a subnet of a class C), the following command will do the necessary filtering:


```
# ipfw add deny tcp from any to my.org/28 6000 setup
```

To see the accounting records:

```
# ipfw -a list
```

or in the short form

```
# ipfw -a 1
```

You can also see the last time a chain entry was matched with:

```
# ipfw -at 1
```

10.7.6 Building a Packet Filtering Firewall

Note: The following suggestions are just that: suggestions. The requirements of each firewall are different and we cannot tell you how to build a firewall to meet your particular requirements.

When initially setting up your firewall, unless you have a test bench setup where you can configure your firewall host in a controlled environment, it is strongly recommend you use the logging version of the commands and enable logging in the kernel. This will allow you to quickly identify problem areas and cure them without too much disruption. Even after the initial setup phase is complete, I recommend using the logging for ‘deny’ as it allows tracing of possible attacks and also modification of the firewall rules if your requirements alter.

Note: If you use the logging versions of the `accept` command, be aware that it can generate *large* amounts of log data. One log entry will be generated for every packet that passes through the firewall, so large FTP/http transfers, etc, will really slow the system down. It also increases the latencies on those packets as it requires more work to be done by the kernel before the packet can be passed on. **syslogd** will also start using up a lot more processor time as it logs all the extra data to disk, and it could quite easily fill the partition `/var/log` is located on.

You should enable your firewall from `/etc/rc.conf.local` or `/etc/rc.conf`. The associated manual page explains which knobs to fiddle and lists some preset firewall configurations. If you do not use a preset configuration, `ipfw list` will output the current ruleset into a file that you can pass to `rc.conf`. If you do not use `/etc/rc.conf.local` or `/etc/rc.conf` to enable your firewall, it is important to make sure your firewall is enabled before any IP interfaces are configured.

The next problem is what your firewall should actually *do*! This is largely dependent on what access to your network you want to allow from the outside, and how much access to the outside world you want to allow from the inside. Some general rules are:

- Block all incoming access to ports below 1024 for TCP. This is where most of the security sensitive services are, like finger, SMTP (mail) and telnet.

- Block *all* incoming UDP traffic. There are very few useful services that travel over UDP, and what useful traffic there is, is normally a security threat (e.g. Suns RPC and NFS protocols). This has its disadvantages also, since UDP is a connectionless protocol, denying incoming UDP traffic also blocks the replies to outgoing UDP traffic. This can cause a problem for people (on the inside) using external archie (prospero) servers. If you want to allow access to archie, you will have to allow packets coming from ports 191 and 1525 to any internal UDP port through the firewall. **ntp** is another service you may consider allowing through, which comes from port 123.
- Block traffic to port 6000 from the outside. Port 6000 is the port used for access to X11 servers, and can be a security threat (especially if people are in the habit of doing `xhost +` on their workstations). X11 can actually use a range of ports starting at 6000, the upper limit being how many X displays you can run on the machine. The upper limit as defined by RFC 1700 (Assigned Numbers) is 6063.
- Check what ports any internal servers use (e.g. SQL servers, etc). It is probably a good idea to block those as well, as they normally fall outside the 1-1024 range specified above.

Another checklist for firewall configuration is available from CERT at http://www.cert.org/tech_tips/packet_filtering.html

As stated above, these are only *guidelines*. You will have to decide what filter rules you want to use on your firewall yourself. We cannot accept ANY responsibility if someone breaks into your network, even if you follow the advice given above.

10.7.7 IPFW Overhead and Optimization

Many people want to know how much overhead IPFW adds to a system. The answer to this depends mostly on your rule set and processor speed. For most applications dealing with Ethernet and small rule sets, the answer is “negligible”. For those of you that need actual measurements to satisfy your curiosity, read on.

The following measurements were made using FreeBSD 2.2.5-STABLE on a 486-66. (While IPFW has changed slightly in later releases of DragonFly, it still performs with similar speed.) IPFW was modified to measure the time spent within the `ip_fw_chk` routine, displaying the results to the console every 1000 packets.

Two rule sets, each with 1000 rules, were tested. The first set was designed to demonstrate a worst case scenario by repeating the rule:

```
# ipfw add deny tcp from any to any 55555
```

This demonstrates a worst case scenario by causing most of IPFW’s packet check routine to be executed before finally deciding that the packet does not match the rule (by virtue of the port number). Following the 999th iteration of this rule was an `allow ip from any to any`.

The second set of rules were designed to abort the rule check quickly:

```
# ipfw add deny ip from 1.2.3.4 to 1.2.3.4
```

The non-matching source IP address for the above rule causes these rules to be skipped very quickly. As before, the 1000th rule was an `allow ip from any to any`.

The per-packet processing overhead in the former case was approximately 2.703 ms/packet, or roughly 2.7 microseconds per rule. Thus the theoretical packet processing limit with these rules is around 370 packets per second. Assuming 10 Mbps Ethernet and a ~1500 byte packet size, we would only be able to achieve 55.5% bandwidth utilization.

For the latter case each packet was processed in approximately 1.172 ms, or roughly 1.2 microseconds per rule. The theoretical packet processing limit here would be about 853 packets per second, which could consume 10 Mbps Ethernet bandwidth.

The excessive number of rules tested and the nature of those rules do not provide a real-world scenario -- they were used only to generate the timing information presented here. Here are a few things to keep in mind when building an efficient rule set:

- Place an *established* rule early on to handle the majority of TCP traffic. Do not put any `allow tcp` statements before this rule.
- Place heavily triggered rules earlier in the rule set than those rarely used (*without changing the permissiveness of the firewall*, of course). You can see which rules are used most often by examining the packet counting statistics with `ipfw -a 1`.

10.8 OpenSSL

OpenSSL (<http://www.openssl.org/>) provides a general-purpose cryptography library, as well as the Secure Sockets Layer v2/v3 (SSLv2/SSLv3) and Transport Layer Security v1 (TLSv1) network security protocols.

However, one of the algorithms (specifically IDEA) included in OpenSSL is protected by patents in the USA and elsewhere, and is not available for unrestricted use. IDEA is included in the OpenSSL sources in DragonFly, but it is not built by default. If you wish to use it, and you comply with the license terms, enable the `MAKE_IDEA` switch in `/etc/make.conf` and rebuild your sources using `make world`.

Today, the RSA algorithm is free for use in USA and other countries. In the past it was protected by a patent.

10.8.1 Source Code Installations

OpenSSL is part of the `src-crypto` and `src-secure` **CVSup** collections. See the **Obtaining DragonFly** section for more information about obtaining and updating DragonFly source code.

10.9 VPN over IPsec

Written by Nik Clayton.

Creating a VPN between two networks, separated by the Internet, using DragonFly gateways.

10.9.1 Understanding IPsec

Written by Hiten M. Pandya.

This section will guide you through the process of setting up IPsec, and to use it in an environment which consists of DragonFly and **Microsoft Windows 2000/XP** machines, to make them communicate securely. In order to set up IPsec, it is necessary that you are familiar with the concepts of building a custom kernel (see Chapter 9).

IPsec is a protocol which sits on top of the Internet Protocol (IP) layer. It allows two or more hosts to communicate in a secure manner (hence the name). The DragonFly IPsec “network stack” is based on the KAME (<http://www.kame.net/>) implementation, which has support for both protocol families, IPv4 and IPv6.

IPsec consists of two sub-protocols:

- *Encapsulated Security Payload (ESP)*, protects the IP packet data from third party interference, by encrypting the contents using symmetric cryptography algorithms (like Blowfish, 3DES).
- *Authentication Header (AH)*, protects the IP packet header from third party interference and spoofing, by computing a cryptographic checksum and hashing the IP packet header fields with a secure hashing function. This is then followed by an additional header that contains the hash, to allow the information in the packet to be authenticated.

ESP and AH can either be used together or separately, depending on the environment.

IPsec can either be used to directly encrypt the traffic between two hosts (known as *Transport Mode*); or to build “virtual tunnels” between two subnets, which could be used for secure communication between two corporate networks (known as *Tunnel Mode*). The latter is more commonly known as a *Virtual Private Network (VPN)*. The `ipsec(4)` manual page should be consulted for detailed information on the IPsec subsystem in DragonFly.

To add IPsec support to your kernel, add the following options to your kernel configuration file:

```
options IPSEC          #IP security
options IPSEC_ESP     #IP security (crypto; define w/ IPSEC)
```

If IPsec debugging support is desired, the following kernel option should also be added:

```
options IPSEC_DEBUG   #debug for IP security
```

10.9.2 The Problem

There’s no standard for what constitutes a VPN. VPNs can be implemented using a number of different technologies, each of which have their own strengths and weaknesses. This article presents a number of scenarios, and strategies for implementing a VPN for each scenario.

10.9.3 Scenario #1: Two networks, connected to the Internet, to behave as one

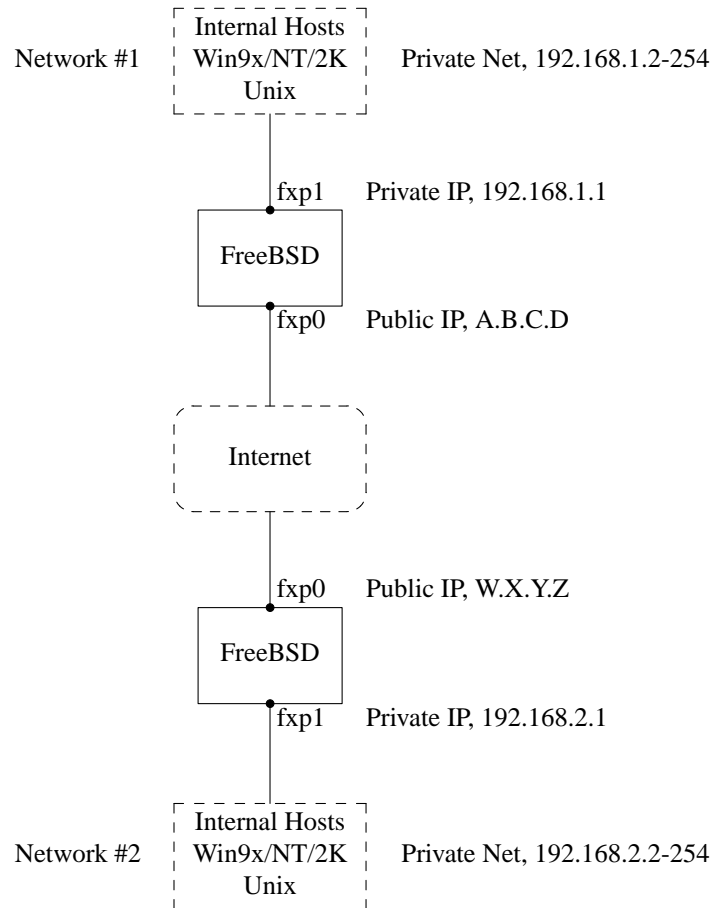
This is the scenario that caused me to first investigating VPNs. The premise is as follows:

- You have at least two sites
- Both sites are using IP internally
- Both sites are connected to the Internet, through a gateway that is running DragonFly.
- The gateway on each network has at least one public IP address.
- The internal addresses of the two networks can be public or private IP addresses, it doesn’t matter. You can be running NAT on the gateway machine if necessary.

- The internal IP addresses of the two networks *do not collide*. While I expect it is theoretically possible to use a combination of VPN technology and NAT to get this to work, I expect it to be a configuration nightmare.

If you find that you are trying to connect two networks, both of which, internally, use the same private IP address range (e.g., both of them use 192.168.1.x), then one of the networks will have to be renumbered.

The network topology might look something like this:



Notice the two public IP addresses. I'll use the letters to refer to them in the rest of this article. Anywhere you see those letters in this article, replace them with your own public IP addresses. Note also that internally, the two gateway machines have .1 IP addresses, and that the two networks have different private IP addresses (192.168.1.x and 192.168.2.x respectively). All the machines on the private networks have been configured to use the .1 machine as their default gateway.

The intention is that, from a network point of view, each network should view the machines on the other network as though they were directly attached the same router -- albeit a slightly slow router with an occasional tendency to drop packets.

This means that (for example), machine 192.168.1.20 should be able to run

```
ping 192.168.2.34
```

and have it work, transparently. Windows machines should be able to see the machines on the other network, browse file shares, and so on, in exactly the same way that they can browse machines on the local network.

And the whole thing has to be secure. This means that traffic between the two networks has to be encrypted.

Creating a VPN between these two networks is a multi-step process. The stages are as follows:

1. Create a “virtual” network link between the two networks, across the Internet. Test it, using tools like `ping(8)`, to make sure it works.
2. Apply security policies to ensure that traffic between the two networks is transparently encrypted and decrypted as necessary. Test this, using tools like `tcpdump(1)`, to ensure that traffic is encrypted.
3. Configure additional software on the DragonFly gateways, to allow Windows machines to see one another across the VPN.

10.9.3.1 Step 1: Creating and testing a “virtual” network link

Suppose that you were logged in to the gateway machine on network #1 (with public IP address `A.B.C.D`, private IP address `192.168.1.1`), and you ran `ping 192.168.2.1`, which is the private address of the machine with IP address `W.X.Y.Z`. What needs to happen in order for this to work?

1. The gateway machine needs to know how to reach `192.168.2.1`. In other words, it needs to have a route to `192.168.2.1`.
2. Private IP addresses, such as those in the `192.168.x` range are not supposed to appear on the Internet at large. Instead, each packet you send to `192.168.2.1` will need to be wrapped up inside another packet. This packet will need to appear to be from `A.B.C.D`, and it will have to be sent to `W.X.Y.Z`. This process is called *encapsulation*.
3. Once this packet arrives at `W.X.Y.Z` it will need to “unencapsulated”, and delivered to `192.168.2.1`.

You can think of this as requiring a “tunnel” between the two networks. The two “tunnel mouths” are the IP addresses `A.B.C.D` and `W.X.Y.Z`, and the tunnel must be told the addresses of the private IP addresses that will be allowed to pass through it. The tunnel is used to transfer traffic with private IP addresses across the public Internet.

This tunnel is created by using the generic interface, or `gif` devices on DragonFly. As you can imagine, the `gif` interface on each gateway host must be configured with four IP addresses; two for the public IP addresses, and two for the private IP addresses.

Support for the `gif` device must be compiled in to the DragonFly kernel on both machines. You can do this by adding the line:

```
pseudo-device gif
```

to the kernel configuration files on both machines, and then compile, install, and reboot as normal.

Configuring the tunnel is a two step process. First the tunnel must be told what the outside (or public) IP addresses are, using `gifconfig(8)`. Then the private IP addresses must be configured using `ifconfig(8)`.

On the gateway machine on network #1 you would run the following two commands to configure the tunnel.

```
gifconfig gif0 A.B.C.D W.X.Y.Z
ifconfig gif0 inet 192.168.1.1 192.168.2.1 netmask 0xffffffff
```

On the other gateway machine you run the same commands, but with the order of the IP addresses reversed.

```
gifconfig gif0 W.X.Y.Z A.B.C.D
```

```
ifconfig gif0 inet 192.168.2.1 192.168.1.1 netmask 0xffffffff
```

You can then run:

```
gifconfig gif0
```

to see the configuration. For example, on the network #1 gateway, you would see this:

```
# gifconfig gif0
gif0: flags=8011<UP,POINTTOPOINT,MULTICAST> mtu 1280
inet 192.168.1.1 --> 192.168.2.1 netmask 0xffffffff
physical address inet A.B.C.D --> W.X.Y.Z
```

As you can see, a tunnel has been created between the physical addresses A.B.C.D and W.X.Y.Z, and the traffic allowed through the tunnel is that between 192.168.1.1 and 192.168.2.1.

This will also have added an entry to the routing table on both machines, which you can examine with the command `netstat -rn`. This output is from the gateway host on network #1.

```
# netstat -rn
Routing tables

Internet:
Destination      Gateway          Flags    Refs    Use    Netif    Expire
...
192.168.2.1      192.168.1.1    UH        0        0     gif0
...
```

As the “Flags” value indicates, this is a host route, which means that each gateway knows how to reach the other gateway, but they do not know how to reach the rest of their respective networks. That problem will be fixed shortly.

It is likely that you are running a firewall on both machines. This will need to be circumvented for your VPN traffic. You might want to allow all traffic between both networks, or you might want to include firewall rules that protect both ends of the VPN from one another.

It greatly simplifies testing if you configure the firewall to allow all traffic through the VPN. You can always tighten things up later. If you are using `ipfw(8)` on the gateway machines then a command like

```
ipfw add 1 allow ip from any to any via gif0
```

will allow all traffic between the two end points of the VPN, without affecting your other firewall rules. Obviously you will need to run this command on both gateway hosts.

This is sufficient to allow each gateway machine to ping the other. On 192.168.1.1, you should be able to run

```
ping 192.168.2.1
```

and get a response, and you should be able to do the same thing on the other gateway machine.

However, you will not be able to reach internal machines on either network yet. This is because of the routing -- although the gateway machines know how to reach one another, they do not know how to reach the network behind each one.

To solve this problem you must add a static route on each gateway machine. The command to do this on the first gateway would be:

```
route add 192.168.2.0 192.168.2.1 netmask 0xffffffff00
```

This says “In order to reach the hosts on the network 192.168.2.0, send the packets to the host 192.168.2.1”. You will need to run a similar command on the other gateway, but with the 192.168.1.x addresses instead.

IP traffic from hosts on one network will now be able to reach hosts on the other network.

That has now created two thirds of a VPN between the two networks, in as much as it is “virtual” and it is a “network”. It is not private yet. You can test this using ping(8) and tcpdump(1). Log in to the gateway host and run

```
tcpdump dst host 192.168.2.1
```

In another log in session on the same host run

```
ping 192.168.2.1
```

You will see output that looks something like this:

```
16:10:24.018080 192.168.1.1 > 192.168.2.1: icmp: echo request
16:10:24.018109 192.168.1.1 > 192.168.2.1: icmp: echo reply
16:10:25.018814 192.168.1.1 > 192.168.2.1: icmp: echo request
16:10:25.018847 192.168.1.1 > 192.168.2.1: icmp: echo reply
16:10:26.028896 192.168.1.1 > 192.168.2.1: icmp: echo request
16:10:26.029112 192.168.1.1 > 192.168.2.1: icmp: echo reply
```

As you can see, the ICMP messages are going back and forth unencrypted. If you had used the `-s` parameter to `tcpdump(1)` to grab more bytes of data from the packets you would see more information.

Obviously this is unacceptable. The next section will discuss securing the link between the two networks so that it all traffic is automatically encrypted.

Summary:

- Configure both kernels with “pseudo-device gif”.
- Edit `/etc/rc.conf` on gateway host #1 and add the following lines (replacing IP addresses as necessary).

```
gifconfig_gif0="A.B.C.D W.X.Y.Z"
ifconfig_gif0="inet 192.168.1.1 192.168.2.1 netmask 0xffffffff"
static_routes="vpn"
route_vpn="192.168.2.0 192.168.2.1 netmask 0xffffffff00"
```

- Edit your firewall script (`/etc/rc.firewall`, or similar) on both hosts, and add


```
ipfw add 1 allow ip from any to any via gif0
```
- Make similar changes to `/etc/rc.conf` on gateway host #2, reversing the order of IP addresses.

10.9.3.2 Step 2: Securing the link

To secure the link we will be using IPsec. IPsec provides a mechanism for two hosts to agree on an encryption key, and to then use this key in order to encrypt data between the two hosts.

There are two areas of configuration to be considered here.

1. There must be a mechanism for two hosts to agree on the encryption mechanism to use. Once two hosts have agreed on this mechanism there is said to be a “security association” between them.
2. There must be a mechanism for specifying which traffic should be encrypted. Obviously, you don’t want to encrypt all your outgoing traffic -- you only want to encrypt the traffic that is part of the VPN. The rules that you put in place to determine what traffic will be encrypted are called “security policies”.

Security associations and security policies are both maintained by the kernel, and can be modified by userland programs. However, before you can do this you must configure the kernel to support IPsec and the Encapsulated Security Payload (ESP) protocol. This is done by configuring a kernel with:

```
options IPSEC
options IPSEC_ESP
```

and recompiling, reinstalling, and rebooting. As before you will need to do this to the kernels on both of the gateway hosts.

You have two choices when it comes to setting up security associations. You can configure them by hand between two hosts, which entails choosing the encryption algorithm, encryption keys, and so forth, or you can use daemons that implement the Internet Key Exchange protocol (IKE) to do this for you.

I recommend the latter. Apart from anything else, it is easier to set up.

Editing and displaying security policies is carried out using `setkey(8)`. By analogy, `setkey` is to the kernel’s security policy tables as `route(8)` is to the kernel’s routing tables. `setkey` can also display the current security associations, and to continue the analogy further, is akin to `netstat -r` in that respect.

There are a number of choices for daemons to manage security associations with DragonFly. This article will describe how to use one of these, `racoon`. `racoon` is in the FreeBSD ports collection, in the `security/` category, and is installed in the usual way.

`racoon` must be run on both gateway hosts. On each host it is configured with the IP address of the other end of the VPN, and a secret key (which you choose, and must be the same on both gateways).

The two daemons then contact one another, confirm that they are who they say they are (by using the secret key that you configured). The daemons then generate a new secret key, and use this to encrypt the traffic over the VPN. They periodically change this secret, so that even if an attacker were to crack one of the keys (which is as theoretically close to unfeasible as it gets) it won’t do them much good -- by the time they’ve cracked the key the two daemons have chosen another one.

`racoon`’s configuration is stored in `/${PREFIX}/etc/racoon`. You should find a configuration file there, which should not need to be changed too much. The other component of `racoon`’s configuration, which you will need to change, is the “pre-shared key”.

The default `racoon` configuration expects to find this in the file `/${PREFIX}/etc/racoon/psk.txt`. It is important to note that the pre-shared key is *not* the key that will be used to encrypt your traffic across the VPN link, it is simply a token that allows the key management daemons to trust one another.

`psk.txt` contains a line for each remote site you are dealing with. In this example, where there are two sites, each `psk.txt` file will contain one line (because each end of the VPN is only dealing with one other end).

On gateway host #1 this line should look like this:

```
W.X.Y.Z          secret
```

That is, the *public* IP address of the remote end, whitespace, and a text string that provides the secret. Obviously, you shouldn't use "secret" as your key -- the normal rules for choosing a password apply.

On gateway host #2 the line would look like this

```
A.B.C.D          secret
```

That is, the public IP address of the remote end, and the same secret key. `psk.txt` must be mode `0600` (i.e., only read/write to `root`) before `racoon` will run.

You must run `racoon` on both gateway machines. You will also need to add some firewall rules to allow the IKE traffic, which is carried over UDP to the ISAKMP (Internet Security Association Key Management Protocol) port. Again, this should be fairly early in your firewall ruleset.

```
ipfw add 1 allow udp from A.B.C.D to W.X.Y.Z isakmp
ipfw add 1 allow udp from W.X.Y.Z to A.B.C.D isakmp
```

Once `racoon` is running you can try pinging one gateway host from the other. The connection is still not encrypted, but `racoon` will then set up the security associations between the two hosts -- this might take a moment, and you may see this as a short delay before the ping commands start responding.

Once the security association has been set up you can view it using `setkey(8)`. Run

```
setkey -D
```

on either host to view the security association information.

That's one half of the problem. The other half is setting your security policies.


To create a sensible security policy, let's review what's been set up so far. This discussion holds for both ends of the link.

Each IP packet that you send out has a header that contains data about the packet. The header includes the IP addresses of both the source and destination. As we already know, private IP addresses, such as the `192.168.x.y` range are not supposed to appear on the public Internet. Instead, they must first be encapsulated inside another packet. This packet must have the public source and destination IP addresses substituted for the private addresses.

So if your outgoing packet started looking like this:

Src: 192.168.1.1 Dst: 192.168.2.1 <other header info>
<packet data>

Then it will be encapsulated inside another packet, looking something like this:



This encapsulation is carried out by the `gif` device. As you can see, the packet now has real IP addresses on the outside, and our original packet has been wrapped up as data inside the packet that will be put out on the Internet.

Obviously, we want all traffic between the VPNs to be encrypted. You might try putting this in to words, as:

“If a packet leaves from `A.B.C.D`, and it is destined for `W.X.Y.Z`, then encrypt it, using the necessary security associations.”

“If a packet arrives from `W.X.Y.Z`, and it is destined for `A.B.C.D`, then decrypt it, using the necessary security associations.”

That’s close, but not quite right. If you did this, all traffic to and from `W.X.Y.Z`, even traffic that was not part of the VPN, would be encrypted. That’s not quite what you want. The correct policy is as follows

“If a packet leaves from `A.B.C.D`, and that packet is encapsulating another packet, and it is destined for `W.X.Y.Z`, then encrypt it, using the necessary security associations.”

“If a packet arrives from `W.X.Y.Z`, and that packet is encapsulating another packet, and it is destined for `A.B.C.D`, then decrypt it, using the necessary security associations.”

A subtle change, but a necessary one.

Security policies are also set using `setkey(8)`. `setkey(8)` features a configuration language for defining the policy. You can either enter configuration instructions via `stdin`, or you can use the `-f` option to specify a filename that contains configuration instructions.

The configuration on gateway host #1 (which has the public IP address `A.B.C.D`) to force all outbound traffic to `W.X.Y.Z` to be encrypted is:

```
spdadd A.B.C.D/32 W.X.Y.Z/32 ipencap -P out ipsec esp/tunnel/A.B.C.D-W.X.Y.Z/require;
```

Put these commands in a file (e.g., `/etc/ipsec.conf`) and then run

```
# setkey -f /etc/ipsec.conf
```

`spdadd` tells `setkey(8)` that we want to add a rule to the secure policy database. The rest of this line specifies which packets will match this policy. `A.B.C.D/32` and `W.X.Y.Z/32` are the IP addresses and netmasks that identify the network or hosts that this policy will apply to. In this case, we want it to apply to traffic between these two hosts. `ipencap` tells the kernel that this policy should only apply to packets that encapsulate other packets. `-P out` says that this policy applies to outgoing packets, and `ipsec` says that the packet will be secured.

The second line specifies how this packet will be encrypted. `esp` is the protocol that will be used, while `tunnel` indicates that the packet will be further encapsulated in an IPsec packet. The repeated use of `A.B.C.D` and `W.X.Y.Z`

is used to select the security association to use, and the final `require` mandates that packets must be encrypted if they match this rule.

This rule only matches outgoing packets. You will need a similar rule to match incoming packets.

```
spdadd W.X.Y.Z/32 A.B.C.D/32 ipencap -P in ipsec esp/tunnel/W.X.Y.Z-A.B.C.D/require;
```

Note the `in` instead of `out` in this case, and the necessary reversal of the IP addresses.

The other gateway host (which has the public IP address `W.X.Y.Z`) will need similar rules.


```
spdadd W.X.Y.Z/32 A.B.C.D/32 ipencap -P out ipsec esp/tunnel/W.X.Y.Z-A.B.C.D/require;
spdadd A.B.C.D/32 W.X.Y.Z/32 ipencap -P in ipsec esp/tunnel/A.B.C.D-W.X.Y.Z/require;
```

Finally, you need to add firewall rules to allow ESP and IPENCAP packets back and forth. These rules will need to be added to both hosts.

```
ipfw add 1 allow esp from A.B.C.D to W.X.Y.Z
ipfw add 1 allow esp from W.X.Y.Z to A.B.C.D
ipfw add 1 allow ipencap from A.B.C.D to W.X.Y.Z
ipfw add 1 allow ipencap from W.X.Y.Z to A.B.C.D
```

Because the rules are symmetric you can use the same rules on each gateway host.

Outgoing packets will now look something like this:



When they are received by the far end of the VPN they will first be decrypted (using the security associations that have been negotiated by `racoon`). Then they will enter the `gif` interface, which will unwrap the second layer, until you are left with the innermost packet, which can then travel in to the inner network.

You can check the security using the same `ping(8)` test from earlier. First, log in to the `A.B.C.D` gateway machine, and run:

```
tcpdump dst host 192.168.2.1
```

In another log in session on the same host run

```
ping 192.168.2.1
```

This time you should see output like the following:

```
XXX tcpdump output
```

Now, as you can see, `tcpdump(1)` shows the ESP packets. If you try to examine them with the `-s` option you will see (apparently) gibberish, because of the encryption.

Congratulations. You have just set up a VPN between two remote sites.

Summary

- Configure both kernels with:

```
options IPSEC
options IPSEC_ESP
```

- Install `security/racoon`. Edit `/${PREFIX}/etc/racoon/psk.txt` on both gateway hosts, adding an entry for the remote host's IP address and a secret key that they both know. Make sure this file is mode 0600.
- Add the following lines to `/etc/rc.conf` on each host:

```
ipsec_enable="YES"
ipsec_file="/etc/ipsec.conf"
```

- Create an `/etc/ipsec.conf` on each host that contains the necessary `spdadd` lines. On gateway host #1 this would be:

```
spdadd A.B.C.D/32 W.X.Y.Z/32 ipencap -P out ipsec
    esp/tunnel/A.B.C.D-W.X.Y.Z/require;
spdadd W.X.Y.Z/32 A.B.C.D/32 ipencap -P in ipsec
    esp/tunnel/W.X.Y.Z-A.B.C.D/require;
```

On gateway host #2 this would be:

```
spdadd W.X.Y.Z/32 A.B.C.D/32 ipencap -P out ipsec
    esp/tunnel/W.X.Y.Z-A.B.C.D/require;
spdadd A.B.C.D/32 W.X.Y.Z/32 ipencap -P in ipsec
    esp/tunnel/A.B.C.D-W.X.Y.Z/require;
```

- Add firewall rules to allow IKE, ESP, and IPENCAP traffic to both hosts:

```
ipfw add 1 allow udp from A.B.C.D to W.X.Y.Z isakmp
ipfw add 1 allow udp from W.X.Y.Z to A.B.C.D isakmp
ipfw add 1 allow esp from A.B.C.D to W.X.Y.Z
ipfw add 1 allow esp from W.X.Y.Z to A.B.C.D
ipfw add 1 allow ipencap from A.B.C.D to W.X.Y.Z
ipfw add 1 allow ipencap from W.X.Y.Z to A.B.C.D
```

The previous two steps should suffice to get the VPN up and running. Machines on each network will be able to refer to one another using IP addresses, and all traffic across the link will be automatically and securely encrypted.

10.10 OpenSSH

Contributed by Chern Lee.

OpenSSH is a set of network connectivity tools used to access remote machines securely. It can be used as a direct replacement for `rlogin`, `rsh`, `rcp`, and `telnet`. Additionally, any other TCP/IP connections can be tunneled/forwarded securely through SSH. **OpenSSH** encrypts all traffic to effectively eliminate eavesdropping, connection hijacking, and other network-level attacks.

OpenSSH is maintained by the OpenBSD project, and is based upon SSH v1.2.12 with all the recent bug fixes and updates. It is compatible with both SSH protocols 1 and 2.

10.10.1 Advantages of Using OpenSSH

Normally, when using `telnet(1)` or `rlogin(1)`, data is sent over the network in a clear, un-encrypted form. Network sniffers anywhere in between the client and server can steal your user/password information or data transferred in your session. **OpenSSH** offers a variety of authentication and encryption methods to prevent this from happening.

10.10.2 Enabling sshd

Be sure to make the following addition to your `rc.conf` file:

```
sshd_enable="YES"
```

This will load `sshd(8)`, the daemon program for **OpenSSH**, the next time your system initializes. Alternatively, you can simply run directly the **sshd** daemon by typing `sshd` on the command line.

10.10.3 SSH Client

The `ssh(1)` utility works similarly to `rlogin(1)`.

```
# ssh user@example.com
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes
Host 'example.com' added to the list of known hosts.
user@example.com's password: *****
```

The login will continue just as it would have if a session was created using `rlogin` or `telnet`. SSH utilizes a key fingerprint system for verifying the authenticity of the server when the client connects. The user is prompted to enter `yes` only when connecting for the first time. Future attempts to login are all verified against the saved fingerprint key. The SSH client will alert you if the saved fingerprint differs from the received fingerprint on future login attempts. The fingerprints are saved in `~/.ssh/known_hosts`, or `~/.ssh/known_hosts2` for SSH v2 fingerprints.

By default, **OpenSSH** servers are configured to accept both SSH v1 and SSH v2 connections. The client, however, can choose between the two. Version 2 is known to be more robust and secure than its predecessor.

The `ssh(1)` command can be forced to use either protocol by passing it the `-1` or `-2` argument for v1 and v2, respectively.

10.10.4 Secure Copy

The `scp(1)` command works similarly to `rcp(1)`; it copies a file to or from a remote machine, except in a secure fashion.

```
# scp user@example.com:/COPYRIGHT COPYRIGHT
user@example.com's password: *****
COPYRIGHT          100% |*****| 4735
00:00
#
```

Since the fingerprint was already saved for this host in the previous example, it is verified when using `scp(1)` here.

The arguments passed to `scp(1)` are similar to `cp(1)`, with the file or files in the first argument, and the destination in the second. Since the file is fetched over the network, through SSH, one or more of the file arguments takes on the form `user@host:<path_to_remote_file>`.

10.10.5 Configuration

The system-wide configuration files for both the **OpenSSH** daemon and client reside within the `/etc/ssh` directory. `ssh_config` configures the client settings, while `sshd_config` configures the daemon.

Additionally, the `sshd_program` (`/usr/sbin/sshd` by default), and `sshd_flags rc.conf` options can provide more levels of configuration.

10.10.6 ssh-keygen

Instead of using passwords, `ssh-keygen(1)` can be used to generate RSA keys to authenticate a user:

```
% ssh-keygen -t rsa1
Initializing random number generator...
Generating p:  .++ (distance 66)
Generating q:  .....++ (distance 498)
Computing the keys...
Key generation complete.
Enter file in which to save the key (/home/user/.ssh/identity):
Enter passphrase:
Enter the same passphrase again:
Your identification has been saved in /home/user/.ssh/identity.
...
```

`ssh-keygen(1)` will create a public and private key pair for use in authentication. The private key is stored in `~/.ssh/identity`, whereas the public key is stored in `~/.ssh/identity.pub`. The public key must be placed in `~/.ssh/authorized_keys` of the remote machine in order for the setup to work.

This will allow connection to the remote machine based upon RSA authentication instead of passwords.

Note: The `-t rsa1` option will create RSA keys for use by SSH protocol version 1. If you want to use RSA keys with the SSH protocol version 2, you have to use the command `ssh-keygen -t rsa`.

If a passphrase is used in `ssh-keygen(1)`, the user will be prompted for a password each time in order to use the private key.

A SSH protocol version 2 DSA key can be created for the same purpose by using the `ssh-keygen -t dsa` command. This will create a public/private DSA key for use in SSH protocol version 2 sessions only. The public key is stored in `~/.ssh/id_dsa.pub`, while the private key is in `~/.ssh/id_dsa`.

DSA public keys are also placed in `~/.ssh/authorized_keys` on the remote machine.

`ssh-agent(1)` and `ssh-add(1)` are utilities used in managing multiple passworded private keys.

Warning: The various options and files can be different according to the **OpenSSH** version you have on your system, to avoid problems you should consult the `ssh-keygen(1)` manual page.

10.10.7 SSH Tunneling

OpenSSH has the ability to create a tunnel to encapsulate another protocol in an encrypted session.

The following command tells `ssh(1)` to create a tunnel for **telnet**:

```
% ssh -2 -N -f -L 5023:localhost:23 user@foo.example.com
%
```

The `ssh` command is used with the following options:

`-2`

Forces `ssh` to use version 2 of the protocol. (Do not use if you are working with older SSH servers)

`-N`

Indicates no command, or tunnel only. If omitted, `ssh` would initiate a normal session.

`-f`

Forces `ssh` to run in the background.

`-L`

Indicates a local tunnel in `localport:remotehost:remoteport` fashion.

```
user@foo.example.com
```

The remote SSH server.

An SSH tunnel works by creating a listen socket on `localhost` on the specified port. It then forwards any connection received on the local host/port via the SSH connection to the specified remote host and port.

In the example, port `5023` on `localhost` is being forwarded to port `23` on `localhost` of the remote machine. Since `23` is **telnet**, this would create a secure **telnet** session through an SSH tunnel.

This can be used to wrap any number of insecure TCP protocols such as SMTP, POP3, FTP, etc.

Example 10-1. Using SSH to Create a Secure Tunnel for SMTP

```
% ssh -2 -N -f -L 5025:localhost:25 user@mailserver.example.com
user@mailserver.example.com's password: *****
% telnet localhost 5025
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
220 mailserver.example.com ESMTTP
```

This can be used in conjunction with an `ssh-keygen(1)` and additional user accounts to create a more seamless/hassle-free SSH tunneling environment. Keys can be used in place of typing a password, and the tunnels can be run as a separate user.

10.10.7.1 Practical SSH Tunneling Examples*10.10.7.1.1 Secure Access of a POP3 Server*

At work, there is an SSH server that accepts connections from the outside. On the same office network resides a mail server running a POP3 server. The network, or network path between your home and office may or may not be completely trustable. Because of this, you need to check your e-mail in a secure manner. The solution is to create an SSH connection to your office's SSH server, and tunnel through to the mail server.

```
% ssh -2 -N -f -L 2110:mail.example.com:110 user@ssh-server.example.com
user@ssh-server.example.com's password: *****
```

When the tunnel is up and running, you can point your mail client to send POP3 requests to `localhost` port 2110. A connection here will be forwarded securely across the tunnel to `mail.example.com`.

10.10.7.1.2 Bypassing a Draconian Firewall

Some network administrators impose extremely draconian firewall rules, filtering not only incoming connections, but outgoing connections. You may be only given access to contact remote machines on ports 22 and 80 for SSH and web surfing.

You may wish to access another (perhaps non-work related) service, such as an Ogg Vorbis server to stream music. If this Ogg Vorbis server is streaming on some other port than 22 or 80, you will not be able to access it.

The solution is to create an SSH connection to a machine outside of your network's firewall, and use it to tunnel to the Ogg Vorbis server.

```
% ssh -2 -N -f -L 8888:music.example.com:8000 user@unfirewalled-system.example.org
user@unfirewalled-system.example.org's password: *****
```

Your streaming client can now be pointed to `localhost` port 8888, which will be forwarded over to `music.example.com` port 8000, successfully evading the firewall.

10.10.8 Further Reading

OpenSSH (<http://www.openssh.com/>)

ssh(1) scp(1) ssh-keygen(1) ssh-agent(1) ssh-add(1)

sshd(8) sftp-server(8)

Notes

1. Under DragonFly the standard login password may be up to 128 characters in length.

Chapter 11 Printing

Contributed by Sean Kelly. Restructured and updated by Jim Mock.

11.1 Synopsis

DragonFly can be used to print to a wide variety of printers, from the oldest impact printer to the latest laser printers, and everything in between, allowing you to produce high quality printed output from the applications you run.

DragonFly can also be configured to act as a print server on a network; in this capacity DragonFly can receive print jobs from a variety of other computers, including other DragonFly computers, Windows and Mac OS hosts.

DragonFly will ensure that one job at a time is printed, and can keep statistics on which users and machines are doing the most printing, produce “banner” pages showing who’s printout is who’s, and more.

After reading this chapter, you will know:

- How to configure the DragonFly print spooler.
- How to install print filters, to handle special print jobs differently, including converting incoming documents to print formats that your printers understand.
- How to enable header, or banner pages on your printout.
- How to print to printers connected to other computers.
- How to print to printers connected directly to the network.
- How to control printer restrictions, including limiting the size of print jobs, and preventing certain users from printing.
- How to keep printer statistics, and account for printer usage.
- How to troubleshoot printing problems.

Before reading this chapter, you should:

- Know how to configure and install a new kernel (Chapter 9).

11.2 Introduction

In order to use printers with DragonFly, you will need to set them up to work with the Berkeley line printer spooling system, also known as the **LPD** spooling system. It is the standard printer control system in DragonFly. This chapter introduces the **LPD** spooling system, often simply called **LPD**, and will guide you through its configuration.

If you are already familiar with **LPD** or another printer spooling system, you may wish to skip to section **Setting up the spooling system**.

LPD controls everything about a host’s printers. It is responsible for a number of things:

- It controls access to attached printers and printers attached to other hosts on the network.
- It enables users to submit files to be printed; these submissions are known as *jobs*.
- It prevents multiple users from accessing a printer at the same time by maintaining a *queue* for each printer.

- It can print *header pages* (also known as *banner* or *burst* pages) so users can easily find jobs they have printed in a stack of printouts.
- It takes care of communications parameters for printers connected on serial ports.
- It can send jobs over the network to a **LPD** spooler on another host.
- It can run special filters to format jobs to be printed for various printer languages or printer capabilities.
- It can account for printer usage.

Through a configuration file (`/etc/printcap`), and by providing the special filter programs, you can enable the **LPD** system to do all or some subset of the above for a great variety of printer hardware.

11.2.1 Why You Should Use the Spooler

If you are the sole user of your system, you may be wondering why you should bother with the spooler when you do not need access control, header pages, or printer accounting. While it is possible to enable direct access to a printer, you should use the spooler anyway since:

- **LPD** prints jobs in the background; you do not have to wait for data to be copied to the printer.
- **LPD** can conveniently run a job to be printed through filters to add date/time headers or convert a special file format (such as a TeX DVI file) into a format the printer will understand. You will not have to do these steps manually.
- Many free and commercial programs that provide a print feature usually expect to talk to the spooler on your system. By setting up the spooling system, you will more easily support other software you may later add or already have.

11.3 Basic Setup

To use printers with the **LPD** spooling system, you will need to set up both your printer hardware and the **LPD** software. This document describes two levels of setup:

- See section *Simple Printer Setup* to learn how to connect a printer, tell **LPD** how to communicate with it, and print plain text files to the printer.
- See section *Advanced Printer Setup* to find out how to print a variety of special file formats, to print header pages, to print across a network, to control access to printers, and to do printer accounting.

11.3.1 Simple Printer Setup

This section tells how to configure printer hardware and the **LPD** software to use the printer. It teaches the basics:

- Section *Hardware Setup* gives some hints on connecting the printer to a port on your computer.
- Section *Software Setup* shows how to set up the **LPD** spooler configuration file (`/etc/printcap`).

If you are setting up a printer that uses a network protocol to accept data to print instead of a serial or parallel interface, see *Printers With Networked Data Stream Interfaces*.

Although this section is called “Simple Printer Setup”, it is actually fairly complex. Getting the printer to work with your computer and the **LPD** spooler is the hardest part. The advanced options like header pages and accounting are fairly easy once you get the printer working.

11.3.1.1 Hardware Setup

This section tells about the various ways you can connect a printer to your PC. It talks about the kinds of ports and cables, and also the kernel configuration you may need to enable DragonFly to speak to the printer.

If you have already connected your printer and have successfully printed with it under another operating system, you can probably skip to section Software Setup.

11.3.1.1.1 Ports and Cables

Printers sold for use on PC's today generally come with one or more of the following three interfaces:

- *Serial* interfaces, also known as RS232C or RS232D, or COM ports, use a serial port on your computer to send data to the printer. Serial interfaces are common in the computer industry and cables are readily available and also easy to construct. Serial interfaces sometimes need special cables and might require you to configure somewhat complex communications options. Most PC serial ports have a maximum transmission rate of 115200 bps, which makes printing large graphic print jobs with them impractical.
- *Parallel* interfaces use a parallel port on your computer to send data to the printer. Parallel interfaces are common in the PC market and are faster than RS232 serial. Cables are readily available but more difficult to construct by hand. There are usually no communications options with parallel interfaces, making their configuration exceedingly simple.

Parallel interfaces are sometimes known as “Centronics” interfaces, named after the connector type on the printer.

- *USB* interfaces, named for the Universal Serial Bus, can run at even faster speeds than parallel or RS232 serial interfaces. Cables are simple and cheap. USB is superior to RS232 Serial and to Parallel for printing, but it is not as well supported under UNIX systems. A way to avoid this problem is to purchase a printer that has both a USB interface and a Parallel interface, as many printers do.

In general, Parallel interfaces usually offer just one-way communication (computer to printer) while serial and USB gives you two-way. Newer parallel ports (EPP and ECP) and printers can communicate in both directions under DragonFly when a IEEE1284 compliant cable is used.

Two-way communication to the printer over a parallel port is generally done one of two ways. The first method uses a custom built printer driver for DragonFly that speaks the proprietary language used by the printer. This is common with inkjet printers and can be used for reporting ink levels and other status information. The second method is used when the printer supports PostScript.

PostScript jobs are actually programs sent to the printer; they need not produce paper at all and may return results directly to the computer. PostScript also uses two-way communication to tell the computer about problems, such as errors in the PostScript program or paper jams. Your users may be appreciative of such information. Furthermore, the best way to do effective accounting with a PostScript printer requires two-way communication: you ask the printer for its page count (how many pages it has printed in its lifetime), then send the user's job, then ask again for its page count. Subtract the two values and you know how much paper to charge the user.

11.3.1.1.2 Parallel Ports

To hook up a printer using a parallel interface, connect the Centronics cable between the printer and the computer. The instructions that came with the printer, the computer, or both should give you complete guidance.

Remember which parallel port you used on the computer. The first parallel port is `/dev/ppc0` to DragonFly; the second is `/dev/ppc1`, and so on. The printer device name uses the same scheme: `/dev/lpt0` for the printer on the first parallel ports etc.

11.3.1.1.3 Serial Ports

To hook up a printer using a serial interface, connect the proper serial cable between the printer and the computer. The instructions that came with the printer, the computer, or both should give you complete guidance.

If you are unsure what the “proper serial cable” is, you may wish to try one of the following alternatives:

- A *modem* cable connects each pin of the connector on one end of the cable straight through to its corresponding pin of the connector on the other end. This type of cable is also known as a “DTE-to-DCE” cable.
- A *null-modem* cable connects some pins straight through, swaps others (send data to receive data, for example), and shorts some internally in each connector hood. This type of cable is also known as a “DTE-to-DTE” cable.
- A *serial printer* cable, required for some unusual printers, is like the null-modem cable, but sends some signals to their counterparts instead of being internally shorted.

You should also set up the communications parameters for the printer, usually through front-panel controls or DIP switches on the printer. Choose the highest `bps` (bits per second, sometimes *baud rate*) rate that both your computer and the printer can support. Choose 7 or 8 data bits; none, even, or odd parity; and 1 or 2 stop bits. Also choose a flow control protocol: either none, or XON/XOFF (also known as “in-band” or “software”) flow control. Remember these settings for the software configuration that follows.

11.3.1.2 Software Setup

This section describes the software setup necessary to print with the **LPD** spooling system in DragonFly.

Here is an outline of the steps involved:

1. Configure your kernel, if necessary, for the port you are using for the printer; section **Kernel Configuration** tells you what you need to do.
2. Set the communications mode for the parallel port, if you are using a parallel port; section **Setting the Communication Mode for the Parallel Port** gives details.
3. Test if the operating system can send data to the printer. Section **Checking Printer Communications** gives some suggestions on how to do this.
4. Set up **LPD** for the printer by modifying the file `/etc/printcap`. You will find out how to do this later in this chapter.

11.3.1.2.1 Kernel Configuration

The operating system kernel is compiled to work with a specific set of devices. The serial or parallel interface for your printer is a part of that set. Therefore, it might be necessary to add support for an additional serial or parallel port if your kernel is not already configured for one.

To find out if the kernel you are currently using supports a serial interface, type:

```
# grep sioN /var/run/dmesg.boot
```

Where *N* is the number of the serial port, starting from zero. If you see output similar to the following:

```
sio2 at port 0x3e8-0x3ef irq 5 on isa
sio2: type 16550A
```

then the kernel supports the port.

To find out if the kernel supports a parallel interface, type:

```
# grep ppcN /var/run/dmesg.boot
```

Where *N* is the number of the parallel port, starting from zero. If you see output similar to the following:

```
ppc0: <Parallel port> at port 0x378-0x37f irq 7 on isa0
ppc0: SMC-like chipset (ECP/EPP/PS2/NIBBLE) in COMPATIBLE mode
ppc0: FIFO with 16/16/8 bytes threshold
```

then the kernel supports the port.

You might have to reconfigure your kernel in order for the operating system to recognize and use the parallel or serial port you are using for the printer.

To add support for a serial port, see the section on kernel configuration. To add support for a parallel port, see that section *and* the section that follows.

11.3.1.3 Adding /dev Entries for the Ports

Even though the kernel may support communication along a serial or parallel port, you will still need a software interface through which programs running on the system can send and receive data. That is what entries in the /dev directory are for.

To add a /dev entry for a port:

1. Become `root` with the `su(1)` command. Enter the `root` password when prompted.
2. Change to the /dev directory:

```
# cd /dev
```

3. Type:

```
# ./MAKEDEV port
```

Where *port* is the device entry for the port you want to make. Use `lpt0` for the printer on the first parallel port, `lpt1` for the printer on the second port, and so on; use `ttyd0` for the first serial port, `ttyd1` for the second, and so on.

4. Type:

```
# ls -l port
```

to make sure the device entry got created.

11.3.1.3.1 Setting the Communication Mode for the Parallel Port

When you are using the parallel interface, you can choose whether DragonFly should use interrupt-driven or polled communication with the printer. The generic printer device driver (`lpt(4)`) on DragonFly uses the `ppbus(4)` system, which controls the port chipset with the `ppc(4)` driver.

- The *interrupt-driven* method is the default with the GENERIC kernel. With this method, the operating system uses an IRQ line to determine when the printer is ready for data.
- The *polled* method directs the operating system to repeatedly ask the printer if it is ready for more data. When it responds ready, the kernel sends more data.

The interrupt-driven method is usually somewhat faster but uses up a precious IRQ line. Some newer HP printers are claimed not to work correctly in interrupt mode, apparently due to some (not yet exactly understood) timing problem. These printers need polled mode. You should use whichever one works. Some printers will work in both modes, but are painfully slow in interrupt mode.

You can set the communications mode in two ways: by configuring the kernel or by using the `lptcontrol(8)` program.

To set the communications mode by configuring the kernel:

1. Edit your kernel configuration file. Look for an `ppc0` entry. If you are setting up the second parallel port, use `ppc1` instead. Use `ppc2` for the third port, and so on.
 - If you want interrupt-driven mode, add the `irq` specifier:


```
device ppc0 at isa? irq N
```

 Where *N* is the IRQ number for your computer's parallel port.
 - If you want polled mode, do not add the `irq` specifier. Use the following line in your kernel configuration file:


```
device ppc0 at isa?
```
2. Save the file. Then configure, build, and install the kernel, then reboot. See kernel configuration for more details.

To set the communications mode with `lptcontrol(8)`:

1. Type:

```
# lptcontrol -i -d /dev/lptN
```

to set interrupt-driven mode for `lptN`.

2. Type:

```
# lptcontrol -p -d /dev/lptN
```

to set polled-mode for `lptN`.

You could put these commands in your `/etc/rc.local` file to set the mode each time your system boots. See `lptcontrol(8)` for more information.

11.3.1.3.2 Checking Printer Communications

Before proceeding to configure the spooling system, you should make sure the operating system can successfully send data to your printer. It is a lot easier to debug printer communication and the spooling system separately.

To test the printer, we will send some text to it. For printers that can immediately print characters sent to them, the program `lptest(1)` is perfect: it generates all 96 printable ASCII characters in 96 lines.

For a PostScript (or other language-based) printer, we will need a more sophisticated test. A small PostScript program, such as the following, will suffice:

```
%!PS
100 100 moveto 300 300 lineto stroke
310 310 moveto /Helvetica findfont 12 scalefont setfont
(Is this thing working?) show
showpage
```

The above PostScript code can be placed into a file and used as shown in the examples appearing in the following sections.

Note: When this document refers to a printer language, it is assuming a language like PostScript, and not Hewlett Packard's PCL. Although PCL has great functionality, you can intermingle plain text with its escape sequences. PostScript cannot directly print plain text, and that is the kind of printer language for which we must make special accommodations.

11.3.1.3.2.1 Checking a Parallel Printer

This section tells you how to check if DragonFly can communicate with a printer connected to a parallel port.

To test a printer on a parallel port:

1. Become `root` with `su(1)`.
2. Send data to the printer.
 - If the printer can print plain text, then use `lptest(1)`. Type:


```
# lptest > /dev/lptN
```

 Where *N* is the number of the parallel port, starting from zero.
 - If the printer understands PostScript or other printer language, then send a small program to the printer. Type:


```
# cat > /dev/lptN
```

 Then, line by line, type the program *carefully* as you cannot edit a line once you have pressed `RETURN` or `ENTER`. When you have finished entering the program, press `CONTROL+D`, or whatever your end of file key is.
 Alternatively, you can put the program in a file and type:


```
# cat file > /dev/lptN
```

 Where *file* is the name of the file containing the program you want to send to the printer.

You should see something print. Do not worry if the text does not look right; we will fix such things later.

11.3.1.3.2.2 Checking a Serial Printer

This section tells you how to check if DragonFly can communicate with a printer on a serial port.

To test a printer on a serial port:

1. Become root with `su(1)`.
2. Edit the file `/etc/remote`. Add the following entry:

```
printer:dv=/dev/port:br#bps-rate:pa=parity
```

Where `port` is the device entry for the serial port (`tttyd0`, `tttyd1`, etc.), `bps-rate` is the bits-per-second rate at which the printer communicates, and `parity` is the parity required by the printer (either even, odd, none, or zero).

Here is a sample entry for a printer connected via a serial line to the third serial port at 19200 bps with no parity:

```
printer:dv=/dev/ttyd2:br#19200:pa=none
```

3. Connect to the printer with `tip(1)`. Type:

```
# tip printer
```

If this step does not work, edit the file `/etc/remote` again and try using `/dev/cuaaN` instead of `/dev/ttydN`.

4. Send data to the printer.
 - If the printer can print plain text, then use `lptest(1)`. Type:

```
% $lptest
```

- If the printer understands PostScript or other printer language, then send a small program to the printer. Type the program, line by line, *very carefully* as backspacing or other editing keys may be significant to the printer. You may also need to type a special end-of-file key for the printer so it knows it received the whole program. For PostScript printers, press `CONTROL+D`.

Alternatively, you can put the program in a file and type:

```
% >file
```

Where `file` is the name of the file containing the program. After `tip(1)` sends the file, press any required end-of-file key.

You should see something print. Do not worry if the text does not look right; we will fix that later.

11.3.1.4 Enabling the Spooler: the `/etc/printcap` File

At this point, your printer should be hooked up, your kernel configured to communicate with it (if necessary), and you have been able to send some simple data to the printer. Now, we are ready to configure **LPD** to control access to your printer.

You configure **LPD** by editing the file `/etc/printcap`. The **LPD** spooling system reads this file each time the spooler is used, so updates to the file take immediate effect.

The format of the `printcap(5)` file is straightforward. Use your favorite text editor to make changes to `/etc/printcap`. The format is identical to other capability files like `/usr/share/misc/termcap` and `/etc/remote`. For complete information about the format, see the `cgetent(3)`.

The simple spooler configuration consists of the following steps:

1. Pick a name (and a few convenient aliases) for the printer, and put them in the `/etc/printcap` file; see the Naming the Printer section for more information on naming.
2. Turn off header pages (which are on by default) by inserting the `sh` capability; see the Suppressing Header Pages section for more information.
3. Make a spooling directory, and specify its location with the `sd` capability; see the Making the Spooling Directory section for more information.
4. Set the `/dev` entry to use for the printer, and note it in `/etc/printcap` with the `lp` capability; see the Identifying the Printer Device for more information. Also, if the printer is on a serial port, set up the communication parameters with the `ms#` capability which is discussed in the Configuring Spooler Communications Parameters section.
5. Install a plain text input filter; see the Installing the Text Filter section for details.
6. Test the setup by printing something with the `lpr(1)` command. More details are available in the Trying It Out and Troubleshooting sections.

Note: Language-based printers, such as PostScript printers, cannot directly print plain text. The simple setup outlined above and described in the following sections assumes that if you are installing such a printer you will print only files that the printer can understand.

Users often expect that they can print plain text to any of the printers installed on your system. Programs that interface to **LPD** to do their printing usually make the same assumption. If you are installing such a printer and want to be able to print jobs in the printer language *and* print plain text jobs, you are strongly urged to add an additional step to the simple setup outlined above: install an automatic plain-text-to-PostScript (or other printer language) conversion program. The section entitled Accommodating Plain Text Jobs on PostScript Printers tells how to do this.

11.3.1.4.1 Naming the Printer

The first (easy) step is to pick a name for your printer. It really does not matter whether you choose functional or whimsical names since you can also provide a number of aliases for the printer.

At least one of the printers specified in the `/etc/printcap` should have the alias `lp`. This is the default printer's name. If users do not have the `PRINTER` environment variable nor specify a printer name on the command line of any of the **LPD** commands, then `lp` will be the default printer they get to use.

Also, it is common practice to make the last alias for a printer be a full description of the printer, including make and model.

Once you have picked a name and some common aliases, put them in the `/etc/printcap` file. The name of the printer should start in the leftmost column. Separate each alias with a vertical bar and put a colon after the last alias.

In the following example, we start with a skeletal `/etc/printcap` that defines two printers (a Diablo 630 line printer and a Panasonic KX-P4455 PostScript laser printer):

```
#
# /etc/printcap for host rose
#
rattan|line|diablo|lp|Diablo 630 Line Printer:
```

```
bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:
```

In this example, the first printer is named `rattan` and has as aliases `line`, `diablo`, `lp`, and `Diablo 630 Line Printer`. Since it has the alias `lp`, it is also the default printer. The second is named `bamboo`, and has as aliases `ps`, `PS`, `S`, `panasonic`, and `Panasonic KX-P4455 PostScript v51.4`.

11.3.1.4.2 Suppressing Header Pages

The **LPD** spooling system will by default print a *header page* for each job. The header page contains the user name who requested the job, the host from which the job came, and the name of the job, in nice large letters. Unfortunately, all this extra text gets in the way of debugging the simple printer setup, so we will suppress header pages.

To suppress header pages, add the `sh` capability to the entry for the printer in `/etc/printcap`. Here is an example `/etc/printcap` with `sh` added:

```
#
# /etc/printcap for host rose - no header pages anywhere
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:
```

Note how we used the correct format: the first line starts in the leftmost column, and subsequent lines are indented with a single TAB. Every line in an entry except the last ends in a backslash character.

11.3.1.4.3 Making the Spooling Directory

The next step in the simple spooler setup is to make a *spooling directory*, a directory where print jobs reside until they are printed, and where a number of other spooler support files live.

Because of the variable nature of spooling directories, it is customary to put these directories under `/var/spool`. It is not necessary to backup the contents of spooling directories, either. Recreating them is as simple as running `mkdir(1)`.

It is also customary to make the directory with a name that is identical to the name of the printer, as shown below:

```
# mkdir /var/spool/printer-name
```

However, if you have a lot of printers on your network, you might want to put the spooling directories under a single directory that you reserve just for printing with **LPD**. We will do this for our two example printers `rattan` and `bamboo`:

```
# mkdir /var/spool/lpd
# mkdir /var/spool/lpd/rattan
# mkdir /var/spool/lpd/bamboo
```

Note: If you are concerned about the privacy of jobs that users print, you might want to protect the spooling directory so it is not publicly accessible. Spooling directories should be owned and be readable, writable, and searchable by user daemon and group daemon, and no one else. We will do this for our example printers:

```
# chown daemon:daemon /var/spool/lpd/rattan
# chown daemon:daemon /var/spool/lpd/bamboo
# chmod 770 /var/spool/lpd/rattan
# chmod 770 /var/spool/lpd/bamboo
```

Finally, you need to tell **LPD** about these directories using the `/etc/printcap` file. You specify the pathname of the spooling directory with the `sd` capability:

```
#
# /etc/printcap for host rose - added spooling directories
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:sd=/var/spool/lpd/rattan:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:
```

Note that the name of the printer starts in the first column but all other entries describing the printer should be indented with a tab and each line escaped with a backslash.

If you do not specify a spooling directory with `sd`, the spooling system will use `/var/spool/lpd` as a default.

11.3.1.4.4 Identifying the Printer Device

In the Adding `/dev` Entries for the Ports section, we identified which entry in the `/dev` directory DragonFly will use to communicate with the printer. Now, we tell **LPD** that information. When the spooling system has a job to print, it will open the specified device on behalf of the filter program (which is responsible for passing data to the printer).

List the `/dev` entry pathname in the `/etc/printcap` file using the `lp` capability.

In our running example, let us assume that `rattan` is on the first parallel port, and `bamboo` is on a sixth serial port; here are the additions to `/etc/printcap`:

```
#
# /etc/printcap for host rose - identified what devices to use
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:sd=/var/spool/lpd/rattan:\
    :lp=/dev/lpt0:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:\
    :lp=/dev/ttyd5:
```

If you do not specify the `lp` capability for a printer in your `/etc/printcap` file, **LPD** uses `/dev/lp` as a default. `/dev/lp` currently does not exist in DragonFly.

If the printer you are installing is connected to a parallel port, skip to the section entitled, Installing the Text Filter. Otherwise, be sure to follow the instructions in the next section.

11.3.1.4.5 Configuring Spooler Communication Parameters

For printers on serial ports, **LPD** can set up the bps rate, parity, and other serial communication parameters on behalf of the filter program that sends data to the printer. This is advantageous since:

- It lets you try different communication parameters by simply editing the `/etc/printcap` file; you do not have to recompile the filter program.
- It enables the spooling system to use the same filter program for multiple printers which may have different serial communication settings.

The following `/etc/printcap` capabilities control serial communication parameters of the device listed in the `lp` capability:

`br#bps-rate`

Sets the communications speed of the device to `bps-rate`, where `bps-rate` can be 50, 75, 110, 134, 150, 200, 300, 600, 1200, 1800, 2400, 4800, 9600, 19200, 38400, 57600, or 115200 bits-per-second.

`ms#stty-mode`

Sets the options for the terminal device after opening the device. `stty(1)` explains the available options.

When **LPD** opens the device specified by the `lp` capability, it sets the characteristics of the device to those specified with the `ms#` capability. Of particular interest will be the `parenb`, `parodd`, `cs5`, `cs6`, `cs7`, `cs8`, `cstopb`, `crtcts`, and `ixon` modes, which are explained in the `stty(1)` manual page.

Let us add to our example printer on the sixth serial port. We will set the bps rate to 38400. For the mode, we will set no parity with `-parenb`, 8-bit characters with `cs8`, no modem control with `cllocal` and hardware flow control with `crtcts`:

```
bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:\
    :lp=/dev/ttyd5:ms#-parenb cs8 cllocal crtcts:
```

11.3.1.4.6 Installing the Text Filter

We are now ready to tell **LPD** what text filter to use to send jobs to the printer. A *text filter*, also known as an *input filter*, is a program that **LPD** runs when it has a job to print. When **LPD** runs the text filter for a printer, it sets the filter's standard input to the job to print, and its standard output to the printer device specified with the `lp` capability. The filter is expected to read the job from standard input, perform any necessary translation for the printer, and write the results to standard output, which will get printed. For more information on the text filter, see the Filters section.

For our simple printer setup, the text filter can be a small shell script that just executes `/bin/cat` to send the job to the printer. DragonFly comes with another filter called `lpf` that handles backspacing and underlining for printers that might not deal with such character streams well. And, of course, you can use any other filter program you want. The filter `lpf` is described in detail in section entitled `lpf: a Text Filter`.

First, let us make the shell script `/usr/local/libexec/if-simple` be a simple text filter. Put the following text into that file with your favorite text editor:

```
#!/bin/sh
#
# if-simple - Simple text input filter for lpd
# Installed in /usr/local/libexec/if-simple
#
# Simply copies stdin to stdout. Ignores all filter arguments.

/bin/cat && exit 0
exit 2
```

Make the file executable:

```
# chmod 555 /usr/local/libexec/if-simple
```

And then tell LPD to use it by specifying it with the `if` capability in `/etc/printcap`. We will add it to the two printers we have so far in the example `/etc/printcap`:

```
#
# /etc/printcap for host rose - added text filter
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:sd=/var/spool/lpd/rattan:\ :lp=/dev/lpt0:\
    :if=/usr/local/libexec/if-simple:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:\
    :lp=/dev/ttyd5:ms#-parenb cs8 clocal crtscts:\
    :if=/usr/local/libexec/if-simple:
```

11.3.1.4.7 Turn on **LPD**

`lpd(8)` is run from `/etc/rc`, controlled by the `lpd_enable` variable. This variable defaults to `NO`. If you have not done so already, add the line:

```
lpd_enable="YES"
```

to `/etc/rc.conf`, and then either restart your machine, or just run `lpd(8)`.

```
# lpd
```

11.3.1.4.8 Trying It Out

You have reached the end of the simple **LPD** setup. Unfortunately, congratulations are not quite yet in order, since we still have to test the setup and correct any problems. To test the setup, try printing something. To print with the **LPD** system, you use the command `lpr(1)`, which submits a job for printing.

You can combine `lpr(1)` with the `lptest(1)` program, introduced in section [Checking Printer Communications](#) to generate some test text.

To test the simple **LPD** setup:

Type:

```
# lptest 20 5 | lpr -Pprinter-name
```

Where *printer-name* is the name of a printer (or an alias) specified in `/etc/printcap`. To test the default printer, type `lpr(1)` without any `-P` argument. Again, if you are testing a printer that expects PostScript, send a PostScript program in that language instead of using `lptest(1)`. You can do so by putting the program in a file and typing `lpr file`.

For a PostScript printer, you should get the results of the program. If you are using `lptest(1)`, then your results should look like the following:

```
! "#$%&'()*+,-./01234
"#$%&'()*+,-./012345
#$%&'()*+,-./0123456
$%&'()*+,-./01234567
%&'()*+,-./012345678
```

To further test the printer, try downloading larger programs (for language-based printers) or running `lptest(1)` with different arguments. For example, `lptest 80 60` will produce 60 lines of 80 characters each.

If the printer did not work, see the Troubleshooting section.

11.4 Advanced Printer Setup

This section describes filters for printing specially formatted files, header pages, printing across networks, and restricting and accounting for printer usage.

11.4.1 Filters

Although **LPD** handles network protocols, queuing, access control, and other aspects of printing, most of the *real* work happens in the *filters*. Filters are programs that communicate with the printer and handle its device dependencies and special requirements. In the simple printer setup, we installed a plain text filter—an extremely simple one that should work with most printers (section [Installing the Text Filter](#)).

However, in order to take advantage of format conversion, printer accounting, specific printer quirks, and so on, you should understand how filters work. It will ultimately be the filter's responsibility to handle these aspects. And the bad news is that most of the time *you* have to provide filters yourself. The good news is that many are generally available; when they are not, they are usually easy to write.

Also, DragonFly comes with one, `/usr/libexec/lpr/lpf`, that works with many printers that can print plain text. (It handles backspacing and tabs in the file, and does accounting, but that is about all it does.) There are also several filters and filter components in `pkgsrc`.

Here is what you will find in this section:

- Section **How Filters Work**, tries to give an overview of a filter's role in the printing process. You should read this section to get an understanding of what is happening "under the hood" when **LPD** uses filters. This knowledge could help you anticipate and debug problems you might encounter as you install more and more filters on each of your printers.
- **LPD** expects every printer to be able to print plain text by default. This presents a problem for PostScript (or other language-based printers) which cannot directly print plain text. Section **Accommodating Plain Text Jobs on PostScript Printers** tells you what you should do to overcome this problem. You should read this section if you have a PostScript printer.
- PostScript is a popular output format for many programs. Even some people (myself included) write PostScript code directly. But PostScript printers are expensive. Section **Simulating PostScript on Non PostScript Printers** tells how you can further modify a printer's text filter to accept and print PostScript data on a *non PostScript* printer. You should read this section if you do not have a PostScript printer.
- Section **Conversion Filters** tells about a way you can automate the conversion of specific file formats, such as graphic or typesetting data, into formats your printer can understand. After reading this section, you should be able to set up your printers such that users can type `lpr -t` to print troff data, or `lpr -d` to print TeX DVI data, or `lpr -v` to print raster image data, and so forth. I recommend reading this section.
- Section **Output Filters** tells all about a not often used feature of **LPD**: output filters. Unless you are printing header pages (see **Header Pages**), you can probably skip that section altogether.
- Section **lpf: a Text Filter** describes `lpf`, a fairly complete if simple text filter for line printers (and laser printers that act like line printers) that comes with DragonFly. If you need a quick way to get printer accounting working for plain text, or if you have a printer which emits smoke when it sees backspace characters, you should definitely consider `lpf`.

11.4.1.1 How Filters Work

As mentioned before, a filter is an executable program started by **LPD** to handle the device-dependent part of communicating with the printer.

When **LPD** wants to print a file in a job, it starts a filter program. It sets the filter's standard input to the file to print, its standard output to the printer, and its standard error to the error logging file (specified in the `lf` capability in `/etc/printcap`, or `/dev/console` by default).

Which filter **LPD** starts and the filter's arguments depend on what is listed in the `/etc/printcap` file and what arguments the user specified for the job on the `lpr(1)` command line. For example, if the user typed `lpr -t`, **LPD** would start the troff filter, listed in the `tf` capability for the destination printer. If the user wanted to print plain text, it would start the `if` filter (this is mostly true: see **Output Filters** for details).

There are three kinds of filters you can specify in `/etc/printcap`:

- The *text filter*, confusingly called the *input filter* in **LPD** documentation, handles regular text printing. Think of it as the default filter. **LPD** expects every printer to be able to print plain text by default, and it is the text filter's job to make sure backspaces, tabs, or other special characters do not confuse the printer. If you are in an environment where you have to account for printer usage, the text filter must also account for pages printed, usually by counting the number of lines printed and comparing that to the number of lines per page the printer supports. The text filter is started with the following argument list:

```
filter-name [-c] -wwidth -llength -iindent -n login -h host acct-file
```

where

`-c`

appears if the job is submitted with `lpr -l`

`width`

is the value from the `pw` (page width) capability specified in `/etc/printcap`, default 132

`length`

is the value from the `pl` (page length) capability, default 66

`indent`

is the amount of the indentation from `lpr -i`, default 0

`login`

is the account name of the user printing the file

`host`

is the host name from which the job was submitted

`acct-file`

is the name of the accounting file from the `af` capability.

- A *conversion filter* converts a specific file format into one the printer can render onto paper. For example, ditroff typesetting data cannot be directly printed, but you can install a conversion filter for ditroff files to convert the ditroff data into a form the printer can digest and print. Section Conversion Filters tells all about them. Conversion filters also need to do accounting, if you need printer accounting. Conversion filters are started with the following arguments:

```
filter-name -xpixel-width -ypixel-height -n login -h host acct-file
```

where `pixel-width` is the value from the `px` capability (default 0) and `pixel-height` is the value from the `py` capability (default 0).

- The *output filter* is used only if there is no text filter, or if header pages are enabled. In my experience, output filters are rarely used. Section Output Filters describe them. There are only two arguments to an output filter:

```
filter-name -width -length
```

which are identical to the text filters `-w` and `-l` arguments.

Filters should also *exit* with the following exit status:

exit 0

If the filter printed the file successfully.

exit 1

If the filter failed to print the file but wants **LPD** to try to print the file again. **LPD** will restart a filter if it exits with this status.

exit 2

If the filter failed to print the file and does not want **LPD** to try again. **LPD** will throw out the file.

The text filter that comes with the DragonFly release, `/usr/libexec/lpr/lpf`, takes advantage of the page width and length arguments to determine when to send a form feed and how to account for printer usage. It uses the login, host, and accounting file arguments to make the accounting entries.

If you are shopping for filters, see if they are LPD-compatible. If they are, they must support the argument lists described above. If you plan on writing filters for general use, then have them support the same argument lists and exit codes.

11.4.1.2 Accommodating Plain Text Jobs on PostScript® Printers

If you are the only user of your computer and PostScript (or other language-based) printer, and you promise to never send plain text to your printer and to never use features of various programs that will want to send plain text to your printer, then you do not need to worry about this section at all.

But, if you would like to send both PostScript and plain text jobs to the printer, then you are urged to augment your printer setup. To do so, we have the text filter detect if the arriving job is plain text or PostScript. All PostScript jobs must start with `%!` (for other printer languages, see your printer documentation). If those are the first two characters in the job, we have PostScript, and can pass the rest of the job directly. If those are not the first two characters in the file, then the filter will convert the text into PostScript and print the result.

How do we do this?

11.4.1.3 Simulating PostScript on Non PostScript Printers

PostScript is the *de facto* standard for high quality typesetting and printing. PostScript is, however, an *expensive* standard. Thankfully, Aladdin Enterprises has a free PostScript work-alike called **Ghostscript** that runs with DragonFly. Ghostscript can read most PostScript files and can render their pages onto a variety of devices, including many brands of non-PostScript printers. By installing Ghostscript and using a special text filter for your printer, you can make your non PostScript printer act like a real PostScript printer.

Ghostscript is in `pkgsrc`, if you would like to install it from there. You can fetch, build, and install it quite easily yourself, as well.

To simulate PostScript, we have the text filter detect if it is printing a PostScript file. If it is not, then the filter will pass the file directly to the printer; otherwise, it will use Ghostscript to first convert the file into a format the printer will understand.

Here is an example: the following script is a text filter for Hewlett Packard DeskJet 500 printers. For other printers, substitute the `-sDEVICE` argument to the `gs` (Ghostscript) command. (Type `gs -h` to get a list of devices the current installation of Ghostscript supports.)

```
#!/bin/sh
#
# ifhp - Print Ghostscript-simulated PostScript on a DeskJet 500
# Installed in /usr/local/libexec/ifhp
```

```

#
# Treat LF as CR+LF:
#
printf "\033&k2G" || exit 2

#
# Read first two characters of the file
#
IFS="" read -r first_line
first_two_chars=`expr "$first_line" : '\(..\)`

if [ "$first_two_chars" = "%!" ]; then
#
# It is PostScript; use Ghostscript to scan-convert and print it.
#
# Note that PostScript files are actually interpreted programs,
# and those programs are allowed to write to stdout, which will
# mess up the printed output. So, we redirect stdout to stderr
# and then make descriptor 3 go to stdout, and have Ghostscript
# write its output there. Exercise for the clever reader:
# capture the stderr output from Ghostscript and mail it back to
# the user originating the print job.
#
exec 3>&1 1>&2
/usr/local/bin/gs -dSAFER -dNOPAUSE -q -sDEVICE=djet500 \
-sOutputFile=/dev/fd/3 - && exit 0
else
#
# Plain text or HP/PCL, so just print it directly; print a form feed
# at the end to eject the last page.
#
echo "$first_line" && cat && printf "\033&l0H" &&
exit 0
fi

exit 2

```

Finally, you need to notify **LPD** of the filter via the `if` capability:

```
:if=/usr/local/libexec/ifhp:
```

That is it. You can type `lpr plain.text` and `lpr whatever.ps` and both should print successfully.

11.4.1.4 Conversion Filters

After completing the simple setup described in *Simple Printer Setup*, the first thing you will probably want to do is install conversion filters for your favorite file formats (besides plain ASCII text).

11.4.1.4.1 Why Install Conversion Filters?

Conversion filters make printing various kinds of files easy. As an example, suppose we do a lot of work with the TeX typesetting system, and we have a PostScript printer. Every time we generate a DVI file from TeX, we cannot print it directly until we convert the DVI file into PostScript. The command sequence goes like this:

```
% dvips seaweed-analysis.dvi
% lpr seaweed-analysis.ps
```

By installing a conversion filter for DVI files, we can skip the hand conversion step each time by having **LPD** do it for us. Now, each time we get a DVI file, we are just one step away from printing it:

```
% lpr -d seaweed-analysis.dvi
```

We got **LPD** to do the DVI file conversion for us by specifying the `-d` option. Section [Formatting and Conversion Options](#) lists the conversion options.

For each of the conversion options you want a printer to support, install a *conversion filter* and specify its pathname in `/etc/printcap`. A conversion filter is like the text filter for the simple printer setup (see section [Installing the Text Filter](#)) except that instead of printing plain text, the filter converts the file into a format the printer can understand.

11.4.1.4.2 Which Conversions Filters Should I Install?

You should install the conversion filters you expect to use. If you print a lot of DVI data, then a DVI conversion filter is in order. If you have got plenty of troff to print out, then you probably want a troff filter.

The following table summarizes the filters that **LPD** works with, their capability entries for the `/etc/printcap` file, and how to invoke them with the `lpr` command:

File type	<code>/etc/printcap</code> capability	<code>lpr</code> option
cifplot	cf	-c
DVI	df	-d
plot	gf	-g
ditroff	nf	-n
FORTTRAN text	rf	-f
troff	tf	-f
raster	vf	-v
plain text	if	none, -p, or -l

In our example, using `lpr -d` means the printer needs a `df` capability in its entry in `/etc/printcap`.

Despite what others might contend, formats like FORTRAN text and plot are probably obsolete. At your site, you can give new meanings to these or any of the formatting options just by installing custom filters. For example, suppose you would like to directly print Printerleaf files (files from the Interleaf desktop publishing program), but will never print plot files. You could install a Printerleaf conversion filter under the `gf` capability and then educate your users that `lpr -g` mean “print Printerleaf files.”

11.4.1.4.3 Installing Conversion Filters

Since conversion filters are programs you install outside of the base DragonFly installation, they should probably go under `/usr/local`. The directory `/usr/local/libexec` is a popular location, since they are specialized programs that only **LPD** will run; regular users should not ever need to run them.

To enable a conversion filter, specify its pathname under the appropriate capability for the destination printer in `/etc/printcap`.

In our example, we will add the DVI conversion filter to the entry for the printer named `bamboo`. Here is the example `/etc/printcap` file again, with the new `df` capability for the printer `bamboo`.

```
#
# /etc/printcap for host rose - added df filter for bamboo
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:sd=/var/spool/lpd/rattan:\
    :lp=/dev/lpt0:\
    :if=/usr/local/libexec/if-simple:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:\
    :lp=/dev/ttyd5:ms#-parenb cs8 clocal crtscts:rw:\
    :if=/usr/local/libexec/psif:\
    :df=/usr/local/libexec/psdf:
```

The DVI filter is a shell script named `/usr/local/libexec/psdf`. Here is that script:

```
#!/bin/sh
#
# psdf - DVI to PostScript printer filter
# Installed in /usr/local/libexec/psdf
#
# Invoked by lpd when user runs lpr -d
#
exec /usr/local/bin/dvips -f | /usr/local/libexec/lprps "$@"
```

This script runs `dvips` in filter mode (the `-f` argument) on standard input, which is the job to print. It then starts the PostScript printer filter `lprps` (see section Accommodating Plain Text Jobs on PostScript Printers) with the arguments **LPD** passed to this script. `lprps` will use those arguments to account for the pages printed.

11.4.1.4.4 More Conversion Filter Examples

Since there is no fixed set of steps to install conversion filters, let me instead provide more examples. Use these as guidance to making your own filters. Use them directly, if appropriate.

This example script is a raster (well, GIF file, actually) conversion filter for a Hewlett Packard LaserJet III-Si printer:

```
#!/bin/sh
#
# hpvf - Convert GIF files into HP/PCL, then print
# Installed in /usr/local/libexec/hpvf

PATH=/usr/X11R6/bin:$PATH; export PATH
```

```
giftopnm | ppmtopgm | pgmtopbm | pbmtolj -resolution 300 \
    && exit 0 \
    || exit 2
```

It works by converting the GIF file into a portable anymap, converting that into a portable graymap, converting that into a portable bitmap, and converting that into LaserJet/PCL-compatible data.

Here is the `/etc/printcap` file with an entry for a printer using the above filter:

```
#
# /etc/printcap for host orchid
#
teak|hp|laserjet|Hewlett Packard LaserJet 3Si:\
    :lp=/dev/lpt0:sh:sd=/var/spool/lpd/teak:mx#0:\
    :if=/usr/local/libexec/hpif:\
    :vf=/usr/local/libexec/hpvf:
```

The following script is a conversion filter for troff data from the groff typesetting system for the PostScript printer named bamboo:

```
#!/bin/sh
#
# pstf - Convert groff's troff data into PS, then print.
# Installed in /usr/local/libexec/pstf
#
exec grops | /usr/local/libexec/lprps "$@"
```

The above script makes use of `lprps` again to handle the communication with the printer. If the printer were on a parallel port, we would use this script instead:

```
#!/bin/sh
#
# pstf - Convert groff's troff data into PS, then print.
# Installed in /usr/local/libexec/pstf
#
exec grops
```

That is it. Here is the entry we need to add to `/etc/printcap` to enable the filter:

```
:tf=/usr/local/libexec/pstf:
```

Here is an example that might make old hands at FORTRAN blush. It is a FORTRAN-text filter for any printer that can directly print plain text. We will install it for the printer `teak`:

```
#!/bin/sh
#
# hprf - FORTRAN text filter for LaserJet 3si:
# Installed in /usr/local/libexec/hprf
#

printf "\033&k2G" && fpr && printf "\033&l0H" &&
    exit 0
exit 2
```

And we will add this line to the `/etc/printcap` for the printer `teak` to enable this filter:

```
:rf=/usr/local/libexec/hprf:
```

11.4.1.4.5 Automated Conversion: an Alternative to Conversion Filters

All these conversion filters accomplish a lot for your printing environment, but at the cost forcing the user to specify (on the `lpr(1)` command line) which one to use. If your users are not particularly computer literate, having to specify a filter option will become annoying. What is worse, though, is that an incorrectly specified filter option may run a filter on the wrong type of file and cause your printer to spew out hundreds of sheets of paper.

Rather than install conversion filters at all, you might want to try having the text filter (since it is the default filter) detect the type of file it has been asked to print and then automatically run the right conversion filter. Tools such as `file` can be of help here. Of course, it will be hard to determine the differences between *some* file types—and, of course, you can still provide conversion filters just for them.

The `pkgsrc` collection has a text filter that performs automatic conversion called `apsfilter`. It can detect plain text, PostScript, and DVI files, run the proper conversions, and print.

11.4.1.5 Output Filters

The **LPD** spooling system supports one other type of filter that we have not yet explored: an output filter. An output filter is intended for printing plain text only, like the text filter, but with many simplifications. If you are using an output filter but no text filter, then:

- **LPD** starts an output filter once for the entire job instead of once for each file in the job.
- **LPD** does not make any provision to identify the start or the end of files within the job for the output filter.
- **LPD** does not pass the user's login or host to the filter, so it is not intended to do accounting. In fact, it gets only two arguments:

```
filter-name -wwidth -llength
```

Where *width* is from the `pw` capability and *length* is from the `pl` capability for the printer in question.

Do not be seduced by an output filter's simplicity. If you would like each file in a job to start on a different page an output filter *will not work*. Use a text filter (also known as an input filter); see section **Installing the Text Filter**. Furthermore, an output filter is actually *more complex* in that it has to examine the byte stream being sent to it for special flag characters and must send signals to itself on behalf of **LPD**.

However, an output filter is *necessary* if you want header pages and need to send escape sequences or other initialization strings to be able to print the header page. (But it is also *futile* if you want to charge header pages to the requesting user's account, since **LPD** does not give any user or host information to the output filter.)

On a single printer, **LPD** allows both an output filter and text or other filters. In such cases, **LPD** will start the output filter to print the header page (see section **Header Pages**) only. **LPD** then expects the output filter to *stop itself* by sending two bytes to the filter: ASCII 031 followed by ASCII 001. When an output filter sees these two bytes (031, 001), it should stop by sending `SIGSTOP` to itself. When **LPD**'s done running other filters, it will restart the output filter by sending `SIGCONT` to it.

If there is an output filter but *no* text filter and **LPD** is working on a plain text job, **LPD** uses the output filter to do the job. As stated before, the output filter will print each file of the job in sequence with no intervening form feeds or other paper advancement, and this is probably *not* what you want. In almost all cases, you need a text filter.

The program `lpf`, which we introduced earlier as a text filter, can also run as an output filter. If you need a quick-and-dirty output filter but do not want to write the byte detection and signal sending code, try `lpf`. You can also wrap `lpf` in a shell script to handle any initialization codes the printer might require.

11.4.1.6 `lpf`: a Text Filter

The program `/usr/libexec/lpr/lpf` that comes with DragonFly is a text filter (input filter) that can indent output (job submitted with `lpr -i`), allow literal characters to pass (job submitted with `lpr -l`), adjust the printing position for backspaces and tabs in the job, and account for pages printed. It can also act like an output filter.

`lpf` is suitable for many printing environments. And although it has no capability to send initialization sequences to a printer, it is easy to write a shell script to do the needed initialization and then execute `lpf`.

In order for `lpf` to do page accounting correctly, it needs correct values filled in for the `pw` and `p1` capabilities in the `/etc/printcap` file. It uses these values to determine how much text can fit on a page and how many pages were in a user's job. For more information on printer accounting, see [Accounting for Printer Usage](#).

11.4.2 Header Pages

If you have *lots* of users, all of them using various printers, then you probably want to consider *header pages* as a necessary evil.

Header pages, also known as *banner* or *burst pages* identify to whom jobs belong after they are printed. They are usually printed in large, bold letters, perhaps with decorative borders, so that in a stack of printouts they stand out from the real documents that comprise users' jobs. They enable users to locate their jobs quickly. The obvious drawback to a header page is that it is yet one more sheet that has to be printed for every job, their ephemeral usefulness lasting not more than a few minutes, ultimately finding themselves in a recycling bin or rubbish heap. (Note that header pages go with each job, not each file in a job, so the paper waste might not be that bad.)

The **LPD** system can provide header pages automatically for your printouts *if* your printer can directly print plain text. If you have a PostScript printer, you will need an external program to generate the header page; see [Header Pages on PostScript Printers](#).

11.4.2.1 Enabling Header Pages

In the [Simple Printer Setup](#) section, we turned off header pages by specifying `sh` (meaning "suppress header") in the `/etc/printcap` file. To enable header pages for a printer, just remove the `sh` capability.

Sounds too easy, right?

You are right. You *might* have to provide an output filter to send initialization strings to the printer. Here is an example output filter for Hewlett Packard PCL-compatible printers:

```
#!/bin/sh
#
# hpof - Output filter for Hewlett Packard PCL-compatible printers
# Installed in /usr/local/libexec/hpof
```

```
printf "\033&k2G" || exit 2
exec /usr/libexec/lpr/lpf
```

Specify the path to the output filter in the `of` capability. See the `Output Filters` section for more information.

Here is an example `/etc/printcap` file for the printer `teak` that we introduced earlier; we enabled header pages and added the above output filter:

```
#
# /etc/printcap for host orchid
#
teak|hp|laserjet|Hewlett Packard LaserJet 3Si:\
    :lp=/dev/lpt0:sd=/var/spool/lpd/teak:mx#0:\
    :if=/usr/local/libexec/hpif:\
    :vf=/usr/local/libexec/hpvf:\
    :of=/usr/local/libexec/hpof:
```

Now, when users print jobs to `teak`, they get a header page with each job. If users want to spend time searching for their printouts, they can suppress header pages by submitting the job with `lpr -h`; see the `Header Page Options` section for more `lpr(1)` options.

Note: **LPD** prints a form feed character after the header page. If your printer uses a different character or sequence of characters to eject a page, specify them with the `ff` capability in `/etc/printcap`.

11.4.2.2 Controlling Header Pages

By enabling header pages, **LPD** will produce a *long header*, a full page of large letters identifying the user, host, and job. Here is an example (kelly printed the job named `outline` from host `rose`):

```

k          ll      ll
k          l       l
k          l       l
k  k      eeee    l       l       y   y
k  k      e  e    l       l       y   y
k  k      eeeee   l       l       y   y
kk  k     e       l       l       y   y
k  k      e  e    l       l       y  yy
k    k    eeee   lll     lll     yyy y
                    y
                    Y   Y
                    YYYY

                    ll
                    t       l       i
                    t       l
oooo  u  u  ttttt  l       ii    n nnn    eeee
o  o  u  u  t       l       i    nn  n    e  e
o  o  u  u  t       l       i    n   n    eeeee
```

```

o   o   u   u   t       l       i   n   n   e
o   o   u   uu  t   t   l       i   n   n   e   e
oooo   uuu u   tt     lll     iii   n   n   eeee

```

```

r rrr      oooo      ssss      eeee
rr  r     o   o     s   s     e   e
r         o   o     ss       eeeeee
r         o   o       ss      e
r         o   o     s   s     e   e
r         oooo      ssss      eeee

```

```

Job:  outline
Date: Sun Sep 17 11:04:58 1995

```

LPD appends a form feed after this text so the job starts on a new page (unless you have `sf` (suppress form feeds) in the destination printer's entry in `/etc/printcap`).

If you prefer, **LPD** can make a *short header*; specify `sb` (short banner) in the `/etc/printcap` file. The header page will look like this:

```
rose:kelly Job: outline Date: Sun Sep 17 11:07:51 1995
```

Also by default, **LPD** prints the header page first, then the job. To reverse that, specify `h1` (header last) in `/etc/printcap`.

11.4.2.3 Accounting for Header Pages

Using **LPD**'s built-in header pages enforces a particular paradigm when it comes to printer accounting: header pages must be *free of charge*.

Why?

Because the output filter is the only external program that will have control when the header page is printed that could do accounting, and it is not provided with any *user or host* information or an accounting file, so it has no idea whom to charge for printer use. It is also not enough to just “add one page” to the text filter or any of the conversion filters (which do have user and host information) since users can suppress header pages with `lpr -h`. They could still be charged for header pages they did not print. Basically, `lpr -h` will be the preferred option of environmentally-minded users, but you cannot offer any incentive to use it.

It is *still not enough* to have each of the filters generate their own header pages (thereby being able to charge for them). If users wanted the option of suppressing the header pages with `lpr -h`, they will still get them and be charged for them since **LPD** does not pass any knowledge of the `-h` option to any of the filters.

So, what are your options?

You can:

- Accept **LPD**'s paradigm and make header pages free.
- Install an alternative to **LPD**, such as **LPRng**. Section Alternatives to the Standard Spooler tells more about other spooling software you can substitute for **LPD**.
- Write a *smart* output filter. Normally, an output filter is not meant to do anything more than initialize a printer or do some simple character conversion. It is suited for header pages and plain text jobs (when there is no text (input) filter). But, if there is a text filter for the plain text jobs, then **LPD** will start the output filter only for the header pages. And the output filter can parse the header page text that **LPD** generates to determine what user and host to charge for the header page. The only other problem with this method is that the output filter still does not know what accounting file to use (it is not passed the name of the file from the `af` capability), but if you have a well-known accounting file, you can hard-code that into the output filter. To facilitate the parsing step, use the `sh` (short header) capability in `/etc/printcap`. Then again, all that might be too much trouble, and users will certainly appreciate the more generous system administrator who makes header pages free.

11.4.2.4 Header Pages on PostScript Printers

As described above, **LPD** can generate a plain text header page suitable for many printers. Of course, PostScript cannot directly print plain text, so the header page feature of **LPD** is useless—or mostly so.

One obvious way to get header pages is to have every conversion filter and the text filter generate the header page. The filters should use the user and host arguments to generate a suitable header page. The drawback of this method is that users will always get a header page, even if they submit jobs with `lpr -h`.

Let us explore this method. The following script takes three arguments (user login name, host name, and job name) and makes a simple PostScript header page:

```
#!/bin/sh
#
# make-ps-header - make a PostScript header page on stdout
# Installed in /usr/local/libexec/make-ps-header
#
#
# These are PostScript units (72 to the inch).  Modify for A4 or
# whatever size paper you are using:
#
page_width=612
page_height=792
border=72
#
# Check arguments
#
if [ $# -ne 3 ]; then
```

```

    echo "Usage: `basename $0` <user> <host> <job>" 1>&2
    exit 1
fi

#
# Save these, mostly for readability in the PostScript, below.
#
user=$1
host=$2
job=$3
date=`date`

#
# Send the PostScript code to stdout.
#
exec cat <<EOF
%!PS

%
% Make sure we do not interfere with user's job that will follow
%
save

%
% Make a thick, unpleasant border around the edge of the paper.
%
$border $border moveto
$page_width $border 2 mul sub 0 rlineto
0 $page_height $border 2 mul sub rlineto
currentscreen 3 -1 roll pop 100 3 1 roll setscreen
$border 2 mul $page_width sub 0 rlineto closepath
0.8 setgray 10 setlinewidth stroke 0 setgray

%
% Display user's login name, nice and large and prominent
%
/Helvetica-Bold findfont 64 scalefont setfont
$page_width ($user) stringwidth pop sub 2 div $page_height 200 sub moveto
($user) show

%
% Now show the boring particulars
%
/Helvetica findfont 14 scalefont setfont
/y 200 def
[ (Job:) (Host:) (Date:) ] {
200 y moveto show /y y 18 sub def }
forall

/Helvetica-Bold findfont 14 scalefont setfont
/y 200 def
[ ($job) ($host) ($date) ] {
270 y moveto show /y y 18 sub def

```

```

} forall

%
% That is it
%
restore
showpage
EOF

```

Now, each of the conversion filters and the text filter can call this script to first generate the header page, and then print the user's job. Here is the DVI conversion filter from earlier in this document, modified to make a header page:

```

#!/bin/sh
#
# psdf - DVI to PostScript printer filter
# Installed in /usr/local/libexec/psdf
#
# Invoked by lpd when user runs lpr -d
#

orig_args="$@"

fail() {
    echo "$@" 1>&2
    exit 2
}

while getopts "x:y:n:h:" option; do
    case $option in
        x|y) ;; # Ignore
        n)   login=$OPTARG ;;
        h)   host=$OPTARG ;;
        *)   echo "LPD started `basename $0` wrong." 1>&2
            exit 2
            ;;
    esac
done

[ "$login" ] || fail "No login name"
[ "$host" ] || fail "No host name"

( /usr/local/libexec/make-ps-header $login $host "DVI File"
  /usr/local/bin/dvips -f ) | eval /usr/local/libexec/lprps $orig_args

```

Notice how the filter has to parse the argument list in order to determine the user and host name. The parsing for the other conversion filters is identical. The text filter takes a slightly different set of arguments, though (see section [How Filters Work](#)).

As we have mentioned before, the above scheme, though fairly simple, disables the “suppress header page” option (the `-h` option) to `lpr`. If users wanted to save a tree (or a few pennies, if you charge for header pages), they would not be able to do so, since every filter's going to print a header page with every job.

To allow users to shut off header pages on a per-job basis, you will need to use the trick introduced in section Accounting for Header Pages: write an output filter that parses the LPD-generated header page and produces a PostScript version. If the user submits the job with `lpr -h`, then **LPD** will not generate a header page, and neither will your output filter. Otherwise, your output filter will read the text from **LPD** and send the appropriate header page PostScript code to the printer.

If you have a PostScript printer on a serial line, you can make use of `lprps`, which comes with an output filter, `psof`, which does the above. Note that `psof` does not charge for header pages.

11.4.3 Networked Printing

DragonFly supports networked printing: sending jobs to remote printers. Networked printing generally refers to two different things:

- Accessing a printer attached to a remote host. You install a printer that has a conventional serial or parallel interface on one host. Then, you set up **LPD** to enable access to the printer from other hosts on the network. Section Printers Installed on Remote Hosts tells how to do this.
- Accessing a printer attached directly to a network. The printer has a network interface in addition (or in place of) a more conventional serial or parallel interface. Such a printer might work as follows:
 - It might understand the **LPD** protocol and can even queue jobs from remote hosts. In this case, it acts just like a regular host running **LPD**. Follow the same procedure in section Printers Installed on Remote Hosts to set up such a printer.
 - It might support a data stream network connection. In this case, you “attach” the printer to one host on the network by making that host responsible for spooling jobs and sending them to the printer. Section Printers with Networked Data Stream Interfaces gives some suggestions on installing such printers.

11.4.3.1 Printers Installed on Remote Hosts

The **LPD** spooling system has built-in support for sending jobs to other hosts also running **LPD** (or are compatible with **LPD**). This feature enables you to install a printer on one host and make it accessible from other hosts. It also works with printers that have network interfaces that understand the **LPD** protocol.

To enable this kind of remote printing, first install a printer on one host, the *printer host*, using the simple printer setup described in the Simple Printer Setup section. Do any advanced setup in Advanced Printer Setup that you need. Make sure to test the printer and see if it works with the features of **LPD** you have enabled. Also ensure that the *local host* has authorization to use the **LPD** service in the *remote host* (see Restricting Jobs from Remote Printers).

If you are using a printer with a network interface that is compatible with **LPD**, then the *printer host* in the discussion below is the printer itself, and the *printer name* is the name you configured for the printer. See the documentation that accompanied your printer and/or printer-network interface.

Tip: If you are using a Hewlett Packard Laserjet then the printer name `text` will automatically perform the LF to CRLF conversion for you, so you will not require the `hpif` script.

Then, on the other hosts you want to have access to the printer, make an entry in their `/etc/printcap` files with the following:

1. Name the entry anything you want. For simplicity, though, you probably want to use the same name and aliases as on the printer host.
2. Leave the `lp` capability blank, explicitly (`:lp=:`).
3. Make a spooling directory and specify its location in the `sd` capability. **LPD** will store jobs here before they get sent to the printer host.
4. Place the name of the printer host in the `rm` capability.
5. Place the printer name on the *printer host* in the `rp` capability.

That is it. You do not need to list conversion filters, page dimensions, or anything else in the `/etc/printcap` file.

Here is an example. The host `rose` has two printers, `bamboo` and `rattan`. We will enable users on the host `orchid` to print to those printers. Here is the `/etc/printcap` file for `orchid` (back from section **Enabling Header Pages**). It already had the entry for the printer `teak`; we have added entries for the two printers on the host `rose`:

```
#
# /etc/printcap for host orchid - added (remote) printers on rose
#
#
# teak is local; it is connected directly to orchid:
#
teak|hp|laserjet|Hewlett Packard LaserJet 3Si:\
    :lp=/dev/lpt0:sd=/var/spool/lpd/teak:mx#0:\
    :if=/usr/local/libexec/ifhp:\
    :vf=/usr/local/libexec/vfhp:\
    :of=/usr/local/libexec/ofhp:
#
# rattan is connected to rose; send jobs for rattan to rose:
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :lp=:rm=rose:rp=rattan:sd=/var/spool/lpd/rattan:
#
# bamboo is connected to rose as well:
#
bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :lp=:rm=rose:rp=bamboo:sd=/var/spool/lpd/bamboo:
```

Then, we just need to make spooling directories on `orchid`:

```
# mkdir -p /var/spool/lpd/rattan /var/spool/lpd/bamboo
# chmod 770 /var/spool/lpd/rattan /var/spool/lpd/bamboo
# chown daemon:daemon /var/spool/lpd/rattan /var/spool/lpd/bamboo
```

Now, users on `orchid` can print to `rattan` and `bamboo`. If, for example, a user on `orchid` typed

```
% lpr -P bamboo -d sushi-review.dvi
```


the **LPD** system on *orchid* would copy the job to the spooling directory `/var/spool/lpd/bamboo` and note that it was a DVI job. As soon as the host *rose* has room in its *bamboo* spooling directory, the two **LPDs** would transfer the file to *rose*. The file would wait in *rose*'s queue until it was finally printed. It would be converted from DVI to PostScript (since *bamboo* is a PostScript printer) on *rose*.

11.4.3.2 Printers with Networked Data Stream Interfaces

Often, when you buy a network interface card for a printer, you can get two versions: one which emulates a spooler (the more expensive version), or one which just lets you send data to it as if you were using a serial or parallel port (the cheaper version). This section tells how to use the cheaper version. For the more expensive one, see the previous section *Printers Installed on Remote Hosts*.

The format of the `/etc/printcap` file lets you specify what serial or parallel interface to use, and (if you are using a serial interface), what baud rate, whether to use flow control, delays for tabs, conversion of newlines, and more. But there is no way to specify a connection to a printer that is listening on a TCP/IP or other network port.

To send data to a networked printer, you need to develop a communications program that can be called by the text and conversion filters. Here is one such example: the script `netprint` takes all data on standard input and sends it to a network-attached printer. We specify the hostname of the printer as the first argument and the port number to which to connect as the second argument to `netprint`. Note that this supports one-way communication only (DragonFly to printer); many network printers support two-way communication, and you might want to take advantage of that (to get printer status, perform accounting, etc.).

```
#!/usr/bin/perl
#
# netprint - Text filter for printer attached to network
# Installed in /usr/local/libexec/netprint
#
$#ARGV eq 1 || die "Usage: $0 <printer-hostname> <port-number>";

$printer_host = $ARGV[0];
$printer_port = $ARGV[1];

require 'sys/socket.ph';

($ignore, $ignore, $protocol) = getprotobyname('tcp');
($ignore, $ignore, $ignore, $ignore, $address)
    = gethostbyname($printer_host);

$sockaddr = pack('S n a4 x8', &AF_INET, $printer_port, $address);

socket(PRINTER, &PF_INET, &SOCK_STREAM, $protocol)
    || die "Can't create TCP/IP stream socket: $!";
connect(PRINTER, $sockaddr) || die "Can't contact $printer_host: $!";
while (<STDIN>) { print PRINTER; }
exit 0;
```

We can then use this script in various filters. Suppose we had a Diablo 750-N line printer connected to the network. The printer accepts data to print on port number 5100. The host name of the printer is *scrivener*. Here is the text filter for the printer:

```
#!/bin/sh
```

```
#
# diablo-if-net - Text filter for Diablo printer 'scrivener' listening
# on port 5100.  Installed in /usr/local/libexec/diablo-if-net
#
exec /usr/libexec/lpr/lpf "$@" | /usr/local/libexec/netprint scrivener 5100
```

11.4.4 Restricting Printer Usage

This section gives information on restricting printer usage. The **LPD** system lets you control who can access a printer, both locally or remotely, whether they can print multiple copies, how large their jobs can be, and how large the printer queues can get.

11.4.4.1 Restricting Multiple Copies

The **LPD** system makes it easy for users to print multiple copies of a file. Users can print jobs with `lpr -#5` (for example) and get five copies of each file in the job. Whether this is a good thing is up to you.

If you feel multiple copies cause unnecessary wear and tear on your printers, you can disable the `-#` option to `lpr(1)` by adding the `sc` capability to the `/etc/printcap` file. When users submit jobs with the `-#` option, they will see:

```
lpr: multiple copies are not allowed
```

Note that if you have set up access to a printer remotely (see section **Printers Installed on Remote Hosts**), you need the `sc` capability on the remote `/etc/printcap` files as well, or else users will still be able to submit multiple-copy jobs by using another host.

Here is an example. This is the `/etc/printcap` file for the host `rose`. The printer `rattan` is quite hearty, so we will allow multiple copies, but the laser printer `bamboo` is a bit more delicate, so we will disable multiple copies by adding the `sc` capability:

```
#
# /etc/printcap for host rose - restrict multiple copies on bamboo
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:sd=/var/spool/lpd/rattan:\
    :lp=/dev/lpt0:\
    :if=/usr/local/libexec/if-simple:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:sc:\
    :lp=/dev/ttyd5:ms#-parenb cs8 clocal crtscts:rw:\
    :if=/usr/local/libexec/psif:\
    :df=/usr/local/libexec/psdf:
```

Now, we also need to add the `sc` capability on the host `orchid`'s `/etc/printcap` (and while we are at it, let us disable multiple copies for the printer `teak`):

```
#
# /etc/printcap for host orchid - no multiple copies for local
# printer teak or remote printer bamboo
teak|hp|laserjet|Hewlett Packard LaserJet 3Si:\
```

```

:lp=/dev/lpt0:sd=/var/spool/lpd/teak:mx#0:sc:\
:if=/usr/local/libexec/ifhp:\
:vf=/usr/local/libexec/vfhp:\
:of=/usr/local/libexec/ofhp:

rattan|line|diablo|lp|Diablo 630 Line Printer:\
:lp=:rm=rose:rp=rattan:sd=/var/spool/lpd/rattan:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
:lp=:rm=rose:rp=bamboo:sd=/var/spool/lpd/bamboo:sc:

```

By using the `sc` capability, we prevent the use of `lpr -#`, but that still does not prevent users from running `lpr(1)` multiple times, or from submitting the same file multiple times in one job like this:

```
% lpr forsale.sign forsale.sign forsale.sign forsale.sign forsale.sign
```

There are many ways to prevent this abuse (including ignoring it) which you are free to explore.

11.4.4.2 Restricting Access to Printers

You can control who can print to what printers by using the UNIX group mechanism and the `rg` capability in `/etc/printcap`. Just place the users you want to have access to a printer in a certain group, and then name that group in the `rg` capability.

Users outside the group (including `root`) will be greeted with `lpr: Not a member of the restricted group` if they try to print to the controlled printer.

As with the `sc` (suppress multiple copies) capability, you need to specify `rg` on remote hosts that also have access to your printers, if you feel it is appropriate (see section [Printers Installed on Remote Hosts](#)).

For example, we will let anyone access the printer `rattan`, but only those in group `artists` can use `bamboo`. Here is the familiar `/etc/printcap` for host `rose`:

```

#
# /etc/printcap for host rose - restricted group for bamboo
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
:sh:sd=/var/spool/lpd/rattan:\
:lp=/dev/lpt0:\
:if=/usr/local/libexec/if-simple:

bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
:sh:sd=/var/spool/lpd/bamboo:sc:rg=artists:\
:lp=/dev/ttyd5:ms#-parenb cs8 clocal crtscts:rw:\
:if=/usr/local/libexec/psif:\
:df=/usr/local/libexec/psdf:

```

Let us leave the other example `/etc/printcap` file (for the host `orchid`) alone. Of course, anyone on `orchid` can print to `bamboo`. It might be the case that we only allow certain logins on `orchid` anyway, and want them to have access to the printer. Or not.

Note: There can be only one restricted group per printer.

11.4.4.3 Controlling Sizes of Jobs Submitted

If you have many users accessing the printers, you probably need to put an upper limit on the sizes of the files users can submit to print. After all, there is only so much free space on the filesystem that houses the spooling directories, and you also need to make sure there is room for the jobs of other users.

LPD enables you to limit the maximum byte size a file in a job can be with the `mx` capability. The units are in `BUFSIZ` blocks, which are 1024 bytes. If you put a zero for this capability, there will be no limit on file size; however, if no `mx` capability is specified, then a default limit of 1000 blocks will be used.

Note: The limit applies to *files* in a job, and *not* the total job size.

LPD will not refuse a file that is larger than the limit you place on a printer. Instead, it will queue as much of the file up to the limit, which will then get printed. The rest will be discarded. Whether this is correct behavior is up for debate.

Let us add limits to our example printers `rattan` and `bamboo`. Since those artists' PostScript files tend to be large, we will limit them to five megabytes. We will put no limit on the plain text line printer:

```
#
# /etc/printcap for host rose
#

#
# No limit on job size:
#
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:mx#0:sd=/var/spool/lpd/rattan:\
    :lp=/dev/lpt0:\
    :if=/usr/local/libexec/if-simple:

#
# Limit of five megabytes:
#
bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:sc:rg=artists:mx#5000:\
    :lp=/dev/ttyd5:ms#-parenb cs8 clocal crtscts:rw:\
    :if=/usr/local/libexec/psif:\
    :df=/usr/local/libexec/psdf:
```

Again, the limits apply to the local users only. If you have set up access to your printers remotely, remote users will not get those limits. You will need to specify the `mx` capability in the remote `/etc/printcap` files as well. See section [Printers Installed on Remote Hosts](#) for more information on remote printing.

There is another specialized way to limit job sizes from remote printers; see section [Restricting Jobs from Remote Printers](#).

11.4.4.4 Restricting Jobs from Remote Printers

The **LPD** spooling system provides several ways to restrict print jobs submitted from remote hosts:

Host restrictions

You can control from which remote hosts a local **LPD** accepts requests with the files `/etc/hosts.equiv` and `/etc/hosts.lpd`. **LPD** checks to see if an incoming request is from a host listed in either one of these files. If not, **LPD** refuses the request.

The format of these files is simple: one host name per line. Note that the file `/etc/hosts.equiv` is also used by the `ruserok(3)` protocol, and affects programs like `rsh(1)` and `rcp(1)`, so be careful.

For example, here is the `/etc/hosts.lpd` file on the host `rose`:

```
orchid
violet
madrival.fishbaum.de
```

This means `rose` will accept requests from the hosts `orchid`, `violet`, and `madrival.fishbaum.de`. If any other host tries to access `rose`'s **LPD**, the job will be refused.

Size restrictions

You can control how much free space there needs to remain on the filesystem where a spooling directory resides. Make a file called `minfree` in the spooling directory for the local printer. Insert in that file a number representing how many disk blocks (512 bytes) of free space there has to be for a remote job to be accepted.

This lets you insure that remote users will not fill your filesystem. You can also use it to give a certain priority to local users: they will be able to queue jobs long after the free disk space has fallen below the amount specified in the `minfree` file.

For example, let us add a `minfree` file for the printer `bamboo`. We examine `/etc/printcap` to find the spooling directory for this printer; here is `bamboo`'s entry:

```
bamboo|ps|PS|S|panasonic|Panasonic KX-P4455 PostScript v51.4:\
    :sh:sd=/var/spool/lpd/bamboo:sc:rg=artists:mx#5000:\
    :lp=/dev/ttyd5:ms#-parenb cs8 clocal crtscts:rw:mx#5000:\
    :if=/usr/local/libexec/psif:\
    :df=/usr/local/libexec/psdf:
```

The spooling directory is given in the `sd` capability. We will make three megabytes (which is 6144 disk blocks) the amount of free disk space that must exist on the filesystem for **LPD** to accept remote jobs:

```
# echo 6144 > /var/spool/lpd/bamboo/minfree
```

User restrictions

You can control which remote users can print to local printers by specifying the `rs` capability in `/etc/printcap`. When `rs` appears in the entry for a locally-attached printer, **LPD** will accept jobs from remote hosts *if* the user submitting the job also has an account of the same login name on the local host. Otherwise, **LPD** refuses the job.

This capability is particularly useful in an environment where there are (for example) different departments sharing a network, and some users transcend departmental boundaries. By giving them accounts on your systems, they can use your printers from their own departmental systems. If you would rather allow them to use

only your printers and not your computer resources, you can give them “token” accounts, with no home directory and a useless shell like `/usr/bin/false`.

11.4.5 Accounting for Printer Usage

So, you need to charge for printouts. And why not? Paper and ink cost money. And then there are maintenance costs—printers are loaded with moving parts and tend to break down. You have examined your printers, usage patterns, and maintenance fees and have come up with a per-page (or per-foot, per-meter, or per-whatever) cost. Now, how do you actually start accounting for printouts?

Well, the bad news is the **LPD** spooling system does not provide much help in this department. Accounting is highly dependent on the kind of printer in use, the formats being printed, and *your* requirements in charging for printer usage.

To implement accounting, you have to modify a printer’s text filter (to charge for plain text jobs) and the conversion filters (to charge for other file formats), to count pages or query the printer for pages printed. You cannot get away with using the simple output filter, since it cannot do accounting. See section Filters.

Generally, there are two ways to do accounting:

- *Periodic accounting* is the more common way, possibly because it is easier. Whenever someone prints a job, the filter logs the user, host, and number of pages to an accounting file. Every month, semester, year, or whatever time period you prefer, you collect the accounting files for the various printers, tally up the pages printed by users, and charge for usage. Then you truncate all the logging files, starting with a clean slate for the next period.
- *Timely accounting* is less common, probably because it is more difficult. This method has the filters charge users for printouts as soon as they use the printers. Like disk quotas, the accounting is immediate. You can prevent users from printing when their account goes in the red, and might provide a way for users to check and adjust their “print quotas.” But this method requires some database code to track users and their quotas.

The **LPD** spooling system supports both methods easily: since you have to provide the filters (well, most of the time), you also have to provide the accounting code. But there is a bright side: you have enormous flexibility in your accounting methods. For example, you choose whether to use periodic or timely accounting. You choose what information to log: user names, host names, job types, pages printed, square footage of paper used, how long the job took to print, and so forth. And you do so by modifying the filters to save this information.

11.4.5.1 Quick and Dirty Printer Accounting

DragonFly comes with two programs that can get you set up with simple periodic accounting right away. They are the text filter `lpf`, described in section `lpf: a Text Filter`, and `pac(8)`, a program to gather and total entries from printer accounting files.

As mentioned in the section on filters (Filters), **LPD** starts the text and the conversion filters with the name of the accounting file to use on the filter command line. The filters can use this argument to know where to write an accounting file entry. The name of this file comes from the `af` capability in `/etc/printcap`, and if not specified as an absolute path, is relative to the spooling directory.

LPD starts `lpf` with page width and length arguments (from the `pw` and `pl` capabilities). `lpf` uses these arguments to determine how much paper will be used. After sending the file to the printer, it then writes an accounting entry in the accounting file. The entries look like this:

```

2.00 rose:andy
3.00 rose:kelly
3.00 orchid:mary
5.00 orchid:mary
2.00 orchid:zhang

```

You should use a separate accounting file for each printer, as `lpf` has no file locking logic built into it, and two `lpfs` might corrupt each other's entries if they were to write to the same file at the same time. An easy way to insure a separate accounting file for each printer is to use `af=acct` in `/etc/printcap`. Then, each accounting file will be in the spooling directory for a printer, in a file named `acct`.

When you are ready to charge users for printouts, run the `pac(8)` program. Just change to the spooling directory for the printer you want to collect on and type `pac`. You will get a dollar-centric summary like the following:

Login	pages/feet	runs	price
orchid:kelly	5.00	1	\$ 0.10
orchid:mary	31.00	3	\$ 0.62
orchid:zhang	9.00	1	\$ 0.18
rose:andy	2.00	1	\$ 0.04
rose:kelly	177.00	104	\$ 3.54
rose:mary	87.00	32	\$ 1.74
rose:root	26.00	12	\$ 0.52
total	337.00	154	\$ 6.74

These are the arguments `pac(8)` expects:

`-Pprinter`

Which *printer* to summarize. This option works only if there is an absolute path in the `af` capability in `/etc/printcap`.

`-c`

Sort the output by cost instead of alphabetically by user name.

`-m`

Ignore host name in the accounting files. With this option, user `smith` on host `alpha` is the same user `smith` on host `gamma`. Without, they are different users.

`-pprice`

Compute charges with *price* dollars per page or per foot instead of the price from the `pc` capability in `/etc/printcap`, or two cents (the default). You can specify *price* as a floating point number.

`-r`

Reverse the sort order.

`-s`

Make an accounting summary file and truncate the accounting file.

name . . .

Print accounting information for the given user *names* only.

In the default summary that `pac(8)` produces, you see the number of pages printed by each user from various hosts. If, at your site, host does not matter (because users can use any host), run `pac -m`, to produce the following summary:

Login	pages/foot	runs	price
andy	2.00	1	\$ 0.04
kelly	182.00	105	\$ 3.64
mary	118.00	35	\$ 2.36
root	26.00	12	\$ 0.52
zhang	9.00	1	\$ 0.18
total	337.00	154	\$ 6.74

To compute the dollar amount due, `pac(8)` uses the `pc` capability in the `/etc/printcap` file (default of 200, or 2 cents per page). Specify, in hundredths of cents, the price per page or per foot you want to charge for printouts in this capability. You can override this value when you run `pac(8)` with the `-p` option. The units for the `-p` option are in dollars, though, not hundredths of cents. For example,

```
# pac -p1.50
```

makes each page cost one dollar and fifty cents. You can really rake in the profits by using this option.

Finally, running `pac -s` will save the summary information in a summary accounting file, which is named the same as the printer's accounting file, but with `_sum` appended to the name. It then truncates the accounting file. When you run `pac(8)` again, it rereads the summary file to get starting totals, then adds information from the regular accounting file.

11.4.5.2 How Can You Count Pages Printed?

In order to perform even remotely accurate accounting, you need to be able to determine how much paper a job uses. This is the essential problem of printer accounting.

For plain text jobs, the problem is not that hard to solve: you count how many lines are in a job and compare it to how many lines per page your printer supports. Do not forget to take into account backspaces in the file which overprint lines, or long logical lines that wrap onto one or more additional physical lines.

The text filter `lpf` (introduced in `lpf: a Text Filter`) takes into account these things when it does accounting. If you are writing a text filter which needs to do accounting, you might want to examine `lpf`'s source code.

How do you handle other file formats, though?

Well, for DVI-to-LaserJet or DVI-to-PostScript conversion, you can have your filter parse the diagnostic output of `dvi1j` or `dvi2ps` and look to see how many pages were converted. You might be able to do similar things with other file formats and conversion programs.

But these methods suffer from the fact that the printer may not actually print all those pages. For example, it could jam, run out of toner, or explode—and the user would still get charged.

So, what can you do?

There is only one *sure* way to do *accurate* accounting. Get a printer that can tell you how much paper it uses, and attach it via a serial line or a network connection. Nearly all PostScript printers support this notion. Other makes and

models do as well (networked Imagen laser printers, for example). Modify the filters for these printers to get the page usage after they print each job and have them log accounting information based on that value *only*. There is no line counting nor error-prone file examination required.

Of course, you can always be generous and make all printouts free.

11.5 Using Printers

This section tells you how to use printers you have set up with DragonFly. Here is an overview of the user-level commands:

`lpr(1)`

Print jobs

`lpq(1)`

Check printer queues

`lprm(1)`

Remove jobs from a printer's queue

There is also an administrative command, `lpc(8)`, described in the section **Administering the LPD Spooler**, used to control printers and their queues.

All three of the commands `lpr(1)`, `lprm(1)`, and `lpq(1)` accept an option `-P printer-name` to specify on which printer/queue to operate, as listed in the `/etc/printcap` file. This enables you to submit, remove, and check on jobs for various printers. If you do not use the `-P` option, then these commands use the printer specified in the `PRINTER` environment variable. Finally, if you do not have a `PRINTER` environment variable, these commands default to the printer named `lp`.

Hereafter, the terminology *default printer* means the printer named in the `PRINTER` environment variable, or the printer named `lp` when there is no `PRINTER` environment variable.

11.5.1 Printing Jobs

To print files, type:

```
% lpr filename ...
```

This prints each of the listed files to the default printer. If you list no files, `lpr(1)` reads data to print from standard input. For example, this command prints some important system files:

```
% lpr /etc/host.conf /etc/hosts.equiv
```

To select a specific printer, type:

```
% lpr -P printer-name filename ...
```

This example prints a long listing of the current directory to the printer named `rattan`:

```
% ls -l | lpr -P rattan
```

Because no files were listed for the `lpr(1)` command, `lpr` read the data to print from standard input, which was the output of the `ls -l` command.

The `lpr(1)` command can also accept a wide variety of options to control formatting, apply file conversions, generate multiple copies, and so forth. For more information, see the section [Printing Options](#).

11.5.2 Checking Jobs

When you print with `lpr(1)`, the data you wish to print is put together in a package called a “print job”, which is sent to the **LPD** spooling system. Each printer has a queue of jobs, and your job waits in that queue along with other jobs from yourself and from other users. The printer prints those jobs in a first-come, first-served order.

To display the queue for the default printer, type `lpq(1)`. For a specific printer, use the `-P` option. For example, the command

```
% lpq -P bamboo
```

shows the queue for the printer named `bamboo`. Here is an example of the output of the `lpq` command:

```
bamboo is ready and printing
Rank  Owner   Job  Files                               Total Size
active kelly   9    /etc/host.conf, /etc/hosts.equiv    88 bytes
2nd   kelly   10   (standard input)                   1635 bytes
3rd   mary    11   ...                                  78519 bytes
```

This shows three jobs in the queue for `bamboo`. The first job, submitted by user `kelly`, got assigned “job number” 9. Every job for a printer gets a unique job number. Most of the time you can ignore the job number, but you will need it if you want to cancel the job; see section [Removing Jobs](#) for details.

Job number nine consists of two files; multiple files given on the `lpr(1)` command line are treated as part of a single job. It is the currently active job (note the word `active` under the “Rank” column), which means the printer should be currently printing that job. The second job consists of data passed as the standard input to the `lpr(1)` command. The third job came from user `mary`; it is a much larger job. The pathname of the file she is trying to print is too long to fit, so the `lpq(1)` command just shows three dots.

The very first line of the output from `lpq(1)` is also useful: it tells what the printer is currently doing (or at least what **LPD** thinks the printer is doing).

The `lpq(1)` command also support a `-l` option to generate a detailed long listing. Here is an example of `lpq -l`:

```
waiting for bamboo to become ready (offline ?)
kelly: 1st  [job 009rose]
        /etc/host.conf                73 bytes
        /etc/hosts.equiv            15 bytes

kelly: 2nd  [job 010rose]
        (standard input)             1635 bytes

mary: 3rd                                     [job 011rose]
        /home/orchid/mary/research/venus/alpha-regio/mapping 78519 bytes
```

11.5.3 Removing Jobs

If you change your mind about printing a job, you can remove the job from the queue with the `lprm(1)` command. Often, you can even use `lprm(1)` to remove an active job, but some or all of the job might still get printed.

To remove a job from the default printer, first use `lpq(1)` to find the job number. Then type:

```
% lprm job-number
```

To remove the job from a specific printer, add the `-P` option. The following command removes job number 10 from the queue for the printer `bamboo`:

```
% lprm -P bamboo 10
```

The `lprm(1)` command has a few shortcuts:

`lprm -`

Removes all jobs (for the default printer) belonging to you.

`lprm user`

Removes all jobs (for the default printer) belonging to `user`. The superuser can remove other users' jobs; you can remove only your own jobs.

`lprm`

With no job number, user name, or `-` appearing on the command line, `lprm(1)` removes the currently active job on the default printer, if it belongs to you. The superuser can remove any active job.

Just use the `-P` option with the above shortcuts to operate on a specific printer instead of the default. For example, the following command removes all jobs for the current user in the queue for the printer named `rattan`:

```
% lprm -P rattan -
```

Note: If you are working in a networked environment, `lprm(1)` will let you remove jobs only from the host from which the jobs were submitted, even if the same printer is available from other hosts. The following command sequence demonstrates this:

```
% lpr -P rattan myfile
% rlogin orchid
% lpq -P rattan
Rank  Owner   Job  Files                Total Size
active seeyan  12  ...                49123 bytes
2nd   kelly    13  myfile                12 bytes
% lprm -P rattan 13
rose: Permission denied
% logout
% lprm -P rattan 13
dfA013rose dequeued
cfA013rose dequeued
```

11.5.4 Beyond Plain Text: Printing Options

The `lpr(1)` command supports a number of options that control formatting text, converting graphic and other file formats, producing multiple copies, handling of the job, and more. This section describes the options.

11.5.4.1 Formatting and Conversion Options

The following `lpr(1)` options control formatting of the files in the job. Use these options if the job does not contain plain text or if you want plain text formatted through the `pr(1)` utility.

For example, the following command prints a DVI file (from the TeX typesetting system) named `fish-report.dvi` to the printer named `bamboo`:

```
% lpr -P bamboo -d fish-report.dvi
```

These options apply to every file in the job, so you cannot mix (say) DVI and ditroff files together in a job. Instead, submit the files as separate jobs, using a different conversion option for each job.

Note: All of these options except `-p` and `-T` require conversion filters installed for the destination printer. For example, the `-d` option requires the DVI conversion filter. Section Conversion Filters gives details.

`-c`

Print cifplot files.

`-d`

Print DVI files.

`-f`

Print FORTRAN text files.

`-g`

Print plot data.

`-i number`

Indent the output by *number* columns; if you omit *number*, indent by 8 columns. This option works only with certain conversion filters.

Note: Do not put any space between the `-i` and the number.

`-l`

Print literal text data, including control characters.

-n

Print ditroff (device independent troff) data.

-P

Format plain text with `pr(1)` before printing. See `pr(1)` for more information.

-T *title*

Use *title* on the `pr(1)` header instead of the file name. This option has effect only when used with the `-p` option.

-t

Print troff data.

-v

Print raster data.

Here is an example: this command prints a nicely formatted version of the `ls(1)` manual page on the default printer:

```
% zcat /usr/share/man/man1/ls.1.gz | troff -t -man | lpr -t
```

The `zcat(1)` command uncompresses the source of the `ls(1)` manual page and passes it to the `troff(1)` command, which formats that source and makes GNU troff output and passes it to `lpr(1)`, which submits the job to the **LPD** spooler. Because we used the `-t` option to `lpr(1)`, the spooler will convert the GNU troff output into a format the default printer can understand when it prints the job.

11.5.4.2 Job Handling Options

The following options to `lpr(1)` tell **LPD** to handle the job specially:

-# *copies*

Produce a number of *copies* of each file in the job instead of just one copy. An administrator may disable this option to reduce printer wear-and-tear and encourage photocopier usage. See section **Restricting Multiple Copies**.

This example prints three copies of `parser.c` followed by three copies of `parser.h` to the default printer:

```
% lpr -#3 parser.c parser.h
```

-m

Send mail after completing the print job. With this option, the **LPD** system will send mail to your account when it finishes handling your job. In its message, it will tell you if the job completed successfully or if there was an error, and (often) what the error was.

-s

Do not copy the files to the spooling directory, but make symbolic links to them instead.

If you are printing a large job, you probably want to use this option. It saves space in the spooling directory (your job might overflow the free space on the filesystem where the spooling directory resides). It saves time as well since **LPD** will not have to copy each and every byte of your job to the spooling directory.

There is a drawback, though: since **LPD** will refer to the original files directly, you cannot modify or remove them until they have been printed.

Note: If you are printing to a remote printer, **LPD** will eventually have to copy files from the local host to the remote host, so the `-s` option will save space only on the local spooling directory, not the remote. It is still useful, though.

`-r`

Remove the files in the job after copying them to the spooling directory, or after printing them with the `-s` option. Be careful with this option!

11.5.4.3 Header Page Options

These options to `lpr(1)` adjust the text that normally appears on a job's header page. If header pages are suppressed for the destination printer, these options have no effect. See section Header Pages for information about setting up header pages.

`-C text`

Replace the hostname on the header page with *text*. The hostname is normally the name of the host from which the job was submitted.

`-J text`

Replace the job name on the header page with *text*. The job name is normally the name of the first file of the job, or `stdin` if you are printing standard input.

`-h`

Do not print any header page.

Note: At some sites, this option may have no effect due to the way header pages are generated. See Header Pages for details.

11.5.5 Administering Printers

As an administrator for your printers, you have had to install, set up, and test them. Using the `lpc(8)` command, you can interact with your printers in yet more ways. With `lpc(8)`, you can

- Start and stop the printers
- Enable and disable their queues

- Rearrange the order of the jobs in each queue.

First, a note about terminology: if a printer is *stopped*, it will not print anything in its queue. Users can still submit jobs, which will wait in the queue until the printer is *started* or the queue is cleared.

If a queue is *disabled*, no user (except `root`) can submit jobs for the printer. An *enabled* queue allows jobs to be submitted. A printer can be *started* for a disabled queue, in which case it will continue to print jobs in the queue until the queue is empty.

In general, you have to have `root` privileges to use the `lpc(8)` command. Ordinary users can use the `lpc(8)` command to get printer status and to restart a hung printer only.

Here is a summary of the `lpc(8)` commands. Most of the commands take a *printer-name* argument to tell on which printer to operate. You can use `all` for the *printer-name* to mean all printers listed in `/etc/printcap`.

`abort printer-name`

Cancel the current job and stop the printer. Users can still submit jobs if the queue is enabled.

`clean printer-name`

Remove old files from the printer's spooling directory. Occasionally, the files that make up a job are not properly removed by **LPD**, particularly if there have been errors during printing or a lot of administrative activity. This command finds files that do not belong in the spooling directory and removes them.

`disable printer-name`

Disable queuing of new jobs. If the printer is running, it will continue to print any jobs remaining in the queue. The superuser (`root`) can always submit jobs, even to a disabled queue.

This command is useful while you are testing a new printer or filter installation: disable the queue and submit jobs as `root`. Other users will not be able to submit jobs until you complete your testing and re-enable the queue with the `enable` command.

`down printer-name message`

Take a printer down. Equivalent to `disable` followed by `stop`. The *message* appears as the printer's status whenever a user checks the printer's queue with `lpq(1)` or status with `lpc status`.

`enable printer-name`

Enable the queue for a printer. Users can submit jobs but the printer will not print anything until it is started.

`help command-name`

Print help on the command *command-name*. With no *command-name*, print a summary of the commands available.

`restart printer-name`

Start the printer. Ordinary users can use this command if some extraordinary circumstance hangs **LPD**, but they cannot start a printer stopped with either the `stop` or `down` commands. The `restart` command is equivalent to `abort` followed by `start`.

`start printer-name`

Start the printer. The printer will print jobs in its queue.

```
stop printer-name
```

Stop the printer. The printer will finish the current job and will not print anything else in its queue. Even though the printer is stopped, users can still submit jobs to an enabled queue.

```
topq printer-name job-or-username
```

Rearrange the queue for *printer-name* by placing the jobs with the listed *job* numbers or the jobs belonging to *username* at the top of the queue. For this command, you cannot use `all` as the *printer-name*.

```
up printer-name
```

Bring a printer up; the opposite of the `down` command. Equivalent to `start` followed by `enable`.

`lpc(8)` accepts the above commands on the command line. If you do not enter any commands, `lpc(8)` enters an interactive mode, where you can enter commands until you type `exit`, `quit`, or end-of-file.

11.6 Alternatives to the Standard Spooler

If you have been reading straight through this manual, by now you have learned just about everything there is to know about the **LPD** spooling system that comes with DragonFly. You can probably appreciate many of its shortcomings, which naturally leads to the question: “What other spooling systems are out there (and work with DragonFly)?”

LPRng

LPRng, which purportedly means “LPR: the Next Generation” is a complete rewrite of PLP. Patrick Powell and Justin Mason (the principal maintainer of PLP) collaborated to make **LPRng**. The main site for **LPRng** is <http://www.lprng.org/>.

CUPS

CUPS, the Common UNIX Printing System, provides a portable printing layer for UNIX-based operating systems. It has been developed by Easy Software Products to promote a standard printing solution for all UNIX vendors and users.

CUPS uses the Internet Printing Protocol (IPP) as the basis for managing print jobs and queues. The Line Printer Daemon (LPD) Server Message Block (SMB), and AppSocket (a.k.a. JetDirect) protocols are also supported with reduced functionality. **CUPS** adds network printer browsing and PostScript Printer Description (PPD) based printing options to support real-world printing under UNIX.

The main site for **CUPS** is <http://www.cups.org/>.

11.7 Troubleshooting

After performing the simple test with `lptest(1)`, you might have gotten one of the following results instead of the correct printout:

It worked, after awhile; or, it did not eject a full sheet.

The printer printed the above, but it sat for awhile and did nothing. In fact, you might have needed to press a **PRINT REMAINING** or **FORM FEED** button on the printer to get any results to appear.

If this is the case, the printer was probably waiting to see if there was any more data for your job before it printed anything. To fix this problem, you can have the text filter send a FORM FEED character (or whatever is necessary) to the printer. This is usually sufficient to have the printer immediately print any text remaining in its internal buffer. It is also useful to make sure each print job ends on a full sheet, so the next job does not start somewhere on the middle of the last page of the previous job.

The following replacement for the shell script `/usr/local/libexec/if-simple` prints a form feed after it sends the job to the printer:

```
#!/bin/sh
#
# if-simple - Simple text input filter for lpd
# Installed in /usr/local/libexec/if-simple
#
# Simply copies stdin to stdout. Ignores all filter arguments.
# Writes a form feed character (\f) after printing job.

/bin/cat && printf "\f" && exit 0
exit 2
```

It produced the “staircase effect.”

You got the following on paper:

```
! "#$%&'()*+,-./01234
      "$%&'()*+,-./012345
                #$%&'()*+,-./0123456
```

You have become another victim of the *staircase effect*, caused by conflicting interpretations of what characters should indicate a new line. UNIX style operating systems use a single character: ASCII code 10, the line feed (LF). MS-DOS, OS/2®, and others uses a pair of characters, ASCII code 10 *and* ASCII code 13 (the carriage return or CR). Many printers use the MS-DOS convention for representing new-lines.

When you print with DragonFly, your text used just the line feed character. The printer, upon seeing a line feed character, advanced the paper one line, but maintained the same horizontal position on the page for the next character to print. That is what the carriage return is for: to move the location of the next character to print to the left edge of the paper.

Here is what DragonFly wants your printer to do:

Printer received CR	Printer prints CR
Printer received LF	Printer prints CR + LF

Here are some ways to achieve this:

- Use the printer’s configuration switches or control panel to alter its interpretation of these characters. Check your printer’s manual to find out how to do this.

Note: If you boot your system into other operating systems besides DragonFly, you may have to *reconfigure* the printer to use a an interpretation for CR and LF characters that those other operating systems use. You might prefer one of the other solutions, below.

- Have DragonFly's serial line driver automatically convert LF to CR+LF. Of course, this works with printers on serial ports *only*. To enable this feature, use the `ms#` capability and set the `onlcr` mode in the `/etc/printcap` file for the printer.
- Send an *escape code* to the printer to have it temporarily treat LF characters differently. Consult your printer's manual for escape codes that your printer might support. When you find the proper escape code, modify the text filter to send the code first, then send the print job.

Here is an example text filter for printers that understand the Hewlett-Packard PCL escape codes. This filter makes the printer treat LF characters as a LF and CR; then it sends the job; then it sends a form feed to eject the last page of the job. It should work with nearly all Hewlett Packard printers.

```
#!/bin/sh
#
# hpif - Simple text input filter for lpd for HP-PCL based printers
# Installed in /usr/local/libexec/hpif
#
# Simply copies stdin to stdout. Ignores all filter arguments.
# Tells printer to treat LF as CR+LF. Ejects the page when done.

printf "\033&k2G" && cat && printf "\033&l0H" && exit 0
exit 2
```

Here is an example `/etc/printcap` from a host called `orchid`. It has a single printer attached to its first parallel port, a Hewlett Packard LaserJet 3Si named `teak`. It is using the above script as its text filter:

```
#
# /etc/printcap for host orchid
#
teak|hp|laserjet|Hewlett Packard LaserJet 3Si:\
    :lp=/dev/lpt0:sh:sd=/var/spool/lpd/teak:mx#0:\
    :if=/usr/local/libexec/hpif:
```

It overprinted each line.

The printer never advanced a line. All of the lines of text were printed on top of each other on one line.

This problem is the “opposite” of the staircase effect, described above, and is much rarer. Somewhere, the LF characters that DragonFly uses to end a line are being treated as CR characters to return the print location to the left edge of the paper, but not also down a line.

Use the printer's configuration switches or control panel to enforce the following interpretation of LF and CR characters:

Printer receives	Printer prints
CR	CR
LF	CR + LF

The printer lost characters.

While printing, the printer did not print a few characters in each line. The problem might have gotten worse as the printer ran, losing more and more characters.

The problem is that the printer cannot keep up with the speed at which the computer sends data over a serial line (this problem should not occur with printers on parallel ports). There are two ways to overcome the problem:

- If the printer supports XON/XOFF flow control, have DragonFly use it by specifying the `ixon` mode in the `ms#` capability.
- If the printer supports carrier flow control, specify the `crtsets` mode in the `ms#` capability. Make sure the cable connecting the printer to the computer is correctly wired for carrier flow control.

It printed garbage.

The printer printed what appeared to be random garbage, but not the desired text.

This is usually another symptom of incorrect communications parameters with a serial printer. Double-check the bps rate in the `br` capability, and the parity setting in the `ms#` capability; make sure the printer is using the same settings as specified in the `/etc/printcap` file.

Nothing happened.

If nothing happened, the problem is probably within DragonFly and not the hardware. Add the log file (`lf`) capability to the entry for the printer you are debugging in the `/etc/printcap` file. For example, here is the entry for `rattan`, with the `lf` capability:

```
rattan|line|diablo|lp|Diablo 630 Line Printer:\
    :sh:sd=/var/spool/lpd/rattan:\
    :lp=/dev/lpt0:\
    :if=/usr/local/libexec/if-simple:\
    :lf=/var/log/rattan.log
```

Then, try printing again. Check the log file (in our example, `/var/log/rattan.log`) to see any error messages that might appear. Based on the messages you see, try to correct the problem.

If you do not specify a `lf` capability, **LPD** uses `/dev/console` as a default.

Chapter 12 Storage

12.1 Synopsis

This chapter covers the use of disks in DragonFly. This includes memory-backed disks, network-attached disks, and standard SCSI/IDE storage devices.

After reading this chapter, you will know:

- The terminology DragonFly uses to describe the organization of data on a physical disk (partitions and slices).
- How to add additional hard disks to your system.
- How to set up virtual file systems, such as memory disks.
- How to use quotas to limit disk space usage.
- How to encrypt disks to secure them against attackers.
- How to create and burn CDs and DVDs on DragonFly.
- The various storage media options for backups.
- How to use backup programs available under DragonFly.
- How to backup to floppy disks.
- What snapshots are and how to use them efficiently.

12.2 Device Names

The following is a list of physical storage devices supported in DragonFly, and the device names associated with them.

Table 12-1. Physical Disk Naming Conventions

Drive type	Drive device name
IDE hard drives	ad
IDE CDROM drives	acd
SCSI hard drives and USB Mass storage devices	da
SCSI CDROM drives	cd
Assorted non-standard CDROM drives	mcd for Mitsumi CD-ROM, scd for Sony CD-ROM,
Floppy drives	fd
SCSI tape drives	sa
IDE tape drives	ast
Flash drives	fla for DiskOnChip® Flash device
RAID drives	aacd for Adaptec® AdvancedRAID, mlx and mlyd for Mylex®, amrd for AMI MegaRAID®, idad for Compaq Smart RAID, twed for 3ware® RAID.

12.3 Adding Disks

Originally contributed by David O'Brien.

Lets say we want to add a new SCSI disk to a machine that currently only has a single drive. First turn off the computer and install the drive in the computer following the instructions of the computer, controller, and drive manufacturer. Due to the wide variations of procedures to do this, the details are beyond the scope of this document.

Login as user `root`. After you have installed the drive, inspect `/var/run/dmesg.boot` to ensure the new disk was found. Continuing with our example, the newly added drive will be `da1` and we want to mount it on `/1` (if you are adding an IDE drive, the device name will be `ad1`).

Because DragonFly runs on IBM-PC compatible computers, it must take into account the PC BIOS partitions. These are different from the traditional BSD partitions. A PC disk has up to four BIOS partition entries. If the disk is going to be truly dedicated to DragonFly, you can use the *dedicated* mode. Otherwise, DragonFly will have to live within one of the PC BIOS partitions. DragonFly calls the PC BIOS partitions *slices* so as not to confuse them with traditional BSD partitions. You may also use slices on a disk that is dedicated to DragonFly, but used in a computer that also has another operating system installed. This is to not confuse the `fdisk` utility of the other operating system.

In the slice case the drive will be added as `/dev/da1s1e`. This is read as: SCSI disk, unit number 1 (second SCSI disk), slice 1 (PC BIOS partition 1), and `e` BSD partition. In the dedicated case, the drive will be added simply as `/dev/da1e`.

12.3.1 Using Command Line Utilities

12.3.1.1 Using Slices

This setup will allow your disk to work correctly with other operating systems that might be installed on your computer and will not confuse other operating systems' `fdisk` utilities. It is recommended to use this method for new disk installs. Only use *dedicated* mode if you have a good reason to do so!

```
# dd if=/dev/zero of=/dev/da1 bs=1k count=1
# fdisk -BI da1 #Initialize your new disk
# disklabel -B -w -r da1s1 auto #Label it.
# disklabel -e da1s1 # Edit the disklabel just created and add any partitions.
# mkdir -p /1
# newfs /dev/da1s1e # Repeat this for every partition you created.
# mount /dev/da1s1e /1 # Mount the partition(s)
# vi /etc/fstab # Add the appropriate entry/entries to your /etc/fstab.
```

If you have an IDE disk, substitute `ad` for `da`.

12.3.1.2 Dedicated

If you will not be sharing the new drive with another operating system, you may use the *dedicated* mode. Remember this mode can confuse Microsoft operating systems; however, no damage will be done by them. IBM's OS/2 however, will "appropriate" any partition it finds which it does not understand.

```
# dd if=/dev/zero of=/dev/da1 bs=1k count=1
# disklabel -Brw da1 auto
# disklabel -e da1 # create the 'e' partition
```

```
# newfs -d0 /dev/da1e
# mkdir -p /1
# vi /etc/fstab # add an entry for /dev/da1e
# mount /1
```

An alternate method is:

```
# dd if=/dev/zero of=/dev/da1 count=2
# disklabel /dev/da1 | disklabel -BrR da1 /dev/stdin
# newfs /dev/da1e
# mkdir -p /1
# vi /etc/fstab # add an entry for /dev/da1e
# mount /1
```

12.4 RAID

12.4.1 Software RAID

12.4.1.1 Concatenated Disk Driver (CCD) Configuration

Original work by Christopher Shumway. Revised by Jim Brown.

When choosing a mass storage solution the most important factors to consider are speed, reliability, and cost. It is rare to have all three in balance; normally a fast, reliable mass storage device is expensive, and to cut back on cost either speed or reliability must be sacrificed.

In designing the system described below, cost was chosen as the most important factor, followed by speed, then reliability. Data transfer speed for this system is ultimately constrained by the network. And while reliability is very important, the CCD drive described below serves online data that is already fully backed up on CD-R's and can easily be replaced.

Defining your own requirements is the first step in choosing a mass storage solution. If your requirements prefer speed or reliability over cost, your solution will differ from the system described in this section.

12.4.1.1.1 Installing the Hardware

In addition to the IDE system disk, three Western Digital 30GB, 5400 RPM IDE disks form the core of the CCD disk described below providing approximately 90GB of online storage. Ideally, each IDE disk would have its own IDE controller and cable, but to minimize cost, additional IDE controllers were not used. Instead the disks were configured with jumpers so that each IDE controller has one master, and one slave.

Upon reboot, the system BIOS was configured to automatically detect the disks attached. More importantly, DragonFly detected them on reboot:

```
ad0: 19574MB <WDC WD205BA> [39770/16/63] at ata0-master UDMA33
ad1: 29333MB <WDC WD307AA> [59598/16/63] at ata0-slave UDMA33
ad2: 29333MB <WDC WD307AA> [59598/16/63] at ata1-master UDMA33
ad3: 29333MB <WDC WD307AA> [59598/16/63] at ata1-slave UDMA33
```

Note: If DragonFly does not detect all the disks, ensure that you have jumpered them correctly. Most IDE drives also have a “Cable Select” jumper. This is *not* the jumper for the master/slave relationship. Consult the drive documentation for help in identifying the correct jumper.

Next, consider how to attach them as part of the file system. You should research both `vinum(8)` (Chapter 13) and `ccd(4)`. In this particular configuration, `ccd(4)` was chosen.

12.4.1.1.2 Setting Up the CCD

The driver `ccd(4)` allows you to take several identical disks and concatenate them into one logical file system. In order to use `ccd(4)`, you need a kernel with `ccd(4)` support built in. Add this line to your kernel configuration file, rebuild, and reinstall the kernel:

```
pseudo-device    ccd      4
```

The `ccd(4)` support can also be loaded as a kernel loadable module.

To set up `ccd(4)`, you must first use `disklabel(8)` to label the disks:

```
disklabel -r -w ad1 auto
disklabel -r -w ad2 auto
disklabel -r -w ad3 auto
```

This creates a disklabel for `ad1c`, `ad2c` and `ad3c` that spans the entire disk.

The next step is to change the disk label type. You can use `disklabel(8)` to edit the disks:

```
disklabel -e ad1
disklabel -e ad2
disklabel -e ad3
```

This opens up the current disk label on each disk with the editor specified by the `EDITOR` environment variable, typically `vi(1)`.

An unmodified disk label will look something like this:

```
16 partitions:
#          size  offset  fstype  [fsize bsize bps/cpg]
  c: 60074784      0  unused      0     0     0 # (Cyl.  0 - 59597)
```

Add a new `e` partition for `ccd(4)` to use. This can usually be copied from the `c` partition, but the `fstype` *must* be **4.2BSD**. The disk label should now look something like this:

```
16 partitions:
#          size  offset  fstype  [fsize bsize bps/cpg]
  c: 60074784      0  unused      0     0     0 # (Cyl.  0 - 59597)
  e: 60074784      0  4.2BSD      0     0     0 # (Cyl.  0 - 59597)
```

12.4.1.1.3 Building the File System

The device node for `ccd0c` may not exist yet, so to create it, perform the following commands:

```
cd /dev
sh MAKEDEV ccd0
```

Now that you have all of the disks labeled, you must build the `ccd(4)`. To do that, use `ccdconfig(8)`, with options similar to the following:

```
ccdconfig ccd0❶ 32❷ 0❸ /dev/ad1e❹ /dev/ad2e /dev/ad3e
```

The use and meaning of each option is shown below:

- ❶ The first argument is the device to configure, in this case, `/dev/ccd0c`. The `/dev/` portion is optional.
- ❷ The interleave for the file system. The interleave defines the size of a stripe in disk blocks, each normally 512 bytes. So, an interleave of 32 would be 16,384 bytes.
- ❸ Flags for `ccdconfig(8)`. If you want to enable drive mirroring, you can specify a flag here. This configuration does not provide mirroring for `ccd(4)`, so it is set at 0 (zero).
- ❹ The final arguments to `ccdconfig(8)` are the devices to place into the array. Use the complete pathname for each device.

After running `ccdconfig(8)` the `ccd(4)` is configured. A file system can be installed. Refer to `newfs(8)` for options, or simply run:

```
newfs /dev/ccd0c
```

12.4.1.1.4 Making it All Automatic

Generally, you will want to mount the `ccd(4)` upon each reboot. To do this, you must configure it first. Write out your current configuration to `/etc/ccd.conf` using the following command:

```
ccdconfig -g > /etc/ccd.conf
```

During reboot, the script `/etc/rc` runs `ccdconfig -C` if `/etc/ccd.conf` exists. This automatically configures the `ccd(4)` so it can be mounted.

Note: If you are booting into single user mode, before you can mount(8) the `ccd(4)`, you need to issue the following command to configure the array:

```
ccdconfig -C
```

To automatically mount the `ccd(4)`, place an entry for the `ccd(4)` in `/etc/fstab` so it will be mounted at boot time:

```
/dev/ccd0c          /media             ufs      rw      2       2
```


12.4.1.2 The Vinum Volume Manager

The Vinum Volume Manager is a block device driver which implements virtual disk drives. It isolates disk hardware from the block device interface and maps data in ways which result in an increase in flexibility, performance and reliability compared to the traditional slice view of disk storage. `vinum(8)` implements the RAID-0, RAID-1 and RAID-5 models, both individually and in combination.

See Chapter 13 for more information about `vinum(8)`.

12.4.2 Hardware RAID

DragonFly also supports a variety of hardware RAID controllers. These devices control a RAID subsystem without the need for DragonFly specific software to manage the array.

Using an on-card BIOS, the card controls most of the disk operations itself. The following is a brief setup description using a Promise IDE RAID controller. When this card is installed and the system is started up, it displays a prompt requesting information. Follow the instructions to enter the card's setup screen. From here, you have the ability to combine all the attached drives. After doing so, the disk(s) will look like a single drive to DragonFly. Other RAID levels can be set up accordingly.

12.4.3 Rebuilding ATA RAID1 Arrays

DragonFly allows you to hot-replace a failed disk in an array. This requires that you catch it before you reboot.

You will probably see something like the following in `/var/log/messages` or in the `dmesg(8)` output:

```
ad6 on monster1 suffered a hard error.
ad6: READ command timeout tag=0 serv=0 - resetting
ad6: trying fallback to PIO mode
ata3: resetting devices .. done
ad6: hard error reading fsbn 1116119 of 0-7 (ad6 bn 1116119; cn 1107 tn 4 sn 11) status=59 error=
ar0: WARNING - mirror lost
```

Using `atacontrol(8)`, check for further information:

```
# atacontrol list
ATA channel 0:
Master:      no device present
Slave:      acd0 <HL-DT-ST CD-ROM GCR-8520B/1.00> ATA/ATAPI rev 0

ATA channel 1:
Master:      no device present
Slave:      no device present

ATA channel 2:
Master:      ad4 <MAXTOR 6L080J4/A93.0500> ATA/ATAPI rev 5
Slave:      no device present

ATA channel 3:
Master:      ad6 <MAXTOR 6L080J4/A93.0500> ATA/ATAPI rev 5
Slave:      no device present
```

```
# atacontrol status ar0
ar0: ATA RAID1 subdisks: ad4 ad6 status: DEGRADED
```

1. You will first need to detach the disk from the array so that you can safely remove it:

```
# atacontrol detach 3
```

2. Replace the disk.

3. Reattach the disk as a spare:

```
# atacontrol attach 3
Master: ad6 <MAXTOR 6L080J4/A93.0500> ATA/ATAPI rev 5
Slave: no device present
```

4. Rebuild the array:

```
# atacontrol rebuild ar0
```

5. The rebuild command hangs until complete. However, it is possible to open another terminal (using **Alt+Fn**) and check on the progress by issuing the following command:

```
# dmesg | tail -10
[output removed]
ad6: removed from configuration
ad6: deleted from ar0 disk1
ad6: inserted into ar0 disk1 as spare

# atacontrol status ar0
ar0: ATA RAID1 subdisks: ad4 ad6 status: REBUILDING 0% completed
```

6. Wait until this operation completes.

12.5 Creating and Using Optical Media (CDs)

Contributed by Mike Meyer.

12.5.1 Introduction

CDs have a number of features that differentiate them from conventional disks. Initially, they were not writable by the user. They are designed so that they can be read continuously without delays to move the head between tracks. They are also much easier to transport between systems than similarly sized media were at the time.

CDs do have tracks, but this refers to a section of data to be read continuously and not a physical property of the disk. To produce a CD on DragonFly, you prepare the data files that are going to make up the tracks on the CD, then write the tracks to the CD.

The ISO 9660 file system was designed to deal with these differences. It unfortunately codifies file system limits that were common then. Fortunately, it provides an extension mechanism that allows properly written CDs to exceed those limits while still working with systems that do not support those extensions.

The `sysutils/mkisofs` program is used to produce a data file containing an ISO 9660 file system. It has options that support various extensions, and is described below. It is installed by default.

Which tool to use to burn the CD depends on whether your CD burner is ATAPI or something else. ATAPI CD burners use the `burncd` program that is part of the base system. SCSI and USB CD burners should use `cdrecord` from the `sysutils/cdrtools` port.

`burncd` has a limited number of supported drives. To find out if a drive is supported, see the CD-R/RW supported drives (<http://www.freebsd.dk/ata/>) list.

Note: It is possible to use `cdrecord` and other tools for SCSI drives on an ATAPI hardware with the ATAPI/CAM module.

12.5.2 mkisofs

`sysutils/mkisofs` produces an ISO 9660 file system that is an image of a directory tree in the UNIX file system name space. The simplest usage is:

```
# mkisofs -o imagefile.iso /path/to/tree
```

This command will create an `imagefile.iso` containing an ISO 9660 file system that is a copy of the tree at `/path/to/tree`. In the process, it will map the file names to names that fit the limitations of the standard ISO 9660 file system, and will exclude files that have names uncharacteristic of ISO file systems.

A number of options are available to overcome those restrictions. In particular, `-R` enables the Rock Ridge extensions common to UNIX systems, `-J` enables Joliet extensions used by Microsoft systems, and `-hfs` can be used to create HFS file systems used by Mac OS.

For CDs that are going to be used only on DragonFly systems, `-U` can be used to disable all filename restrictions. When used with `-R`, it produces a file system image that is identical to the DragonFly tree you started from, though it may violate the ISO 9660 standard in a number of ways.

The last option of general use is `-b`. This is used to specify the location of the boot image for use in producing an “El Torito” bootable CD. This option takes an argument which is the path to a boot image from the top of the tree being written to the CD. So, given that `/tmp/myboot` holds a bootable DragonFly system with the boot image in `/tmp/myboot/boot/cdboot`, you could produce the image of an ISO 9660 file system in `/tmp/bootable.iso` like so:

```
# mkisofs -U -R -b boot/cdboot -o /tmp/bootable.iso /tmp/myboot
```

Having done that, if you have `vn` configured in your kernel, you can mount the file system with:

```
# vnconfig -e vn0c /tmp/bootable.iso
# mount -t cd9660 /dev/vn0c /mnt
```

At which point you can verify that `/mnt` and `/tmp/myboot` are identical.

There are many other options you can use with `sysutils/mkisofs` to fine-tune its behavior. In particular: modifications to an ISO 9660 layout and the creation of Joliet and HFS discs. See the `mkisofs(8)` manual page for details.

12.5.3 burncd

If you have an ATAPI CD burner, you can use the `burncd` command to burn an ISO image onto a CD. `burncd` is part of the base system, installed as `/usr/sbin/burncd`. Usage is very simple, as it has few options:

```
# burncd -f cddevice data imagefile.iso fixate
```

Will burn a copy of `imagefile.iso` on `cddevice`. The default device is `/dev/acd0c`. See `burncd(8)` for options to set the write speed, eject the CD after burning, and write audio data.

12.5.4 cdrecord

If you do not have an ATAPI CD burner, you will have to use `cdrecord` to burn your CDs. `cdrecord` is not part of the base system; you must install it from either the port at `sysutils/cdrtools` or the appropriate package. Changes to the base system can cause binary versions of this program to fail, possibly resulting in a “coaster”. You should therefore either upgrade the port when you upgrade your system.

While `cdrecord` has many options, basic usage is even simpler than `burncd`. Burning an ISO 9660 image is done with:

```
# cdrecord dev=device imagefile.iso
```

The tricky part of using `cdrecord` is finding the `dev` to use. To find the proper setting, use the `-scanbus` flag of `cdrecord`, which might produce results like this:

```
# cdrecord -scanbus
Cdrecord 1.9 (i386-unknown-freebsd4.2) Copyright (C) 1995-2000 Jörg Schilling
Using libscg version 'schily-0.1'
scsibus0:
  0,0,0   0) 'SEAGATE ' 'ST39236LW      ' '0004' Disk
  0,1,0   1) 'SEAGATE ' 'ST39173W      ' '5958' Disk
  0,2,0   2) *
  0,3,0   3) 'iomega  ' 'jaz 1GB       ' 'J.86' Removable Disk
  0,4,0   4) 'NEC      ' 'CD-ROM DRIVE:466' '1.26' Removable CD-ROM
  0,5,0   5) *
  0,6,0   6) *
  0,7,0   7) *
scsibus1:
  1,0,0  100) *
  1,1,0  101) *
  1,2,0  102) *
  1,3,0  103) *
  1,4,0  104) *
  1,5,0  105) 'YAMAHA ' 'CRW4260      ' '1.0q' Removable CD-ROM
  1,6,0  106) 'ARTEC  ' 'AM12S        ' '1.06' Scanner
  1,7,0  107) *
```

This lists the appropriate `dev` value for the devices on the list. Locate your CD burner, and use the three numbers separated by commas as the value for `dev`. In this case, the CRW device is 1,5,0, so the appropriate input would be `dev=1,5,0`. There are easier ways to specify this value; see `cdrecord(1)` for details. That is also the place to look for information on writing audio tracks, controlling the speed, and other things.

12.5.5 Duplicating Audio CDs

You can duplicate an audio CD by extracting the audio data from the CD to a series of files, and then writing these files to a blank CD. The process is slightly different for ATAPI and SCSI drives.

SCSI Drives

1. Use `cdda2wav` to extract the audio.


```
% cdda2wav -v255 -D2,0 -B -Owav
```
2. Use `cdrecord` to write the `.wav` files.

```
% cdrecord -v dev=2,0 -dao -useinfo *.wav
```

Make sure that `2.0` is set appropriately, as described in Section 12.5.4.

ATAPI Drives

1. The ATAPI CD driver makes each track available as `/dev/acd0t nn` , where d is the drive number, and nn is the track number written with two decimal digits, prefixed with zero as needed. So the first track on the first disk is `/dev/acd0t01`, the second is `/dev/acd0t02`, the third is `/dev/acd0t03`, and so on.

Make sure the appropriate files exist in `/dev`.

```
# cd /dev
# sh MAKEDEV acd0t99
```

2. Extract each track using `dd(1)`. You must also use a specific block size when extracting the files.

```
# dd if=/dev/acd0t01 of=track1.cdr bs=2352
# dd if=/dev/acd0t02 of=track2.cdr bs=2352
...
```

3. Burn the extracted files to disk using `burncd`. You must specify that these are audio files, and that `burncd` should fixate the disk when finished.

```
# burncd -f /dev/acd0c audio track1.cdr track2.cdr ... fixate
```

12.5.6 Duplicating Data CDs

You can copy a data CD to a image file that is functionally equivalent to the image file created with `sysutils/mkisofs`, and you can use it to duplicate any data CD. The example given here assumes that your CDROM device is `acd0c`. Substitute your correct CDROM device. A `c` must be appended to the end of the device name to indicate the entire partition or, in the case of CDROMs, the entire disc.

```
# dd if=/dev/acd0c of=file.iso bs=2048
```

Now that you have an image, you can burn it to CD as described above.

12.5.7 Using Data CDs

Now that you have created a standard data CDROM, you probably want to mount it and read the data on it. By default, `mount(8)` assumes that a file system is of type `ufs`. If you try something like:

```
# mount /dev/cd0 /mnt
```

you will get a complaint about `Incorrect super block`, and no mount. The CDROM is not a UFS file system, so attempts to mount it as such will fail. You just need to tell `mount(8)` that the file system is of type `ISO9660`, and everything will work. You do this by specifying the `-t cd9660` option `mount(8)`. For example, if you want to mount the CDROM device, `/dev/cd0`, under `/mnt`, you would execute:

```
# mount -t cd9660 /dev/cd0 /mnt
```

Note that your device name (`/dev/cd0` in this example) could be different, depending on the interface your CDROM uses. Also, the `-t cd9660` option just executes `mount_cd9660(8)`. The above example could be shortened to:

```
# mount_cd9660 /dev/cd0 /mnt
```

You can generally use data CDROMs from any vendor in this way. Disks with certain ISO 9660 extensions might behave oddly, however. For example, Joliet disks store all filenames in two-byte Unicode characters. The DragonFly kernel does not speak Unicode (yet!), so non-English characters show up as question marks. (The CD9660 driver includes hooks to load an appropriate Unicode conversion table on the fly. Modules for some of the common encodings are available via the `sysutils/cd9660_unicode` port.)

Occasionally, you might get `Device not configured` when trying to mount a CDROM. This usually means that the CDROM drive thinks that there is no disk in the tray, or that the drive is not visible on the bus. It can take a couple of seconds for a CDROM drive to realize that it has been fed, so be patient.

Sometimes, a SCSI CDROM may be missed because it did not have enough time to answer the bus reset. If you have a SCSI CDROM please add the following option to your kernel configuration and rebuild your kernel.

```
options SCSI_DELAY=15000
```

This tells your SCSI bus to pause 15 seconds during boot, to give your CDROM drive every possible chance to answer the bus reset.

12.5.8 Burning Raw Data CDs

You can choose to burn a file directly to CD, without creating an ISO 9660 file system. Some people do this for backup purposes. This runs more quickly than burning a standard CD:

```
# burncd -f /dev/acd1 -s 12 data archive.tar.gz fixate
```

In order to retrieve the data burned to such a CD, you must read data from the raw device node:

```
# tar xzvf /dev/acd1
```

You cannot mount this disk as you would a normal CDROM. Such a CDROM cannot be read under any operating system except DragonFly. If you want to be able to mount the CD, or share data with another operating system, you must use `sysutils/mkisofs` as described above.

12.5.9 Using the ATAPI/CAM Driver

Contributed by Marc Fonvieille.

This driver allows ATAPI devices (CD-ROM, CD-RW, DVD drives etc...) to be accessed through the SCSI subsystem, and so allows the use of applications like `sysutils/cdrdao` or `cdrecord(1)`.

To use this driver, you will need to add the following lines to your kernel configuration file:

```
device atapicam
device scbus
device cd
device pass
```

You also need the following line in your kernel configuration file:

```
device ata
```

which should already be present.

Then rebuild, install your new kernel, and reboot your machine. During the boot process, your burner should show up, like so:

```
acd0: CD-RW <MATSHITA CD-RW/DVD-ROM UJDA740> at ata1-master PIO4
cd0 at ata1 bus 0 target 0 lun 0
cd0: <MATSHITA CDRW/DVD UJDA740 1.00> Removable CD-ROM SCSI-0 device
cd0: 16.000MB/s transfers
cd0: Attempt to query device size failed: NOT READY, Medium not present - tray closed
```

The drive could now be accessed via the `/dev/cd0` device name, for example to mount a CD-ROM on `/mnt`, just type the following:

```
# mount -t cd9660 /dev/cd0 /mnt
```

As root, you can run the following command to get the SCSI address of the burner:

```
# camcontrol devlist
<MATSHITA CDRW/DVD UJDA740 1.00> at scbus1 target 0 lun 0 (pass0,cd0)
```

So `1,0,0` will be the SCSI address to use with `cdrecord(1)` and other SCSI application.

For more information about ATAPI/CAM and SCSI system, refer to the `atapicam(4)` and `cam(4)` manual pages.

12.6 Creating and Using Optical Media (DVDs)

Contributed by Marc Fonvieille. With inputs from Andy Polyakov.

12.6.1 Introduction

Compared to the CD, the DVD is the next generation of optical media storage technology. The DVD can hold more data than any CD and is nowadays the standard for video publishing.

Five physical recordable formats can be defined for what we will call a recordable DVD:

- DVD-R: This was the first DVD recordable format available. The DVD-R standard is defined by the DVD Forum (<http://www.dvdforum.com/forum.shtml>). This format is write once.
- DVD-RW: This is the rewriteable version of the DVD-R standard. A DVD-RW can be rewritten about 1000 times.
- DVD-RAM: This is also a rewriteable format supported by the DVD Forum. A DVD-RAM can be seen as a removable hard drive. However, this media is not compatible with most DVD-ROM drives and DVD-Video players; only a few DVD writers support the DVD-RAM format.
- DVD+RW: This is a rewriteable format defined by the DVD+RW Alliance (<http://www.dvdrw.com/>). A DVD+RW can be rewritten about 1000 times.
- DVD+R: This format is the write once variation of the DVD+RW format.

A single layer recordable DVD can hold up to 4,700,000,000 bytes which is actually 4.38 GB or 4485 MB (1 kilobyte is 1024 bytes).

Note: A distinction must be made between the physical media and the application. For example, a DVD-Video is a specific file layout that can be written on any recordable DVD physical media: DVD-R, DVD+R, DVD-RW etc. Before choosing the type of media, you must be sure that both the burner and the DVD-Video player (a standalone player or a DVD-ROM drive on a computer) are compatible with the media under consideration.

12.6.2 Configuration

The program `growisofs(1)` will be used to perform DVD recording. This command is part of the **dvd+rw-tools** utilities (`sysutils/dvd+rw-tools`). The **dvd+rw-tools** support all DVD media types.

These tools use the SCSI subsystem to access to the devices, therefore the ATAPI/CAM support must be added to your kernel.

You also have to enable DMA access for ATAPI devices, this can be done in adding the following line to the `/boot/loader.conf` file:

```
hw.ata.atapi_dma="1"
```

Before attempting to use the **dvd+rw-tools** you should consult the dvd+rw-tools' hardware compatibility notes (<http://fy.chalmers.se/~appro/linux/DVD+RW/hcn.html>) for any information related to your DVD burner.

12.6.3 Burning Data DVDs

The `growisofs(1)` command is a frontend to `mkisofs`, it will invoke `mkisofs(8)` to create the file system layout and will perform the write on the DVD. This means you do not need to create an image of the data before the burning process.

To burn onto a DVD+R or a DVD-R the data from the `/path/to/data` directory, use the following command:

```
# growisofs -dvd-compat -Z /dev/cd0 -J -R /path/to/data
```

The options `-J -R` are passed to `mkisofs(8)` for the file system creation (in this case: an ISO 9660 file system with Joliet and Rock Ridge extensions), consult the `mkisofs(8)` manual page for more details.

The option `-z` is used for the initial session recording in any case: multiple sessions or not. The DVD device, `/dev/cd0`, must be changed according to your configuration. The `-dvd-compat` parameter will close the disk, the recording will be unappendable. In return this should provide better media compatibility with DVD-ROM drives.

It is also possible to burn a pre-mastered image, for example to burn the image `imagefile.iso`, we will run:

```
# growisofs -dvd-compat -Z /dev/cd0=imagefile.iso
```

The write speed should be detected and automatically set according to the media and the drive being used. If you want to force the write speed, use the `-speed=` parameter. For more information, read the `growisofs(1)` manual page.

12.6.4 Burning a DVD-Video

A DVD-Video is a specific file layout based on ISO 9660 and the micro-UDF (M-UDF) specifications. The DVD-Video also presents a specific data structure hierarchy, it is the reason why you need a particular program such as `multimedia/dvdauthor` to author the DVD.

If you already have an image of the DVD-Video file system, just burn it in the same way as for any image, see the previous section for an example. If you have made the DVD authoring and the result is in, for example, the directory `/path/to/video`, the following command should be used to burn the DVD-Video:

```
# growisofs -Z /dev/cd0 -dvd-video /path/to/video
```

The `-dvd-video` option will be passed down to `mkisofs(8)` and will instruct it to create a DVD-Video file system layout. Beside this, the `-dvd-video` option implies `-dvd-compat growisofs(1)` option.

12.6.5 Using a DVD+RW

Unlike CD-RW, a virgin DVD+RW needs to be formatted before first use. The `growisofs(1)` program will take care of it automatically whenever appropriate, which is the *recommended* way. However you can use the `dvd+rw-format` command to format the DVD+RW:

```
# dvd+rw-format /dev/cd0
```

You need to perform this operation just once, keep in mind that only virgin DVD+RW medias need to be formatted. Then you can burn the DVD+RW in the way seen in previous sections.

If you want to burn new data (burn a totally new file system not append some data) onto a DVD+RW, you do not need to blank it, you just have to write over the previous recording (in performing a new initial session), like this:

```
# growisofs -Z /dev/cd0 -J -R /path/to/newdata
```

DVD+RW format offers the possibility to easily append data to a previous recording. The operation consists in merging a new session to the existing one, it is not multisession writing, `growisofs(1)` will *grow* the ISO 9660 file system present on the media.

For example, if we want to append data to our previous DVD+RW, we have to use the following:

```
# growisofs -M /dev/cd0 -J -R /path/to/nextdata
```

The same `mkisofs(8)` options we used to burn the initial session should be used during next writes.

Note: You may want to use the `-dvd-compat` option if you want better media compatibility with DVD-ROM drives. In the DVD+RW case, this will not prevent you from adding data.

If for any reason you really want to blank the media, do the following:

```
# growisofs -Z /dev/cd0=/dev/zero
```

12.6.6 Using a DVD-RW

A DVD-RW accepts two disc formats: the incremental sequential one and the restricted overwrite. By default DVD-RW discs are in sequential format.

A virgin DVD-RW can be directly written without the need of a formatting operation, however a non-virgin DVD-RW in sequential format needs to be blanked before to be able to write a new initial session.

To blank a DVD-RW in sequential mode, run:

```
# dvd+rw-format -blank=full /dev/cd0
```

Note: A full blanking (`-blank=full`) will take about one hour on a 1x media. A fast blanking can be performed using the `-blank` option if the DVD-RW will be recorded in Disk-At-Once (DAO) mode. To burn the DVD-RW in DAO mode, use the command:

```
# growisofs -use-the-force-luke=dao -Z /dev/cd0=imagefile.iso
```

The `-use-the-force-luke=dao` option should not be required since `growisofs(1)` attempts to detect minimally (fast blanked) media and engage DAO write.

In fact one should use restricted overwrite mode with any DVD-RW, this format is more flexible than the default incremental sequential one.

To write data on a sequential DVD-RW, use the same instructions as for the other DVD formats:

```
# growisofs -Z /dev/cd0 -J -R /path/to/data
```

If you want to append some data to your previous recording, you will have to use the `growisofs(1)` `-M` option.

However, if you perform data addition on a DVD-RW in incremental sequential mode, a new session will be created on the disc and the result will be a multi-session disc.

A DVD-RW in restricted overwrite format does not need to be blanked before a new initial session, you just have to overwrite the disc with the `-Z` option, this is similar to the DVD+RW case. It is also possible to grow an existing ISO 9660 file system written on the disc in a same way as for a DVD+RW with the `-M` option. The result will be a one-session DVD.

To put a DVD-RW in the restricted overwrite format, the following command must be used:

```
# dvd+rw-format /dev/cd0
```

To change back to the sequential format use:

```
# dvd+rw-format -blank=full /dev/cd0
```

12.6.7 Multisession

Very few DVD-ROM and DVD-Video players support multisession DVDs, they will most of time, hopefully, only read the first session. DVD+R, DVD-R and DVD-RW in sequential format can accept multiple sessions, the notion of multiple sessions does not exist for the DVD+RW and the DVD-RW restricted overwrite formats.

Using the following command after an initial (non-closed) session on a DVD+R, DVD-R, or DVD-RW in sequential format, will add a new session to the disc:

```
# growisofs -M /dev/cd0 -J -R /path/to/nextdata
```

Using this command line with a DVD+RW or a DVD-RW in restricted overwrite mode, will append data in merging the new session to the existing one. The result will be a single-session disc. This is the way used to add data after an initial write on these medias.

Note: Some space on the media is used between each session for end and start of sessions. Therefore, one should add sessions with large amount of data to optimize media space. The number of sessions is limited to 154 for a DVD+R and about 2000 for a DVD-R.

12.6.8 For More Information

To obtain more information about a DVD, the `dvd+rw-mediainfo /dev/cd0` command can be ran with the disc in the drive.

More information about the **dvd+rw-tools** can be found in the `growisofs(1)` manual page, on the `dvd+rw-tools` web site (<http://fy.chalmers.se/~appro/linux/DVD+RW/>) and in the `cdwrite` mailing list (<http://lists.debian.org/cdwrite/>) archives.

Note: The `dvd+rw-mediainfo` output of the resulting recording or the media with issues is mandatory for any problem report. Without this output, it will be quite impossible to help you.

12.7 Creating and Using Floppy Disks

Original work by Julio Merino. Rewritten by Martin Karlsson.

Storing data on floppy disks is sometimes useful, for example when one does not have any other removable storage media or when one needs to transfer small amounts of data to another computer.

This section will explain how to use floppy disks in DragonFly. It will primarily cover formatting and usage of 3.5inch DOS floppies, but the concepts are similar for other floppy disk formats.

12.7.1 Formatting Floppies

12.7.1.1 The Device

Floppy disks are accessed through entries in `/dev`, just like other devices. To access the raw floppy disk, one uses `/dev/fdN`, where *N* stands for the drive number, usually 0, or `/dev/fdNX`, where *X* stands for a letter.

There are also `/dev/fdN.size` devices, where *size* is a floppy disk size in kilobytes. These entries are used at low-level format time to determine the disk size. 1440kB is the size that will be used in the following examples.

Sometimes the entries under `/dev` will have to be (re)created. To do that, issue:

```
# cd /dev && ./MAKEDEV "fd*"
```

12.7.1.2 Formatting

A floppy disk needs to be low-level formatted before it can be used. This is usually done by the vendor, but formatting is a good way to check media integrity. Although it is possible to force larger (or smaller) disk sizes, 1440kB is what most floppy disks are designed for.

To low-level format the floppy disk you need to use `fdformat(1)`. This utility expects the device name as an argument.

Make note of any error messages, as these can help determine if the disk is good or bad.

Use the `/dev/fdN.size` devices to format the floppy. Insert a new 3.5inch floppy disk in your drive and issue:

```
# /usr/sbin/fdformat /dev/fd0.1440
```

12.7.2 The Disk Label

After low-level formatting the disk, you will need to place a disk label on it. This disk label will be destroyed later, but it is needed by the system to determine the size of the disk and its geometry later.

The new disk label will take over the whole disk, and will contain all the proper information about the geometry of the floppy. The geometry values for the disk label are listed in `/etc/disktab`.

You can run `nowdisklabel(8)` like so:

```
# /sbin/disklabel -B -r -w /dev/fd0 fd1440
```

12.7.3 The File System

Now the floppy is ready to be high-level formatted. This will place a new file system on it, which will let DragonFly read and write to the disk. After creating the new file system, the disk label is destroyed, so if you want to reformat the disk, you will have to recreate the disk label.

The floppy's file system can be either UFS or FAT. FAT is generally a better choice for floppies.

To put a new file system on the floppy, issue:

```
# /sbin/newfs_msdos /dev/fd0
```

The disk is now ready for use.

12.7.4 Using the Floppy

To use the floppy, mount it with `mount_msdos(8)`. One can also use `sysutils/mttools` from `pkgsrc`.

12.8 Creating and Using Data Tapes

The major tape media are the 4mm, 8mm, QIC, mini-cartridge and DLT.

12.8.1 4mm (DDS: Digital Data Storage)

4mm tapes are replacing QIC as the workstation backup media of choice. This trend accelerated greatly when Conner purchased Archive, a leading manufacturer of QIC drives, and then stopped production of QIC drives. 4mm drives are small and quiet but do not have the reputation for reliability that is enjoyed by 8mm drives. The cartridges are less expensive and smaller (3 x 2 x 0.5 inches, 76 x 51 x 12 mm) than 8mm cartridges. 4mm, like 8mm, has comparatively short head life for the same reason, both use helical scan.

Data throughput on these drives starts ~150 kB/s, peaking at ~500 kB/s. Data capacity starts at 1.3 GB and ends at 2.0 GB. Hardware compression, available with most of these drives, approximately doubles the capacity. Multi-drive tape library units can have 6 drives in a single cabinet with automatic tape changing. Library capacities reach 240 GB.

The DDS-3 standard now supports tape capacities up to 12 GB (or 24 GB compressed).

4mm drives, like 8mm drives, use helical-scan. All the benefits and drawbacks of helical-scan apply to both 4mm and 8mm drives.

Tapes should be retired from use after 2,000 passes or 100 full backups.

12.8.2 8mm (Exabyte)

8mm tapes are the most common SCSI tape drives; they are the best choice of exchanging tapes. Nearly every site has an Exabyte 2 GB 8mm tape drive. 8mm drives are reliable, convenient and quiet. Cartridges are inexpensive and small (4.8 x 3.3 x 0.6 inches; 122 x 84 x 15 mm). One downside of 8mm tape is relatively short head and tape life due to the high rate of relative motion of the tape across the heads.

Data throughput ranges from ~250 kB/s to ~500 kB/s. Data sizes start at 300 MB and go up to 7 GB. Hardware compression, available with most of these drives, approximately doubles the capacity. These drives are available as single units or multi-drive tape libraries with 6 drives and 120 tapes in a single cabinet. Tapes are changed automatically by the unit. Library capacities reach 840+ GB.

The Exabyte “Mammoth” model supports 12 GB on one tape (24 GB with compression) and costs approximately twice as much as conventional tape drives.

Data is recorded onto the tape using helical-scan, the heads are positioned at an angle to the media (approximately 6 degrees). The tape wraps around 270 degrees of the spool that holds the heads. The spool spins while the tape slides over the spool. The result is a high density of data and closely packed tracks that angle across the tape from one edge to the other.

12.8.3 QIC

QIC-150 tapes and drives are, perhaps, the most common tape drive and media around. QIC tape drives are the least expensive “serious” backup drives. The downside is the cost of media. QIC tapes are expensive compared to 8mm or 4mm tapes, up to 5 times the price per GB data storage. But, if your needs can be satisfied with a half-dozen tapes, QIC may be the correct choice. QIC is the *most* common tape drive. Every site has a QIC drive of some density or another. Therein lies the rub, QIC has a large number of densities on physically similar (sometimes identical) tapes. QIC drives are not quiet. These drives audibly seek before they begin to record data and are clearly audible whenever reading, writing or seeking. QIC tapes measure (6 x 4 x 0.7 inches; 15.2 x 10.2 x 1.7 mm). Mini-cartridges, which also use 1/4" wide tape are discussed separately. Tape libraries and changers are not available.

Data throughput ranges from ~150 kB/s to ~500 kB/s. Data capacity ranges from 40 MB to 15 GB. Hardware compression is available on many of the newer QIC drives. QIC drives are less frequently installed; they are being supplanted by DAT drives.

Data is recorded onto the tape in tracks. The tracks run along the long axis of the tape media from one end to the other. The number of tracks, and therefore the width of a track, varies with the tape’s capacity. Most if not all newer drives provide backward-compatibility at least for reading (but often also for writing). QIC has a good reputation regarding the safety of the data (the mechanics are simpler and more robust than for helical scan drives).

Tapes should be retired from use after 5,000 backups.

12.8.4 XXX* Mini-Cartridge

12.8.5 DLT

DLT has the fastest data transfer rate of all the drive types listed here. The 1/2" (12.5mm) tape is contained in a single spool cartridge (4 x 4 x 1 inches; 100 x 100 x 25 mm). The cartridge has a swinging gate along one entire side of the cartridge. The drive mechanism opens this gate to extract the tape leader. The tape leader has an oval hole in it which the drive uses to “hook” the tape. The take-up spool is located inside the tape drive. All the other tape cartridges listed here (9 track tapes are the only exception) have both the supply and take-up spools located inside the tape cartridge itself.

Data throughput is approximately 1.5 MB/s, three times the throughput of 4mm, 8mm, or QIC tape drives. Data capacities range from 10 GB to 20 GB for a single drive. Drives are available in both multi-tape changers and multi-tape, multi-drive tape libraries containing from 5 to 900 tapes over 1 to 20 drives, providing from 50 GB to 9 TB of storage.

With compression, DLT Type IV format supports up to 70 GB capacity.

Data is recorded onto the tape in tracks parallel to the direction of travel (just like QIC tapes). Two tracks are written at once. Read/write head lifetimes are relatively long; once the tape stops moving, there is no relative motion between the heads and the tape.

12.8.6 AIT

AIT is a new format from Sony, and can hold up to 50 GB (with compression) per tape. The tapes contain memory chips which retain an index of the tape’s contents. This index can be rapidly read by the tape drive to determine the position of files on the tape, instead of the several minutes that would be required for other tapes. Software such as

SAMS:Alexandria can operate forty or more AIT tape libraries, communicating directly with the tape's memory chip to display the contents on screen, determine what files were backed up to which tape, locate the correct tape, load it, and restore the data from the tape.

Libraries like this cost in the region of \$20,000, pricing them a little out of the hobbyist market.

12.8.7 Using a New Tape for the First Time

The first time that you try to read or write a new, completely blank tape, the operation will fail. The console messages should be similar to:

```
sa0(ncr1:4:0): NOT READY asc:4,1
sa0(ncr1:4:0): Logical unit is in process of becoming ready
```

The tape does not contain an Identifier Block (block number 0). All QIC tape drives since the adoption of QIC-525 standard write an Identifier Block to the tape. There are two solutions:

- `mt fsf 1` causes the tape drive to write an Identifier Block to the tape.
- Use the front panel button to eject the tape.

Re-insert the tape and `dump` data to the tape.

`dump` will report `DUMP: End of tape detected` and the console will show: `HARDWARE FAILURE info:280 asc:80,96`.

rewind the tape using: `mt rewind`.

Subsequent tape operations are successful.

12.9 Backups to Floppies

12.9.1 Can I Use Floppies for Backing Up My Data?

Floppy disks are not really a suitable media for making backups as:

- The media is unreliable, especially over long periods of time.
- Backing up and restoring is very slow.
- They have a very limited capacity (the days of backing up an entire hard disk onto a dozen or so floppies has long since passed).

However, if you have no other method of backing up your data then floppy disks are better than no backup at all.

If you do have to use floppy disks then ensure that you use good quality ones. Floppies that have been lying around the office for a couple of years are a bad choice. Ideally use new ones from a reputable manufacturer.

12.9.2 So How Do I Backup My Data to Floppies?

The best way to backup to floppy disk is to use `tar(1)` with the `-M` (multi volume) option, which allows backups to span multiple floppies.

To backup all the files in the current directory and sub-directory use this (as `root`):

```
# tar Mcvf /dev/fd0 *
```

When the first floppy is full `tar(1)` will prompt you to insert the next volume (because `tar(1)` is media independent it refers to volumes; in this context it means floppy disk).

Prepare volume #2 for `/dev/fd0` and hit return:

This is repeated (with the volume number incrementing) until all the specified files have been archived.

12.9.3 Can I Compress My Backups?

Unfortunately, `tar(1)` will not allow the `-z` option to be used for multi-volume archives. You could, of course, `gzip(1)` all the files, `tar(1)` them to the floppies, then `gunzip(1)` the files again!

12.9.4 How Do I Restore My Backups?

To restore the entire archive use:

```
# tar Mxvf /dev/fd0
```

There are two ways that you can use to restore only specific files. First, you can start with the first floppy and use:

```
# tar Mxvf /dev/fd0 filename
```

The utility `tar(1)` will prompt you to insert subsequent floppies until it finds the required file.

Alternatively, if you know which floppy the file is on then you can simply insert that floppy and use the same command as above. Note that if the first file on the floppy is a continuation from the previous one then `tar(1)` will warn you that it cannot restore it, even if you have not asked it to!

12.10 Backup Basics

The three major backup programs are `dump(8)`, `tar(1)`, and `cpio(1)`.

12.10.1 Dump and Restore

The traditional UNIX backup programs are `dump` and `restore`. They operate on the drive as a collection of disk blocks, below the abstractions of files, links and directories that are created by the file systems. `dump` backs up an entire file system on a device. It is unable to backup only part of a file system or a directory tree that spans more than one file system. `dump` does not write files and directories to tape, but rather writes the raw data blocks that comprise files and directories.

Note: If you use `dump` on your root directory, you would not back up `/home`, `/usr` or many other directories since these are typically mount points for other file systems or symbolic links into those file systems.

`dump` has quirks that remain from its early days in Version 6 of AT&T UNIX (circa 1975). The default parameters are suitable for 9-track tapes (6250 bpi), not the high-density media available today (up to 62,182 ftpi). These defaults must be overridden on the command line to utilize the capacity of current tape drives.

It is also possible to backup data across the network to a tape drive attached to another computer with `rdump` and `rrestore`. Both programs rely upon `rcmd` and `ruserok` to access the remote tape drive. Therefore, the user performing the backup must be listed in the `.rhosts` file on the remote computer. The arguments to `rdump` and `rrestore` must be suitable to use on the remote computer. When `rdumping` from a DragonFly computer to an Exabyte tape drive connected to a Sun called `komodo`, use:

```
# /sbin/rdump 0dsbfu 54000 13000 126 komodo:/dev/nsa8 /dev/da0a 2>&1
```

Beware: there are security implications to allowing `.rhosts` authentication. Evaluate your situation carefully.

It is also possible to use `dump` and `restore` in a more secure fashion over `ssh`.

Example 12-1. Using `dump` over `ssh`

```
# /sbin/dump -0uan -f - /usr | gzip -2 | ssh1 -c blowfish \
    targetuser@targetmachine.example.com dd of=/mybigfiles/dump-usr-10.gz
```

Or using `dump`'s built-in method, setting the environment variable `RSH`:

Example 12-2. Using `dump` over `ssh` with `RSH` set

```
# RSH=/usr/bin/ssh /sbin/dump -0uan -f targetuser@targetmachine.example.com:/dev/sa0
```

12.10.2 tar

`tar(1)` also dates back to Version 6 of AT&T UNIX (circa 1975). `tar` operates in cooperation with the file system; `tar` writes files and directories to tape. `tar` does not support the full range of options that are available from `cpio(1)`, but `tar` does not require the unusual command pipeline that `cpio` uses.

Most versions of `tar` do not support backups across the network. The GNU version of `tar`, which DragonFly utilizes, supports remote devices using the same syntax as `rdump`. To `tar` to an Exabyte tape drive connected to a Sun called `komodo`, use:

```
# /usr/bin/tar cf komodo:/dev/nsa8 . 2>&1
```

For versions without remote device support, you can use a pipeline and `rsh` to send the data to a remote tape drive.

```
# tar cf - . | rsh hostname dd of=tape-device obs=20b
```

If you are worried about the security of backing up over a network you should use the `ssh` command instead of `rsh`.

12.10.3 cpio

`cpio(1)` is the original UNIX file interchange tape program for magnetic media. `cpio` has options (among many others) to perform byte-swapping, write a number of different archive formats, and pipe the data to other programs. This last feature makes `cpio` an excellent choice for installation media. `cpio` does not know how to walk the directory tree and a list of files must be provided through `stdin`.

`cpio` does not support backups across the network. You can use a pipeline and `rsh` to send the data to a remote tape drive.

```
# for f in directory_list; do
find $f >> backup.list
done
# cpio -v -o --format=newc < backup.list | ssh user@host "cat > backup_device"
```

Where `directory_list` is the list of directories you want to back up, `user@host` is the user/hostname combination that will be performing the backups, and `backup_device` is where the backups should be written to (e.g., `/dev/nsa0`).

12.10.4 pax

`pax(1)` is IEEE/POSIX's answer to `tar` and `cpio`. Over the years the various versions of `tar` and `cpio` have gotten slightly incompatible. So rather than fight it out to fully standardize them, POSIX created a new archive utility. `pax` attempts to read and write many of the various `cpio` and `tar` formats, plus new formats of its own. Its command set more resembles `cpio` than `tar`.

12.10.5 Amanda

Amanda (Advanced Maryland Network Disk Archiver) is a client/server backup system, rather than a single program. An **Amanda** server will backup to a single tape drive any number of computers that have **Amanda** clients and a network connection to the **Amanda** server. A common problem at sites with a number of large disks is that the length of time required to backup to data directly to tape exceeds the amount of time available for the task. **Amanda** solves this problem. **Amanda** can use a "holding disk" to backup several file systems at the same time. **Amanda** creates "archive sets": a group of tapes used over a period of time to create full backups of all the file systems listed in **Amanda**'s configuration file. The "archive set" also contains nightly incremental (or differential) backups of all the file systems. Restoring a damaged file system requires the most recent full backup and the incremental backups.

The configuration file provides fine control of backups and the network traffic that **Amanda** generates. **Amanda** will use any of the above backup programs to write the data to tape. **Amanda** is available as either a port or a package, it is not installed by default.

12.10.6 Do Nothing

"Do nothing" is not a computer program, but it is the most widely used backup strategy. There are no initial costs. There is no backup schedule to follow. Just say no. If something happens to your data, grin and bear it!

If your time and your data is worth little to nothing, then "Do nothing" is the most suitable backup program for your computer. But beware, UNIX is a useful tool, you may find that within six months you have a collection of files that are valuable to you.

“Do nothing” is the correct backup method for `/usr/obj` and other directory trees that can be exactly recreated by your computer. An example is the files that comprise the HTML or PostScript version of this Handbook. These document formats have been created from SGML input files. Creating backups of the HTML or PostScript files is not necessary. The SGML files are backed up regularly.

12.10.7 Which Backup Program Is Best?

`dump(8)` *Period*. Elizabeth D. Zwicky torture tested all the backup programs discussed here. The clear choice for preserving all your data and all the peculiarities of UNIX file systems is `dump`. Elizabeth created file systems containing a large variety of unusual conditions (and some not so unusual ones) and tested each program by doing a backup and restore of those file systems. The peculiarities included: files with holes, files with holes and a block of nulls, files with funny characters in their names, unreadable and unwritable files, devices, files that change size during the backup, files that are created/deleted during the backup and more. She presented the results at LISA V in Oct. 1991. Read the “Torture-testing Backup and Archive Programs” (<http://berdmann.dyndns.org/zwicky/testdump.doc.html>) study for more information.

12.10.8 Emergency Restore Procedure

12.10.8.1 Before the Disaster

There are only four steps that you need to perform in preparation for any disaster that may occur.

First, print the disklabel from each of your disks (e.g. `disklabel da0 | lpr`), your file system table (`/etc/fstab`) and all boot messages, two copies of each.

Second, determine that the boot and fix-it floppies (`boot.flp` and `fixit.flp`) have all your devices. The easiest way to check is to reboot your machine with the boot floppy in the floppy drive and check the boot messages. If all your devices are listed and functional, skip on to step three.

Otherwise, you have to create two custom bootable floppies which have a kernel that can mount all of your disks and access your tape drive. These floppies must contain: `fdisk`, `disklabel`, `newfs`, `mount`, and whichever backup program you use. These programs must be statically linked. If you use `dump`, the floppy must contain `restore`.

Third, create backup tapes regularly. Any changes that you make after your last backup may be irretrievably lost. Write-protect the backup tapes.

Fourth, test the floppies (either `boot.flp` and `fixit.flp` or the two custom bootable floppies you made in step two.) and backup tapes. Make notes of the procedure. Store these notes with the bootable floppy, the printouts and the backup tapes. You will be so distraught when restoring that the notes may prevent you from destroying your backup tapes (How? In place of `tar xvf /dev/sa0`, you might accidentally type `tar cvf /dev/sa0` and over-write your backup tape).

For an added measure of security, make bootable floppies and two backup tapes each time. Store one of each at a remote location. A remote location is NOT the basement of the same office building. A number of firms in the World Trade Center learned this lesson the hard way. A remote location should be physically separated from your computers and disk drives by a significant distance.

Example 12-3. A Script for Creating a Bootable Floppy

```
#!/bin/sh
```

```

#
# create a restore floppy
#
# format the floppy
#
PATH=/bin:/sbin:/usr/sbin:/usr/bin

fdformat -q fd0
if [ $? -ne 0 ]
then
    echo "Bad floppy, please use a new one"
    exit 1
fi

# place boot blocks on the floppy
#
disklabel -w -B /dev/fd0c fd1440

#
# newfs the one and only partition
#
newfs -t 2 -u 18 -l 1 -c 40 -i 5120 -m 5 -o space /dev/fd0a

#
# mount the new floppy
#
mount /dev/fd0a /mnt

#
# create required directories
#
mkdir /mnt/dev
mkdir /mnt/bin
mkdir /mnt/sbin
mkdir /mnt/etc
mkdir /mnt/root
mkdir /mnt/mnt # for the root partition
mkdir /mnt/tmp
mkdir /mnt/var

#
# populate the directories
#
if [ ! -x /sys/compile/MINI/kernel ]
then
    cat << EOM
The MINI kernel does not exist, please create one.
Here is an example config file:
#
# MINI -- A kernel to get &os; onto a disk.
#
machine          "i386"
cpu              "I486_CPU"

```

```

ident      MINI
maxusers   5

options    INET                # needed for _tcp _icmpstat _ipstat
           #                   #           _udpstat _tcpstat _udb
options    FFS                 #Berkeley Fast File System
options    FAT_CURSOR          #block cursor in syscons or pcons
options    SCSI_DELAY=15      #Be pessimistic about Joe SCSI device
options    NCONS=2            #1 virtual consoles
options    USERCONFIG         #Allow user configuration with -c XXX

config     kernel root on da0 swap on da0 and da1 dumps on da0

device     isa0
device     pci0

device     fdc0 at isa? port "IO_FD1" bio irq 6 drq 2 vector fdintr
device     fd0 at fdc0 drive 0

device     ncr0

device     scbus0

device     sc0 at isa? port "IO_KBD" tty irq 1 vector scintr
device     npx0 at isa? port "IO_NPX" irq 13 vector npxintr

device     da0
device     da1
device     da2

device     sa0

pseudo-device loop            # required by INET
pseudo-device gzip           # Exec gzipped a.out's
EOM
  exit 1
fi

cp -f /sys/compile/MINI/kernel /mnt

gzip -c -best /sbin/init > /mnt/sbin/init
gzip -c -best /sbin/fsck > /mnt/sbin/fsck
gzip -c -best /sbin/mount > /mnt/sbin/mount
gzip -c -best /sbin/halt > /mnt/sbin/halt
gzip -c -best /sbin/restore > /mnt/sbin/restore

gzip -c -best /bin/sh > /mnt/bin/sh
gzip -c -best /bin/sync > /mnt/bin/sync

cp /root/.profile /mnt/root

cp -f /dev/MAKEDEV /mnt/dev
chmod 755 /mnt/dev/MAKEDEV

```

```

chmod 500 /mnt/sbin/init
chmod 555 /mnt/sbin/fsck /mnt/sbin/mount /mnt/sbin/halt
chmod 555 /mnt/bin/sh /mnt/bin/sync
chmod 6555 /mnt/sbin/restore

#
# create the devices nodes
#
cd /mnt/dev
./MAKEDEV std
./MAKEDEV da0
./MAKEDEV da1
./MAKEDEV da2
./MAKEDEV sa0
./MAKEDEV pty0
cd /

#
# create minimum file system table
#
cat > /mnt/etc/fstab <<EOM
/dev/fd0a    /      ufs    rw  1  1
EOM

#
# create minimum passwd file
#
cat > /mnt/etc/passwd <<EOM
root:*:0:0:Charlie &:/root:/bin/sh
EOM

cat > /mnt/etc/master.passwd <<EOM
root::0:0::0:0:Charlie &:/root:/bin/sh
EOM

chmod 600 /mnt/etc/master.passwd
chmod 644 /mnt/etc/passwd
/usr/sbin/pwd_mkdb -d/mnt/etc /mnt/etc/master.passwd

#
# umount the floppy and inform the user
#
/sbin/umount /mnt
echo "The floppy has been unmounted and is now ready."

```

12.10.8.2 After the Disaster

The key question is: did your hardware survive? You have been doing regular backups so there is no need to worry about the software.

If the hardware has been damaged, the parts should be replaced before attempting to use the computer.

If your hardware is okay, check your floppies. If you are using a custom boot floppy, boot single-user (type `-s` at the `boot :` prompt). Skip the following paragraph.

If you are using the `boot.flp` and `fixit.flp` floppies, keep reading. Insert the `boot.flp` floppy in the first floppy drive and boot the computer. The original install menu will be displayed on the screen. Select the `Fixit--Repair` mode with `CDROM` or `floppy.` option. Insert the `fixit.flp` when prompted. `restore` and the other programs that you need are located in `/mnt2/stand`.

Recover each file system separately.

Try to `mount` (e.g. `mount /dev/da0a /mnt`) the root partition of your first disk. If the `disklabel` was damaged, use `disklabel` to re-partition and label the disk to match the label that you printed and saved. Use `newfs` to re-create the file systems. Re-mount the root partition of the floppy read-write (`mount -u -o rw /mnt`). Use your backup program and backup tapes to recover the data for this file system (e.g. `restore vrf /dev/sa0`). Unmount the file system (e.g. `umount /mnt`). Repeat for each file system that was damaged.

Once your system is running, backup your data onto new tapes. Whatever caused the crash or data loss may strike again. Another hour spent now may save you from further distress later.

12.11 Network, Memory, and File-Backed File Systems

Reorganized and enhanced by Marc Fonvieille.

Aside from the disks you physically insert into your computer: floppies, CDs, hard drives, and so forth; other forms of disks are understood by DragonFly - the *virtual disks*.

These include network file systems such as the Network File System, memory-based file systems and file-backed file systems.

12.11.1 File-Backed File System

The utility `vnconfig(8)` configures and enables `vnode` pseudo-disk devices. A *vnode* is a representation of a file, and is the focus of file activity. This means that `vnconfig(8)` uses files to create and operate a file system. One possible use is the mounting of floppy or CD images kept in files.

To use `vnconfig(8)`, you need `vn(4)` support in your kernel configuration file:

```
pseudo-device vn
```

To mount an existing file system image:

Example 12-4. Using `vnconfig` to Mount an Existing File System Image

```
# vnconfig vn0 diskimage
# mount /dev/vn0c /mnt
```

To create a new file system image with `vnconfig(8)`:

Example 12-5. Creating a New File-Backed Disk with `vnconfig`

```
# dd if=/dev/zero of=newimage bs=1k count=5k
5120+0 records in
5120+0 records out
# vnconfig -s labels -c vn0 newimage
# disklabel -r -w vn0 auto
# newfs vn0c
Warning: 2048 sector(s) in last cylinder unallocated
/dev/vn0c:      10240 sectors in 3 cylinders of 1 tracks, 4096 sectors
              5.0MB in 1 cyl groups (16 c/g, 32.00MB/g, 1280 i/g)
super-block backups (for fsck -b #) at:
 32
# mount /dev/vn0c /mnt
# df /mnt
Filesystem 1K-blocks      Used    Avail Capacity  Mounted on
/dev/vn0c      4927          1     4532      0%    /mnt
```

12.11.2 Memory-Based File System

The `md(4)` driver is a simple, efficient means to create memory file systems. `malloc(9)` is used to allocate the memory.

Simply take a file system you have prepared with, for example, `vnconfig(8)`, and:

Example 12-6. `md` Memory Disk

```
# dd if=newimage of=/dev/md0
5120+0 records in
5120+0 records out
# mount /dev/md0c /mnt
# df /mnt
Filesystem 1K-blocks      Used    Avail Capacity  Mounted on
/dev/md0c      4927          1     4532      0%    /mnt
```

For more details, please refer to `md(4)` manual page.

12.11.3 Detaching a Memory Disk from the System

When a memory-based or file-based file system is not used, you should release all resources to the system. The first thing to do is to unmount the file system, then use `mdconfig(8)` to detach the disk from the system and release the resources.

For example to detach and free all resources used by `/dev/md4`:

```
# mdconfig -d -u 4
```

It is possible to list information about configured `md(4)` devices in using the command `mdconfig -l`.

`vnconfig(8)` is used to detach the device. For example to detach and free all resources used by `/dev/vn4`:


```
# vnconfig -u vn4
```

12.12 File System Quotas

Quotas are an optional feature of the operating system that allow you to limit the amount of disk space and/or the number of files a user or members of a group may allocate on a per-file system basis. This is used most often on timesharing systems where it is desirable to limit the amount of resources any one user or group of users may allocate. This will prevent one user or group of users from consuming all of the available disk space.

12.12.1 Configuring Your System to Enable Disk Quotas

Before attempting to use disk quotas, it is necessary to make sure that quotas are configured in your kernel. This is done by adding the following line to your kernel configuration file:

```
options QUOTA
```

The stock `GENERIC` kernel does not have this enabled by default, so you will have to configure, build and install a custom kernel in order to use disk quotas. Please refer to Chapter 9 for more information on kernel configuration.

Next you will need to enable disk quotas in `/etc/rc.conf`. This is done by adding the line:

```
enable_quotas="YES"
```

For finer control over your quota startup, there is an additional configuration variable available. Normally on bootup, the quota integrity of each file system is checked by the `quotacheck(8)` program. The `quotacheck(8)` facility insures that the data in the quota database properly reflects the data on the file system. This is a very time consuming process that will significantly affect the time your system takes to boot. If you would like to skip this step, a variable in `/etc/rc.conf` is made available for the purpose:

```
check_quotas="NO"
```

Finally you will need to edit `/etc/fstab` to enable disk quotas on a per-file system basis. This is where you can either enable user or group quotas or both for all of your file systems.

To enable per-user quotas on a file system, add the `userquota` option to the options field in the `/etc/fstab` entry for the file system you want to enable quotas on. For example:

```
/dev/dals2g /home ufs rw,userquota 1 2
```

Similarly, to enable group quotas, use the `groupquota` option instead of `userquota`. To enable both user and group quotas, change the entry as follows:

```
/dev/dals2g /home ufs rw,userquota,groupquota 1 2
```

By default, the quota files are stored in the root directory of the file system with the names `quota.user` and `quota.group` for user and group quotas respectively. See `fstab(5)` for more information. Even though the `fstab(5)` manual page says that you can specify an alternate location for the quota files, this is not recommended because the various quota utilities do not seem to handle this properly.

At this point you should reboot your system with your new kernel. `/etc/rc` will automatically run the appropriate commands to create the initial quota files for all of the quotas you enabled in `/etc/fstab`, so there is no need to manually create any zero length quota files.

In the normal course of operations you should not be required to run the `quotacheck(8)`, `quotaon(8)`, or `quotaoff(8)` commands manually. However, you may want to read their manual pages just to be familiar with their operation.

12.12.2 Setting Quota Limits

Once you have configured your system to enable quotas, verify that they really are enabled. An easy way to do this is to run:

```
# quota -v
```

You should see a one line summary of disk usage and current quota limits for each file system that quotas are enabled on.

You are now ready to start assigning quota limits with the `edquota(8)` command.

You have several options on how to enforce limits on the amount of disk space a user or group may allocate, and how many files they may create. You may limit allocations based on disk space (block quotas) or number of files (inode quotas) or a combination of both. Each of these limits are further broken down into two categories: hard and soft limits.

A hard limit may not be exceeded. Once a user reaches his hard limit he may not make any further allocations on the file system in question. For example, if the user has a hard limit of 500 blocks on a file system and is currently using 490 blocks, the user can only allocate an additional 10 blocks. Attempting to allocate an additional 11 blocks will fail.

Soft limits, on the other hand, can be exceeded for a limited amount of time. This period of time is known as the grace period, which is one week by default. If a user stays over his or her soft limit longer than the grace period, the soft limit will turn into a hard limit and no further allocations will be allowed. When the user drops back below the soft limit, the grace period will be reset.

The following is an example of what you might see when you run the `edquota(8)` command. When the `edquota(8)` command is invoked, you are placed into the editor specified by the `EDITOR` environment variable, or in the `vi` editor if the `EDITOR` variable is not set, to allow you to edit the quota limits.

```
# edquota -u test
```

```
Quotas for user test:
```

```
/usr: blocks in use: 65, limits (soft = 50, hard = 75)
      inodes in use: 7, limits (soft = 50, hard = 60)
/usr/var: blocks in use: 0, limits (soft = 50, hard = 75)
          inodes in use: 0, limits (soft = 50, hard = 60)
```

You will normally see two lines for each file system that has quotas enabled. One line for the block limits, and one line for inode limits. Simply change the value you want updated to modify the quota limit. For example, to raise this user's block limit from a soft limit of 50 and a hard limit of 75 to a soft limit of 500 and a hard limit of 600, change:

```
/usr: blocks in use: 65, limits (soft = 50, hard = 75)
```

to:

```
/usr: blocks in use: 65, limits (soft = 500, hard = 600)
```

The new quota limits will be in place when you exit the editor.

Sometimes it is desirable to set quota limits on a range of UIDs. This can be done by use of the `-p` option on the `edquota(8)` command. First, assign the desired quota limit to a user, and then run `edquota -p protouser startuid-enduid`. For example, if user `test` has the desired quota limits, the following command can be used to duplicate those quota limits for UIDs 10,000 through 19,999:

```
# edquota -p test 10000-19999
```

For more information see `edquota(8)` manual page.

12.12.3 Checking Quota Limits and Disk Usage

You can use either the `quota(1)` or the `repquota(8)` commands to check quota limits and disk usage. The `quota(1)` command can be used to check individual user or group quotas and disk usage. A user may only examine his own quota, and the quota of a group he is a member of. Only the super-user may view all user and group quotas. The `repquota(8)` command can be used to get a summary of all quotas and disk usage for file systems with quotas enabled.

The following is some sample output from the `quota -v` command for a user that has quota limits on two file systems.

```
Disk quotas for user test (uid 1002):
  Filesystem  blocks  quota  limit  grace  files  quota  limit  grace
    /usr      65*    50     75   5days    7     50     60
  /usr/var    0      50     75           0     50     60
```

On the `/usr` file system in the above example, this user is currently 15 blocks over the soft limit of 50 blocks and has 5 days of the grace period left. Note the asterisk `*` which indicates that the user is currently over his quota limit.

Normally file systems that the user is not using any disk space on will not show up in the output from the `quota(1)` command, even if he has a quota limit assigned for that file system. The `-v` option will display those file systems, such as the `/usr/var` file system in the above example.

12.12.4 Quotas over NFS

Quotas are enforced by the quota subsystem on the NFS server. The `rpc.rquotad(8)` daemon makes quota information available to the `quota(1)` command on NFS clients, allowing users on those machines to see their quota statistics.

Enable `rpc.rquotad` in `/etc/inetd.conf` like so:

```
rquotad/1      dgram rpc/udp wait root /usr/libexec/rpc.rquotad rpc.rquotad
```

Now restart `inetd`:

```
# kill -HUP `cat /var/run/inetd.pid`
```

Chapter 13 The Vinum Volume Manager

Contributed by Greg Lehey.

13.1 Synopsis

No matter what disks you have, there will always be limitations:

- They can be too small.
- They can be too slow.
- They can be too unreliable.

13.2 Disks Are Too Small

Originally written by Greg Lehey.

Vinum is a so-called *Volume Manager*, a virtual disk driver that addresses these three problems. Let us look at them in more detail. Various solutions to these problems have been proposed and implemented:

Disks are getting bigger, but so are data storage requirements. Often you will find you want a file system that is bigger than the disks you have available. Admittedly, this problem is not as acute as it was ten years ago, but it still exists. Some systems have solved this by creating an abstract device which stores its data on a number of disks.

13.3 Access Bottlenecks

Modern systems frequently need to access data in a highly concurrent manner. For example, large FTP or HTTP servers can maintain thousands of concurrent sessions and have multiple 100 Mbit/s connections to the outside world, well beyond the sustained transfer rate of most disks.

Current disk drives can transfer data sequentially at up to 70 MB/s, but this value is of little importance in an environment where many independent processes access a drive, where they may achieve only a fraction of these values. In such cases it is more interesting to view the problem from the viewpoint of the disk subsystem: the important parameter is the load that a transfer places on the subsystem, in other words the time for which a transfer occupies the drives involved in the transfer.

In any disk transfer, the drive must first position the heads, wait for the first sector to pass under the read head, and then perform the transfer. These actions can be considered to be atomic: it does not make any sense to interrupt them.


Consider a typical transfer of about 10 kB: the current generation of high-performance disks can position the heads in an average of 3.5 ms. The fastest drives spin at 15,000 rpm, so the average rotational latency (half a revolution) is 2 ms. At 70 MB/s, the transfer itself takes about 150 μ s, almost nothing compared to the positioning time. In such a case, the effective transfer rate drops to a little over 1 MB/s and is clearly highly dependent on the transfer size.

The traditional and obvious solution to this bottleneck is “more spindles”: rather than using one large disk, it uses several smaller disks with the same aggregate storage space. Each disk is capable of positioning and transferring independently, so the effective throughput increases by a factor close to the number of disks used.

The exact throughput improvement is, of course, smaller than the number of disks involved: although each drive is capable of transferring in parallel, there is no way to ensure that the requests are evenly distributed across the drives. Inevitably the load on one drive will be higher than on another.

The evenness of the load on the disks is strongly dependent on the way the data is shared across the drives. In the following discussion, it is convenient to think of the disk storage as a large number of data sectors which are addressable by number, rather like the pages in a book. The most obvious method is to divide the virtual disk into groups of consecutive sectors the size of the individual physical disks and store them in this manner, rather like taking a large book and tearing it into smaller sections. This method is called *concatenation* and has the advantage that the disks are not required to have any specific size relationships. It works well when the access to the virtual disk is spread evenly about its address space. When access is concentrated on a smaller area, the improvement is less marked. Figure 13-1 illustrates the sequence in which storage units are allocated in a concatenated organization.

Figure 13-1. Concatenated Organization



An alternative mapping is to divide the address space into smaller, equal-sized components and store them sequentially on different devices. For example, the first 256 sectors may be stored on the first disk, the next 256 sectors on the next disk and so on. After filling the last disk, the process repeats until the disks are full. This mapping is called *striping* or RAID-0¹. Striping requires somewhat more effort to locate the data, and it can cause additional I/O load where a transfer is spread over multiple disks, but it can also provide a more constant load across the disks. Figure 13-2 illustrates the sequence in which storage units are allocated in a striped organization.

Figure 13-2. Striped Organization

Disk 1	Disk 2	Disk 3	Disk 4
0	1	2	3
4	5	6	7
8	9	10	11
12	13	14	15
16	17	18	19
20	21	22	23

13.4 Data Integrity

The final problem with current disks is that they are unreliable. Although disk drive reliability has increased tremendously over the last few years, they are still the most likely core component of a server to fail. When they do, the results can be catastrophic: replacing a failed disk drive and restoring data to it can take days.

The traditional way to approach this problem has been *mirroring*, keeping two copies of the data on different physical hardware. Since the advent of the RAID levels, this technique has also been called RAID level 1 or RAID-1. Any write to the volume writes to both locations; a read can be satisfied from either, so if one drive fails, the data is still available on the other drive.

Mirroring has two problems:

- The price. It requires twice as much disk storage as a non-redundant solution.
- The performance impact. Writes must be performed to both drives, so they take up twice the bandwidth of a non-mirrored volume. Reads do not suffer from a performance penalty: it even looks as if they are faster.

An alternative solution is *parity*, implemented in the RAID levels 2, 3, 4 and 5. Of these, RAID-5 is the most interesting. As implemented in Vinum, it is a variant on a striped organization which dedicates one block of each stripe to parity of the other blocks. As implemented by Vinum, a RAID-5 plex is similar to a striped plex, except that it implements RAID-5 by including a parity block in each stripe. As required by RAID-5, the location of this parity block changes from one stripe to the next. The numbers in the data blocks indicate the relative block numbers.

Figure 13-3. RAID-5 Organization

Disk 1	Disk 2	Disk 3	Disk 4
0	1	2	Parity
3	4	Parity	5
6	Parity	7	8
Parity	9	10	11
12	13	14	Parity
15	16	Parity	17

Compared to mirroring, RAID-5 has the advantage of requiring significantly less storage space. Read access is similar to that of striped organizations, but write access is significantly slower, approximately 25% of the read performance. If one drive fails, the array can continue to operate in degraded mode: a read from one of the remaining accessible drives continues normally, but a read from the failed drive is recalculated from the corresponding block from all the remaining drives.

13.5 Vinum Objects

In order to address these problems, Vinum implements a four-level hierarchy of objects:

- The most visible object is the virtual disk, called a *volume*. Volumes have essentially the same properties as a UNIX disk drive, though there are some minor differences. They have no size limitations.
- Volumes are composed of *plexes*, each of which represent the total address space of a volume. This level in the hierarchy thus provides redundancy. Think of plexes as individual disks in a mirrored array, each containing the same data.
- Since Vinum exists within the UNIX disk storage framework, it would be possible to use UNIX partitions as the building block for multi-disk plexes, but in fact this turns out to be too inflexible: UNIX disks can have only a limited number of partitions. Instead, Vinum subdivides a single UNIX partition (the *drive*) into contiguous areas called *subdisks*, which it uses as building blocks for plexes.
- Subdisks reside on Vinum *drives*, currently UNIX partitions. Vinum drives can contain any number of subdisks. With the exception of a small area at the beginning of the drive, which is used for storing configuration and state information, the entire drive is available for data storage.

The following sections describe the way these objects provide the functionality required of Vinum.

13.5.1 Volume Size Considerations

Plexes can include multiple subdisks spread over all drives in the Vinum configuration. As a result, the size of an individual drive does not limit the size of a plex, and thus of a volume.

13.5.2 Redundant Data Storage

Vinum implements mirroring by attaching multiple plexes to a volume. Each plex is a representation of the data in a volume. A volume may contain between one and eight plexes.

Although a plex represents the complete data of a volume, it is possible for parts of the representation to be physically missing, either by design (by not defining a subdisk for parts of the plex) or by accident (as a result of the failure of a drive). As long as at least one plex can provide the data for the complete address range of the volume, the volume is fully functional.

13.5.3 Performance Issues

Vinum implements both concatenation and striping at the plex level:

- A *concatenated plex* uses the address space of each subdisk in turn.
- A *striped plex* stripes the data across each subdisk. The subdisks must all have the same size, and there must be at least two subdisks in order to distinguish it from a concatenated plex.

13.5.4 Which Plex Organization?

Vinum implements two kinds of plex:

- Concatenated plexes are the most flexible: they can contain any number of subdisks, and the subdisks may be of different length. The plex may be extended by adding additional subdisks. They require less CPU time than striped plexes, though the difference in CPU overhead is not measurable. On the other hand, they are most susceptible to hot spots, where one disk is very active and others are idle.
- The greatest advantage of striped (RAID-0) plexes is that they reduce hot spots: by choosing an optimum sized stripe (about 256 kB), you can even out the load on the component drives. The disadvantages of this approach are (fractionally) more complex code and restrictions on subdisks: they must be all the same size, and extending a plex by adding new subdisks is so complicated that Vinum currently does not implement it. Vinum imposes an additional, trivial restriction: a striped plex must have at least two subdisks, since otherwise it is indistinguishable from a concatenated plex.

Table 13-1 summarizes the advantages and disadvantages of each plex organization.

Table 13-1. Vinum Plex Organizations

Plex type	Minimum subdisks	Can add subdisks	Must be equal size	Application
-----------	------------------	------------------	--------------------	-------------

Plex type	Minimum subdisks	Can add subdisks	Must be equal size	Application
concatenated	1	yes	no	Large data storage with maximum placement flexibility and moderate performance
striped	2	no	yes	High performance in combination with highly concurrent access

13.6 Some Examples

Vinum maintains a *configuration database* which describes the objects known to an individual system. Initially, the user creates the configuration database from one or more configuration files with the aid of the `vinum(8)` utility program. Vinum stores a copy of its configuration database on each disk slice (which Vinum calls a *device*) under its control. This database is updated on each state change, so that a restart accurately restores the state of each Vinum object.

13.6.1 The Configuration File

The configuration file describes individual Vinum objects. The definition of a simple volume might be:

```
drive a device /dev/da3h
volume myvol
  plex org concat
    sd length 512m drive a
```

This file describes four Vinum objects:

- The *drive* line describes a disk partition (*drive*) and its location relative to the underlying hardware. It is given the symbolic name *a*. This separation of the symbolic names from the device names allows disks to be moved from one location to another without confusion.
- The *volume* line describes a volume. The only required attribute is the name, in this case *myvol*.
- The *plex* line defines a plex. The only required parameter is the organization, in this case *concat*. No name is necessary: the system automatically generates a name from the volume name by adding the suffix *.px*, where *x* is the number of the plex in the volume. Thus this plex will be called *myvol.p0*.
- The *sd* line describes a subdisk. The minimum specifications are the name of a drive on which to store it, and the length of the subdisk. As with plexes, no name is necessary: the system automatically assigns names derived from the plex name by adding the suffix *.sx*, where *x* is the number of the subdisk in the plex. Thus Vinum gives this subdisk the name *myvol.p0.s0*.

After processing this file, `vinum(8)` produces the following output:

```
# vinum -> create config1
Configuration summary
Drives:          1 (4 configured)
```

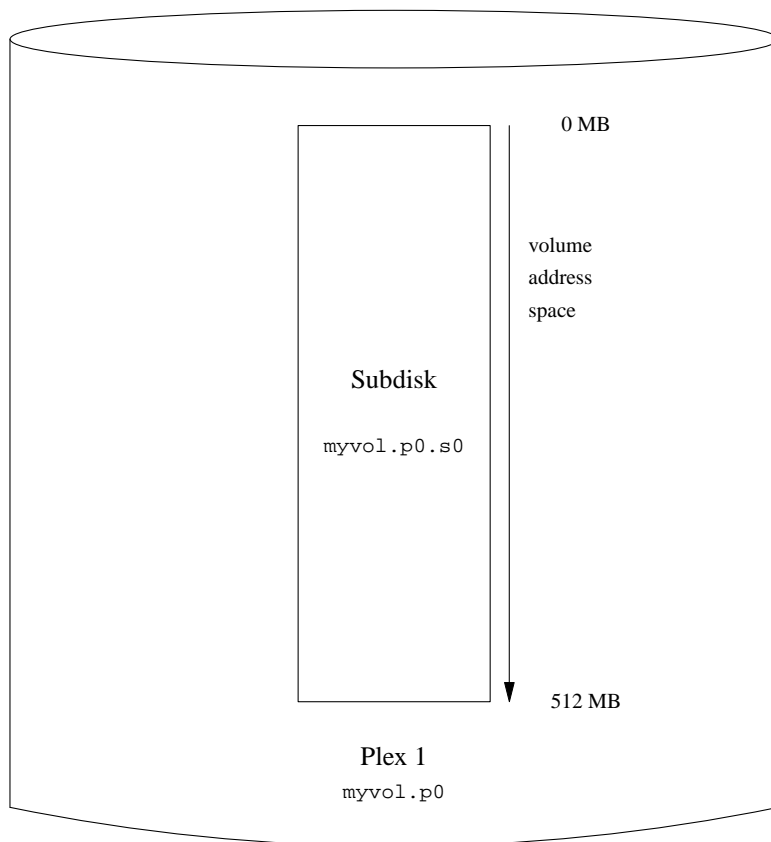
```

Volumes:      1 (4 configured)
Plexes:      1 (8 configured)
Subdisks:    1 (16 configured)

D a          State: up      Device /dev/da3h      Avail: 2061/2573 MB (80%)
V myvol     State: up      Plexes:      1 Size:      512 MB
P myvol.p0  C State: up      Subdisks:    1 Size:      512 MB
S myvol.p0.s0 State: up      PO:         0 B Size:      512 MB
    
```

This output shows the brief listing format of vinum(8). It is represented graphically in Figure 13-4.

Figure 13-4. A Simple Vinum Volume



This figure, and the ones which follow, represent a volume, which contains the plexes, which in turn contain the subdisks. In this trivial example, the volume contains one plex, and the plex contains one subdisk.

This particular volume has no specific advantage over a conventional disk partition. It contains a single plex, so it is not redundant. The plex contains a single subdisk, so there is no difference in storage allocation from a conventional disk partition. The following sections illustrate various more interesting configuration methods.

13.6.2 Increased Resilience: Mirroring

The resilience of a volume can be increased by mirroring. When laying out a mirrored volume, it is important to ensure that the subdisks of each plex are on different drives, so that a drive failure will not take down both plexes. The following configuration mirrors a volume:

```
drive b device /dev/da4h
volume mirror
    plex org concat
        sd length 512m drive a
    plex org concat
        sd length 512m drive b
```


In this example, it was not necessary to specify a definition of drive *a* again, since Vinum keeps track of all objects in its configuration database. After processing this definition, the configuration looks like:

```
Drives:          2 (4 configured)
Volumes:         2 (4 configured)
Plexes:          3 (8 configured)
Subdisks:        3 (16 configured)
```

D a	State: up	Device /dev/da3h	Avail: 1549/2573 MB (60%)
D b	State: up	Device /dev/da4h	Avail: 2061/2573 MB (80%)
V myvol	State: up	Plexes: 1	Size: 512 MB
V mirror	State: up	Plexes: 2	Size: 512 MB
P myvol.p0	C State: up	Subdisks: 1	Size: 512 MB
P mirror.p0	C State: up	Subdisks: 1	Size: 512 MB
P mirror.p1	C State: initializing	Subdisks: 1	Size: 512 MB
S myvol.p0.s0	State: up	PO: 0	B Size: 512 MB
S mirror.p0.s0	State: up	PO: 0	B Size: 512 MB
S mirror.p1.s0	State: empty	PO: 0	B Size: 512 MB

Figure 13-5 shows the structure graphically.

Figure 13-5. A Mirrored Vinum Volume



In this example, each plex contains the full 512 MB of address space. As in the previous example, each plex contains only a single subdisk.

13.6.3 Optimizing Performance

The mirrored volume in the previous example is more resistant to failure than an unmirrored volume, but its performance is less: each write to the volume requires a write to both drives, using up a greater proportion of the total disk bandwidth. Performance considerations demand a different approach: instead of mirroring, the data is striped across as many disk drives as possible. The following configuration shows a volume with a plex striped across four disk drives:

```
drive c device /dev/da5h
drive d device /dev/da6h
volume stripe
plex org striped 512k
  sd length 128m drive a
```

```
sd length 128m drive b
sd length 128m drive c
sd length 128m drive d
```

As before, it is not necessary to define the drives which are already known to Vinum. After processing this definition, the configuration looks like:

```
Drives:          4 (4 configured)
Volumes:         3 (4 configured)
Plexes:          4 (8 configured)
Subdisks:        7 (16 configured)

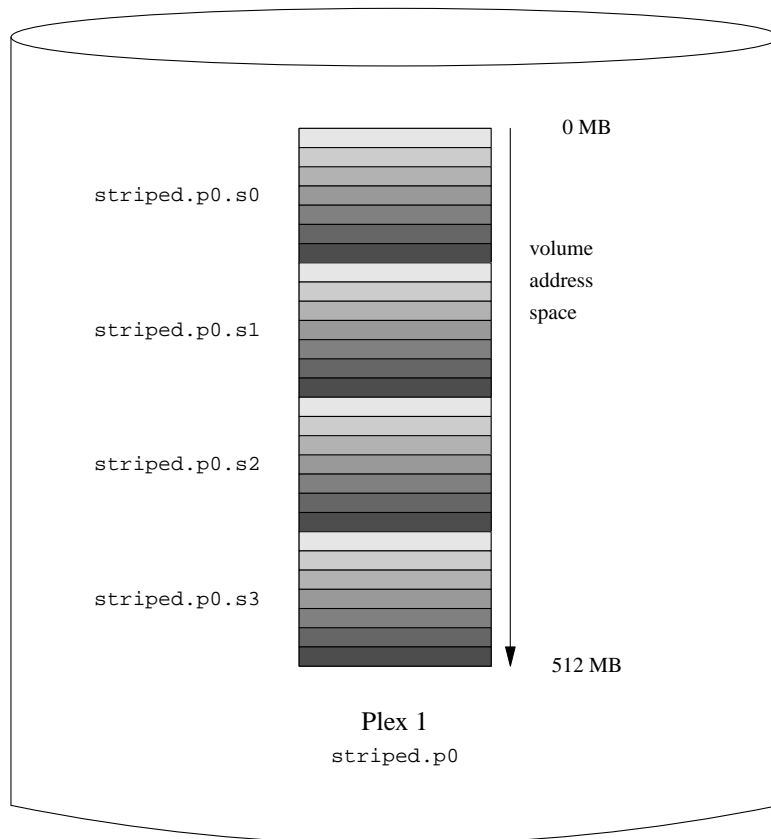
D a              State: up      Device /dev/da3h      Avail: 1421/2573 MB (55%)
D b              State: up      Device /dev/da4h      Avail: 1933/2573 MB (75%)
D c              State: up      Device /dev/da5h      Avail: 2445/2573 MB (95%)
D d              State: up      Device /dev/da6h      Avail: 2445/2573 MB (95%)

V myvol          State: up      Plexes:      1 Size:      512 MB
V mirror         State: up      Plexes:      2 Size:      512 MB
V striped        State: up      Plexes:      1 Size:      512 MB

P myvol.p0       C State: up      Subdisks:    1 Size:      512 MB
P mirror.p0      C State: up      Subdisks:    1 Size:      512 MB
P mirror.p1      C State: initializing Subdisks:    1 Size:      512 MB
P striped.p1     State: up      Subdisks:    1 Size:      512 MB

S myvol.p0.s0    State: up      PO:          0 B Size:      512 MB
S mirror.p0.s0   State: up      PO:          0 B Size:      512 MB
S mirror.p1.s0   State: empty    PO:          0 B Size:      512 MB
S striped.p0.s0  State: up      PO:          0 B Size:      128 MB
S striped.p0.s1  State: up      PO:          512 kB Size:      128 MB
S striped.p0.s2  State: up      PO:          1024 kB Size:      128 MB
S striped.p0.s3  State: up      PO:          1536 kB Size:      128 MB
```

Figure 13-6. A Striped Vinum Volume



This volume is represented in Figure 13-6. The darkness of the stripes indicates the position within the plex address space: the lightest stripes come first, the darkest last.

13.6.4 Resilience and Performance

With sufficient hardware, it is possible to build volumes which show both increased resilience and increased performance compared to standard UNIX partitions. A typical configuration file might be:

```

volume raid10
  plex org striped 512k
    sd length 102480k drive a
    sd length 102480k drive b
    sd length 102480k drive c
    sd length 102480k drive d
    sd length 102480k drive e
  plex org striped 512k
    sd length 102480k drive c


```

```
sd length 102480k drive d
sd length 102480k drive e
sd length 102480k drive a
sd length 102480k drive b
```

The subdisks of the second plex are offset by two drives from those of the first plex: this helps ensure that writes do not go to the same subdisks even if a transfer goes over two drives.

Figure 13-7 represents the structure of this volume.

Figure 13-7. A Mirrored, Striped Vinum Volume



13.7 Object Naming

As described above, Vinum assigns default names to plexes and subdisks, although they may be overridden. Overriding the default names is not recommended: experience with the VERITAS volume manager, which allows

arbitrary naming of objects, has shown that this flexibility does not bring a significant advantage, and it can cause confusion.

Names may contain any non-blank character, but it is recommended to restrict them to letters, digits and the underscore characters. The names of volumes, plexes and subdisks may be up to 64 characters long, and the names of drives may be up to 32 characters long.

Vinum objects are assigned device nodes in the hierarchy `/dev/vinum`. The configuration shown above would cause Vinum to create the following device nodes:

- The control devices `/dev/vinum/control` and `/dev/vinum/controld`, which are used by `vinum(8)` and the Vinum daemon respectively.
- Block and character device entries for each volume. These are the main devices used by Vinum. The block device names are the name of the volume, while the character device names follow the BSD tradition of prepending the letter `r` to the name. Thus the configuration above would include the block devices `/dev/vinum/myvol`, `/dev/vinum/mirror`, `/dev/vinum/striped`, `/dev/vinum/raid5` and `/dev/vinum/raid10`, and the character devices `/dev/vinum/rmyvol`, `/dev/vinum/rmirror`, `/dev/vinum/rstriped`, `/dev/vinum/rraid5` and `/dev/vinum/rraid10`. There is obviously a problem here: it is possible to have two volumes called `r` and `rr`, but there will be a conflict creating the device node `/dev/vinum/rr`: is it a character device for volume `r` or a block device for volume `rr`? Currently Vinum does not address this conflict: the first-defined volume will get the name.
- A directory `/dev/vinum/drive` with entries for each drive. These entries are in fact symbolic links to the corresponding disk nodes.
- A directory `/dev/vinum/volume` with entries for each volume. It contains subdirectories for each plex, which in turn contain subdirectories for their component subdisks.
- The directories `/dev/vinum/plex`, `/dev/vinum/sd`, and `/dev/vinum/rsd`, which contain block device nodes for each plex and block and character device nodes respectively for each subdisk.

For example, consider the following configuration file:

```
drive drive1 device /dev/sd1h
drive drive2 device /dev/sd2h
drive drive3 device /dev/sd3h
drive drive4 device /dev/sd4h
  volume s64 setupstate
    plex org striped 64k
      sd length 100m drive drive1
      sd length 100m drive drive2
      sd length 100m drive drive3
      sd length 100m drive drive4
```

After processing this file, `vinum(8)` creates the following structure in `/dev/vinum`:

```
brwx----- 1 root wheel 25, 0x40000001 Apr 13 16:46 Control
brwx----- 1 root wheel 25, 0x40000002 Apr 13 16:46 control
brwx----- 1 root wheel 25, 0x40000000 Apr 13 16:46 controld
drwxr-xr-x 2 root wheel 512 Apr 13 16:46 drive
drwxr-xr-x 2 root wheel 512 Apr 13 16:46 plex
crwxr-xr-- 1 root wheel 91, 2 Apr 13 16:46 rs64
drwxr-xr-x 2 root wheel 512 Apr 13 16:46 rsd
```



```

drwxr-xr-x 2 root wheel      512 Apr 13 16:46 rvol
brwxr-xr-- 1 root wheel    25,  2 Apr 13 16:46 s64
drwxr-xr-x 2 root wheel      512 Apr 13 16:46 sd
drwxr-xr-x 3 root wheel      512 Apr 13 16:46 vol

/dev/vinum/drive:
total 0
lrwxr-xr-x 1 root wheel  9 Apr 13 16:46 drive1 -> /dev/sd1h
lrwxr-xr-x 1 root wheel  9 Apr 13 16:46 drive2 -> /dev/sd2h
lrwxr-xr-x 1 root wheel  9 Apr 13 16:46 drive3 -> /dev/sd3h
lrwxr-xr-x 1 root wheel  9 Apr 13 16:46 drive4 -> /dev/sd4h

/dev/vinum/plex:
total 0
brwxr-xr-- 1 root wheel    25, 0x10000002 Apr 13 16:46 s64.p0

/dev/vinum/rsd:
total 0
crwxr-xr-- 1 root wheel  91, 0x20000002 Apr 13 16:46 s64.p0.s0
crwxr-xr-- 1 root wheel  91, 0x20100002 Apr 13 16:46 s64.p0.s1
crwxr-xr-- 1 root wheel  91, 0x20200002 Apr 13 16:46 s64.p0.s2
crwxr-xr-- 1 root wheel  91, 0x20300002 Apr 13 16:46 s64.p0.s3

/dev/vinum/rvol:
total 0
crwxr-xr-- 1 root wheel  91,  2 Apr 13 16:46 s64

/dev/vinum/sd:
total 0
brwxr-xr-- 1 root wheel    25, 0x20000002 Apr 13 16:46 s64.p0.s0
brwxr-xr-- 1 root wheel    25, 0x20100002 Apr 13 16:46 s64.p0.s1
brwxr-xr-- 1 root wheel    25, 0x20200002 Apr 13 16:46 s64.p0.s2
brwxr-xr-- 1 root wheel    25, 0x20300002 Apr 13 16:46 s64.p0.s3

/dev/vinum/vol:
total 1
brwxr-xr-- 1 root wheel    25,  2 Apr 13 16:46 s64
drwxr-xr-x 3 root wheel      512 Apr 13 16:46 s64.plex

/dev/vinum/vol/s64.plex:
total 1
brwxr-xr-- 1 root wheel    25, 0x10000002 Apr 13 16:46 s64.p0
drwxr-xr-x 2 root wheel      512 Apr 13 16:46 s64.p0.sd

/dev/vinum/vol/s64.plex/s64.p0.sd:
total 0
brwxr-xr-- 1 root wheel    25, 0x20000002 Apr 13 16:46 s64.p0.s0
brwxr-xr-- 1 root wheel    25, 0x20100002 Apr 13 16:46 s64.p0.s1
brwxr-xr-- 1 root wheel    25, 0x20200002 Apr 13 16:46 s64.p0.s2
brwxr-xr-- 1 root wheel    25, 0x20300002 Apr 13 16:46 s64.p0.s3

```

Although it is recommended that plexes and subdisks should not be allocated specific names, Vinum drives must be named. This makes it possible to move a drive to a different location and still recognize it automatically. Drive names may be up to 32 characters long.

13.7.1 Creating File Systems

Volumes appear to the system to be identical to disks, with one exception. Unlike UNIX drives, Vinum does not partition volumes, which thus do not contain a partition table. This has required modification to some disk utilities, notably `newfs(8)`, which previously tried to interpret the last letter of a Vinum volume name as a partition identifier. For example, a disk drive may have a name like `/dev/ad0a` or `/dev/da2h`. These names represent the first partition (a) on the first (0) IDE disk (ad) and the eighth partition (h) on the third (2) SCSI disk (da) respectively. By contrast, a Vinum volume might be called `/dev/vinum/concat`, a name which has no relationship with a partition name.

Normally, `newfs(8)` interprets the name of the disk and complains if it cannot understand it. For example:

```
# newfs /dev/vinum/concat
newfs: /dev/vinum/concat: can't figure out file system partition
```

13.8 Configuring Vinum

The `GENERIC` kernel does not contain Vinum. It is possible to build a special kernel which includes Vinum, but this is not recommended. The standard way to start Vinum is as a kernel module (kld). You do not even need to use `kldload(8)` for Vinum: when you start `vinum(8)`, it checks whether the module has been loaded, and if it is not, it loads it automatically.

13.8.1 Startup

Vinum stores configuration information on the disk slices in essentially the same form as in the configuration files. When reading from the configuration database, Vinum recognizes a number of keywords which are not allowed in the configuration files. For example, a disk configuration might contain the following text:

```
volume myvol state up
volume bigraid state down
plex name myvol.p0 state up org concat vol myvol
plex name myvol.p1 state up org concat vol myvol
plex name myvol.p2 state init org striped 512b vol myvol
plex name bigraid.p0 state initializing org raid5 512b vol bigraid
sd name myvol.p0.s0 drive a plex myvol.p0 state up len 1048576b driveoffset 265b plexoffset 0b
sd name myvol.p0.s1 drive b plex myvol.p0 state up len 1048576b driveoffset 265b plexoffset 10485
sd name myvol.p1.s0 drive c plex myvol.p1 state up len 1048576b driveoffset 265b plexoffset 0b
sd name myvol.p1.s1 drive d plex myvol.p1 state up len 1048576b driveoffset 265b plexoffset 10485
sd name myvol.p2.s0 drive a plex myvol.p2 state init len 524288b driveoffset 1048841b plexoffset 0
sd name myvol.p2.s1 drive b plex myvol.p2 state init len 524288b driveoffset 1048841b plexoffset 1
sd name myvol.p2.s2 drive c plex myvol.p2 state init len 524288b driveoffset 1048841b plexoffset 1
sd name myvol.p2.s3 drive d plex myvol.p2 state init len 524288b driveoffset 1048841b plexoffset 1
sd name bigraid.p0.s0 drive a plex bigraid.p0 state initializing len 4194304b driveoff set 157312
sd name bigraid.p0.s1 drive b plex bigraid.p0 state initializing len 4194304b driveoff set 157312
sd name bigraid.p0.s2 drive c plex bigraid.p0 state initializing len 4194304b driveoff set 157312
```

```
sd name bigraid.p0.s3 drive d plex bigraid.p0 state initializing len 4194304b driveoff set 157312
sd name bigraid.p0.s4 drive e plex bigraid.p0 state initializing len 4194304b driveoff set 157312
```

The obvious differences here are the presence of explicit location information and naming (both of which are also allowed, but discouraged, for use by the user) and the information on the states (which are not available to the user). Vinum does not store information about drives in the configuration information: it finds the drives by scanning the configured disk drives for partitions with a Vinum label. This enables Vinum to identify drives correctly even if they have been assigned different UNIX drive IDs.

13.8.1.1 Automatic Startup

In order to start Vinum automatically when you boot the system, ensure that you have the following line in your `/etc/rc.conf`:

```
start_vinum="YES" # set to YES to start vinum
```

If you do not have a file `/etc/rc.conf`, create one with this content. This will cause the system to load the Vinum `kld` at startup, and to start any objects mentioned in the configuration. This is done before mounting file systems, so it is possible to automatically `fsck(8)` and mount file systems on Vinum volumes.

When you start Vinum with the `vinum start` command, Vinum reads the configuration database from one of the Vinum drives. Under normal circumstances, each drive contains an identical copy of the configuration database, so it does not matter which drive is read. After a crash, however, Vinum must determine which drive was updated most recently and read the configuration from this drive. It then updates the configuration if necessary from progressively older drives.

13.9 Using Vinum for the Root Filesystem

For a machine that has fully-mirrored filesystems using Vinum, it is desirable to also mirror the root filesystem. Setting up such a configuration is less trivial than mirroring an arbitrary filesystem because:

- The root filesystem must be available very early during the boot process, so the Vinum infrastructure must already be available at this time.
- The volume containing the root filesystem also contains the system bootstrap and the kernel, which must be read using the host system's native utilities (e. g. the BIOS on PC-class machines) which often cannot be taught about the details of Vinum.

In the following sections, the term "root volume" is generally used to describe the Vinum volume that contains the root filesystem. It is probably a good idea to use the name `"root"` for this volume, but this is not technically required in any way. All command examples in the following sections assume this name though.

13.9.1 Starting up Vinum Early Enough for the Root Filesystem

There are several measures to take for this to happen:

- Vinum must be available in the kernel at boot-time. Thus, the method to start Vinum automatically described in Section 13.8.1.1 is not applicable to accomplish this task, and the `start_vinum` parameter must actually *not* be

set when the following setup is being arranged. The first option would be to compile Vinum statically into the kernel, so it is available all the time, but this is usually not desirable. There is another option as well, to have `/boot/loader` (Section 7.3.3) load the vinum kernel module early, before starting the kernel. This can be accomplished by putting the line

```
vinum_load="YES"
```

into the file `/boot/loader.conf`.

- Vinum must be initialized early since it needs to supply the volume for the root filesystem. By default, the Vinum kernel part is not looking for drives that might contain Vinum volume information until the administrator (or one of the startup scripts) issues a `vinum start` command.

Vinum must explicitly be told which disks to scan, using a line like the following one in `/boot/loader.conf`:

```
vinum.drives="/dev/da0 /dev/da1"
```

It is important that all drives are mentioned that could possibly contain Vinum data. It does not harm if *more* drives are listed, nor is it necessary to add each slice and/or partition explicitly, since Vinum will scan all slices and partitions of the named drives for valid Vinum headers.

Since the routines used to parse the name of the root filesystem, and derive the device ID (major/minor number) are only prepared to handle “classical” device names like `/dev/ad0s1a`, they cannot make any sense out of a root volume name like `/dev/vinum/root`. For that reason, Vinum itself needs to pre-setup the internal kernel parameter that holds the ID of the root device during its own initialization. This is requested by passing the name of the root volume in the loader variable `vinum.root`. The entry in `/boot/loader.conf` to accomplish this looks like:

```
vinum.root="root"
```

Now, when the kernel initialization tries to find out the root device to mount, it sees whether some kernel module has already pre-initialized the kernel parameter for it. If that is the case, *and* the device claiming the root device matches the major number of the driver as figured out from the name of the root device string being passed (that is, “vinum” in our case), it will use the pre-allocated device ID, instead of trying to figure out one itself. That way, during the usual automatic startup, it can continue to mount the Vinum root volume for the root filesystem.

However, when `boot -a` has been requesting to ask for entering the name of the root device manually, it must be noted that this routine still cannot actually parse a name entered there that refers to a Vinum volume. If any device name is entered that does not refer to a Vinum device, the mismatch between the major numbers of the pre-allocated root parameter and the driver as figured out from the given name will make this routine enter its normal parser, so entering a string like `ufs:da0d` will work as expected. Note that if this fails, it is however no longer possible to re-enter a string like `ufs:vinum/root` again, since it cannot be parsed. The only way out is to reboot again, and start over then. (At the “askroot” prompt, the initial `/dev/` can always be omitted.)

By placing the line:

```
vinum.drives="/dev/da0 /dev/da1"
```

```
vinum.autostart="YES"
```

into `/boot/loader.conf`, Vinum is instructed to automatically scan all drives for Vinum information as part of the kernel startup.

It is important that all drives are mentioned that could possibly contain Vinum data. It does not harm if *more* drives are listed, nor is it necessary to add each slice and/or partition explicitly, since Vinum will scan all slices and partitions of the named drives for valid Vinum headers.

Vinum itself needs to pre-setup the internal kernel parameter that holds the ID of the root device during its own initialization. This is requested by passing the name of the root volume in the loader variable `vinum.root`. The entry in `/boot/loader.conf` to accomplish this looks like:

```
vinum.root="root"
```

Now, when the kernel initialization tries to find out the root device to mount, it sees whether some kernel module has already pre-initialized the kernel parameter for it. If that is the case, *and* the device claiming the root device matches the major number of the driver as figured out from the name of the root device string being passed (that is, "vinum" in our case), it will use the pre-allocated device ID, instead of trying to figure out one itself. That way, during the usual automatic startup, it can continue to mount the Vinum root volume for the root filesystem.

However, when `boot -a` has been requesting to ask for entering the name of the root device manually, it must be noted that this routine still cannot actually parse a name entered there that refers to a Vinum volume. If any device name is entered that does not refer to a Vinum device, the mismatch between the major numbers of the pre-allocated root parameter and the driver as figured out from the given name will make this routine enter its normal parser, so entering a string like `ufs:da0d` will work as expected. Note that if this fails, it is however no longer possible to re-enter a string like `ufs:vinum/root` again, since it cannot be parsed. The only way out is to reboot again, and start over then. (At the "askroot" prompt, the initial `/dev/` can always be omitted.)

13.9.2 Making a Vinum-based Root Volume Accessible to the Bootstrap

Since the current DragonFly bootstrap is very small, and already has the burden of reading files (like `/boot/loader`) from the UFS filesystem, it is sheer impossible to also teach it about internal Vinum structures so it could parse the Vinum configuration data, and figure out about the elements of a boot volume itself. Thus, some tricks are necessary to provide the bootstrap code with the illusion of a standard "a" partition that contains the root filesystem.

For this to be possible at all, the following requirements must be met for the root volume:

- The root volume must not be striped or RAID-5.
- The root volume must not contain more than one concatenated subdisk per plex.

Note that it is desirable and possible that there are multiple plexes, each containing one replica of the root filesystem. The bootstrap process will, however, only use one of these replica for finding the bootstrap and all the files, until the kernel will eventually mount the root filesystem itself. Each single subdisk within these plexes will then need its own "a" partition illusion, for the respective device to become bootable. It is not strictly needed that each of these faked "a" partitions is located at the same offset within its device, compared with other devices containing plexes of the root volume. However, it is probably a good idea to create the Vinum volumes that way so the resulting mirrored devices are symmetric, to avoid confusion.

In order to set up these "a" partitions, for each device containing part of the root volume, the following needs to be done:

1. The location (offset from the beginning of the device) and size of this device's subdisk that is part of the root volume need to be examined, using the command

```
vinum l -rv root
```

Note that Vinum offsets and sizes are measured in bytes. They must be divided by 512 in order to obtain the block numbers that are to be used in the `disklabel` command.

2. Run the command

```
disklabel -e devname
```

for each device that participates in the root volume. *devname* must be either the name of the disk (like *da0*) for disks without a slice (aka. *fdisk*) table, or the name of the slice (like *ad0s1*).

If there is already an "a" partition on the device (presumably, containing a pre-Vinum root filesystem), it should be renamed to something else, so it remains accessible (just in case), but will no longer be used by default to bootstrap the system. Note that active partitions (like a root filesystem currently mounted) cannot be renamed, so this must be executed either when being booted from a "Fixit" medium, or in a two-step process, where (in a mirrored situation) the disk that has not been currently booted is being manipulated first.

Then, the offset the Vinum partition on this device (if any) must be added to the offset of the respective root volume subdisk on this device. The resulting value will become the "offset" value for the new "a" partition. The "size" value for this partition can be taken verbatim from the calculation above. The "fstype" should be 4.2BSD. The "fsize", "bsize", and "cpg" values should best be chosen to match the actual filesystem, though they are fairly unimportant within this context.

That way, a new "a" partition will be established that overlaps the Vinum partition on this device. Note that the *disklabel* will only allow for this overlap if the Vinum partition has properly been marked using the "vinum" fstype.

3. That's all! A faked "a" partition does exist now on each device that has one replica of the root volume. It is highly recommendable to verify the result again, using a command like

```
fsck -n /dev/devnamea
```

It should be remembered that all files containing control information must be relative to the root filesystem in the Vinum volume which, when setting up a new Vinum root volume, might not match the root filesystem that is currently active. So in particular, the files */etc/fstab* and */boot/loader.conf* need to be taken care of.

At next reboot, the bootstrap should figure out the appropriate control information from the new Vinum-based root filesystem, and act accordingly. At the end of the kernel initialization process, after all devices have been announced, the prominent notice that shows the success of this setup is a message like:

```
Mounting root from ufs:/dev/vinum/root
```

13.9.3 Example of a Vinum-based Root Setup

After the Vinum root volume has been set up, the output of `vinum l -rv root` could look like:

```
...
Subdisk root.p0.s0:
Size:          125829120 bytes (120 MB)
State: up
Plex root.p0 at offset 0 (0 B)
Drive disk0 (/dev/da0h) at offset 135680 (132 kB)

Subdisk root.pl.s0:
Size:          125829120 bytes (120 MB)
```

```
State: up
Plex root.pl at offset 0 (0 B)
Drive disk1 (/dev/da1h) at offset 135680 (132 kB)
```

The values to note are 135680 for the offset (relative to partition `/dev/da0h`). This translates to 265 512-byte disk blocks in `disklabel`'s terms. Likewise, the size of this root volume is 245760 512-byte blocks. `/dev/da1h`, containing the second replica of this root volume, has a symmetric setup.

The `disklabel` for these devices might look like:

```
...
8 partitions:
#      size  offset  fstype  [fsize bsize bps/cpg]
a:   245760    281   4.2BSD   2048 16384    0 # (Cyl.  0*- 15*)
c:  71771688     0  unused     0     0    # (Cyl.  0 - 4467*)
h:  71771672    16   vinum                # (Cyl.  0*- 4467*)
```

It can be observed that the "size" parameter for the faked "a" partition matches the value outlined above, while the "offset" parameter is the sum of the offset within the Vinum partition "h", and the offset of this partition within the device (or slice). This is a typical setup that is necessary to avoid the problem described in Section 13.9.4.3. It can also be seen that the entire "a" partition is completely within the "h" partition containing all the Vinum data for this device.

Note that in the above example, the entire device is dedicated to Vinum, and there is no leftover pre-Vinum root partition, since this has been a newly set-up disk that was only meant to be part of a Vinum configuration, ever.

13.9.4 Troubleshooting

If something goes wrong, a way is needed to recover from the situation. The following list contains few known pitfalls and solutions.

13.9.4.1 System Bootstrap Loads, but System Does Not Boot

If for any reason the system does not continue to boot, the bootstrap can be interrupted with by pressing the **space** key at the 10-seconds warning. The loader variables (like `vinum.autostart`) can be examined using the `show`, and manipulated using `set` or `unset` commands.

If the only problem was that the Vinum kernel module was not yet in the list of modules to load automatically, a simple `load vinum` will help.

When ready, the boot process can be continued with a `boot -as`. The options `-as` will request the kernel to ask for the root filesystem to mount (`-a`), and make the boot process stop in single-user mode (`-s`), where the root filesystem is mounted read-only. That way, even if only one plex of a multi-plex volume has been mounted, no data inconsistency between plexes is being risked.

At the prompt asking for a root filesystem to mount, any device that contains a valid root filesystem can be entered. If `/etc/fstab` had been set up correctly, the default should be something like `ufs:/dev/vinum/root`. A typical alternate choice would be something like `ufs:da0d` which could be a hypothetical partition that contains the pre-Vinum root filesystem. Care should be taken if one of the alias "a" partitions are entered here that are actually

reference to the subdisks of the Vinum root device, because in a mirrored setup, this would only mount one piece of a mirrored root device. If this filesystem is to be mounted read-write later on, it is necessary to remove the other plex(es) of the Vinum root volume since these plexes would otherwise carry inconsistent data.

13.9.4.2 Only Primary Bootstrap Loads

If `/boot/loader` fails to load, but the primary bootstrap still loads (visible by a single dash in the left column of the screen right after the boot process starts), an attempt can be made to interrupt the primary bootstrap at this point, using the **space** key. This will make the bootstrap stop in stage two, see Section 7.3.2. An attempt can be made here to boot off an alternate partition, like the partition containing the previous root filesystem that has been moved away from "a" above.

13.9.4.3 Nothing Boots, the Bootstrap Panics

This situation will happen if the bootstrap had been destroyed by the Vinum installation. Unfortunately, Vinum accidentally currently leaves only 4 KB at the beginning of its partition free before starting to write its Vinum header information. However, the stage one and two bootstraps plus the `disklabel` embedded between them currently require 8 KB. So if a Vinum partition was started at offset 0 within a slice or disk that was meant to be bootable, the Vinum setup will trash the bootstrap.

Similarly, if the above situation has been recovered, for example by booting from a "Fixit" medium, and the bootstrap has been re-installed using `disklabel -B` as described in Section 7.3.2, the bootstrap will trash the Vinum header, and Vinum will no longer find its disk(s). Though no actual Vinum configuration data or data in Vinum volumes will be trashed by this, and it would be possible to recover all the data by entering exact the same Vinum configuration data again, the situation is hard to fix at all. It would be necessary to move the entire Vinum partition by at least 4 KB off, in order to have the Vinum header and the system bootstrap no longer collide.

Notes

1. RAID stands for *Redundant Array of Inexpensive Disks* and offers various forms of fault tolerance, though the latter term is somewhat misleading: it provides no redundancy.

Chapter 14 Localization - I18N/L10N Usage and Setup

Contributed by Andrey A. Chernov. Rewritten by Michael C. Wu.

14.1 Synopsis

DragonFly is a very distributed project with users and contributors located all over the world. This chapter discusses the internationalization and localization features of DragonFly that allow non-English speaking users to get real work done. There are many aspects of the i18n implementation in both the system and application levels, so where applicable we refer the reader to more specific sources of documentation.

After reading this chapter, you will know:

- How different languages and locales are encoded on modern operating systems.
- How to set the locale for your login shell.
- How to configure your console for non-English languages.
- How to use X Window System effectively with different languages.
- Where to find more information about writing i18n-compliant applications.

Before reading this chapter, you should:

- Know how to install additional third-party applications (Chapter 4).

14.2 The Basics

14.2.1 What Is I18N/L10N?

Developers shortened internationalization into the term I18N, counting the number of letters between the first and the last letters of internationalization. L10N uses the same naming scheme, coming from “localization”. Combined together, I18N/L10N methods, protocols, and applications allow users to use languages of their choice.

I18N applications are programmed using I18N kits under libraries. It allows for developers to write a simple file and translate displayed menus and texts to each language. We strongly encourage programmers to follow this convention.

14.2.2 Why Should I Use I18N/L10N?

I18N/L10N is used whenever you wish to either view, input, or process data in non-English languages.

14.2.3 What Languages Are Supported in the I18N Effort?

I18N and L10N are not DragonFly specific. Currently, one can choose from most of the major languages of the World, including but not limited to: Chinese, German, Japanese, Korean, French, Russian, Vietnamese and others.

14.3 Using Localization

In all its splendor, I18N is not DragonFly-specific and is a convention. We encourage you to help DragonFly in following this convention.

Localization settings are based on three main terms: Language Code, Country Code, and Encoding. Locale names are constructed from these parts as follows:

LanguageCode_CountryCode.Encoding

14.3.1 Language and Country Codes

In order to localize a DragonFly system to a specific language (or any other I18N-supporting UNIX like systems), the user needs to find out the codes for the specify country and language (country codes tell applications what variation of given language to use). In addition, web browsers, SMTP/POP servers, web servers, etc. make decisions based on them. The following are examples of language/country codes:

Language/Country Code	Description
en_US	English - United States
ru_RU	Russian for Russia
zh_TW	Traditional Chinese for Taiwan

14.3.2 Encodings

Some languages use non-ASCII encodings that are 8-bit, wide or multibyte characters, see `multibyte(3)` for more details. Older applications do not recognize them and mistake them for control characters. Newer applications usually do recognize 8-bit characters. Depending on the implementation, users may be required to compile an application with wide or multibyte characters support, or configure it correctly. To be able to input and process wide or multibyte characters, the FreeBSD Ports collection ([../ports/index.html](http://www.freebsd.org/ports/index.html)) has provided each language with different programs. Refer to the I18N documentation in the respective FreeBSD Port.

Specifically, the user needs to look at the application documentation to decide on how to configure it correctly or to pass correct values into the `configure/Makefile/compiler`.

Some things to keep in mind are:

- Language specific single C chars character sets (see `multibyte(3)`), e.g. ISO-8859-1, ISO-8859-15, KOI8-R, CP437.
- Wide or multibyte encodings, e.g. EUC, Big5.

You can check the active list of character sets at the IANA Registry (<http://www.iana.org/assignments/character-sets>).

Note: DragonFly uses X11-compatible locale encodings instead.

14.3.3 I18N Applications

In the FreeBSD Ports and Package system, I18N applications have been named with I18N in their names for easy identification. However, they do not always support the language needed.

14.3.4 Setting Locale

Usually it is sufficient to export the value of the locale name as `LANG` in the login shell. This could be done in the user's `~/.login_conf` file or in the startup file of the user's shell (`~/.profile`, `~/.bashrc`, `~/.cshrc`). There is no need to set the locale subsets such as `LC_CTYPE`, `LC_TIME`. Please refer to language-specific DragonFly documentation for more information.

You should set the following two environment variables in your configuration files:

- `LANG` for POSIX `setlocale(3)` family functions
- `MM_CHARSET` for applications' MIME character set

This includes the user shell configuration, the specific application configuration, and the X11 configuration.

14.3.4.1 Setting Locale Methods

There are two methods for setting locale, and both are described below. The first (recommended one) is by assigning the environment variables in login class, and the second is by adding the environment variable assignments to the system's shell startup file.

14.3.4.1.1 Login Classes Method

This method allows environment variables needed for locale name and MIME character sets to be assigned once for every possible shell instead of adding specific shell assignments to each shell's startup file. User Level Setup can be done by an user himself and Administrator Level Setup require superuser privileges.

14.3.4.1.1.1 User Level Setup

Here is a minimal example of a `.login_conf` file in user's home directory which has both variables set for Latin-1 encoding:

```
me:\
:charset=ISO-8859-1:\
:lang=de_DE.ISO8859-1:
```

Here is an example of a `.login_conf` that sets the variables for Traditional Chinese in BIG-5 encoding. Notice the many more variables set because some software does not respect locale variables correctly for Chinese, Japanese, and Korean.

```
#Users who do not wish to use monetary units or time formats
#of Taiwan can manually change each variable
me:\
:lang=zh_TW.Big5:\
:lc_all=zh_TW.Big5:\
:lc_collate=zh_TW.Big5:\
:lc_ctype=zh_TW.Big5:\
```

```
:lc_messages=zh_TW.Big5:\
:lc_monetary=zh_TW.Big5:\
:lc_numeric=zh_TW.Big5:\
:lc_time=zh_TW.Big5:\
:charset=big5:\
:xmodifiers="@im=xcin": #Setting the XIM Input Server
```

See Administrator Level Setup and `login.conf(5)` for more details.

14.3.4.1.1.2 Administrator Level Setup

Verify that the user's login class in `/etc/login.conf` sets the correct language. Make sure these settings appear in `/etc/login.conf`:

```
language_name:accounts_title:\
:charset=MIME_charset:\
:lang=locale_name:\
:tc=default:
```

So sticking with our previous example using Latin-1, it would look like this:

```
german:German Users Accounts:\
:charset=ISO-8859-1:\
:lang=de_DE.ISO8859-1:\
:tc=default:
```

Changing Login Classes with `vipw(8)`

Use `vipw` to add new users, and make the entry look like this:

```
user:password:1111:11:language:0:0:User Name:/home/user:/bin/sh
```

Changing Login Classes with `adduser(8)`

Use `adduser` to add new users, and do the following:

- Set `defaultclass = language` in `/etc/adduser.conf`. Keep in mind you must enter a `default` class for all users of other languages in this case.
- An alternative variant is answering the specified language each time that


```
Enter login class: default []:
```

 appears from `adduser(8)`.
- Another alternative is to use the following for each user of a different language that you wish to add:

```
# adduser -class language
```

Changing Login Classes with `pw(8)`

If you use `pw(8)` for adding new users, call it in this form:

```
# pw useradd user_name -L language
```

14.3.4.1.2 Shell Startup File Method

Note: This method is not recommended because it requires a different setup for each possible shell program chosen. Use the Login Class Method instead.

To add the locale name and MIME character set, just set the two environment variables shown below in the `/etc/profile` and/or `/etc/csh.login` shell startup files. We will use the German language as an example below:

In `/etc/profile`:

```
LANG=de_DE.ISO8859-1; export LANG
MM_CHARSET=ISO-8859-1; export MM_CHARSET
```

Or in `/etc/csh.login`:

```
setenv LANG de_DE.ISO8859-1
setenv MM_CHARSET ISO-8859-1
```

Alternatively, you can add the above instructions to `/usr/share/skel/dot.profile` (similar to what was used in `/etc/profile` above), or `/usr/share/skel/dot.login` (similar to what was used in `/etc/csh.login` above).

For X11:

In `$HOME/.xinitrc`:

```
LANG=de_DE.ISO8859-1; export LANG
```

Or:

```
setenv LANG de_DE.ISO8859-1
```

Depending on your shell (see above).

14.3.5 Console Setup

For all single C chars character sets, set the correct console fonts in `/etc/rc.conf` for the language in question with:

```
font8x16=font_name
font8x14=font_name
font8x8=font_name
```

The `font_name` here is taken from the `/usr/share/syscons/fonts` directory, without the `.fnt` suffix.

Also be sure to set the correct keymap and screenmap for your single C chars character set. You can add the following to `/etc/rc.conf`:

```
scrnmap=screenmap_name
keymap=keymap_name
keychange="fkey_number sequence"
```

The *screenmap_name* here is taken from the `/usr/share/syscons/scrnmaps` directory, without the `.scm` suffix. A screenmap with a corresponding mapped font is usually needed as a workaround for expanding bit 8 to bit 9 on a VGA adapter's font character matrix in pseudographics area, i.e., to move letters out of that area if screen font uses a bit 8 column.

If you have the **moused** daemon enabled by setting the following in your `/etc/rc.conf`:

```
moused_enable="YES"
```

then examine the mouse cursor information in the next paragraph.

By default the mouse cursor of the `syscons(4)` driver occupies the `0xd0-0xd3` range in the character set. If your language uses this range, you need to move the cursor's range outside of it. To enable this workaround, insert the following line into your kernel configuration:

```
options SC_MOUSE_CHAR=0x03
```

Insert the following line into `/etc/rc.conf`:

```
mousechar_start=3
```

The *keymap_name* here is taken from the `/usr/share/syscons/keymaps` directory, without the `.kbd` suffix. If you're uncertain which keymap to use, you can use `kbdmap(1)` to test keymaps without rebooting.

The *keychange* is usually needed to program function keys to match the selected terminal type because function key sequences cannot be defined in the key map.

Also be sure to set the correct console terminal type in `/etc/ttys` for all `ttys*` entries. Current pre-defined correspondences are:

Character Set	Terminal Type
ISO-8859-1 or ISO-8859-15	cons2511
ISO-8859-2	cons2512
ISO-8859-7	cons2517
KOI8-R	cons25r
KOI8-U	cons25u
CP437 (VGA default)	cons25
US-ASCII	cons25w

For wide or multibyte characters languages, use the correct FreeBSD port in your `/usr/ports/language` directory. Some ports appear as console while the system sees it as serial vtty's, hence you must reserve enough vtty's for both X11 and the pseudo-serial console. Here is a partial list of applications for using other languages in console:

Language	Location
Traditional Chinese (BIG-5)	chinese/big5con
Japanese	japanese/ja-kon2-* or japanese/Mule_Wnn
Korean	korean/ko-han

14.3.6 X11 Setup

Although X11 is not part of DragonFly, we have included some information here for DragonFly users. For more details, refer to the XFree86 web site (<http://www.xfree86.org/>) or whichever X11 Server you use.

In `~/ .Xresources`, you can additionally tune application specific I18N settings (e.g., fonts, menus, etc.).

14.3.6.1 Displaying Fonts

Install the X11 TrueType Common server (`x11-servers/XttxF86srv-common`) and install the language TrueType fonts. Setting the correct locale should allow you to view your selected language in menus and such.

14.3.6.2 Inputting Non-English Characters

The X11 Input Method (XIM) Protocol is a new standard for all X11 clients. All X11 applications should be written as XIM clients that take input from XIM Input servers. There are several XIM servers available for different languages.

14.3.7 Printer Setup

Some single C chars character sets are usually hardware coded into printers. Wide or multibyte character sets require special setup and we recommend using **apsfilter**. You may also convert the document to PostScript or PDF formats using language specific converters.

14.3.8 Kernel and File Systems

The DragonFly fast filesystem (FFS) is 8-bit clean, so it can be used with any single C chars character set (see `multibyte(3)`), but there is no character set name stored in the filesystem; i.e., it is raw 8-bit and does not know anything about encoding order. Officially, FFS does not support any form of wide or multibyte character sets yet. However, some wide or multibyte character sets have independent patches for FFS enabling such support. They are only temporary unportable solutions or hacks and we have decided to not include them in the source tree. Refer to respective languages' web sites for more informations and the patch files.

The DragonFly MS-DOS filesystem has the configurable ability to convert between MS-DOS, Unicode character sets and chosen DragonFly filesystem character sets. See `mount_msdos(8)` for details.

14.4 Compiling I18N Programs

Many FreeBSD Ports have been ported with I18N support. Some of them are marked with `-I18N` in the port name. These and many other programs have built in support for I18N and need no special consideration.

However, some applications such as **MySQL** need to be have the `Makefile` configured with the specific charset. This is usually done in the `Makefile` or done by passing a value to **configure** in the source.

14.5 Localizing DragonFly to Specific Languages

14.5.1 Russian Language (KOI8-R Encoding)

Originally contributed by Andrey A. Chernov.

For more information about KOI8-R encoding, see the KOI8-R References (Russian Net Character Set) (<http://koi8.pp.ru/>).

14.5.1.1 Locale Setup

Put the following lines into your `~/ .login_conf` file:

```
me:My Account:\
:charset=KOI8-R:\
:lang=ru_RU.KOI8-R:
```

See earlier in this chapter for examples of setting up the locale.

14.5.1.2 Console Setup

- Add the following line to your kernel configuration file:

```
options SC_MOUSE_CHAR=0x03
```

Insert the following line into `/etc/rc.conf`:

```
mousechar_start=3
```

- Use following settings in `/etc/rc.conf`:

```
keymap="ru.koi8-r"
scrnmap="koi8-r2cp866"
font8x16="cp866b-8x16"
font8x14="cp866-8x14"
font8x8="cp866-8x8"
```

- For each `ttyv*` entry in `/etc/ttys`, use `cons25r` as the terminal type.

See earlier in this chapter for examples of setting up the console.

14.5.1.3 Printer Setup

Since most printers with Russian characters come with hardware code page CP866, a special output filter is needed to convert from KOI8-R to CP866. Such a filter is installed by default as `/usr/libexec/lpr/ru/koi2alt`. A Russian printer `/etc/printcap` entry should look like:

```
lp|Russian local line printer:\
:sh:of=/usr/libexec/lpr/ru/koi2alt:\
:lp=/dev/lpt0:sd=/var/spool/output/lpd:lf=/var/log/lpd-errs:
```

See `printcap(5)` for a detailed description.

14.5.1.4 MS-DOS® FS and Russian Filenames

The following example `fstab(5)` entry enables support for Russian filenames in mounted MS-DOS filesystems:

```
/dev/ad0s2      /dos/c  msdos   rw,-Wkoi2dos,-Lru_RU.KOI8-R 0 0
```

The option `-L` selects the locale name used, and `-w` sets the character conversion table. To use the `-w` option, be sure to mount `/usr` before the MS-DOS partition because the conversion tables are located in `/usr/libdata/msdosfs`. For more information, see the `mount_msdos(8)` manual page.

14.5.1.5 X11 Setup

1. Do non-X locale setup first as described.

Note: The Russian KOI8-R locale may not work with old **XFree86** releases (lower than 3.3). **XFree86 4.X** is now the default version of the X Window System on DragonFly. This should not be an issue unless you explicitly install an older version of **XFree86**.

2. Go to the `russian/X.language` directory and issue the following command:

```
# make install
```

The above port installs the latest version of the KOI8-R fonts. **XFree86 3.3** already has some KOI8-R fonts, but these are scaled better.

Check the "Files" section in your `/etc/XF86Config` file. The following lines must be added *before* any other `FontPath` entries:

```
FontPath "/usr/X11R6/lib/X11/fonts/cyrillic/misc"
FontPath "/usr/X11R6/lib/X11/fonts/cyrillic/75dpi"
FontPath "/usr/X11R6/lib/X11/fonts/cyrillic/100dpi"
```

If you use a high resolution video mode, swap the 75 dpi and 100 dpi lines.

3. To activate a Russian keyboard, add the following to the "Keyboard" section of your `XF86Config` file.

For **XFree86 3.X**:

```
XkbLayout "ru"
XkbOptions "grp:caps_toggle"
```

For **XFree86 4.X**:

```
Option "XkbLayout" "ru"
Option "XkbOptions" "grp:caps_toggle"
```

Also make sure that `XkbDisable` is turned off (commented out) there.

The RUS/LAT switch will be **CapsLock**. The old **CapsLock** function is still available via **Shift+CapsLock** (in LAT mode only).

If you have "Windows" keys on your keyboard, and notice that some non-alphabetical keys are mapped incorrectly in RUS mode, add the following line in your `XF86Config` file.

For **XFree86 3.X**:

```
XkbVariant "winkeys"
```

For **XFree86 4.X**:

```
Option "XkbVariant" "winkeys"
```

Note: The Russian XKB keyboard may not work with old **XFree86** versions, see the above note for more information. The Russian XKB keyboard may also not work with non-localized applications as well. Minimally localized applications should call a `XtSetLanguageProc (NULL, NULL, NULL);` function early in the program. See KOI8-R for X Window (<http://koi8.pp.ru/xwin.html>) for more instructions on localizing X11 applications.

14.5.2 Traditional Chinese Localization for Taiwan

The FreeBSD-Taiwan Project has an I18N/L10N tutorial for FreeBSD at <http://freebsd.sinica.edu.tw/~ncvs/zh-l10n-tut/> using many Chinese ports. Much of that project can apply to DragonFly. The editor for the zh-L10N-tut is Clive Lin <Clive@CirX.org>. You can also cvsup the following collections at freebsd.sinica.edu.tw:

Collection	Description
outta-port tag=.	Beta-quality ports collection for Chinese
zh-L10N-tut tag=.	Localizing FreeBSD Tutorial in BIG-5 Traditional Chinese
zh-doc tag=.	FreeBSD Documentation Translation to BIG-5 Traditional Chinese

Chuan-Hsing Shen <s874070@mail.yzu.edu.tw> has created the Chinese FreeBSD Collection (CFC) (<http://cnpa.yzu.edu.tw/~cfc/>) using FreeBSD-Taiwan's zh-L10N-tut. The packages and the script files are available at [ftp://ftp.csie.ncu.edu.tw/OS/FreeBSD/taiwan/CFC/](http://ftp.csie.ncu.edu.tw/OS/FreeBSD/taiwan/CFC/).

14.5.3 German Language Localization (for All ISO 8859-1 Languages)

Slaven Rezic <eserte@cs.tu-berlin.de> wrote a tutorial how to use umlauts on a FreeBSD machine. The tutorial is written in German and available at <http://www.de.FreeBSD.org/de/umlaute/>.

14.5.4 Japanese and Korean Language Localization

For Japanese, refer to <http://www.jp.FreeBSD.org/>, and for Korean, refer to <http://www.kr.FreeBSD.org/>.

14.5.5 Non-English DragonFly Documentation

Non-English documentation will be made available as it is created, at the main site ([../..../index.html](http://www.dragonflybsd.org/..../index.html)) or in `/usr/share/doc`.

Chapter 15 Desktop Applications

Contributed by Christophe Juniet.

15.1 Synopsis

DragonFly can run a wide variety of desktop applications, such as browsers and word processors. Most of these are available as packages or can be automatically built from the ports collection. Many new users expect to find these kinds of applications on their desktop. This chapter will show you how to install some popular desktop applications effortlessly, either from their packages or from the ports collection.

Warning: This chapter contains a number of outdated references to the FreeBSD ports collection. Most instructions still apply to pkgsrc, but proceed with caution until this chapter is updated.

Note that when installing programs from the ports, they are compiled from source. This can take a very long time, depending on what you are compiling and the processing power of your machine(s). If building from source takes a prohibitively long amount of time for you, you can install most of the programs of the ports collection from pre-built packages.

As DragonFly features Linux binary compatibility, many applications originally developed for Linux are available for your desktop. It is strongly recommended that you read [Chapter 22](#) before installing any of the Linux applications. Many of the ports using the Linux binary compatibility start with “linux-”. Remember this when you search for a particular port, for instance with `whereis(1)`. In the following text, it is assumed that you have enabled Linux binary compatibility before installing any of the Linux applications.

Here are the categories covered by this chapter:

- Browsers (such as **Mozilla**, **Netscape**, **Opera**)
- Productivity (such as **KOffice**, **AbiWord**, **The GIMP**, **OpenOffice.org**)
- Document Viewers (such as **Acrobat Reader®**, **gv**, **Xpdf**, **GQview**)
- Finance (such as **GnuCash**, **Gnumeric**, **Abacus**)

Before reading this chapter, you should:

- Know how to install additional third-party software ([Chapter 4](#)).
- Know how to install additional Linux software ([Chapter 22](#)).

For information on how to get a multimedia environment, read [Chapter 16](#). If you want to set up and use electronic mail, please refer to [Chapter 20](#).

15.2 Browsers

DragonFly does not come with a particular browser pre-installed. Instead, the `www` (<http://www.FreeBSD.org/ports/www.html>) directory of the ports collection contains a lot of browsers ready to be

installed. If you do not have time to compile everything (this can take a very long time in some cases) many of them are available as packages.

KDE and **GNOME** already provide HTML browsers. Please refer to Section 5.7 for more information on how to set up these complete desktops.

If you are looking for light-weight browsers, you should investigate the ports collection for `www/dillo`, `www/links`, or `www/w3m`.

This section covers these applications:

Application Name	Resources Needed	Installation from Ports	Major Dependencies
Mozilla	heavy	heavy	Gtk+
Netscape	heavy	light	Linux Binary Compatibility
Opera	light	light	FreeBSD version (should work on DragonFly): None. Linux version: Linux Binary Compatibility and linux-openmotif

15.2.1 Mozilla

Mozilla is perhaps the most suitable browser for your DragonFly Desktop. It is modern and stable. It features a very standards-compliant HTML display engine. It provides a mail and news reader. It even has a HTML composer if you plan to write some web pages yourself. Users of **Netscape** will recognize the similarities with **Communicator** suite, as both browsers shared the same basis.

On slow machines, with a CPU speed less than 233MHz or with less than 64MB of RAM, **Mozilla** can be too resource-consuming to be fully usable. You may want to look at the **Opera** browser instead, described a little later in this chapter.

The Mozilla package from the network by:

```
# pkg_add mozilla
```

If the package is not available, and you have enough time and disk space, you can get the source for **Mozilla**, compile it and install it on your system. This is accomplished by:

```
# cd /usr/ports/www/mozilla
# make install clean
```

The **Mozilla** port ensures a correct initialization by running the chrome registry setup with `root` privileges. However, if you want to fetch some add-ons like mouse gestures, you must run **Mozilla** as `root` to get them properly installed.

Once you have completed the installation of **Mozilla**, you do not need to be `root` any longer. You can start **Mozilla** as a browser by typing:

```
% mozilla
```

You can start it directly as a mail and news reader as shown below:

```
% mozilla -mail
```

15.2.2 Mozilla, Java™, and Macromedia® Flash™

Contributed by Tom Rhodes.

Installing **Mozilla** is simple, but unfortunately installing **Mozilla** with support for add-ons like Java™ and Macromedia® Flash™ consumes both time and disk space.

The first thing is to download the files which will be used with **Mozilla**. Take your current web browser up to <http://www.sun.com/software/java2/download.html> and create an account on their website. Remember to save the username and password from here as it may be needed in the future. Download a copy of the file `j2sdk-1_3_1-src.tar.gz` and place this in `/usr/ports/distfiles/` as the port will not fetch it automatically. This is due to license restrictions. While we are here, download the “java environment” from <http://java.sun.com/webapps/download/Display?BundleId=7905>. The filename is `j2sdk-1_3_1_08-linux-i586.bin` and is large (about 25 megabytes!). Like before, this file must be placed into `/usr/ports/distfiles/`. Finally download a copy of the “java patchkit” from <http://www.eyesbeyond.com/freebsd/ports/java/> and place it into `/usr/ports/distfiles/`.

Install the `java/jdk13` port with the standard `make install clean` and then install the `www/flashpluginwrapper` port. This port requires `emulators/linux_base` which is a large port. True that other **Flash** plugins exist, however they have not worked for me.

Install the `www/mozilla` port, if **Mozilla** is not already installed.

Now copy the **Flash** plug-in files with:

```
# cp /usr/local/lib/flash/libflashplayer.so \
/usr/X11R6/lib/browser_plugins/libflashplayer_linux.so

# cp /usr/local/lib/flash/ShockwaveFlash.class \
/usr/X11R6/lib/browser_plugins/
```

Now add the following lines to the top of (but right under `#!/bin/sh`) **Mozilla** startup script:

```
/usr/X11R6/bin/mozilla.
```

```
LD_PRELOAD=/usr/local/lib/libflashplayer.so.1
export LD_PRELOAD
```

This will enable the **Flash** plug-in.

Now just start **Mozilla** with:

```
% mozilla &
```

And access the About Plug-ins option from the Help menu. A list should appear with all the currently available plugins. **Java** and **Shockwave® Flash** should both be listed.

15.2.3 Netscape®

The ports collection contains several versions of the Netscape browser. Since the native FreeBSD ones contain a serious security bug, installing them is strongly discouraged. Instead, use a more recent Linux or DIGITAL UNIX version.

The latest stable release of the Netscape browser is **Netscape 7**. It can be installed from the ports collection:

```
# cd /usr/ports/www/netscape7
```

```
# make install clean
```

There are localized versions in the French, German, and Japanese categories.

Caution: Netscape 4.x versions are not recommended because they are not compliant with today's standards. However, Netscape 7.x and newer versions are only available for the i386 platform.

15.2.4 Opera

Opera is a very fast, full-featured, and standards-compliant browser. It comes in a version that runs under Linux emulation. There is a no-cost version of the browser that displays advertising and an ad-free version that can be purchased on the Opera web site (<http://www.opera.com/>).

To browse the Web with **Opera**, install the package:

```
# pkg_add opera
```

Some FTP sites do not have all the packages, but the same result can be obtained with the ports collection by typing:

```
# cd /usr/ports/www/opera
# make install clean
```

15.3 Productivity

When it comes to productivity, new users often look for a good office suite or a friendly word processor. While some desktop environments like **KDE** already provide an office suite, there is no default application. DragonFly provides all that is needed, regardless of your desktop environment.

This section covers these applications:

Application Name	Resources Needed	Installation from Ports	Major Dependencies
KOffice	light	heavy	KDE
AbiWord	light	light	Gtk+ or GNOME
The Gimp	light	heavy	Gtk+
OpenOffice.org	heavy	huge	GCC 3.1 , JDK™ 1.3 , Mozilla

15.3.1 KOffice

The KDE community has provided its desktop environment with an office suite which can be used outside **KDE**. It includes the four standard components that can be found in other office suites. **KWord** is the word processor, **KSpread** is the spreadsheet program, **KPresenter** manages slide presentations, and **Kontour** lets you draw graphical documents.

Before installing the latest **KOffice**, make sure you have an up-to-date version of **KDE**.

To install **KOffice** as a package, issue the following command:

```
# pkg_add koffice
```

If the package is not available, you can use the ports collection. For instance, to install **KOffice** for **KDE3**, do:

```
# cd /usr/ports/editors/koffice-kde3
# make install clean
```

15.3.2 AbiWord

AbiWord is a free word processing program similar in look and feel to **Microsoft Word**. It is suitable for typing papers, letters, reports, memos, and so forth. It is very fast, contains many features, and is very user-friendly.

AbiWord can import or export many file formats, including some proprietary ones like Microsoft `.doc`.

AbiWord is available as a package. You can install it by:

```
# pkg_add AbiWord
```

If the package is not available, it can be compiled from the ports collection. The ports collection should be more up to date. It can be done as follows:

```
# cd /usr/ports/editors/AbiWord
# make install clean
```

15.3.3 The GIMP

For image authoring or picture retouching, **The GIMP** is a very sophisticated image manipulation program. It can be used as a simple paint program or as a quality photo retouching suite. It supports a large number of plug-ins and features a scripting interface. **The GIMP** can read and write a wide range of file formats. It supports interfaces with scanners and tablets.

You can install the package by issuing this command:

```
# pkg_add -r gimp
```

If your FTP site does not have this package, you can use the ports collection. The graphics (<http://www.FreeBSD.org/ports/graphics.html>) directory of the ports collection also contains **The Gimp Manual**. Here is how to get them installed:

```
# cd /usr/ports/graphics/gimp1
# make install clean
# cd /usr/ports/graphics/gimp-manual-pdf
# make install clean
```

Note: The graphics (<http://www.FreeBSD.org/ports/graphics.html>) directory of the ports collection holds the development version of **The GIMP** in `graphics/gimp-devel`. HTML and PostScript versions of **The Gimp Manual** are in `graphics/gimp-manual-html` and `graphics/gimp-manual-ps`.

15.3.4 OpenOffice.org

OpenOffice.org includes all of the mandatory applications in a complete office productivity suite: a word processor, a spreadsheet, a presentation manager, and a drawing program. Its user interface is very similar to other office suites, and it can import and export in various popular file formats. It is available in a number of different languages including interfaces, spell checkers, and dictionaries.

The word processor of **OpenOffice.org** uses a native XML file format for increased portability and flexibility. The spreadsheet program features a macro language and it can be interfaced with external databases. **OpenOffice.org** is already stable and runs natively on Windows, Solaris™, Linux, FreeBSD, and Mac OS X. More information about **OpenOffice.org** can be found on the OpenOffice web site (<http://www.openoffice.org/>). For FreeBSD specific information, and to directly download packages use the FreeBSD OpenOffice Porting Team (<http://projects.imp.ch/openoffice/>)'s web site. These packages should work on DragonFly.

To install **OpenOffice.org**, do:

```
# pkg_add openoffice
```

Once the package is installed, you must run the setup program and choose a standard workstation installation. Run this command as the user who will use **OpenOffice.org**:

```
% openoffice-setup
```

If the **OpenOffice.org** packages are not available, you still have the option to compile the port. However, you must bear in mind that it requires a lot of disk space and a fairly long time to compile.

```
# cd /usr/ports/editors/openoffice
# make install clean
```

Once this is done, run the setup as the user who will use **OpenOffice.org** and choose a standard workstation installation by:

```
% cd /usr/ports/editors/openoffice
% make install-user
```

If you want to use a localized version, here are the available ports:

Language	Port
Arabic	editors/openoffice-ar
Danish	editors/openoffice-dk
Spanish	editors/openoffice-es
Greek	editors/openoffice-gr
Italian	editors/openoffice-it
Dutch	editors/openoffice-nl
Swedish	editors/openoffice-se
Turkish	editors/openoffice-tr
French	french/openoffice
German	german/openoffice
Japanese	japanese/openoffice

Language	Port
Korean	korean/openoffice
Polish	polish/openoffice
Portuguese	portuguese/openoffice
Russian	russian/openoffice

15.4 Document Viewers

Some new document formats have recently gained popularity. The standard viewers they require may not be available in the base system. We will see how to install them in this section.

This section covers these applications:

Application Name	Resources Needed	Installation from Ports	Major Dependencies
Acrobat Reader	light	light	Linux Binary Compatibility
gv	light	light	Xaw3d
Xpdf	light	light	FreeType
GQview	light	light	Gtk+ or GNOME

15.4.1 Acrobat Reader®

Many documents are now distributed as PDF files, which stands for “Portable Document Format”. One of the recommended viewers for these types of files is **Acrobat Reader**, released by Adobe for Linux. As DragonFly can run Linux binaries, it is also available for DragonFly.

To install the **Acrobat Reader 5** package, do:

```
# pkg_add acroread5
```

As usual, if the package is not available or you want the latest version, you can use the ports collection as well:

```
# cd /usr/ports/print/acroread5
# make install clean
```

Note: **Acrobat Reader** is available in several different versions. At this time of writing, there are: `print/acroread` (version 3.0.2), `print/acroread4` (version 4.0.5), and `print/acroread5` (version 5.0.6). They may not all have been packaged for your version of DragonFly. The ports collection will always contain the latest versions.

15.4.2 gv

gv is a PostScript and PDF viewer. It is originally based on **ghostview** but it has a nicer look thanks to the **Xaw3d** library. It is fast and its interface is clean. **gv** has many features like orientation, paper size, scale, or antialias. Almost

any operation can be done either from the keyboard or the mouse.

To install **gv** as a package, do:

```
# pkg_add gv
```

If you cannot get the package, you can use the ports collection:

```
# cd /usr/ports/print/gv
# make install clean
```

15.4.3 Xpdf

If you want a small DragonFly PDF viewer, **Xpdf** is a light-weight and efficient viewer. It requires very few resources and is very stable. It uses the standard X fonts and does not require **Motif** or any other X toolkit.

To install the **Xpdf** package, issue this command:

```
# pkg_add xpdf
```

If the package is not available or you prefer to use the ports collection, do:

```
# cd /usr/ports/graphics/xpdf
# make install clean
```

Once the installation is complete, you can launch **Xpdf** and use the right mouse button to activate the menu.

15.4.4 GQview

GQview is an image manager. You can view a file with a single click, launch an external editor, get thumbnail previews, and much more. It also features a slideshow mode and some basic file operations. You can manage image collections and easily find duplicates. **GQview** can do full screen viewing and supports internationalization.

If you want to install the **GQview** package, do:

```
# pkg_add gqview
```

If the package is not available or you prefer to use the ports collection, do:

```
# cd /usr/ports/graphics/gqview
# make install clean
```

15.5 Finance

If, for any reason, you would like to manage your personal finances on your DragonFly Desktop, there are some powerful and easy to use applications ready to be installed. Some of them are compatible with widespread file formats like those of **Quicken**® or **Excel** documents.

This section covers these applications:

Application Name	Resources Needed	Installation from Ports	Major Dependencies
GnuCash	light	heavy	GNOME
Gnumeric	light	heavy	GNOME
Abacus	light	light	Tcl/Tk

15.5.1 GnuCash

GnuCash is part of the **GNOME** effort to provide user-friendly yet powerful applications to end-users. With **GnuCash**, you can keep track of your income and expenses, your bank accounts, or your stocks. It features an intuitive interface while remaining very professional.

GnuCash provides a smart register, a hierarchical system of accounts, many keyboard accelerators and auto-completion methods. It can split a single transaction into several more detailed pieces. **GnuCash** can import and merge **Quicken** QIF files. It also handles most international date and currency formats.

To install **GnuCash** on your system, do:

```
# pkg_add gnuCash
```

If the package is not available, you can use the ports collection:

```
# cd /usr/ports/finance/gnuCash
# make install clean
```

15.5.2 Gnumeric

Gnumeric is a spreadsheet, part of the **GNOME** desktop environment. It features convenient automatic “guessing” of user input according to the cell format and an autofill system for many sequences. It can import files in a number of popular formats like those of **Excel**, **Lotus 1-2-3**, or **Quattro Pro**. **Gnumeric** supports graphs through the `math/guppi` graphing program. It has a large number of built-in functions and allows all of the usual cell formats such as number, currency, date, time, and much more.

To install **Gnumeric** as a package, type in:

```
# pkg_add gnumeric
```

If the package is not available, you can use the ports collection by doing:

```
# cd /usr/ports/math/gnumeric
# make install clean
```

15.5.3 Abacus

Abacus is a small and easy to use spreadsheet. It includes many built-in functions useful in several domains such as statistics, finances, and mathematics. It can import and export the **Excel** file format. **Abacus** can produce PostScript output.

To install **Abacus** from its package, do:

```
# pkg_add abacus
```

If the package is not available, you can use the ports collection by doing:

```
# cd /usr/ports/desktoputils/abacus
# make install clean
```

15.6 Summary

DragonFly is quite ready for day-to-day use as a desktop. With several thousand applications available as packages (<http://www.FreeBSD.org/where.html>) or ports (<http://www.FreeBSD.org/ports/index.html>), you can build a perfect desktop that suits all your needs.

Once you have achieved the installation of your desktop, you may want to go one step further with `misc/instant-workstation`. This “meta-port” allows you to build a typical set of ports for a workstation. You can customize it by editing `/usr/ports/misc/instant-workstation/Makefile`. Follow the syntax used for the default set to add or remove ports, and build it with the usual procedure. Eventually, you will be able to create a big package that corresponds to your very own desktop and install it to your other workstations!

Here is a quick review of all the desktop applications covered in this chapter:

Application Name	Package Name	Ports Name
Mozilla	mozilla	www/mozilla
Netscape	linux-netscape7	www/netscape7
Opera	linux-opera	www/linux-opera
KOffice	koffice-kde3	editors/koffice-kde3
AbiWord	AbiWord	editors/AbiWord
The GIMP	gimp	graphics/gimp1
OpenOffice.org	openoffice	editors/openoffice
Acrobat Reader	acroread5	print/acroread5
gv	gv	print/gv
Xpdf	xpdf	graphics/xpdf
GQview	gqview	graphics/gqview
GnuCash	gnucash	finance/gnucash
Gnumeric	gnumeric	math/gnumeric
Abacus	abacus	desktoputils/abacus

Chapter 16 Multimedia

Edited by Ross Lippert.

16.1 Synopsis

DragonFly supports a wide variety of sound cards, allowing you to enjoy high fidelity output from your computer. This includes the ability to record and playback audio in the MPEG Audio Layer 3 (MP3), WAV, and Ogg Vorbis formats as well as many other formats. The pkgsrc tree also contain applications allowing you to edit your recorded audio, add sound effects, and control attached MIDI devices.

With some willingness to experiment, DragonFly can support playback of video files and DVD's. The number of applications to encode, convert, and playback various video media is more limited than the number of sound applications. For example as of this writing, there is no good re-encoding application in the pkgsrc tree, which could be use to convert between formats, as there is with `audio/sox`. However, the software landscape in this area is changing rapidly.

This chapter will describe the necessary steps to configure your sound card. The configuration and installation of X11 (Chapter 5) has already taken care of the hardware issues for your video card, though there may be some tweaks to apply for better playback.

After reading this chapter, you will know:

- How to configure your system so that your sound card is recognized.
- Methods to test that your card is working using sample applications.
- How to troubleshoot your sound setup.
- How to playback and encode MP3s and other audio.
- How video is supported by the X server.
- Some video player/encoder ports which give good results.
- How to playback DVD's, `.mpg` and `.avi` files.
- How to rip CD and DVD information into files.
- How to configure a TV card.

Before reading this chapter, you should:

- Know how to configure and install a new kernel (Chapter 9).

Warning: Trying to mount audio CDs with the `mount(8)` command will result in an error, at least, and a *kernel panic*, at worst. These media have specialized encodings which differ from the usual ISO-filesystem.

16.2 Setting Up the Sound Card

Contributed by Moses Moore.

16.2.1 Locating the Correct Device

Before you begin, you should know the model of the card you have, the chip it uses, and whether it is a PCI or ISA card. DragonFly supports a wide variety of both PCI and ISA cards. If you do not see your card in the following list, check the pcm(4) manual page. This is not a complete list; however, it does list some of the most common cards.

- Crystal 4237, 4236, 4232, 4231
- Yamaha OPL-SA_x
- OPTi931
- Ensoniq AudioPCI 1370/1371
- ESS Solo-1/1E
- NeoMagic 256AV/ZX
- SoundBlaster® Pro, 16, 32, AWE64, AWE128, Live
- Creative ViBRA16
- Advanced Asound 100, 110, and Logic ALS120
- ES 1868, 1869, 1879, 1888
- Gravis UltraSound
- Aureal Vortex 1 or 2

To use your sound device, you will need to load the proper device driver. This may be accomplished in one of two ways. The easiest way is to simply load a kernel module for your sound card with `kldload(8)` which can either be done from the command line:

```
# kldload snd_emu10k1.ko
```

or by adding the appropriate line to the file `/boot/loader.conf` like this:

```
snd_emu10k1_load="YES"
```

These examples are for a Creative SoundBlaster Live! sound card. Other available loadable sound modules are listed in `/boot/defaults/loader.conf`.

Alternatively, you may statically compile in support for your sound card in your kernel. The sections below provide the information you need to add support for your hardware in this manner. For more information about recompiling your kernel, please see Chapter 9.

16.2.1.1 Creative, Advance, and ESS Sound Cards

If you have one of the above cards, you will need to add:

```
device pcm
```

to your kernel configuration file. If you have a PnP ISA card, you will also need to add:

```
device sbc
```

For a non-PnP ISA card, add:

```
device pcm
device sbc0 at isa? port 0x220 irq 5 drq 1 flags 0x15
```

to your kernel configuration file. The settings shown above are the defaults. You may need to change the IRQ or the other settings to match your card. See the sbc(4) manual page for more information.

16.2.1.2 Gravis UltraSound Cards

For a PnP ISA card, you will need to add:

```
device pcm
device gusc
```

to your kernel configuration file. If you have a non-PnP ISA card, you will need to add:

```
device pcm
device gusc0 at isa? port 0x220 irq 5 drq 1 flags 0x13
```

to your kernel configuration file. You may need to change the IRQ or the other settings to match your card. See the gusc(4) manual page for more information.

16.2.1.3 Crystal Sound Cards

For Crystal cards, you will need to add:

```
device pcm
device csa
```

to your kernel configuration file.

16.2.1.4 Generic Support

For PnP ISA or PCI cards, you will need to add:

```
device pcm
```

to your kernel configuration file. If you have a non-PnP ISA sound card that does not have a bridge driver, you will need to add:

```
device pcm0 at isa? irq 10 drq 1 flags 0x0
```

to your kernel configuration file. You may need to change the IRQ or the other settings to match your card.

16.2.1.5 Onboard Sound

Some systems with built-in motherboard sound devices may require the following device in your kernel configuration:

```
device pnpbios
```

16.2.2 Creating and Testing the Device Nodes

After you reboot, log in and check for the device in the `/var/run/dmesg.boot` file, as shown below:

```
# grep pcm /var/run/dmesg.boot
pcm0: <SB16 DSP 4.11> on sbc0
```

The output from your system may look different. If no `pcm` devices show up, something went wrong earlier. If that happens, go through your kernel configuration file again and make sure you chose the correct device. Common problems are listed in Section 16.2.2.1.

If the previous command returned `pcm0`, you will have to run the following as `root`:

```
# cd /dev
# sh MAKEDEV snd0
```

If the command returned `pcm1`, follow the same steps as shown above, replacing `snd0` with `snd1`.

Note: The above commands will *not* create a `/dev/snd` device!

MAKEDEV will create a group of device nodes, including:

Device	Description
<code>/dev/audio</code>	Sparc® compatible audio device
<code>/dev/dsp</code>	Digitized voice device
<code>/dev/dspW</code>	Like <code>/dev/dsp</code> , but 16 bits per sample
<code>/dev/midi</code>	Raw midi access device
<code>/dev/mixer</code>	Control port mixer device
<code>/dev/music</code>	Level 2 sequencer interface
<code>/dev/sequencer</code>	Sequencer device
<code>/dev/pss</code>	Programmable device interface

If all goes well, you should now have a functioning sound card. If your CD-ROM or DVD-ROM drive is properly coupled to your sound card, you can put a CD in the drive and play it with `cdcontrol(1)`:

```
% cdcontrol -f /dev/acd0c play 1
```

Various applications, such as `audio/xmms` offer a better interface. You may want to install an application such as `audio/mpg123` to listen to MP3 audio files.

16.2.2.1 Common Problems

Error	Solution
unsupported subdevice XX	One or more of the device nodes was not created correctly. Repeat the steps above.
sb_dspwr(XX) timed out	The I/O port is not set correctly.
bad irq XX	The IRQ is set incorrectly. Make sure that the set IRQ and the sound IRQ are the same.
xxx: gus pcm not attached, out of memory	There is not enough available memory to use the device.
xxx: can't open /dev/dsp!	Check with <code>fstat grep dsp</code> if another application is holding the device open. Noteworthy troublemakers are esound and KDE's sound support.

16.2.3 Utilizing Multiple Sound Sources

Contributed by Munish Chopra.

It is often desirable to have multiple sources of sound that are able to play simultaneously, such as when **esound** or **artsd** do not support sharing of the sound device with a certain application.

DragonFly lets you do this through *Virtual Sound Channels*, which can be set with the `sysctl(8)` facility. Virtual channels allow you to multiplex your sound card's playback channels by mixing sound in the kernel.

To set the number of virtual channels, there are two `sysctl` knobs which, if you are the `root` user, can be set like this:

```
# sysctl hw.snd.pcm0.vchans=4
# sysctl hw.snd.maxautovchans=4
```

The above example allocates four virtual channels, which is a practical number for everyday use.

`hw.snd.pcm0.vchans` is the number of virtual channels `pcm0` has, and is configurable once a device has been attached. `hw.snd.maxautovchans` is the number of virtual channels a new audio device is given when it is attached using `kldload(8)`. Since the `pcm` module can be loaded independently of the hardware drivers, `hw.snd.maxautovchans` can store how many virtual channels any devices which are attached later will be given.

As DragonFly doesn't support `devfs` you will have to point your applications at `/dev/dsp0.x`, where `x` is 0 to 3 if `hw.snd.pcm0.vchans` is set to 4 as in the above example.

16.3 MP3 Audio

Contributed by Chern Lee.

MP3 (MPEG Layer 3 Audio) accomplishes near CD-quality sound, leaving no reason to let your DragonFly workstation fall short of its offerings.

16.3.1 MP3 Players

A popular X11 MP3 player is **XMMS** (X Multimedia System). **Winamp** skins can be used with **XMMS** since the GUI is almost identical to that of Nullsoft's **Winamp**. **XMMS** also has native plug-in support.

XMMS can be installed from the `audio/xmms` port or package.

XMMS' interface is intuitive, with a playlist, graphic equalizer, and more. Those familiar with **Winamp** will find **XMMS** simple to use.

The `audio/mpg123` port is an alternative, command-line MP3 player.

mpg123 can be run by specifying the sound device and the MP3 file on the command line, as shown below:

```
# mpg123 -a /dev/dsp1.0 Foobar-GreatestHits.mp3
High Performance MPEG 1.0/2.0/2.5 Audio Player for Layer 1, 2 and 3.
Version 0.59r (1999/Jun/15). Written and copyrights by Michael Hipp.
Uses code from various people. See 'README' for more!
THIS SOFTWARE COMES WITH ABSOLUTELY NO WARRANTY! USE AT YOUR OWN RISK!
```

```
Playing MPEG stream from Foobar-GreatestHits.mp3 ...
MPEG 1.0 layer III, 128 kbit/s, 44100 Hz joint-stereo
```

`/dev/dsp1.0` should be replaced with the `dsp` device entry on your system.

16.3.2 Ripping CD Audio Tracks

Before encoding a CD or CD track to MP3, the audio data on the CD must be ripped onto the hard drive. This is done by copying the raw CDDA (CD Digital Audio) data to WAV files.

The `cdda2wav` tool, which is a part of the `sysutils/cdrtools` suite, is used for ripping audio information from CDs and the information associated with them.

With the audio CD in the drive, the following command can be issued (as `root`) to rip an entire CD into individual (per track) WAV files:

```
# cdda2wav -D 0,1,0 -B
```

cdda2wav will support ATAPI (IDE) CDROM drives. To rip from an IDE drive, specify the device name in place of the SCSI unit numbers. For example, to rip track 7 from an IDE drive:

```
# cdda2wav -D /dev/acd0a -t 7
```

The `-D 0,1,0` indicates the SCSI device `0,1,0`, which corresponds to the output of `cdrecord -scanbus`.

To rip individual tracks, make use of the `-t` option as shown:

```
# cdda2wav -D 0,1,0 -t 7
```

This example rips track seven of the audio CDROM. To rip a range of tracks, for example, track one to seven, specify a range:

```
# cdda2wav -D 0,1,0 -t 1+7
```

The utility `dd(1)` can also be used to extract audio tracks on ATAPI drives, read Section 12.5.5 for more information on that possibility.

16.3.3 Encoding MP3s

Nowadays, the mp3 encoder of choice is **lame**. **Lame** can be found at `audio/lame` in the `pkgsrc` tree.

Using the ripped WAV files, the following command will convert `audio01.wav` to `audio01.mp3`:

```
# lame -h -b 128 \
--tt "Foo Song Title" \
--ta "FooBar Artist" \
--tl "FooBar Album" \
--ty "2001" \
--tc "Ripped and encoded by Foo" \
--tg "Genre" \
audio01.wav audio01.mp3
```

128 kbits seems to be the standard MP3 bitrate in use. Many enjoy the higher quality 160, or 192. The higher the bitrate, the more disk space the resulting MP3 will consume--but the quality will be higher. The `-h` option turns on the “higher quality but a little slower” mode. The options beginning with `--t` indicate ID3 tags, which usually contain song information, to be embedded within the MP3 file. Additional encoding options can be found by consulting the lame man page.

16.3.4 Decoding MP3s

In order to burn an audio CD from MP3s, they must be converted to a non-compressed WAV format. Both **XMMS** and **mpg123** support the output of MP3 to an uncompressed file format.

Writing to Disk in **XMMS**:

1. Launch **XMMS**.
2. Right-click on the window to bring up the **XMMS** menu.
3. Select `Preference` under `Options`.
4. Change the Output Plugin to “Disk Writer Plugin”.
5. Press `Configure`.
6. Enter (or choose browse) a directory to write the uncompressed files to.
7. Load the MP3 file into **XMMS** as usual, with volume at 100% and EQ settings turned off.
8. Press `Play` — **XMMS** will appear as if it is playing the MP3, but no music will be heard. It is actually playing the MP3 to a file.
9. Be sure to set the default Output Plugin back to what it was before in order to listen to MP3s again.

Writing to stdout in **mpg123**:

1. Run `mpg123 -s audio01.mp3 > audio01.pcm`

XMMS writes a file in the WAV format, while **mpg123** converts the MP3 into raw PCM audio data. Both of these formats can be used with **cdrecord** to create audio CDs. You have to use raw PCM with `burncd(8)`. If you use WAV files, you will notice a small tick sound at the beginning of each track, this sound is the header of the WAV file. You can simply remove the header of a WAV file with the utility **SoX** (it can be installed from the `audio/sox` port or package):

```
% sox -t wav -r 44100 -s -w -c 2 track.wav track.raw
```

Read Section 12.5 for more information on using a CD burner in DragonFly.

16.4 Video Playback

Contributed by Ross Lippert.

Video playback is a very new and rapidly developing application area. Be patient. Not everything is going to work as smoothly as it did with sound.

Before you begin, you should know the model of the video card you have and the chip it uses. While **XFree86** and **X.org** support a wide variety of video cards, fewer give good playback performance. To obtain a list of extensions supported by the X server using your card use the command `xdpyinfo(1)` while X11 is running.

It is a good idea to have a short MPEG file which can be treated as a test file for evaluating various players and options. Since some DVD players will look for DVD media in `/dev/dvd` by default, or have this device name hardcoded in them, you might find it useful to make symbolic links to the proper devices:

```
# ln -sf /dev/acd0c /dev/dvd
# ln -sf /dev/racd0c /dev/rdvd
```

Additionally, DVD decryption, which requires invoking special DVD-ROM functions, requires write permission on the DVD devices.

Some of the packages discussed rely on the following kernel options to build correctly. Before attempting to build, add these options to the kernel configuration file, build a new kernel, and reboot:

```
option CPU_ENABLE_SSE
option USER_LDT
```

To enhance the shared memory X11 interface, it is recommended that the values of some `sysctl(8)` variables should be increased:

```
kern.ipc.shmmax=67108864
kern.ipc.shmall=32768
```

16.4.1 Determining Video Capabilities

There are several possible ways to display video under X11. What will really work is largely hardware dependent. Each method described below will have varying quality across different hardware. Secondly, the rendering of video in X11 is a topic receiving a lot of attention lately, and with each version of **X.org** or **XFree86** there may be significant improvement.

A list of common video interfaces:

1. X11: normal X11 output using shared memory.
2. XVideo: an extension to the X11 interface which supports video in any X11 drawable.
3. SDL: the Simple Directmedia Layer.
4. DGA: the Direct Graphics Access.
5. SVGAlib: low level console graphics layer.

16.4.1.1 XVideo

Both **XFree86 4.X** and **X.org** have an extension called *XVideo* (aka Xvideo, aka Xv, aka xv) which allows video to be directly displayed in drawable objects through a special acceleration. This extension provides very good quality playback even on low-end machines (for example my PIII 400 Mhz laptop). Unfortunately, the list of cards in which this feature is supported “out of the box” is currently:

1. 3DFX Voodoo 3
2. Intel i810 and i815
3. some S3 chips (such as Savage/IX and Savage/MX)

If your card is not one of these, do not be disappointed yet. **XFree86 4.X** adds new xv capabilities with each release¹. To check whether the extension is running, use `xvinfo`:

```
% xvinfo
```

XVideo is supported for your card if the result looks like:

```
X-Video Extension version 2.2
screen #0
  Adaptor #0: "Savage Streams Engine"
    number of ports: 1
    port base: 43
    operations supported: PutImage
    supported visuals:
      depth 16, visualID 0x22
      depth 16, visualID 0x23
    number of attributes: 5
      "XV_COLORKEY" (range 0 to 16777215)
        client settable attribute
        client gettable attribute (current value is 2110)
      "XV_BRIGHTNESS" (range -128 to 127)
        client settable attribute
        client gettable attribute (current value is 0)
      "XV_CONTRAST" (range 0 to 255)
        client settable attribute
        client gettable attribute (current value is 128)
      "XV_SATURATION" (range 0 to 255)
        client settable attribute
        client gettable attribute (current value is 128)
      "XV_HUE" (range -180 to 180)
```

```

        client settable attribute
        client gettable attribute (current value is 0)
maximum XvImage size: 1024 x 1024
Number of image formats: 7
  id: 0x32595559 (YUY2)
    guid: 59555932-0000-0010-8000-00aa00389b71
    bits per pixel: 16
    number of planes: 1
    type: YUV (packed)
  id: 0x32315659 (YV12)
    guid: 59563132-0000-0010-8000-00aa00389b71
    bits per pixel: 12
    number of planes: 3
    type: YUV (planar)
  id: 0x30323449 (I420)
    guid: 49343230-0000-0010-8000-00aa00389b71
    bits per pixel: 12
    number of planes: 3
    type: YUV (planar)
  id: 0x36315652 (RV16)
    guid: 52563135-0000-0000-0000-000000000000
    bits per pixel: 16
    number of planes: 1
    type: RGB (packed)
    depth: 0
    red, green, blue masks: 0x1f, 0x3e0, 0x7c00
  id: 0x35315652 (RV15)
    guid: 52563136-0000-0000-0000-000000000000
    bits per pixel: 16
    number of planes: 1
    type: RGB (packed)
    depth: 0
    red, green, blue masks: 0x1f, 0x7e0, 0xf800
  id: 0x31313259 (Y211)
    guid: 59323131-0000-0010-8000-00aa00389b71
    bits per pixel: 6
    number of planes: 3
    type: YUV (packed)
  id: 0x0
    guid: 00000000-0000-0000-0000-000000000000
    bits per pixel: 0
    number of planes: 0
    type: RGB (packed)
    depth: 1
    red, green, blue masks: 0x0, 0x0, 0x0

```

Also note that the formats listed (YUV2, YUV12, etc) are not present with every implementation of XVideo and their absence may hinder some players.

If the result looks like:

```

X-Video Extension version 2.2
screen #0
no adaptors present

```

Then XVideo is probably not supported for your card.

If XVideo is not supported for your card, this only means that it will be more difficult for your display to meet the computational demands of rendering video. Depending on your video card and processor, though, you might still be able to have a satisfying experience. You should probably read about ways of improving performance in the advanced reading Section 16.4.3.

16.4.1.2 Simple Directmedia Layer

The Simple Directmedia Layer, SDL, was intended to be a porting layer between Microsoft Windows, BeOS, and UNIX, allowing cross-platform applications to be developed which made efficient use of sound and graphics. The SDL layer provides a low-level abstraction to the hardware which can sometimes be more efficient than the X11 interface.

The SDL can be found at `devel/SDL`.

16.4.1.3 Direct Graphics Access

Direct Graphics Access is an X11 extension which allows a program to bypass the X server and directly alter the framebuffer. Because it relies on a low level memory mapping to effect this sharing, programs using it must be run as `root`.

The DGA extension can be tested and benchmarked by `dga(1)`. When `dga` is running, it changes the colors of the display whenever a key is pressed. To quit, use `q`.

16.4.2 Ports and Packages Dealing with Video

This section discusses the software available from the `pkgsrc` tree which can be used for video playback. Video playback is a very active area of software development, and the capabilities of various applications are bound to diverge somewhat from the descriptions given here.

Firstly, it is important to know that many of the video applications which run on DragonFly were developed as Linux applications. Many of these applications are still beta-quality. Some of the problems that you may encounter with video packages on DragonFly include:

1. An application cannot playback a file which another application produced.
2. An application cannot playback a file which the application itself produced.
3. The same application on two different machines, rebuilt on each machine for that machine, plays back the same file differently.
4. A seemingly trivial filter like rescaling of the image size results in very bad artifacts from a buggy rescaling routine.
5. An application frequently dumps core.
6. Documentation is not installed with the port and can be found either on the web or under the port's `work` directory.

Many of these applications may also exhibit "Linux-isms". That is, there may be issues resulting from the way some standard libraries are implemented in the Linux distributions, or some features of the Linux kernel which have been

assumed by the authors of the applications. These issues are not always noticed and worked around by the port maintainers, which can lead to problems like these:

1. The use of `/proc/cpuinfo` to detect processor characteristics.
2. A misuse of threads which causes a program to hang upon completion instead of truly terminating.
3. Software not yet in the `pkgsrc` tree which is commonly used in conjunction with the application.

So far, these application developers have been cooperative with port maintainers to minimize the work-arounds needed for port-ing.

16.4.2.1 MPlayer

MPlayer is a recently developed and rapidly developing video player. The goals of the **MPlayer** team are speed and flexibility on Linux and other Unices. The project was started when the team founder got fed up with bad playback performance on then available players. Some would say that the graphical interface has been sacrificed for a streamlined design. However, once you get used to the command line options and the key-stroke controls, it works very well.

16.4.2.1.1 Building MPlayer

MPlayer resides in `multimedia/mplayer`. **MPlayer** performs a variety of hardware checks during the build process, resulting in a binary which will not be portable from one system to another. Therefore, it is important to build it from ports and not to use a binary package. Additionally, a number of options can be specified in the `bmake` command line, as described in the `bmake show-options`.

```
# cd /usr/pkgsrc/multimedia/mplayer
# bmake show-options
```

Any of the following general options may be selected:

```
aalib      Enable aalib support.
arts       Use the aRts audio daemon.
dts        Enable DTS Coherent Acoustics support.
dv         Enable usage of the libdv library.
dvdread    Enable reproduction of DVDs.
esound     Enable support for the Enlightenment Sound Daemon.
ggi        Enable GGI support.
gif        Enable GIF support.
jpeg       Enable JPEG support.
mad        Enable usage of the mad library to play MP3 files.
mplayer-menu  Enable support for user-defined menus.
mplayer-real  Enable usage of Real codecs.
mplayer-runtime-cpudetection  Enable CPU detection at run time.
mplayer-win32  Enable usage of Win32 DLLs (codecs).
nas        Enable usage of the Network Audio System.
oss        Enable support for the Open Sound System audio library.
png        Enable PNG support.
sdl        Use SDL as display library.
theora     Use the theora codecs.
vorbis     Enable Ogg Vorbis support.
xvid       Enable usage of Xvid codecs.
```

At most one of the following faadgroup options may be selected:

```
faad       Enable AAC decoding support using faad2.
```



```
mplayer-internal-faad
```

These options are enabled by default:

```
aalib arts dv dvdread esound faad gif jpeg
mad mplayer-menu mplayer-real mplayer-runtime-cpudetection
mplayer-win32 nas oss png sdl theora vorbis
xvid
```

These options are currently enabled:

```
aalib arts dv dvdread esound faad gif jpeg
mad mplayer-menu mplayer-real mplayer-runtime-cpudetection
mplayer-win32 nas oss png sdl theora vorbis
xvid
```

You can select which build options to use by setting `PKG_DEFAULT_OPTIONS` or `PKG_OPTIONS.mplayer`.

If you have `x11/gtk` installed, then you might as well enable the GUI. Otherwise, it is not worth the effort. If it is illegal to play (possibly CSS encoded) DVD's with **MPlayer** on your country you must disable the DVD support option². Otherwise the defaults are ok:

```
# bmake
```

As of this writing, the **MPlayer** port will build its HTML documentation and one executable, `mplayer`. By default it builds an encoder, `mencoder`, which is a tool for re-encoding video. A modification to the `Makefile` can disable it.

The HTML documentation for **MPlayer** is very informative. If the reader finds the information on video hardware and interfaces in this chapter lacking, the **MPlayer** documentation is a very thorough supplement. You should definitely take the time to read the **MPlayer** documentation if you are looking for information about video support in UNIX.

16.4.2.1.2 Using MPlayer

The command options for `mplayer` are listed in the manual page. For even more detail there is HTML documentation. In this section, we will describe only a few common uses.

To play a file, such as `testfile.avi`, through one of the various video interfaces set the `-vo` option:

```
% mplayer -vo xv testfile.avi
% mplayer -vo sdl testfile.avi
% mplayer -vo x11 testfile.avi
# mplayer -vo dga testfile.avi
# mplayer -vo 'sdl:dga' testfile.avi
```

It is worth trying all of these options, as their relative performance depends on many factors and will vary significantly with hardware.

To play from a DVD, replace the `testfile.avi` with `-dvd N DEVICE` where `N` is the title number to play and `DEVICE` is the device node for the DVD-ROM. For example, to play title 3 from `/dev/dvd`:

```
# mplayer -vo dga -dvd 3 /dev/dvd
```

To stop, pause, advance and so on, consult the keybindings, which are output by running `mplayer -h` or read the manual page.

Additional important options for playback are: `-fs` `-zoom` which engages the fullscreen mode and `-framedrop` which helps performance.

In order for the `mplayer` command line to not become too large, the user can create a file `.mplayer/config` and set default options there:

```
vo=xv
fs=yes
zoom=yes
```

Finally, `mplayer` can be used to rip a DVD title into a `.vob` file. To dump out the second title from a DVD, type this:

```
# mplayer -dumpstream -dumpfile out.vob -dvd 2 /dev/dvd
```

The output file, `out.vob`, will be MPEG and can be manipulated by the other packages described in this section.

16.4.2.1.3 mencoder

If you installed `mencoder` when you build **MPlayer**, be forewarned that it is still an experimental component. Before using `mencoder` it is a good idea to familiarize yourself with the options from the HTML documentation. There is a manual page, but it is not very useful without the HTML documentation. There are innumerable ways to improve quality, lower bitrate, and change formats, and some of these tricks may make the difference between good or bad performance. Here are a couple of examples to get you going. First a simple copy:

```
% mencoder input.avi -oac copy -ovc copy -o output.avi
```

Improper combinations of command line options can yield output files that are unplayable even by `mplayer`. Thus, if you just want to rip to a file, stick to the `-dumpfile` in `mplayer`.

To convert `input.avi` to the MPEG4 codec with MPEG3 audio encoding (`audio/lame` is required):

```
% mencoder input.avi -oac mp3lame -lameopts br=192 \
  -ovc lavc -lavcopts vcodec=mpeg4:vhq -o output.avi
```

This has produced output playable by `mplayer` and `xine`.

`input.avi` can be replaced with `-dvd 1 /dev/dvd` and run as `root` to re-encode a DVD title directly. Since you are likely to be dissatisfied with your results the first time around, it is recommended you dump the title to a file and work on the file.

16.4.2.2 The xine Video Player

The **xine** video player is a project of wide scope aiming not only at being an all in one video solution, but also in producing a reusable base library and a modular executable which can be extended with plugins. It comes both as a package and as a port, `multimedia/xine-ui`.

The **xine** player is still very rough around the edges, but it is clearly off to a good start. In practice, **xine** requires either a fast CPU with a fast video card, or support for the XVideo extension. The GUI is usable, but a bit clumsy.

As of this writing, there is no input module shipped with **xine** which will play CSS encoded DVD's. There are third party builds which do have modules for this built in them, but none of these are in the pkgsrc tree.

Compared to **MPlayer**, **xine** does more for the user, but at the same time, takes some of the more fine-grained control away from the user. The **xine** video player performs best on XVideo interfaces.

By default, **xine** player will start up in a graphical user interface. The menus can then be used to open a specific file:

```
% xine
```

Alternatively, it may be invoked to play a file immediately without the GUI with the command:

```
% xine -g -p mymovie.avi
```

16.4.2.3 The transcode Utilities

The software **transcode** is not a player, but a suite of tools for re-encoding .avi and .mpg files. With **transcode**, one has the ability to merge video files, repair broken files, using command line tools with stdin/stdout stream interfaces.

Like **MPlayer**, **transcode** is very experimental software which must be build from pkgsrc at multimedia/transcode. You can see all the available options with the `bmake show-options` command. I recommend the default ones.

```
# bmake
```

Here are two examples of using **transcode** for video conversion which produce rescaled output. The first encodes the output to an openDIVX AVI file, while the second encodes to the much more portable MPEG format.

```
% transcode -i input.vob -x vob -V -Z 320x240 \
-y opendivx -N 0x55 -o output.avi

% transcode -i input.vob -x vob -V -Z 320x240 \
-y mpeg -N 0x55 -o output.tmp
% tcplex -o output.mpg -i output.tmp.mlv -p output.tmp.mpa -m 1
```

There is a manual page for **transcode**, but there is little documentation for the various `tc*` utilities (such as `tcplex`) which are also installed. However, the `-h` command line option can always be given to get curt usage instructions for a command.

In comparison, **transcode** runs significantly slower than **mencoder**, but it has a better chance of producing a more widely playable file. MPEGs created by **transcode** have been known to play on **Windows Media® Player** and Apple's **Quicktime®**, for example.

16.4.3 Further Reading

The various video software packages for DragonFly are developing rapidly. It is quite possible that in the near future many of the problems discussed here will have been resolved. In the mean time, those who want to get the very most

out of DragonFly's A/V capabilities will have to cobble together knowledge from several FAQs and tutorials and use a few different applications. This section exists to give the reader pointers to such additional information.

The MPlayer documentation (<http://www.mplayerhq.hu/DOCS/>) is very technically informative. These documents should probably be consulted by anyone wishing to obtain a high level of expertise with UNIX video. The **MPlayer** mailing list is hostile to anyone who has not bothered to read the documentation, so if you plan on making bug reports to them, RTFM.

The xine HOWTO (http://dvd.sourceforge.net/xine-howto/en_GB/html/howto.html) contains a chapter on performance improvement which is general to all players.

Finally, there are some other promising applications which the reader may try:

- Avifile (<http://avifile.sourceforge.net/>) which is not yet in the official pkgsrc tree, but you can find it on pkgsrc-wip (<http://pkgsrc-wip.sourceforge.net/>).
- Ogle (<http://www.dtek.chalmers.se/groups/dvd/>) which is also a port `multimedia/ogle`.
- Xtheater (<http://xtheater.sourceforge.net/>)
- `multimedia/dvdauthor`, an open source package for authoring DVD content.

16.5 Setting Up TV Cards

Original contribution by Josef El-Rayes. Enhanced and adapted by Marc Fonvieille.

16.5.1 Introduction

TV cards allow you to watch broadcast or cable TV on your computer. Most of them accept composite video via an RCA or S-video input and some of these cards come with a FM radio tuner.

DragonFly provides support for PCI-based TV cards using a Brooktree Bt848/849/878/879 or a Conexant CN-878/Fusion 878a Video Capture Chip with the `bktr(4)` driver. You must also ensure the board comes with a supported tuner, consult the `bktr(4)` manual page for a list of supported tuners.

16.5.2 Adding the Driver

To use your card, you will need to load the `bktr(4)` driver, this can be done by adding the following line to the `/boot/loader.conf` file like this:

```
bktr_load="YES"
```

Alternatively, you may statically compile the support for the TV card in your kernel, in that case add the following lines to your kernel configuration:

```
device bktr
device iicbus
device iicbb
device smbus
```

These additional device drivers are necessary because of the card components being interconnected via an I2C bus. Then build and install a new kernel.

Once the support was added to your system, you have to reboot your machine. During the boot process, your TV card should show up, like this:

```
bktr0: <BrookTree 848A> mem 0xd7000000-0xd7000fff irq 10 at device 10.0 on pci0
iicbb0: <I2C bit-banging driver> on bti2c0
iicbus0: <Philips I2C bus> on iicbb0 master-only
iicbus1: <Philips I2C bus> on iicbb0 master-only
smbus0: <System Management Bus> on bti2c0
bktr0: Pinnacle/Miro TV, Philips SECAM tuner.
```

Of course these messages can differ according to your hardware. However you should check if the tuner is correctly detected; it is still possible to override some of the detected parameters with `sysctl(8)` MIBs and kernel configuration file options. For example, if you want to force the tuner to a Philips SECAM tuner, you should add the following line to your kernel configuration file:

```
options OVERRIDE_TUNER=6
```

or you can directly use `sysctl(8)`:

```
# sysctl hw.bt848.tuner=6
```

See the `bktr(4)` manual page and the `/usr/src/sys/i386/conf/LINT` file for more details on the available options.

16.5.3 Useful Applications

To use your TV card you need to install one of the following applications:

- `multimedia/fxrtv` provides TV-in-a-window and image/audio/video capture capabilities.
- `multimedia/xawrtv` is also a TV application, with the same features as `fxrtv`.
- `multimedia/alevt` decodes and displays Videotext/Teletext.
- `audio/xmradio`, an application to use the FM radio tuner coming with some TV cards.

More applications are available in the `pkgsrc` tree.

16.5.4 Troubleshooting

If you encounter any problem with your TV card, you should check at first if the video capture chip and the tuner are really supported by the `bktr(4)` driver and if you used the right configuration options. For more support and various questions about your TV card you may want to contact and use the archives of the DragonFly User related mailing list (<http://leaf.dragonflybsd.org/mailarchive/>).

Notes

1. A popular familiar graphics card with generally very good **XFree86** performance, nVidia, has yet to release the specifications on their XVideo support to the **XFree86** team. It may be some time before **XFree86** fully support XVideo for these cards.
2. Unauthorized DVD playback is a serious criminal act in some countries. Check local laws before enabling this option.

Chapter 17 Serial Communications

17.1 Synopsis

UNIX has always had support for serial communications. In fact, the very first UNIX machines relied on serial lines for user input and output. Things have changed a lot from the days when the average “terminal” consisted of a 10-character-per-second serial printer and a keyboard. This chapter will cover some of the ways in which DragonFly uses serial communications.

After reading this chapter, you will know:

- How to connect terminals to your DragonFly system.
- How to use a modem to dial out to remote hosts.
- How to allow remote users to login to your system with a modem.
- How to boot your system from a serial console.

Before reading this chapter, you should:

- Know how to configure and install a new kernel (Chapter 9).
- Understand UNIX permissions and processes (Chapter 3).
- Have access to the technical manual for the serial hardware (modem or multi-port card) that you would like to use with DragonFly.

17.2 Introduction

17.2.1 Terminology

bps

Bits per Second — the rate at which data is transmitted

DTE

Data Terminal Equipment — for example, your computer

DCE

Data Communications Equipment — your modem

RS-232

EIA standard for hardware serial communications

When talking about communications data rates, this section does not use the term “baud”. Baud refers to the number of electrical state transitions that may be made in a period of time, while “bps” (bits per second) is the *correct* term to use (at least it does not seem to bother the curmudgeons quite as much).

17.2.2 Cables and Ports

To connect a modem or terminal to your DragonFly system, you will need a serial port on your computer and the proper cable to connect to your serial device. If you are already familiar with your hardware and the cable it requires, you can safely skip this section.

17.2.2.1 Cables

There are several different kinds of serial cables. The two most common types for our purposes are null-modem cables and standard (“straight”) RS-232 cables. The documentation for your hardware should describe the type of cable required.

17.2.2.1.1 Null-modem Cables

A null-modem cable passes some signals, such as “signal ground”, straight through, but switches other signals. For example, the “send data” pin on one end goes to the “receive data” pin on the other end.

If you like making your own cables, you can construct a null-modem cable for use with terminals. This table shows the RS-232C signal names and the pin numbers on a DB-25 connector.

Signal	Pin #		Pin #	Signal
SG	7	connects to	7	SG
TxD	2	connects to	3	RxD
RxD	3	connects to	2	TxD
RTS	4	connects to	5	CTS
CTS	5	connects to	4	RTS
DTR	20	connects to	6	DSR
DCD	8		6	DSR
DSR	6	connects to	20	DTR

Note: Connect “Data Set Ready” (DSR) and “Data Carrier Detect” (DCD) internally in the connector hood, and then to “Data Terminal Ready” (DTR) in the remote hood.

17.2.2.1.2 Standard RS-232C Cables

A standard serial cable passes all the RS-232C signals straight-through. That is, the “send data” pin on one end of the cable goes to the “send data” pin on the other end. This is the type of cable to use to connect a modem to your DragonFly system, and is also appropriate for some terminals.

17.2.2.2 Ports

Serial ports are the devices through which data is transferred between the DragonFly host computer and the terminal. This section describes the kinds of ports that exist and how they are addressed in DragonFly.

17.2.2.2.1 Kinds of Ports

Several kinds of serial ports exist. Before you purchase or construct a cable, you need to make sure it will fit the ports on your terminal and on the DragonFly system.

Most terminals will have DB25 ports. Personal computers, including PCs running DragonFly, will have DB25 or DB9 ports. If you have a multiport serial card for your PC, you may have RJ-12 or RJ-45 ports.

See the documentation that accompanied the hardware for specifications on the kind of port in use. A visual inspection of the port often works too.

17.2.2.2.2 Port Names

In DragonFly, you access each serial port through an entry in the `/dev` directory. There are two different kinds of entries:

- Call-in ports are named `/dev/ttydN` where N is the port number, starting from zero. Generally, you use the call-in port for terminals. Call-in ports require that the serial line assert the data carrier detect (DCD) signal to work correctly.
- Call-out ports are named `/dev/cuaaN`. You usually do not use the call-out port for terminals, just for modems. You may use the call-out port if the serial cable or the terminal does not support the carrier detect signal.

If you have connected a terminal to the first serial port (COM1 in MS-DOS), then you will use `/dev/ttyd0` to refer to the terminal. If the terminal is on the second serial port (also known as COM2), use `/dev/ttyd1`, and so forth.

17.2.3 Kernel Configuration

DragonFly supports four serial ports by default. In the MS-DOS world, these are known as COM1, COM2, COM3, and COM4. DragonFly currently supports “dumb” multiport serial interface cards, such as the BocaBoard 1008 and 2016, as well as more intelligent multi-port cards such as those made by Digiboard and Stallion Technologies. However, the default kernel only looks for the standard COM ports.

To see if your kernel recognizes any of your serial ports, watch for messages while the kernel is booting, or use the `/sbin/dmesg` command to replay the kernel’s boot messages. In particular, look for messages that start with the characters `sio`.

Tip: To view just the messages that have the word `sio`, use the command:

```
# /sbin/dmesg | grep 'sio'
```

For example, on a system with four serial ports, these are the serial-port specific kernel boot messages:

```
sio0 at 0x3f8-0x3ff irq 4 on isa
sio0: type 16550A
sio1 at 0x2f8-0x2ff irq 3 on isa
sio1: type 16550A
sio2 at 0x3e8-0x3ef irq 5 on isa
```

```
sio2: type 16550A
sio3 at 0x2e8-0x2ef irq 9 on isa
sio3: type 16550A
```

If your kernel does not recognize all of your serial ports, you will probably need to configure a custom DragonFly kernel for your system. For detailed information on configuring your kernel, please see Chapter 9.

The relevant device lines for your kernel configuration file would look like this:

```
device sio0 at isa? port IO_COM1 irq 4
device sio1 at isa? port IO_COM2 irq 3
device sio2 at isa? port IO_COM3 irq 5
device sio3 at isa? port IO_COM4 irq 9
```

You can comment-out or completely remove lines for devices you do not have. Please refer to the `sio(4)` manual page for more information on serial ports and multiport boards configuration.

Note: `port IO_COM1` is a substitution for `port 0x3f8`, `IO_COM2` is `0x2f8`, `IO_COM3` is `0x3e8`, and `IO_COM4` is `0x2e8`, which are fairly common port addresses for their respective serial ports; interrupts 4, 3, 5, and 9 are fairly common interrupt request lines. Also note that regular serial ports *cannot* share interrupts on ISA-bus PCs (multiport boards have on-board electronics that allow all the 16550A's on the board to share one or two interrupt request lines).

17.2.4 Device Special Files

Most devices in the kernel are accessed through “device special files”, which are located in the `/dev` directory. The `sio` devices are accessed through the `/dev/ttydN` (dial-in) and `/dev/cuaaN` (call-out) devices. DragonFly also provides initialization devices (`/dev/ttyidN` and `/dev/cuaiaN`) and locking devices (`/dev/ttyldN` and `/dev/cualaN`). The initialization devices are used to initialize communications port parameters each time a port is opened, such as `crtsets` for modems which use RTS/CTS signaling for flow control. The locking devices are used to lock flags on ports to prevent users or programs changing certain parameters; see the manual pages `termios(4)`, `sio(4)`, and `stty(1)` for information on the terminal settings, locking and initializing devices, and setting terminal options, respectively.

17.2.4.1 Making Device Special Files

A shell script called `MAKEDEV` in the `/dev` directory manages the device special files. To use `MAKEDEV` to make dial-up device special files for COM1 (port 0), `cd` to `/dev` and issue the command `MAKEDEV ttyd0`. Likewise, to make dial-up device special files for COM2 (port 1), use `MAKEDEV ttyd1`.

`MAKEDEV` not only creates the `/dev/ttydN` device special files, but also the `/dev/cuaaN`, `/dev/cuaiaN`, `/dev/cualaN`, `/dev/ttyldN`, and `/dev/ttyidN` nodes.

After making new device special files, be sure to check the permissions on the files (especially the `/dev/cua*` files) to make sure that only users who should have access to those device special files can read and write on them — you probably do not want to allow your average user to use your modems to dial-out. The default permissions on the `/dev/cua*` files should be sufficient:

```
crw-rw----  1 uucp      dialer    28, 129 Feb 15 14:38 /dev/cuaa1
crw-rw----  1 uucp      dialer    28, 161 Feb 15 14:38 /dev/cuaia1
```

```
crw-rw----  1 uucp    dialer   28, 193 Feb 15 14:38 /dev/cua1a1
```

These permissions allow the user `uucp` and users in the group `dialer` to use the call-out devices.

17.2.5 Serial Port Configuration

The `ttymN` (or `cuaaN`) device is the regular device you will want to open for your applications. When a process opens the device, it will have a default set of terminal I/O settings. You can see these settings with the command

```
# stty -a -f /dev/ttyd1
```

When you change the settings to this device, the settings are in effect until the device is closed. When it is reopened, it goes back to the default set. To make changes to the default set, you can open and adjust the settings of the “initial state” device. For example, to turn on `CLOCAL` mode, 8 bit communication, and `XON/XOFF` flow control by default for `ttyd5`, type:

```
# stty -f /dev/ttyid5 clocal cs8 ixon ixoff
```

System-wide initialization of the serial devices is controlled in `/etc/rc.serial`. This file affects the default settings of serial devices.

To prevent certain settings from being changed by an application, make adjustments to the “lock state” device. For example, to lock the speed of `ttyd5` to 57600 bps, type:

```
# stty -f /dev/tty1d5 57600
```

Now, an application that opens `ttyd5` and tries to change the speed of the port will be stuck with 57600 bps.

Naturally, you should make the initial state and lock state devices writable only by the `root` account.

17.3 Terminals

Contributed by Sean Kelly.

Terminals provide a convenient and low-cost way to access your DragonFly system when you are not at the computer’s console or on a connected network. This section describes how to use terminals with DragonFly.

17.3.1 Uses and Types of Terminals

The original UNIX systems did not have consoles. Instead, people logged in and ran programs through terminals that were connected to the computer’s serial ports. It is quite similar to using a modem and terminal software to dial into a remote system to do text-only work.

Today’s PCs have consoles capable of high quality graphics, but the ability to establish a login session on a serial port still exists in nearly every UNIX style operating system today; DragonFly is no exception. By using a terminal attached to an unused serial port, you can log in and run any text program that you would normally run on the console or in an `xterm` window in the X Window System.

For the business user, you can attach many terminals to a DragonFly system and place them on your employees' desktops. For a home user, a spare computer such as an older IBM PC or a Macintosh can be a terminal wired into a more powerful computer running DragonFly. You can turn what might otherwise be a single-user computer into a powerful multiple user system.

For DragonFly, there are three kinds of terminals:

- Dumb terminals
- PCs acting as terminals
- X terminals

The remaining subsections describe each kind.

17.3.1.1 Dumb Terminals

Dumb terminals are specialized pieces of hardware that let you connect to computers over serial lines. They are called “dumb” because they have only enough computational power to display, send, and receive text. You cannot run any programs on them. It is the computer to which you connect them that has all the power to run text editors, compilers, email, games, and so forth.

There are hundreds of kinds of dumb terminals made by many manufacturers, including Digital Equipment Corporation's VT-100 and Wyse's WY-75. Just about any kind will work with DragonFly. Some high-end terminals can even display graphics, but only certain software packages can take advantage of these advanced features.

Dumb terminals are popular in work environments where workers do not need access to graphical applications such as those provided by the X Window System.

17.3.1.2 PCs Acting as Terminals

If a dumb terminal has just enough ability to display, send, and receive text, then certainly any spare personal computer can be a dumb terminal. All you need is the proper cable and some *terminal emulation* software to run on the computer.

Such a configuration is popular in homes. For example, if your spouse is busy working on your DragonFly system's console, you can do some text-only work at the same time from a less powerful personal computer hooked up as a terminal to the DragonFly system.

17.3.1.3 X Terminals

X terminals are the most sophisticated kind of terminal available. Instead of connecting to a serial port, they usually connect to a network like Ethernet. Instead of being relegated to text-only applications, they can display any X application.

We introduce X terminals just for the sake of completeness. However, this chapter does *not* cover setup, configuration, or use of X terminals.

17.3.2 Configuration

This section describes what you need to configure on your DragonFly system to enable a login session on a terminal. It assumes you have already configured your kernel to support the serial port to which the terminal is connected—and that you have connected it.

Recall from Chapter 7 that the `init` process is responsible for all process control and initialization at system startup. One of the tasks performed by `init` is to read the `/etc/ttys` file and start a `getty` process on the available terminals. The `getty` process is responsible for reading a login name and starting the `login` program.

Thus, to configure terminals for your DragonFly system the following steps should be taken as `root`:

1. Add a line to `/etc/ttys` for the entry in the `/dev` directory for the serial port if it is not already there.
2. Specify that `/usr/libexec/getty` be run on the port, and specify the appropriate `getty` type from the `/etc/gettytab` file.
3. Specify the default terminal type.
4. Set the port to “on.”
5. Specify whether the port should be “secure.”
6. Force `init` to reread the `/etc/ttys` file.

As an optional step, you may wish to create a custom `getty` type for use in step 2 by making an entry in `/etc/gettytab`. This chapter does not explain how to do so; you are encouraged to see the `gettytab(5)` and the `getty(8)` manual pages for more information.

17.3.2.1 Adding an Entry to `/etc/ttys`

The `/etc/ttys` file lists all of the ports on your DragonFly system where you want to allow logins. For example, the first virtual console `ttysv0` has an entry in this file. You can log in on the console using this entry. This file also contains entries for the other virtual consoles, serial ports, and pseudo-ttys. For a hardwired terminal, just list the serial port's `/dev` entry without the `/dev` part (for example, `/dev/ttysv0` would be listed as `ttysv0`).

A default DragonFly install includes an `/etc/ttys` file with support for the first four serial ports: `ttysd0` through `ttysd3`. If you are attaching a terminal to one of those ports, you do not need to add another entry.

Example 17-1. Adding Terminal Entries to `/etc/ttys`

Suppose we would like to connect two terminals to the system: a Wyse-50 and an old 286 IBM PC running **Procomm** terminal software emulating a VT-100 terminal. We connect the Wyse to the second serial port and the 286 to the sixth serial port (a port on a multiport serial card). The corresponding entries in the `/etc/ttys` file would look like this:

```
ttysd1 ① "/usr/libexec/getty std.38400"② wy50③ on④ insecure⑤
ttysd5  "/usr/libexec/getty std.19200" vt100 on insecure
```

- ① The first field normally specifies the name of the terminal special file as it is found in `/dev`.
- ② The second field is the command to execute for this line, which is usually `getty(8)`. `getty` initializes and opens the line, sets the speed, prompts for a user name and then executes the `login(1)` program.

The `getty` program accepts one (optional) parameter on its command line, the `getty` type. A `getty` type configures characteristics on the terminal line, like bps rate and parity. The `getty` program reads these characteristics from the file `/etc/gettytab`.

The file `/etc/gettytab` contains lots of entries for terminal lines both old and new. In almost all cases, the entries that start with the text `std` will work for hardwired terminals. These entries ignore parity. There is a `std` entry for each bps rate from 110 to 115200. Of course, you can add your own entries to this file. The `gettytab(5)` manual page provides more information.

When setting the `getty` type in the `/etc/ttys` file, make sure that the communications settings on the terminal match.

For our example, the Wyse-50 uses no parity and connects at 38400 bps. The 286 PC uses no parity and connects at 19200 bps.

- ③ The third field is the type of terminal usually connected to that tty line. For dial-up ports, `unknown` or `dialup` is typically used in this field since users may dial up with practically any type of terminal or software. For hardwired terminals, the terminal type does not change, so you can put a real terminal type from the `termcap(5)` database file in this field.

For our example, the Wyse-50 uses the real terminal type while the 286 PC running **Procomm** will be set to emulate at VT-100.

- ④ The fourth field specifies if the port should be enabled. Putting `on` here will have the `init` process start the program in the second field, `getty`. If you put `off` in this field, there will be no `getty`, and hence no logins on the port.
- ⑤ The final field is used to specify whether the port is secure. Marking a port as secure means that you trust it enough to allow the `root` account (or any account with a user ID of 0) to login from that port. Insecure ports do not allow `root` logins. On an insecure port, users must login from unprivileged accounts and then use `su(1)` or a similar mechanism to gain superuser privileges.

It is highly recommended that you use “insecure” even for terminals that are behind locked doors. It is quite easy to login and use `su` if you need superuser privileges.

17.3.2.2 Force `init` to Reread `/etc/ttys`

After making the necessary changes to the `/etc/ttys` file you should send a `SIGHUP` (hangup) signal to the `init` process to force it to re-read its configuration file. For example:

```
# kill -HUP 1
```

Note: `init` is always the first process run on a system, therefore it will always have PID 1.

If everything is set up correctly, all cables are in place, and the terminals are powered up, then a `getty` process should be running on each terminal and you should see login prompts on your terminals at this point.

17.3.3 Troubleshooting Your Connection

Even with the most meticulous attention to detail, something could still go wrong while setting up a terminal. Here is a list of symptoms and some suggested fixes.

17.3.3.1 No Login Prompt Appears

Make sure the terminal is plugged in and powered up. If it is a personal computer acting as a terminal, make sure it is running terminal emulation software on the correct serial port.

Make sure the cable is connected firmly to both the terminal and the DragonFly computer. Make sure it is the right kind of cable.

Make sure the terminal and DragonFly agree on the bps rate and parity settings. If you have a video display terminal, make sure the contrast and brightness controls are turned up. If it is a printing terminal, make sure paper and ink are in good supply.

Make sure that a `getty` process is running and serving the terminal. For example, to get a list of running `getty` processes with `ps`, type:

```
# ps -axww|grep getty
```

You should see an entry for the terminal. For example, the following display shows that a `getty` is running on the second serial port `ttyd1` and is using the `std.38400` entry in `/etc/gettytab`:

```
22189  d1  Is+    0:00.03 /usr/libexec/getty std.38400 ttyd1
```

If no `getty` process is running, make sure you have enabled the port in `/etc/ttys`. Also remember to run `kill -HUP 1` after modifying the `ttys` file.

If the `getty` process is running but the terminal still does not display a login prompt, or if it displays a prompt but will not allow you to type, your terminal or cable may not support hardware handshaking. Try changing the entry in `/etc/ttys` from `std.38400` to `3wire.38400` remember to run `kill -HUP 1` after modifying `/etc/ttys`). The `3wire` entry is similar to `std`, but ignores hardware handshaking. You may need to reduce the baud rate or enable software flow control when using `3wire` to prevent buffer overflows.

17.3.3.2 If Garbage Appears Instead of a Login Prompt

Make sure the terminal and DragonFly agree on the bps rate and parity settings. Check the `getty` processes to make sure the correct `getty` type is in use. If not, edit `/etc/ttys` and run `kill -HUP 1`.

17.3.3.3 Characters Appear Doubled; the Password Appears When Typed

Switch the terminal (or the terminal emulation software) from “half duplex” or “local echo” to “full duplex.”

17.4 Dial-in Service

Contributed by Guy Helmer. Additions by Sean Kelly.

Configuring your DragonFly system for dial-in service is very similar to connecting terminals except that you are dealing with modems instead of terminals.

17.4.1 External vs. Internal Modems

External modems seem to be more convenient for dial-up, because external modems often can be semi-permanently configured via parameters stored in non-volatile RAM and they usually provide lighted indicators that display the state of important RS-232 signals. Blinking lights impress visitors, but lights are also very useful to see whether a modem is operating properly.

Internal modems usually lack non-volatile RAM, so their configuration may be limited only to setting DIP switches. If your internal modem has any signal indicator lights, it is probably difficult to view the lights when the system's cover is in place.

17.4.1.1 Modems and Cables

If you are using an external modem, then you will of course need the proper cable. A standard RS-232C serial cable should suffice as long as all of the normal signals are wired:

- Transmitted Data (SD)
- Received Data (RD)
- Request to Send (RTS)
- Clear to Send (CTS)
- Data Set Ready (DSR)
- Data Terminal Ready (DTR)
- Carrier Detect (CD)
- Signal Ground (SG)

DragonFly needs the RTS and CTS signals for flow-control at speeds above 2400 bps, the CD signal to detect when a call has been answered or the line has been hung up, and the DTR signal to reset the modem after a session is complete. Some cables are wired without all of the needed signals, so if you have problems, such as a login session not going away when the line hangs up, you may have a problem with your cable.

Like other UNIX like operating systems, DragonFly uses the hardware signals to find out when a call has been answered or a line has been hung up and to hangup and reset the modem after a call. DragonFly avoids sending commands to the modem or watching for status reports from the modem. If you are familiar with connecting modems to PC-based bulletin board systems, this may seem awkward.

17.4.2 Serial Interface Considerations

DragonFly supports NS8250-, NS16450-, NS16550-, and NS16550A-based EIA RS-232C (CCITT V.24) communications interfaces. The 8250 and 16450 devices have single-character buffers. The 16550 device provides a

16-character buffer, which allows for better system performance. (Bugs in plain 16550's prevent the use of the 16-character buffer, so use 16550A's if possible). Because single-character-buffer devices require more work by the operating system than the 16-character-buffer devices, 16550A-based serial interface cards are much preferred. If the system has many active serial ports or will have a heavy load, 16550A-based cards are better for low-error-rate communications.

17.4.3 Quick Overview

As with terminals, `init` spawns a `getty` process for each configured serial port for dial-in connections. For example, if a modem is attached to `/dev/ttyd0`, the command `ps ax` might show this:

```
4850 ?? I      0:00.09 /usr/libexec/getty V19200 ttyd0
```

When a user dials the modem's line and the modems connect, the CD (Carrier Detect) line is reported by the modem. The kernel notices that carrier has been detected and completes `getty`'s open of the port. `getty` sends a `login:` prompt at the specified initial line speed. `getty` watches to see if legitimate characters are received, and, in a typical configuration, if it finds junk (probably due to the modem's connection speed being different than `getty`'s speed), `getty` tries adjusting the line speeds until it receives reasonable characters.

After the user enters his/her login name, `getty` executes `/usr/bin/login`, which completes the login by asking for the user's password and then starting the user's shell.

17.4.4 Configuration Files

There are three system configuration files in the `/etc` directory that you will probably need to edit to allow dial-up access to your DragonFly system. The first, `/etc/gettytab`, contains configuration information for the `/usr/libexec/getty` daemon. Second, `/etc/ttys` holds information that tells `/sbin/init` what `tty` devices should have `getty` processes running on them. Lastly, you can place port initialization commands in the `/etc/rc.serial` script.

There are two schools of thought regarding dial-up modems on UNIX. One group likes to configure their modems and systems so that no matter at what speed a remote user dials in, the local computer-to-modem RS-232 interface runs at a locked speed. The benefit of this configuration is that the remote user always sees a system login prompt immediately. The downside is that the system does not know what a user's true data rate is, so full-screen programs like Emacs will not adjust their screen-painting methods to make their response better for slower connections.

The other school configures their modems' RS-232 interface to vary its speed based on the remote user's connection speed. For example, V.32bis (14.4 Kbps) connections to the modem might make the modem run its RS-232 interface at 19.2 Kbps, while 2400 bps connections make the modem's RS-232 interface run at 2400 bps. Because `getty` does not understand any particular modem's connection speed reporting, `getty` gives a `login:` message at an initial speed and watches the characters that come back in response. If the user sees junk, it is assumed that they know they should press the Enter key until they see a recognizable prompt. If the data rates do not match, `getty` sees anything the user types as "junk", tries going to the next speed and gives the `login:` prompt again. This procedure can continue ad nauseam, but normally only takes a keystroke or two before the user sees a good prompt. Obviously, this login sequence does not look as clean as the former "locked-speed" method, but a user on a low-speed connection should receive better interactive response from full-screen programs.

This section will try to give balanced configuration information, but is biased towards having the modem's data rate follow the connection rate.

17.4.4.1 /etc/gettytab

`/etc/gettytab` is a termcap(5)-style file of configuration information for `getty(8)`. Please see the `gettytab(5)` manual page for complete information on the format of the file and the list of capabilities.

17.4.4.1.1 Locked-speed Config

If you are locking your modem's data communications rate at a particular speed, you probably will not need to make any changes to `/etc/gettytab`.

17.4.4.1.2 Matching-speed Config

You will need to set up an entry in `/etc/gettytab` to give `getty` information about the speeds you wish to use for your modem. If you have a 2400 bps modem, you can probably use the existing `D2400` entry.

```
#
# Fast dialup terminals, 2400/1200/300 rotary (can start either way)
#
D2400|d2400|Fast-Dial-2400:\
      :nx=D1200:tc=2400-baud:
3|D1200|Fast-Dial-1200:\
      :nx=D300:tc=1200-baud:
5|D300|Fast-Dial-300:\
      :nx=D2400:tc=300-baud:
```

If you have a higher speed modem, you will probably need to add an entry in `/etc/gettytab`; here is an entry you could use for a 14.4 Kbps modem with a top interface speed of 19.2 Kbps:

```
#
# Additions for a V.32bis Modem
#
um|V300|High Speed Modem at 300,8-bit:\
      :nx=V19200:tc=std.300:
un|V1200|High Speed Modem at 1200,8-bit:\
      :nx=V300:tc=std.1200:
uo|V2400|High Speed Modem at 2400,8-bit:\
      :nx=V1200:tc=std.2400:
up|V9600|High Speed Modem at 9600,8-bit:\
      :nx=V2400:tc=std.9600:
uq|V19200|High Speed Modem at 19200,8-bit:\
      :nx=V9600:tc=std.19200:
```

This will result in 8-bit, no parity connections.

The example above starts the communications rate at 19.2 Kbps (for a V.32bis connection), then cycles through 9600 bps (for V.32), 2400 bps, 1200 bps, 300 bps, and back to 19.2 Kbps. Communications rate cycling is implemented with the `nx=` ("next table") capability. Each of the lines uses a `tc=` ("table continuation") entry to pick up the rest of the "standard" settings for a particular data rate.

If you have a 28.8 Kbps modem and/or you want to take advantage of compression on a 14.4 Kbps modem, you need to use a higher communications rate than 19.2 Kbps. Here is an example of a `gettytab` entry starting a 57.6 Kbps:

```
#
```

```
# Additions for a V.32bis or V.34 Modem
# Starting at 57.6 Kbps
#
vm|VH300|Very High Speed Modem at 300,8-bit:\
    :nx=VH57600:tc=std.300:
vn|VH1200|Very High Speed Modem at 1200,8-bit:\
    :nx=VH300:tc=std.1200:
vo|VH2400|Very High Speed Modem at 2400,8-bit:\
    :nx=VH1200:tc=std.2400:
vp|VH9600|Very High Speed Modem at 9600,8-bit:\
    :nx=VH2400:tc=std.9600:
vq|VH57600|Very High Speed Modem at 57600,8-bit:\
    :nx=VH9600:tc=std.57600:
```

If you have a slow CPU or a heavily loaded system and do not have 16550A-based serial ports, you may receive `sio` “silo” errors at 57.6 Kbps.

17.4.4.2 `/etc/ttys`

Configuration of the `/etc/ttys` file was covered in Example 17-1. Configuration for modems is similar but we must pass a different argument to `getty` and specify a different terminal type. The general format for both locked-speed and matching-speed configurations is:

```
ttyd0  "/usr/libexec/getty xxx"  dialup on
```

The first item in the above line is the device special file for this entry — `ttyd0` means `/dev/ttyd0` is the file that this `getty` will be watching. The second item, `"/usr/libexec/getty xxx"` (`xxx` will be replaced by the initial `gettytab` capability) is the process `init` will run on the device. The third item, `dialup`, is the default terminal type. The fourth parameter, `on`, indicates to `init` that the line is operational. There can be a fifth parameter, `secure`, but it should only be used for terminals which are physically secure (such as the system console).

The default terminal type (`dialup` in the example above) may depend on local preferences. `dialup` is the traditional default terminal type on dial-up lines so that users may customize their login scripts to notice when the terminal is `dialup` and automatically adjust their terminal type. However, the author finds it easier at his site to specify `vt102` as the default terminal type, since the users just use VT102 emulation on their remote systems.

After you have made changes to `/etc/ttys`, you may send the `init` process a HUP signal to re-read the file. You can use the command

```
# kill -HUP 1
```

to send the signal. If this is your first time setting up the system, you may want to wait until your modem(s) are properly configured and connected before signaling `init`.

17.4.4.2.1 Locked-speed Config

For a locked-speed configuration, your `ttys` entry needs to have a fixed-speed entry provided to `getty`. For a modem whose port speed is locked at 19.2 Kbps, the `ttys` entry might look like this:

```
ttyd0  "/usr/libexec/getty std.19200"  dialup on
```

If your modem is locked at a different data rate, substitute the appropriate value for `std.speed` instead of `std.19200`. Make sure that you use a valid type listed in `/etc/gettytab`.

17.4.4.2 Matching-speed Config

In a matching-speed configuration, your `ttys` entry needs to reference the appropriate beginning “auto-baud” (sic) entry in `/etc/gettytab`. For example, if you added the above suggested entry for a matching-speed modem that starts at 19.2 Kbps (the `gettytab` entry containing the `V19200` starting point), your `ttys` entry might look like this:

```
ttyd0  "/usr/libexec/getty V19200"  dialup on
```

17.4.4.3 /etc/rc.serial

High-speed modems, like V.32, V.32bis, and V.34 modems, need to use hardware (RTS/CTS) flow control. You can add `stty` commands to `/etc/rc.serial` to set the hardware flow control flag in the DragonFly kernel for the modem ports.

For example to set the `termios` flag `crtsets` on serial port #1's (COM2) dial-in and dial-out initialization devices, the following lines could be added to `/etc/rc.serial`:

```
# Serial port initial configuration
stty -f /dev/ttyid1 crtsets
stty -f /dev/cuaial crtsets
```

17.4.5 Modem Settings

If you have a modem whose parameters may be permanently set in non-volatile RAM, you will need to use a terminal program (such as `Telx` under MS-DOS or `tip` under DragonFly) to set the parameters. Connect to the modem using the same communications speed as the initial speed `getty` will use and configure the modem's non-volatile RAM to match these requirements:

- CD asserted when connected
- DTR asserted for operation; dropping DTR hangs up line and resets modem
- CTS transmitted data flow control
- Disable XON/XOFF flow control
- RTS received data flow control
- Quiet mode (no result codes)
- No command echo

Please read the documentation for your modem to find out what commands and/or DIP switch settings you need to give it.

For example, to set the above parameters on a U.S. Robotics® Sportster® 14,400 external modem, one could give these commands to the modem:

```
ATZ
AT&C1&D2&H1&I0&R2&W
```

You might also want to take this opportunity to adjust other settings in the modem, such as whether it will use V.42bis and/or MNP5 compression.

The U.S. Robotics Sportster 14,400 external modem also has some DIP switches that need to be set; for other modems, perhaps you can use these settings as an example:

- Switch 1: UP — DTR Normal
- Switch 2: N/A (Verbal Result Codes/Numeric Result Codes)
- Switch 3: UP — Suppress Result Codes
- Switch 4: DOWN — No echo, offline commands
- Switch 5: UP — Auto Answer
- Switch 6: UP — Carrier Detect Normal
- Switch 7: UP — Load NVRAM Defaults
- Switch 8: N/A (Smart Mode/Dumb Mode)

Result codes should be disabled/suppressed for dial-up modems to avoid problems that can occur if `getty` mistakenly gives a `login:` prompt to a modem that is in command mode and the modem echoes the command or returns a result code. This sequence can result in an extended, silly conversation between `getty` and the modem.

17.4.5.1 Locked-speed Config

For a locked-speed configuration, you will need to configure the modem to maintain a constant modem-to-computer data rate independent of the communications rate. On a U.S. Robotics Sportster 14,400 external modem, these commands will lock the modem-to-computer data rate at the speed used to issue the commands:

```
ATZ
AT&B1&W
```

17.4.5.2 Matching-speed Config

For a variable-speed configuration, you will need to configure your modem to adjust its serial port data rate to match the incoming call rate. On a U.S. Robotics Sportster 14,400 external modem, these commands will lock the modem's error-corrected data rate to the speed used to issue the commands, but allow the serial port rate to vary for non-error-corrected connections:

```
ATZ
AT&B2&W
```

17.4.5.3 Checking the Modem's Configuration

Most high-speed modems provide commands to view the modem's current operating parameters in a somewhat human-readable fashion. On the U.S. Robotics Sportster 14,400 external modems, the command `ATI5` displays the

settings that are stored in the non-volatile RAM. To see the true operating parameters of the modem (as influenced by the modem's DIP switch settings), use the commands `ATZ` and then `ATI4`.

If you have a different brand of modem, check your modem's manual to see how to double-check your modem's configuration parameters.

17.4.6 Troubleshooting

Here are a few steps you can follow to check out the dial-up modem on your system.

17.4.6.1 Checking Out the DragonFly System

Hook up your modem to your DragonFly system, boot the system, and, if your modem has status indication lights, watch to see whether the modem's DTR indicator lights when the `login:` prompt appears on the system's console — if it lights up, that should mean that DragonFly has started a `getty` process on the appropriate communications port and is waiting for the modem to accept a call.

If the DTR indicator does not light, login to the DragonFly system through the console and issue a `ps ax` to see if DragonFly is trying to run a `getty` process on the correct port. You should see lines like these among the processes displayed:

```
114 ?? I      0:00.10 /usr/libexec/getty V19200 ttyd0
115 ?? I      0:00.10 /usr/libexec/getty V19200 ttyd1
```

If you see something different, like this:

```
114 d0 I      0:00.10 /usr/libexec/getty V19200 ttyd0
```

and the modem has not accepted a call yet, this means that `getty` has completed its `open` on the communications port. This could indicate a problem with the cabling or a mis-configured modem, because `getty` should not be able to open the communications port until `CD` (carrier detect) has been asserted by the modem.

If you do not see any `getty` processes waiting to open the desired `ttydN` port, double-check your entries in `/etc/ttys` to see if there are any mistakes there. Also, check the log file `/var/log/messages` to see if there are any log messages from `init` or `getty` regarding any problems. If there are any messages, triple-check the configuration files `/etc/ttys` and `/etc/gettytab`, as well as the appropriate device special files `/dev/ttydN`, for any mistakes, missing entries, or missing device special files.

17.4.6.2 Try Dialing In

Try dialing into the system; be sure to use 8 bits, no parity, and 1 stop bit on the remote system. If you do not get a prompt right away, or get garbage, try pressing `Enter` about once per second. If you still do not see a `login:` prompt after a while, try sending a `BREAK`. If you are using a high-speed modem to do the dialing, try dialing again after locking the dialing modem's interface speed (via `AT&B1` on a U.S. Robotics Sportster modem, for example).

If you still cannot get a `login:` prompt, check `/etc/gettytab` again and double-check that

- The initial capability name specified in `/etc/ttys` for the line matches a name of a capability in `/etc/gettytab`

- Each `nx=` entry matches another `gettytab` capability name
- Each `tc=` entry matches another `gettytab` capability name

If you dial but the modem on the DragonFly system will not answer, make sure that the modem is configured to answer the phone when DTR is asserted. If the modem seems to be configured correctly, verify that the DTR line is asserted by checking the modem's indicator lights (if it has any).

If you have gone over everything several times and it still does not work, take a break and come back to it later. If it still does not work, perhaps you can send an electronic mail message to the DragonFly User related mailing list (<http://leaf.dragonflybsd.org/mailarchive/>) describing your modem and your problem, and the good folks on the list will try to help.

17.5 Dial-out Service

The following are tips for getting your host to be able to connect over the modem to another computer. This is appropriate for establishing a terminal session with a remote host.

This is useful to log onto a BBS.

This kind of connection can be extremely helpful to get a file on the Internet if you have problems with PPP. If you need to FTP something and PPP is broken, use the terminal session to FTP it. Then use `zmodem` to transfer it to your machine.

17.5.1 My Stock Hayes Modem Is Not Supported, What Can I Do?

Actually, the manual page for `tip` is out of date. There is a generic Hayes dialer already built in. Just use `at=hayes` in your `/etc/remote` file.

The Hayes driver is not smart enough to recognize some of the advanced features of newer modems—messages like `BUSY`, `NO DIALTONE`, or `CONNECT 115200` will just confuse it. You should turn those messages off when you use `tip` (using `ATX0&W`).

Also, the dial timeout for `tip` is 60 seconds. Your modem should use something less, or else `tip` will think there is a communication problem. Try `ATS7=45&W`.

Note: As shipped, `tip` does not yet support Hayes modems fully. The solution is to edit the file `tipconf.h` in the directory `/usr/src/usr.bin/tip/tip`. Obviously you need the source distribution to do this.

Edit the line `#define HAYES 0` to `#define HAYES 1`. Then `make` and `make install`. Everything works nicely after that.

17.5.2 How Am I Expected to Enter These AT Commands?

Make what is called a “direct” entry in your `/etc/remote` file. For example, if your modem is hooked up to the first serial port, `/dev/cuaa0`, then put in the following line:

```
cuaa0:dv=/dev/cuaa0:br#19200:pa=none
```

Use the highest bps rate your modem supports in the br capability. Then, type `tip cuaa0` and you will be connected to your modem.

If there is no `/dev/cuaa0` on your system, do this:

```
# cd /dev
# sh MAKEDEV cuaa0
```

Or use `cu` as root with the following command:

```
# cu -lline -sspeed
```

line is the serial port (e.g. `/dev/cuaa0`) and *speed* is the speed (e.g. 57600). When you are done entering the AT commands hit `~.` to exit.

17.5.3 The @ Sign for the pn Capability Does Not Work!

The @ sign in the phone number capability tells `tip` to look in `/etc/phones` for a phone number. But the @ sign is also a special character in capability files like `/etc/remote`. Escape it with a backslash:

```
pn=\<@
```

17.5.4 How Can I Dial a Phone Number on the Command Line?

Put what is called a “generic” entry in your `/etc/remote` file. For example:

```
tip115200|Dial any phone number at 115200 bps:\
      :dv=/dev/cuaa0:br#115200:at=hayes:pa=none:du:
tip57600|Dial any phone number at 57600 bps:\
      :dv=/dev/cuaa0:br#57600:at=hayes:pa=none:du:
```

Then you can do things like:

```
# tip -115200 5551234
```

If you prefer `cu` over `tip`, use a generic `cu` entry:

```
cu115200|Use cu to dial any number at 115200bps:\
      :dv=/dev/cuaa1:br#57600:at=hayes:pa=none:du:
```

and type:

```
# cu 5551234 -s 115200
```

17.5.5 Do I Have to Type in the bps Rate Every Time I Do That?

Put in an entry for `tip1200` or `cu1200`, but go ahead and use whatever bps rate is appropriate with the br capability. `tip` thinks a good default is 1200 bps which is why it looks for a `tip1200` entry. You do not have to use 1200 bps, though.

17.5.6 I Access a Number of Hosts Through a Terminal Server

Rather than waiting until you are connected and typing `CONNECT <host>` each time, use `tip`'s `cm` capability. For example, these entries in `/etc/remote`:

```
pain|pain.deep13.com|Forrester's machine:\
    :cm=CONNECT pain\n:tc=deep13:
muffin|muffin.deep13.com|Frank's machine:\
    :cm=CONNECT muffin\n:tc=deep13:
deep13:Gizmonics Institute terminal server:\
    :dv=/dev/cuaa2:br#38400:at=hayes:du:pa=none:pn=5551234:
```

will let you type `tip pain` or `tip muffin` to connect to the hosts `pain` or `muffin`, and `tip deep13` to get to the terminal server.

17.5.7 Can Tip Try More Than One Line for Each Site?

This is often a problem where a university has several modem lines and several thousand students trying to use them.

Make an entry for your university in `/etc/remote` and use `@` for the `pn` capability:

```
big-university:\
    :pn=\@:tc=dialout
dialout:\
    :dv=/dev/cuaa3:br#9600:at=courier:du:pa=none:
```

Then, list the phone numbers for the university in `/etc/phones`:

```
big-university 5551111
big-university 5551112
big-university 5551113
big-university 5551114
```

`tip` will try each one in the listed order, then give up. If you want to keep retrying, run `tip` in a while loop.

17.5.8 Why Do I Have to Hit Ctrl+P Twice to Send Ctrl+P Once?

Ctrl+P is the default “force” character, used to tell `tip` that the next character is literal data. You can set the force character to any other character with the `~s` escape, which means “set a variable.”

Type `~sforce=single-char` followed by a newline. *single-char* is any single character. If you leave out *single-char*, then the force character is the nul character, which you can get by typing **Ctrl+2** or **Ctrl+Space**. A pretty good value for *single-char* is **Shift+Ctrl+6**, which is only used on some terminal servers.

You can have the force character be whatever you want by specifying the following in your `$HOME/.tiprc` file:

```
force=<single-char>
```

17.5.9 Suddenly Everything I Type Is in Upper Case??

You must have pressed **Ctrl+A**, `tip`'s "raise character," specially designed for people with broken caps-lock keys. Use `~s` as above and set the variable `raisechar` to something reasonable. In fact, you can set it to the same as the force character, if you never expect to use either of these features.

Here is a sample `.tiprc` file perfect for **Emacs** users who need to type **Ctrl+2** and **Ctrl+A** a lot:

```
force=^^
raisechar=^^
```

The ^^ is **Shift+Ctrl+6**.

17.5.10 How Can I Do File Transfers with `tip`?

If you are talking to another UNIX system, you can send and receive files with `~p` (put) and `~t` (take). These commands run `cat` and `echo` on the remote system to accept and send files. The syntax is:

```
~p local-file [remote-file]
```

```
~t remote-file [local-file]
```

There is no error checking, so you probably should use another protocol, like `zmodem`.

17.5.11 How Can I Run `zmodem` with `tip`?

To receive files, start the sending program on the remote end. Then, type `~C rz` to begin receiving them locally.

To send files, start the receiving program on the remote end. Then, type `~C sz files` to send them to the remote system.

17.6 Setting Up the Serial Console

Contributed by Kazutaka YOKOTA. Based on a document by Bill Paul.

17.6.1 Introduction

DragonFly has the ability to boot on a system with only a dumb terminal on a serial port as a console. Such a configuration should be useful for two classes of people: system administrators who wish to install DragonFly on machines that have no keyboard or monitor attached, and developers who want to debug the kernel or device drivers.

As described in Chapter 7, DragonFly employs a three stage bootstrap. The first two stages are in the boot block code which is stored at the beginning of the DragonFly slice on the boot disk. The boot block will then load and run the boot loader (`/boot/loader`) as the third stage code.

In order to set up the serial console you must configure the boot block code, the boot loader code and the kernel.

17.6.2 Serial Console Configuration, Terse Version

This section assumes that you are using the default setup, know how to connect serial ports and just want a fast overview of a serial console. If you encounter difficulty with these steps, please see the more extensive explanation of all the options and advanced settings in Section 17.6.3.

1. Connect the serial port. The serial console will be on COM1.
2. `echo -h > /boot.config` to enable the serial console for the boot loader and kernel.
3. Edit `/etc/ttys` and change `off` to `on` for the `ttyd0` entry. This enables a login prompt on the serial console, which mirrors how video consoles are typically setup.
4. `shutdown -r now` will reboot the system with the serial console.

17.6.3 Serial Console Configuration

1. Prepare a serial cable.

You will need either a null-modem cable or a standard serial cable and a null-modem adapter. See Section 17.2.2 for a discussion on serial cables.

2. Unplug your keyboard.

Most PC systems probe for the keyboard during the Power-On Self-Test (POST) and will generate an error if the keyboard is not detected. Some machines complain loudly about the lack of a keyboard and will not continue to boot until it is plugged in.

If your computer complains about the error, but boots anyway, then you do not have to do anything special. (Some machines with Phoenix BIOS installed merely say `Keyboard failed` and continue to boot normally.)

If your computer refuses to boot without a keyboard attached then you will have to configure the BIOS so that it ignores this error (if it can). Consult your motherboard's manual for details on how to do this.

Tip: Setting the keyboard to "Not installed" in the BIOS setup does *not* mean that you will not be able to use your keyboard. All this does is tell the BIOS not to probe for a keyboard at power-on, so it will not complain if the keyboard is not plugged in. You can leave the keyboard plugged in even with this flag set to "Not installed" and the keyboard will still work.

Note: If your system has a PS/2® mouse, chances are very good that you may have to unplug your mouse as well as your keyboard. This is because PS/2 mice share some hardware with the keyboard and leaving the mouse plugged in can fool the keyboard probe into thinking the keyboard is still there. In general, this is not a problem since the mouse is not much good without the keyboard anyway.

3. Plug a dumb terminal into COM1 (`si00`).

If you do not have a dumb terminal, you can use an old PC/XT with a modem program, or the serial port on another UNIX box. If you do not have a COM1 (`si00`), get one. At this time, there is no way to select a port other than COM1 for the boot blocks without recompiling the boot blocks. If you are already using COM1 for another device, you will have to temporarily remove that device and install a new boot block and kernel once you get

DragonFly up and running. (It is assumed that COM1 will be available on a file/compute/terminal server anyway; if you really need COM1 for something else (and you cannot switch that something else to COM2 (`sio1`)), then you probably should not even be bothering with all this in the first place.)

4. Make sure the configuration file of your kernel has appropriate flags set for COM1 (`sio0`).

Relevant flags are:

0x10

Enables console support for this unit. The other console flags are ignored unless this is set. Currently, at most one unit can have console support; the first one (in config file order) with this flag set is preferred. This option alone will not make the serial port the console. Set the following flag or use the `-h` option described below, together with this flag.

0x20

Forces this unit to be the console (unless there is another higher priority console), regardless of the `-h` option discussed below. This flag replaces the `COMCONSOLE` option in DragonFly versions 2.x. The flag `0x20` must be used together with the `0x10` flag.

0x40

Reserves this unit (in conjunction with `0x10`) and makes the unit unavailable for normal access. You should not set this flag to the serial port unit which you want to use as the serial console. This reserves this port for "low-level IO", i.e. kernel debugging.

0x80

This port will be used for remote kernel debugging.

Example:

```
device sio0 at isa? port IO_COM1 flags 0x10 irq 4
```

See the `sio(4)` manual page for more details.

If the flags were not set, you need to run `UserConfig` (on a different console) or recompile the kernel.

5. Create `boot.config` in the root directory of the `a` partition on the boot drive.

This file will instruct the boot block code how you would like to boot the system. In order to activate the serial console, you need one or more of the following options—if you want multiple options, include them all on the same line:

`-h`

Toggles internal and serial consoles. You can use this to switch console devices. For instance, if you boot from the internal (video) console, you can use `-h` to direct the boot loader and the kernel to use the serial port as its console device. Alternatively, if you boot from the serial port, you can use the `-h` to tell the boot loader and the kernel to use the video display as the console instead.

`-D`

Toggles single and dual console configurations. In the single configuration the console will be either the internal console (video display) or the serial port, depending on the state of the `-h` option above. In the dual console configuration, both the video display and the serial port will become the console at the same time,

regardless of the state of the `-h` option. However, note that the dual console configuration takes effect only during the boot block is running. Once the boot loader gets control, the console specified by the `-h` option becomes the only console.

`-P`

Makes the boot block probe the keyboard. If no keyboard is found, the `-D` and `-h` options are automatically set.

Note: Due to space constraints in the current version of the boot blocks, the `-P` option is capable of detecting extended keyboards only. Keyboards with less than 101 keys (and without F11 and F12 keys) may not be detected. Keyboards on some laptop computers may not be properly found because of this limitation. If this is the case with your system, you have to abandon using the `-P` option. Unfortunately there is no workaround for this problem.

Use either the `-P` option to select the console automatically, or the `-h` option to activate the serial console.

You may include other options described in `boot(8)` as well.

The options, except for `-P`, will be passed to the boot loader (`/boot/loader`). The boot loader will determine which of the internal video or the serial port should become the console by examining the state of the `-h` option alone. This means that if you specify the `-D` option but not the `-h` option in `/boot.config`, you can use the serial port as the console only during the boot block; the boot loader will use the internal video display as the console.

6. Boot the machine.

When you start your DragonFly box, the boot blocks will echo the contents of `/boot.config` to the console. For example:

```
/boot.config: -P
Keyboard: no
```

The second line appears only if you put `-P` in `/boot.config` and indicates presence/absence of the keyboard. These messages go to either serial or internal console, or both, depending on the option in `/boot.config`.

Options	Message goes to
none	internal console
<code>-h</code>	serial console
<code>-D</code>	serial and internal consoles
<code>-Dh</code>	serial and internal consoles
<code>-P</code> , keyboard present	internal console
<code>-P</code> , keyboard absent	serial console

After the above messages, there will be a small pause before the boot blocks continue loading the boot loader and before any further messages printed to the console. Under normal circumstances, you do not need to interrupt the boot blocks, but you may want to do so in order to make sure things are set up correctly.

Hit any key, other than Enter, at the console to interrupt the boot process. The boot blocks will then prompt you for further action. You should now see something like:

```
>> DragonFly/i386 BOOT
Default: 0:ad(0,a)/boot/loader
boot:
```

Verify the above message appears on either the serial or internal console or both, according to the options you put in `/boot.config`. If the message appears in the correct console, hit Enter to continue the boot process.

If you want the serial console but you do not see the prompt on the serial terminal, something is wrong with your settings. In the meantime, you enter `-h` and hit Enter/Return (if possible) to tell the boot block (and then the boot loader and the kernel) to choose the serial port for the console. Once the system is up, go back and check what went wrong.

After the boot loader is loaded and you are in the third stage of the boot process you can still switch between the internal console and the serial console by setting appropriate environment variables in the boot loader. See Section 17.6.6.

17.6.4 Summary

Here is the summary of various settings discussed in this section and the console eventually selected.

17.6.4.1 Case 1: You Set the Flags to 0x10 for `sio0`

```
device sio0 at isa? port IO_COM1 flags 0x10 irq 4
```

Options in <code>/boot.config</code>	Console during boot blocks	Console during boot loader	Console in kernel
nothing	internal	internal	internal
<code>-h</code>	serial	serial	serial
<code>-D</code>	serial and internal	internal	internal
<code>-Dh</code>	serial and internal	serial	serial
<code>-P</code> , keyboard present	internal	internal	internal
<code>-P</code> , keyboard absent	serial and internal	serial	serial

17.6.4.2 Case 2: You Set the Flags to 0x30 for `sio0`

```
device sio0 at isa? port IO_COM1 flags 0x30 irq 4
```

Options in <code>/boot.config</code>	Console during boot blocks	Console during boot loader	Console in kernel
nothing	internal	internal	serial
<code>-h</code>	serial	serial	serial
<code>-D</code>	serial and internal	internal	serial
<code>-Dh</code>	serial and internal	serial	serial
<code>-P</code> , keyboard present	internal	internal	serial
<code>-P</code> , keyboard absent	serial and internal	serial	serial

17.6.5 Tips for the Serial Console

17.6.5.1 Setting a Faster Serial Port Speed

By default, the serial port settings are: 9600 baud, 8 bits, no parity, and 1 stop bit. If you wish to change the speed, you need to recompile at least the boot blocks. Add the following line to `/etc/make.conf` and compile new boot blocks:

```
BOOT_COMCONSOLE_SPEED=19200
```

If the serial console is configured in some other way than by booting with `-h`, or if the serial console used by the kernel is different from the one used by the boot blocks, then you must also add the following option to the kernel configuration file and compile a new kernel:

```
options CONSPEED=19200
```

17.6.5.2 Using Serial Port Other Than `si00` for the Console

Using a port other than `si00` as the console requires some recompiling. If you want to use another serial port for whatever reasons, recompile the boot blocks, the boot loader and the kernel as follows.

1. Get the kernel source. (See Section 21.1)
2. Edit `/etc/make.conf` and set `BOOT_COMCONSOLE_PORT` to the address of the port you want to use (0x3F8, 0x2F8, 0x3E8 or 0x2E8). Only `si00` through `si03` (COM1 through COM4) can be used; multiport serial cards will not work. No interrupt setting is needed.
3. Create a custom kernel configuration file and add appropriate flags for the serial port you want to use. For example, if you want to make `si01` (COM2) the console:

```
device si01 at isa? port IO_COM2 flags 0x10 irq 3
```

or

```
device si01 at isa? port IO_COM2 flags 0x30 irq 3
```

The console flags for the other serial ports should not be set.

4. Recompile and install the boot blocks and the boot loader:


```
# cd /sys/boot
# make
# make install
```
5. Rebuild and install the kernel.
6. Write the boot blocks to the boot disk with `disklabel(8)` and boot from the new kernel.

17.6.5.3 Entering the DDB Debugger from the Serial Line

If you wish to drop into the kernel debugger from the serial console (useful for remote diagnostics, but also dangerous if you generate a spurious `BREAK` on the serial port!) then you should compile your kernel with the following options:

```
options BREAK_TO_DEBUGGER
options DDB
```

17.6.5.4 Getting a Login Prompt on the Serial Console

While this is not required, you may wish to get a *login* prompt over the serial line, now that you can see boot messages and can enter the kernel debugging session through the serial console. Here is how to do it.

Open the file `/etc/ttys` with an editor and locate the lines:

```
ttyd0 "/usr/libexec/getty std.9600" unknown off secure
ttyd1 "/usr/libexec/getty std.9600" unknown off secure
ttyd2 "/usr/libexec/getty std.9600" unknown off secure
ttyd3 "/usr/libexec/getty std.9600" unknown off secure
```

`ttyd0` through `ttyd3` corresponds to `COM1` through `COM4`. Change `off` to `on` for the desired port. If you have changed the speed of the serial port, you need to change `std.9600` to match the current setting, e.g. `std.19200`.

You may also want to change the terminal type from `unknown` to the actual type of your serial terminal.

After editing the file, you must `kill -HUP 1` to make this change take effect.

17.6.6 Changing Console from the Boot Loader

Previous sections described how to set up the serial console by tweaking the boot block. This section shows that you can specify the console by entering some commands and environment variables in the boot loader. As the boot loader is invoked at the third stage of the boot process, after the boot block, the settings in the boot loader will override the settings in the boot block.

17.6.6.1 Setting Up the Serial Console

You can easily specify the boot loader and the kernel to use the serial console by writing just one line in `/boot/loader.rc`:

```
set console=comconsole
```

This will take effect regardless of the settings in the boot block discussed in the previous section.

You had better put the above line as the first line of `/boot/loader.rc` so as to see boot messages on the serial console as early as possible.

Likewise, you can specify the internal console as:

```
set console=vidconsole
```

If you do not set the boot loader environment variable `console`, the boot loader, and subsequently the kernel, will use whichever console indicated by the `-h` option in the boot block.

In versions 3.2 or later, you may specify the console in `/boot/loader.conf.local` or `/boot/loader.conf`, rather than in `/boot/loader.rc`. In this method your `/boot/loader.rc` should look like:

```
include /boot/loader.4th
```



```
start
```

Then, create `/boot/loader.conf.local` and put the following line there.

```
console=comconsole
```

or

```
console=vidconsole
```

See `loader.conf(5)` for more information.

Note: At the moment, the boot loader has no option equivalent to the `-P` option in the boot block, and there is no provision to automatically select the internal console and the serial console based on the presence of the keyboard.

17.6.6.2 Using a Serial Port Other Than `si00` for the Console

You need to recompile the boot loader to use a serial port other than `si00` for the serial console. Follow the procedure described in Section 17.6.5.2.

17.6.7 Caveats

The idea here is to allow people to set up dedicated servers that require no graphics hardware or attached keyboards. Unfortunately, while most systems will let you boot without a keyboard, there are quite a few that will not let you boot without a graphics adapter. Machines with AMI BIOSes can be configured to boot with no graphics adapter installed simply by changing the “graphics adapter” setting in the CMOS configuration to “Not installed.”

However, many machines do not support this option and will refuse to boot if you have no display hardware in the system. With these machines, you will have to leave some kind of graphics card plugged in, (even if it is just a junky mono board) although you will not have to attach a monitor. You might also try installing an AMI BIOS.

Chapter 18 PPP and SLIP

Restructured, reorganized, and updated by Jim Mock.

18.1 Synopsis

DragonFly has a number of ways to link one computer to another. To establish a network or Internet connection through a dial-up modem, or to allow others to do so through you, requires the use of PPP or SLIP. This chapter describes setting up these modem-based communication services in detail.

After reading this chapter, you will know:

- How to set up user PPP.
- How to set up kernel PPP.
- How to set up PPPoE (PPP over Ethernet).
- How to set up PPPoA (PPP over ATM).
- How to configure and set up a SLIP client and server.

Before reading this chapter, you should:

- Be familiar with basic network terminology.
- Understand the basics and purpose of a dialup connection and PPP and/or SLIP.

You may be wondering what the main difference is between user PPP and kernel PPP. The answer is simple: user PPP processes the inbound and outbound data in userland rather than in the kernel. This is expensive in terms of copying the data between the kernel and userland, but allows a far more feature-rich PPP implementation. User PPP uses the `tun` device to communicate with the outside world whereas kernel PPP uses the `ppp` device.

Note: Throughout in this chapter, user PPP will simply be referred to as **ppp** unless a distinction needs to be made between it and any other PPP software such as **pppd**. Unless otherwise stated, all of the commands explained in this chapter should be executed as `root`.

18.2 Using User PPP

Updated and enhanced by Tom Rhodes. Originally contributed by Brian Somers. With input from Nik Clayton, Dirk Frömberg, and Peter Childs.

18.2.1 User PPP

18.2.1.1 Assumptions

This document assumes you have the following:

- An account with an Internet Service Provider (ISP) which you connect to using PPP.
- You have a modem or other device connected to your system and configured correctly which allows you to connect to your ISP.
- The dial-up number(s) of your ISP.
- Your login name and password. (Either a regular UNIX style login and password pair, or a PAP or CHAP login and password pair.)
- The IP address of one or more name servers. Normally, you will be given two IP addresses by your ISP to use for this. If they have not given you at least one, then you can use the `enable dns` command in `ppp.conf` and `ppp` will set the name servers for you. This feature depends on your ISP's PPP implementation supporting DNS negotiation.

The following information may be supplied by your ISP, but is not completely necessary:

- The IP address of your ISP's gateway. The gateway is the machine to which you will connect and will be set up as your *default route*. If you do not have this information, we can make one up and your ISP's PPP server will tell us the correct value when we connect.

This IP number is referred to as `HISADDR` by `ppp`.

- The netmask you should use. If your ISP has not provided you with one, you can safely use `255.255.255.255`.
- If your ISP provides you with a static IP address and hostname, you can enter it. Otherwise, we simply let the peer assign whatever IP address it sees fit.

If you do not have any of the required information, contact your ISP.

Note: Throughout this section, many of the examples showing the contents of configuration files are numbered by line. These numbers serve to aid in the presentation and discussion only and are not meant to be placed in the actual file. Proper indentation with tab and space characters is also important.

18.2.1.2 Creating PPP Device Nodes

Under normal circumstances, most users will only need one `tun` device (`/dev/tun0`). References to `tun0` below may be changed to `tunN` where `N` is any unit number corresponding to your system.

The easiest way to make sure that the `tun0` device is configured correctly is to remake the device. To remake the device, do the following:

```
# cd /dev
# sh MAKEDEV tun0
```

If you need 16 tunnel devices in your kernel, you will need to create them. This can be done by executing the following commands:

```
# cd /dev
# sh MAKEDEV tun15
```

18.2.1.3 Automatic PPP Configuration

Both `ppp` and `pppd` (the kernel level implementation of PPP) use the configuration files located in the `/etc/ppp` directory. Examples for user `ppp` can be found in `/usr/share/examples/ppp/`.

Configuring `ppp` requires that you edit a number of files, depending on your requirements. What you put in them depends to some extent on whether your ISP allocates IP addresses statically (i.e., you get given one IP address, and always use that one) or dynamically (i.e., your IP address changes each time you connect to your ISP).

18.2.1.3.1 PPP and Static IP Addresses

You will need to edit the `/etc/ppp/ppp.conf` configuration file. It should look similar to the example below.

Note: Lines that end in a `:` start in the first column (beginning of the line)—all other lines should be indented as shown using spaces or tabs.

```

1  default:
2      set log Phase Chat LCP IPCP CCP tun command
3      ident user-ppp VERSION (built COMPILATIONDATE)
4      set device /dev/cuaa0
5      set speed 115200
6      set dial "ABORT BUSY ABORT NO\\sCARRIER TIMEOUT 5 \
7              \"\" AT OK-AT-OK ATE1Q0 OK \\dATDT\\T TIMEOUT 40 CONNECT"
8      set timeout 180
9      enable dns
10
11  provider:
12      set phone "(123) 456 7890"
13      set authname foo
14      set authkey bar
15      set login "TIMEOUT 10 \"\" \"\" gin:--gin: \\U word: \\P col: ppp"
16      set timeout 300
17      set ifaddr x.x.x.x y.y.y.y 255.255.255.255 0.0.0.0
18      add default HISADDR

```

Line 1:

Identifies the default entry. Commands in this entry are executed automatically when `ppp` is run.

Line 2:

Enables logging parameters. When the configuration is working satisfactorily, this line should be reduced to saying

```
set log phase tun
```

in order to avoid excessive log file sizes.

Line 3:

Tells PPP how to identify itself to the peer. PPP identifies itself to the peer if it has any trouble negotiating and setting up the link, providing information that the peers administrator may find useful when investigating such problems.

Line 4:

Identifies the device to which the modem is connected. COM1 is `/dev/cuaa0` and COM2 is `/dev/cuaa1`.

Line 5:

Sets the speed you want to connect at. If 115200 does not work (it should with any reasonably new modem), try 38400 instead.

Line 6 & 7:

The dial string. User PPP uses an expect-send syntax similar to the `chat(8)` program. Refer to the manual page for information on the features of this language.

Note that this command continues onto the next line for readability. Any command in `ppp.conf` may do this if the last character on the line is a “\” character.

Line 8:

Sets the idle timeout for the link. 180 seconds is the default, so this line is purely cosmetic.

Line 9:

Tells PPP to ask the peer to confirm the local resolver settings. If you run a local name server, this line should be commented out or removed.

Line 10:

A blank line for readability. Blank lines are ignored by PPP.

Line 11:

Identifies an entry for a provider called “provider”. This could be changed to the name of your ISP so that later you can use the `load ISP` to start the connection.

Line 12:

Sets the phone number for this provider. Multiple phone numbers may be specified using the colon (`:`) or pipe character (`|`) as a separator. The difference between the two separators is described in `ppp(8)`. To summarize, if you want to rotate through the numbers, use a colon. If you want to always attempt to dial the first number first and only use the other numbers if the first number fails, use the pipe character. Always quote the entire set of phone numbers as shown.

You must enclose the phone number in quotation marks (`"`) if there is any intention on using spaces in the phone number. This can cause a simple, yet subtle error.

Line 13 & 14:

Identifies the user name and password. When connecting using a UNIX style login prompt, these values are referred to by the `set login` command using the `\U` and `\P` variables. When connecting using PAP or CHAP, these values are used at authentication time.

Line 15:

If you are using PAP or CHAP, there will be no login at this point, and this line should be commented out or removed. See PAP and CHAP authentication for further details.

The login string is of the same chat-like syntax as the dial string. In this example, the string works for a service whose login session looks like this:

```
J. Random Provider
login: foo
password: bar
protocol: ppp
```

You will need to alter this script to suit your own needs. When you write this script for the first time, you should ensure that you have enabled “chat” logging so you can determine if the conversation is going as expected.

Line 16:

Sets the default idle timeout (in seconds) for the connection. Here, the connection will be closed automatically after 300 seconds of inactivity. If you never want to timeout, set this value to zero or use the `-ddial` command line switch.

Line 17:

Sets the interface addresses. The string `x.x.x.x` should be replaced by the IP address that your provider has allocated to you. The string `y.y.y.y` should be replaced by the IP address that your ISP indicated for their gateway (the machine to which you connect). If your ISP has not given you a gateway address, use `10.0.0.2/0`. If you need to use a “guessed” address, make sure that you create an entry in `/etc/ppp/ppp.linkup` as per the instructions for PPP and Dynamic IP addresses. If this line is omitted, `ppp` cannot run in `-auto` mode.

Line 18:

Adds a default route to your ISP’s gateway. The special word `HISADDR` is replaced with the gateway address specified on line 17. It is important that this line appears after line 17, otherwise `HISADDR` will not yet be initialized.

If you do not wish to run `ppp` in `-auto`, this line should be moved to the `ppp.linkup` file.

It is not necessary to add an entry to `ppp.linkup` when you have a static IP address and are running `ppp` in `-auto` mode as your routing table entries are already correct before you connect. You may however wish to create an entry to invoke programs after connection. This is explained later with the `sendmail` example.

Example configuration files can be found in the `/usr/share/examples/ppp/` directory.

18.2.1.3.2 PPP and Dynamic IP Addresses

If your service provider does not assign static IP addresses, `ppp` can be configured to negotiate the local and remote addresses. This is done by “guessing” an IP address and allowing `ppp` to set it up correctly using the IP

Configuration Protocol (IPCP) after connecting. The `ppp.conf` configuration is the same as PPP and Static IP Addresses, with the following change:

```
17      set ifaddr 10.0.0.1/0 10.0.0.2/0 255.255.255.255
```

Again, do not include the line number, it is just for reference. Indentation of at least one space is required.

Line 17:

The number after the / character is the number of bits of the address that ppp will insist on. You may wish to use IP numbers more appropriate to your circumstances, but the above example will always work.

The last argument (0.0.0.0) tells PPP to start negotiations using address 0.0.0.0 rather than 10.0.0.1 and is necessary for some ISPs. Do not use 0.0.0.0 as the first argument to `set ifaddr` as it prevents PPP from setting up an initial route in `-auto` mode.

If you are not running in `-auto` mode, you will need to create an entry in `/etc/ppp/ppp.linkup`. `ppp.linkup` is used after a connection has been established. At this point, ppp will have assigned the interface addresses and it will now be possible to add the routing table entries:

```
1 provider:
2 add default HISADDR
```

Line 1:

On establishing a connection, ppp will look for an entry in `ppp.linkup` according to the following rules: First, try to match the same label as we used in `ppp.conf`. If that fails, look for an entry for the IP address of our gateway. This entry is a four-octet IP style label. If we still have not found an entry, look for the `MYADDR` entry.

Line 2:

This line tells ppp to add a default route that points to `HISADDR`. `HISADDR` will be replaced with the IP number of the gateway as negotiated by the IPCP.

See the `pmdemand` entry in the files `/usr/share/examples/ppp/ppp.conf.sample` and `/usr/share/examples/ppp/ppp.linkup.sample` for a detailed example.

18.2.1.3.3 Receiving Incoming Calls

When you configure **ppp** to receive incoming calls on a machine connected to a LAN, you must decide if you wish to forward packets to the LAN. If you do, you should allocate the peer an IP number from your LAN's subnet, and use the command `enable proxy` in your `/etc/ppp/ppp.conf` file. You should also confirm that the `/etc/rc.conf` file contains the following:

```
gateway_enable="YES"
```

18.2.1.3.4 Which getty?

Configuring DragonFly for Dial-up Services provides a good description on enabling dial-up services using `getty(8)`.

An alternative to `getty` is `mgetty` (<http://www.leo.org/~doering/mgetty/index.html>), a smarter version of `getty` designed with dial-up lines in mind.

The advantages of using `mgetty` is that it actively *talks* to modems, meaning if port is turned off in `/etc/ttys` then your modem will not answer the phone.

Later versions of `mgetty` (from 0.99beta onwards) also support the automatic detection of PPP streams, allowing your clients script-less access to your server.

Refer to `Mgetty and AutoPPP` for more information on `mgetty`.

18.2.1.3.5 PPP Permissions

The `ppp` command must normally be run as the `root` user. If however, you wish to allow `ppp` to run in server mode as a normal user by executing `ppp` as described below, that user must be given permission to run `ppp` by adding them to the `network` group in `/etc/group`.

You will also need to give them access to one or more sections of the configuration file using the `allow` command:

```
allow users fred mary
```

If this command is used in the `default` section, it gives the specified users access to everything.

18.2.1.3.6 PPP Shells for Dynamic-IP Users

Create a file called `/etc/ppp/ppp-shell` containing the following:

```
#!/bin/sh
IDENT=`echo $0 | sed -e 's/^\.*-\(.*\)$/\1/'`
CALLEDAS="$IDENT"
TTY=`tty`

if [ x$IDENT = xdialup ]; then
    IDENT=`basename $TTY`
fi

echo "PPP for $CALLEDAS on $TTY"
echo "Starting PPP for $IDENT"

exec /usr/sbin/ppp -direct $IDENT
```

This script should be executable. Now make a symbolic link called `ppp-dialup` to this script using the following commands:

```
# ln -s ppp-shell /etc/ppp/ppp-dialup
```

You should use this script as the *shell* for all of your dialup users. This is an example from `/etc/passwd` for a dialup PPP user with username `pchilds` (remember do not directly edit the password file, use `vipw`).

```
pchilds:*:1011:300:Peter Childs PPP:/home/ppp:/etc/ppp/ppp-dialup
```

Create a `/home/ppp` directory that is world readable containing the following 0 byte files:

```
-r--r--r--  1 root    wheel      0 May 27  02:23 .hushlogin
-r--r--r--  1 root    wheel      0 May 27  02:22 .rhosts
```

which prevents `/etc/motd` from being displayed.

18.2.1.3.7 PPP Shells for Static-IP Users

Create the `ppp-shell` file as above, and for each account with statically assigned IPs create a symbolic link to `ppp-shell`.

For example, if you have three dialup customers, fred, sam, and mary, that you route class C networks for, you would type the following:

```
# ln -s /etc/ppp/ppp-shell /etc/ppp/ppp-fred
# ln -s /etc/ppp/ppp-shell /etc/ppp/ppp-sam
# ln -s /etc/ppp/ppp-shell /etc/ppp/ppp-mary
```

Each of these users dialup accounts should have their shell set to the symbolic link created above (for example, mary's shell should be `/etc/ppp/ppp-mary`).

18.2.1.3.8 Setting Up `ppp.conf` for Dynamic-IP Users

The `/etc/ppp/ppp.conf` file should contain something along the lines of:

```
default:
  set debug phase lcp chat
  set timeout 0

ttyd0:
  set ifaddr 203.14.100.1 203.14.100.20 255.255.255.255
  enable proxy

ttyd1:
  set ifaddr 203.14.100.1 203.14.100.21 255.255.255.255
  enable proxy
```

Note: The indenting is important.

The `default:` section is loaded for each session. For each dialup line enabled in `/etc/ttys` create an entry similar to the one for `ttyd0:` above. Each line should get a unique IP address from your pool of IP addresses for dynamic users.

18.2.1.3.9 Setting Up `ppp.conf` for Static-IP Users

Along with the contents of the sample `/usr/share/examples/ppp/ppp.conf` above you should add a section for each of the statically assigned dialup users. We will continue with our fred, sam, and mary example.

```
fred:
  set ifaddr 203.14.100.1 203.14.101.1 255.255.255.255

sam:
  set ifaddr 203.14.100.1 203.14.102.1 255.255.255.255

mary:
  set ifaddr 203.14.100.1 203.14.103.1 255.255.255.255
```

The file `/etc/ppp/ppp.linkup` should also contain routing information for each static IP user if required. The line below would add a route for the `203.14.101.0` class C via the client's ppp link.

```
fred:
  add 203.14.101.0 netmask 255.255.255.0 HISADDR

sam:
  add 203.14.102.0 netmask 255.255.255.0 HISADDR

mary:
  add 203.14.103.0 netmask 255.255.255.0 HISADDR
```

18.2.1.3.10 *mgetty and AutoPPP*

Configuring and compiling `mgetty` with the `AUTO_PPP` option enabled allows `mgetty` to detect the LCP phase of PPP connections and automatically spawn off a `ppp` shell. However, since the default login/password sequence does not occur it is necessary to authenticate users using either PAP or CHAP.

This section assumes the user has successfully configured, compiled, and installed a version of `mgetty` with the `AUTO_PPP` option (v0.99beta or later).

Make sure your `/usr/local/etc/mgetty+sendfax/login.config` file has the following in it:

```
/AutoPPP/ - - /etc/ppp/ppp-pap-dialup
```

This will tell `mgetty` to run the `ppp-pap-dialup` script for detected PPP connections.

Create a file called `/etc/ppp/ppp-pap-dialup` containing the following (the file should be executable):

```
#!/bin/sh
exec /usr/sbin/ppp -direct pap$IDENT
```

For each dialup line enabled in `/etc/ttys`, create a corresponding entry in `/etc/ppp/ppp.conf`. This will happily co-exist with the definitions we created above.

```
pap:
  enable pap
  set ifaddr 203.14.100.1 203.14.100.20-203.14.100.40
  enable proxy
```

Each user logging in with this method will need to have a username/password in `/etc/ppp/ppp.secret` file, or alternatively add the following option to authenticate users via PAP from `/etc/passwd` file.

```
enable passwdauth
```

If you wish to assign some users a static IP number, you can specify the number as the third argument in `/etc/ppp/ppp.secret`. See `/usr/share/examples/ppp/ppp.secret.sample` for examples.

18.2.1.3.11 *MS Extensions*

It is possible to configure PPP to supply DNS and NetBIOS nameserver addresses on demand.

To enable these extensions with PPP version 1.x, the following lines might be added to the relevant section of `/etc/ppp/ppp.conf`.

```
enable msxt
```

```
set ns 203.14.100.1 203.14.100.2
set nbns 203.14.100.5
```

And for PPP version 2 and above:

```
accept dns
set dns 203.14.100.1 203.14.100.2
set nbns 203.14.100.5
```

This will tell the clients the primary and secondary name server addresses, and a NetBIOS nameserver host.

In version 2 and above, if the `set dns` line is omitted, PPP will use the values found in `/etc/resolv.conf`.

18.2.1.3.12 PAP and CHAP Authentication

Some ISPs set their system up so that the authentication part of your connection is done using either of the PAP or CHAP authentication mechanisms. If this is the case, your ISP will not give a `login:` prompt when you connect, but will start talking PPP immediately.

PAP is less secure than CHAP, but security is not normally an issue here as passwords, although being sent as plain text with PAP, are being transmitted down a serial line only. There is not much room for crackers to “eavesdrop”.

Referring back to the PPP and Static IP addresses or PPP and Dynamic IP addresses sections, the following alterations must be made:

```
13      set authname MyUserName
14      set authkey MyPassword
15      set login
```

Line 13:

This line specifies your PAP/CHAP user name. You will need to insert the correct value for `MyUserName`.

Line 14:

This line specifies your PAP/CHAP password. You will need to insert the correct value for `MyPassword`. You may want to add an additional line, such as:

```
16      accept PAP
```

or

```
16      accept CHAP
```

to make it obvious that this is the intention, but PAP and CHAP are both accepted by default.

Line 15:

Your ISP will not normally require that you log into the server if you are using PAP or CHAP. You must therefore disable your “set login” string.

18.2.1.3.13 Changing Your `ppp` Configuration on the Fly

It is possible to talk to the `ppp` program while it is running in the background, but only if a suitable diagnostic port has been set up. To do this, add the following line to your configuration:

```
set server /var/run/ppp-tun%d DiagnosticPassword 0177
```

This will tell PPP to listen to the specified UNIX domain socket, asking clients for the specified password before allowing access. The %d in the name is replaced with the tun device number that is in use.

Once a socket has been set up, the pppctl(8) program may be used in scripts that wish to manipulate the running program.

18.2.1.4 Using PPP Network Address Translation Capability

PPP has ability to use internal NAT without kernel diverting capabilities. This functionality may be enabled by the following line in `/etc/ppp/ppp.conf`:

```
nat enable yes
```

Alternatively, PPP NAT may be enabled by command-line option `-nat`. There is also `/etc/rc.conf` knob named `ppp_nat`, which is enabled by default.

If you use this feature, you may also find useful the following `/etc/ppp/ppp.conf` options to enable incoming connections forwarding:

```
nat port tcp 10.0.0.2:ftp ftp
nat port tcp 10.0.0.2:http http
```

or do not trust the outside at all

```
nat deny_incoming yes
```

18.2.1.5 Final System Configuration

You now have `ppp` configured, but there are a few more things to do before it is ready to work. They all involve editing the `/etc/rc.conf` file.

Working from the top down in this file, make sure the `hostname=` line is set, e.g.:

```
hostname="foo.example.com"
```

If your ISP has supplied you with a static IP address and name, it is probably best that you use this name as your host name.

Look for the `network_interfaces` variable. If you want to configure your system to dial your ISP on demand, make sure the `tun0` device is added to the list, otherwise remove it.

```
network_interfaces="lo0 tun0"
ifconfig_tun0=
```

Note: The `ifconfig_tun0` variable should be empty, and a file called `/etc/start_if.tun0` should be created. This file should contain the line:

```
ppp -auto mysystem
```

This script is executed at network configuration time, starting your ppp daemon in automatic mode. If you have a LAN for which this machine is a gateway, you may also wish to use the `-alias` switch. Refer to the manual page for further details.

Make sure the router program set to NO with following line in your `/etc/rc.conf`:

```
router_enable="NO"
```

It is important that the `routed` daemon is not started (it is by default), as `routed` tends to delete the default routing table entries created by `ppp`.

It is probably worth your while ensuring that the `sendmail_flags` line does not include the `-q` option, otherwise `sendmail` will attempt to do a network lookup every now and then, possibly causing your machine to dial out. You may try:

```
sendmail_flags="-bd"
```

The downside of this is that you must force `sendmail` to re-examine the mail queue whenever the ppp link is up by typing:

```
# /usr/sbin/sendmail -q
```

You may wish to use the `!bg` command in `ppp.linkup` to do this automatically:

```
1 provider:
2 delete ALL
3 add 0 0 HISADDR
4 !bg sendmail -bd -q30m
```

If you do not like this, it is possible to set up a “dfilter” to block SMTP traffic. Refer to the sample files for further details.

All that is left is to reboot the machine. After rebooting, you can now either type:

```
# ppp
```

and then `dial provider` to start the PPP session, or, if you want `ppp` to establish sessions automatically when there is outbound traffic (and you have not created the `start_if.tun0` script), type:

```
# ppp -auto provider
```

18.2.1.6 Summary

To recap, the following steps are necessary when setting up ppp for the first time:

Client side:

1. Ensure that the `tun` device is built into your kernel.
2. Ensure that the `tunN` device file is available in the `/dev` directory.
3. Create an entry in `/etc/ppp/ppp.conf`. The `pmdemand` example should suffice for most ISPs.

4. If you have a dynamic IP address, create an entry in `/etc/ppp/ppp.linkup`.
5. Update your `/etc/rc.conf` file.
6. Create a `start_if.tun0` script if you require demand dialing.

Server side:

1. Ensure that the `tun` device is built into your kernel.
2. Ensure that the `tunN` device file is available in the `/dev` directory.
3. Create an entry in `/etc/passwd` (using the `vipw(8)` program).
4. Create a profile in this users home directory that runs `ppp -direct direct-server` or similar.
5. Create an entry in `/etc/ppp/ppp.conf`. The `direct-server` example should suffice.
6. Create an entry in `/etc/ppp/ppp.linkup`.
7. Update your `/etc/rc.conf` file.

18.3 Using Kernel PPP

Parts originally contributed by Gennady B. Sorokopud and Robert Huff.

18.3.1 Setting Up Kernel PPP

Before you start setting up PPP on your machine, make sure that `pppd` is located in `/usr/sbin` and the directory `/etc/ppp` exists.

`pppd` can work in two modes:

1. As a “client” — you want to connect your machine to the outside world via a PPP serial connection or modem line.
2. As a “server” — your machine is located on the network, and is used to connect other computers using PPP.

In both cases you will need to set up an options file (`/etc/ppp/options` or `~/ .ppprc` if you have more than one user on your machine that uses PPP).

You will also need some modem/serial software (preferably `comms/kermit`), so you can dial and establish a connection with the remote host.

18.3.2 Using `pppd` as a Client

Based on information provided by Trev Roydhouse.

The following `/etc/ppp/options` might be used to connect to a Cisco terminal server PPP line.

```
crtscts      # enable hardware flow control
modem       # modem control line
```

```

noipdefault      # remote PPP server must supply your IP address
                  # if the remote host does not send your IP during IPCP
                  # negotiation, remove this option
passive          # wait for LCP packets
domain ppp.foo.com      # put your domain name here

:<remote_ip>     # put the IP of remote PPP host here
                  # it will be used to route packets via PPP link
                  # if you didn't specified the noipdefault option
                  # change this line to <local_ip>:<remote_ip>

defaultroute     # put this if you want that PPP server will be your
                  # default router

```

To connect:

1. Dial to the remote host using **kermit** (or some other modem program), and enter your user name and password (or whatever is needed to enable PPP on the remote host).
2. Exit **kermit** (without hanging up the line).
3. Enter the following:

```
# /usr/src/usr.sbin/pppd.new/pppd /dev/tty01 19200
```

Be sure to use the appropriate speed and device name.

Now your computer is connected with PPP. If the connection fails, you can add the debug option to the `/etc/ppp/options` file, and check console messages to track the problem.

Following `/etc/ppp/pppup` script will make all 3 stages automatic:

```

#!/bin/sh
ps ax |grep pppd |grep -v grep
pid=`ps ax |grep pppd |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing pppd, PID=' ${pid}
    kill ${pid}
fi
ps ax |grep kermit |grep -v grep
pid=`ps ax |grep kermit |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing kermit, PID=' ${pid}
    kill -9 ${pid}
fi

ifconfig ppp0 down
ifconfig ppp0 delete

kermit -y /etc/ppp/kermit.dial
pppd /dev/tty01 19200

```

`/etc/ppp/kermit.dial` is a **kermit** script that dials and makes all necessary authorization on the remote host (an example of such a script is attached to the end of this document).

Use the following `/etc/ppp/pppdown` script to disconnect the PPP line:

```
#!/bin/sh
pid=`ps ax |grep pppd |grep -v grep|awk '{print $1;}'`
if [ X${pid} != "X" ] ; then
    echo 'killing pppd, PID=' ${pid}
    kill -TERM ${pid}
fi

ps ax |grep kermit |grep -v grep
pid=`ps ax |grep kermit |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing kermit, PID=' ${pid}
    kill -9 ${pid}
fi

/sbin/ifconfig ppp0 down
/sbin/ifconfig ppp0 delete
kermit -y /etc/ppp/kermit.hup
/etc/ppp/ppptest
```

Check to see if `pppd` is still running by executing `/usr/etc/ppp/ppptest`, which should look like this:

```
#!/bin/sh
pid=`ps ax| grep pppd |grep -v grep|awk '{print $1;}'`
if [ X${pid} != "X" ] ; then
    echo 'pppd running: PID=' ${pid-NONE}
else
    echo 'No pppd running.'
fi
set -x
netstat -n -I ppp0
ifconfig ppp0
```

To hang up the modem, execute `/etc/ppp/kermit.hup`, which should contain:

```
set line /dev/tty01 ; put your modem device here
set speed 19200
set file type binary
set file names literal
set win 8
set rec pack 1024
set send pack 1024
set block 3
set term bytesize 8
set command bytesize 8
set flow none

pau 1
out +++
inp 5 OK
out ATH0\13
echo \13
```



```
exit
```

Here is an alternate method using chat instead of kermit:

The following two files are sufficient to accomplish a pppd connection.

```
/etc/ppp/options:
```

```
/dev/cuaa1 115200
```

```
crtscts # enable hardware flow control
modem # modem control line
connect "/usr/bin/chat -f /etc/ppp/login.chat.script"
noipdefault # remote PPP server must supply your IP address
            # if the remote host doesn't send your IP during
            # IPCP negotiation, remove this option
passive     # wait for LCP packets
domain <your.domain> # put your domain name here

: # put the IP of remote PPP host here
  # it will be used to route packets via PPP link
  # if you didn't specified the noipdefault option
  # change this line to <local_ip>:<remote_ip>

defaultroute # put this if you want that PPP server will be
             # your default router
```

```
/etc/ppp/login.chat.script:
```

Note: The following should go on a single line.

```
ABORT BUSY ABORT 'NO CARRIER' "" AT OK ATDT<phone.number>
CONNECT "" TIMEOUT 10 ogin:-\\r-ogin: <login-id>
TIMEOUT 5 sword: <password>
```

Once these are installed and modified correctly, all you need to do is run pppd, like so:

```
# pppd
```

18.3.3 Using pppd as a Server

/etc/ppp/options should contain something similar to the following:

```
crtscts # Hardware flow control
netmask 255.255.255.0 # netmask (not required)
192.114.208.20:192.114.208.165 # IP's of local and remote hosts
                                # local ip must be different from one
                                # you assigned to the ethernet (or other)
                                # interface on your machine.
                                # remote IP is IP address that will be
                                # assigned to the remote machine
```

```

domain ppp.foo.com          # your domain
passive                     # wait for LCP
modem                       # modem line

```

The following `/etc/ppp/pppserv` script will tell **pppd** to behave as a server:

```

#!/bin/sh
ps ax |grep pppd |grep -v grep
pid=`ps ax |grep pppd |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing pppd, PID=' ${pid}
    kill ${pid}
fi
ps ax |grep kermi |grep -v grep
pid=`ps ax |grep kermi |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing kermi, PID=' ${pid}
    kill -9 ${pid}
fi

# reset ppp interface
ifconfig ppp0 down
ifconfig ppp0 delete

# enable autoanswer mode
kermi -y /etc/ppp/kermi.ans

# run ppp
pppd /dev/tty01 19200

```

Use this `/etc/ppp/pppservdown` script to stop the server:

```

#!/bin/sh
ps ax |grep pppd |grep -v grep
pid=`ps ax |grep pppd |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing pppd, PID=' ${pid}
    kill ${pid}
fi
ps ax |grep kermi |grep -v grep
pid=`ps ax |grep kermi |grep -v grep|awk '{print $1;}'`
if [ "X${pid}" != "X" ] ; then
    echo 'killing kermi, PID=' ${pid}
    kill -9 ${pid}
fi
ifconfig ppp0 down
ifconfig ppp0 delete

kermi -y /etc/ppp/kermi.noans

```

The following **kermi** script (`/etc/ppp/kermi.ans`) will enable/disable autoanswer mode on your modem. It should look like this:

```

set line /dev/tty01
set speed 19200
set file type binary
set file names literal
set win 8
set rec pack 1024
set send pack 1024
set block 3
set term bytesize 8
set command bytesize 8
set flow none

pau 1
out +++
inp 5 OK
out ATH0\13
inp 5 OK
echo \13
out ATS0=1\13 ; change this to out ATS0=0\13 if you want to disable
                ; autoanswer mode

inp 5 OK
echo \13
exit

```

A script named `/etc/ppp/kermit.dial` is used for dialing and authenticating on the remote host. You will need to customize it for your needs. Put your login and password in this script; you will also need to change the input statement depending on responses from your modem and remote host.

```

;
; put the com line attached to the modem here:
;
set line /dev/tty01
;
; put the modem speed here:
;
set speed 19200
set file type binary ; full 8 bit file xfer
set file names literal
set win 8
set rec pack 1024
set send pack 1024
set block 3
set term bytesize 8
set command bytesize 8
set flow none
set modem hayes
set dial hangup off
set carrier auto ; Then SET CARRIER if necessary,
set dial display on ; Then SET DIAL if necessary,
set input echo on
set input timeout proceed
set input case ignore
def \%x 0 ; login prompt counter

```

```

goto slhup

:slcmd                                ; put the modem in command mode
echo Put the modem in command mode.
clear                                  ; Clear unread characters from input buffer
pause 1
output +++                             ; hayes escape sequence
input 1 OK\13\10                       ; wait for OK
if success goto slhup
output \13
pause 1
output at\13
input 1 OK\13\10
if fail goto slcmd                     ; if modem doesn't answer OK, try again

:slhup                                  ; hang up the phone
clear                                  ; Clear unread characters from input buffer
pause 1
echo Hanging up the phone.
output ath0\13                          ; hayes command for on hook
input 2 OK\13\10
if fail goto slcmd                     ; if no OK answer, put modem in command mode

:sldial                                 ; dial the number
pause 1
echo Dialing.
output atdt9,550311\13\10               ; put phone number here
assign \%x 0                            ; zero the time counter

:look
clear                                  ; Clear unread characters from input buffer
increment \%x                           ; Count the seconds
input 1 {CONNECT }
if success goto sllogin
reinput 1 {NO CARRIER\13\10}
if success goto sldial
reinput 1 {NO DIALTONE\13\10}
if success goto slnodial
reinput 1 {\255}
if success goto slhup
reinput 1 {\127}
if success goto slhup
if < \%x 60 goto look
else goto slhup

:sllogin                                ; login
assign \%x 0                            ; zero the time counter
pause 1
echo Looking for login prompt.

:slloop
increment \%x                           ; Count the seconds
clear                                    ; Clear unread characters from input buffer

```

```

output \13
;
; put your expected login prompt here:
;
input 1 {Username: }
if success goto sluid
reinput 1 {\255}
if success goto slhup
reinput 1 {\127}
if success goto slhup
if < \%x 10 goto slloop           ; try 10 times to get a login prompt
else goto slhup                   ; hang up and start again if 10 failures

:sluid
;
; put your userid here:
;
output ppp-login\13
input 1 {Password: }
;
; put your password here:
;
output ppp-password\13
input 1 {Entering SLIP mode.}
echo
quit

:slnodial
echo \7No dialtone.  Check the telephone line!\7
exit 1

; local variables:
; mode: csh
; comment-start: ";"
; comment-start-skip: ";"
; end:

```

18.4 Troubleshooting PPP Connections

Contributed by Tom Rhodes.

This section covers a few issues which may arise when using PPP over a modem connection. For instance, perhaps you need to know exactly what prompts the system you are dialing into will present. Some ISPs present the `ssword` prompt, and others will present `password`; if the `ppp` script is not written accordingly, the login attempt will fail. The most common way to debug `ppp` connections is by connecting manually. The following information will walk you through a manual connection step by step.

18.4.1 Check the Device Nodes

If you reconfigured your kernel then you recall the `sio` device. If you did not configure your kernel, there is no reason to worry. Just check the `dmesg` output for the modem device with:

```
#dmesg | grep sio
```

You should get some pertinent output about the `sio` devices. These are the COM ports we need. If your modem acts like a standard serial port then you should see it listed on `sio1`, or COM2. If so, you are not required to rebuild the kernel, you just need to make the serial device. You can do this by changing your directory to `/dev` and running the `MAKEDEV` script like above. Now make the serial devices with:

```
# sh MAKEDEV cuaa0 cuaa1 cuaa2 cuaa3
```

which will create the serial devices for your system. When matching up `sio` modem is on `sio1` or COM2 if you are in DOS, then your modem device would be `/dev/cuaa1`.

18.4.2 Connecting Manually

Connecting to the Internet by manually controlling `ppp` is quick, easy, and a great way to debug a connection or just get information on how your ISP treats `ppp` client connections. Lets start **PPP** from the command line. Note that in all of our examples we will use *example* as the hostname of the machine running **PPP**. You start `ppp` by just typing `ppp`:

```
# ppp
```

We have now started `ppp`.

```
ppp ON example> set device /dev/cuaa1
```

We set our modem device, in this case it is `cuaa1`.

```
ppp ON example> set speed 115200
```

Set the connection speed, in this case we are using 115,200 kbps.

```
ppp ON example> enable dns
```

Tell `ppp` to configure our resolver and add the nameserver lines to `/etc/resolv.conf`. If `ppp` cannot determine our hostname, we can set one manually later.

```
ppp ON example> term
```

Switch to “terminal” mode so that we can manually control the modem.

```
deflink: Entering terminal mode on /dev/cuaa1
type '~h' for help
```

```
at
```

```
OK
```

```
atdt123456789
```

Use `at` to initialize the modem, then use `atdt` and the number for your ISP to begin the dial in process.

```
CONNECT
```

Confirmation of the connection, if we are going to have any connection problems, unrelated to hardware, here is where we will attempt to resolve them.

```
ISP Login:myusername
```

Here you are prompted for a username, return the prompt with the username that was provided by the ISP.

```
ISP Pass:mypassword
```

This time we are prompted for a password, just reply with the password that was provided by the ISP. Just like logging into DragonFly, the password will not echo.

```
Shell or PPP:ppp
```

Depending on your ISP this prompt may never appear. Here we are being asked if we wish to use a shell on the provider, or to start ppp. In this example, we have chosen to use ppp as we want an Internet connection.

```
Ppp ON example>
```

Notice that in this example the first p has been capitalized. This shows that we have successfully connected to the ISP.

```
PPp ON example>
```

We have successfully authenticated with our ISP and are waiting for the assigned IP address.

```
PPP ON example>
```

We have made an agreement on an IP address and successfully completed our connection.

```
PPP ON example>add default HISADDR
```

Here we add our default route, we need to do this before we can talk to the outside world as currently the only established connection is with the peer. If this fails due to existing routes you can put a bang character ! in front of the add. Alternatively, you can set this before making the actual connection and it will negotiate a new route accordingly.

If everything went good we should now have an active connection to the Internet, which could be thrown into the background using **CTRL+z** If you notice the PPP return to ppp then we have lost our connection. This is good to know because it shows our connection status. Capital P's show that we have a connection to the ISP and lowercase p's show that the connection has been lost for whatever reason. ppp only has these 2 states.

18.4.2.1 Debugging

If you have a direct line and cannot seem to make a connection, then turn hardware flow CTS/RTS to off with the `set ctsrts off`. This is mainly the case if you are connected to some **PPP** capable terminal servers, where **PPP** hangs when it tries to write data to your communication link, so it would be waiting for a CTS, or Clear To Send signal which may never come. If you use this option however, you should also use the `set accmap` option, which may be required to defeat hardware dependent on passing certain characters from end to end, most of the time XON/XOFF. See the ppp(8) manual page for more information on this option, and how it is used.

If you have an older modem, you may need to use the `set parity even`. Parity is set at none by default, but is used for error checking (with a large increase in traffic) on older modems and some ISPs. You may need this option for the Compuserve ISP.

PPP may not return to the command mode, which is usually a negotiation error where the ISP is waiting for your side to start negotiating. At this point, using the `~p` command will force ppp to start sending the configuration information.

If you never obtain a login prompt, then most likely you need to use PAP or CHAP authentication instead of the UNIX style in the example above. To use PAP or CHAP just add the following options to **PPP** before going into terminal mode:

```
ppp ON example> set authname myusername
```

Where *myusername* should be replaced with the username that was assigned by the ISP.

```
ppp ON example> set authkey mypassword
```

Where *mypassword* should be replaced with the password that was assigned by the ISP.

If you connect fine, but cannot seem to find any domain name, try to use `ping(8)` with an IP address and see if you can get any return information. If you experience 100 percent (100%) packet loss, then it is most likely that you were not assigned a default route. Double check that the option `add default HISADDR` was set during the connection. If you can connect to a remote IP address then it is possible that a resolver address has not been added to the `/etc/resolv.conf`. This file should look like:

```
domain example.com
nameserver x.x.x.x
nameserver y.y.y.y
```

Where *x.x.x.x* and *y.y.y.y* should be replaced with the IP address of your ISP's DNS servers. This information may or may not have been provided when you signed up, but a quick call to your ISP should remedy that.

You could also have `syslog(3)` provide a logging function for your **PPP** connection. Just add:

```
!ppp
*. *      /var/log/ppp.log
```

to `/etc/syslog.conf`. In most cases, this functionality already exists.

18.5 Using PPP over Ethernet (PPPoE)

Contributed (from <http://node.to/freebsd/how-tos/how-to-freebsd-pppoe.html>) by Jim Mock.

This section describes how to set up PPP over Ethernet (PPPoE).

18.5.1 Configuring the Kernel

No kernel configuration is necessary for PPPoE any longer. If the necessary netgraph support is not built into the kernel, it will be dynamically loaded by **ppp**.

18.5.2 Setting Up `ppp.conf`

Here is an example of a working `ppp.conf`:

```
default:
    set log Phase tun command # you can add more detailed logging if you wish
    set ifaddr 10.0.0.1/0 10.0.0.2/0

name_of_service_provider:
    set device PPPoE:x11 # replace x11 with your ethernet device
    set authname YOURLOGINNAME
    set authkey YOURPASSWORD
    set dial
    set login
    add default HISADDR
```

18.5.3 Running `ppp`

As root, you can run:

```
# ppp -ddial name_of_service_provider
```

18.5.4 Starting `ppp` at Boot

Add the following to your `/etc/rc.conf` file:

```
ppp_enable="YES"
ppp_mode="ddial"
ppp_nat="YES" # if you want to enable nat for your local network, otherwise NO
ppp_profile="name_of_service_provider"
```

18.5.5 Using a PPPoE Service Tag

Sometimes it will be necessary to use a service tag to establish your connection. Service tags are used to distinguish between different PPPoE servers attached to a given network.

You should have been given any required service tag information in the documentation provided by your ISP. If you cannot locate it there, ask your ISP's tech support personnel.

As a last resort, you could try the method suggested by the Roaring Penguin PPPoE (<http://www.roaringpenguin.com/pppoe/>) program which can be found in the `pkgsrc` collection. Bear in mind however, this may de-program your modem and render it useless, so think twice before doing it. Simply install the program shipped with the modem by your provider. Then, access the **System** menu from the program. The name of your profile should be listed there. It is usually *ISP*.

The profile name (service tag) will be used in the PPPoE configuration entry in `ppp.conf` as the provider part of the `set device` command (see the `ppp(8)` manual page for full details). It should look like this:

```
set device PPPoE:x11:ISP
```

Do not forget to change `x11` to the proper device for your Ethernet card.

Do not forget to change `ISP` to the profile you have just found above.

For additional information, see:

- Cheaper Broadband with FreeBSD on DSL (<http://renaud.waldura.com/doc/freebsd/pppoe/>) by Renaud Waldura.
- Nutzung von T-DSL und T-Online mit FreeBSD (<http://www.ruhr.de/home/nathan/FreeBSD/tdsl-freebsd.html>) by Udo Erdelhoff (in German).

18.5.6 PPPoE with a 3Com® HomeConnect® ADSL Modem Dual Link

This modem does not follow RFC 2516 (<http://www.faqs.org/rfcs/rfc2516.html>) (*A Method for transmitting PPP over Ethernet (PPPoE)*), written by L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, and R. Wheeler). Instead, different packet type codes have been used for the Ethernet frames. Please complain to 3Com (<http://www.3com.com/>) if you think it should comply with the PPPoE specification.

In order to make DragonFly capable of communicating with this device, a `sysctl` must be set. This can be done automatically at boot time by updating `/etc/sysctl.conf`:

```
net.graph.nonstandard_pppoe=1
```

or can be done for immediate effect with the command `sysctl net.graph.nonstandard_pppoe=1`.

Unfortunately, because this is a system-wide setting, it is not possible to talk to a normal PPPoE client or server and a 3Com® HomeConnect® ADSL Modem at the same time.

18.6 Using SLIP

Originally contributed by Satoshi Asami. With input from Guy Helmer and Piero Serini.

18.6.1 Setting Up a SLIP Client

The following is one way to set up a DragonFly machine for SLIP on a static host network. For dynamic hostname assignments (your address changes each time you dial up), you probably need to have a more complex setup.

First, determine which serial port your modem is connected to. Many people set up a symbolic link, such as `/dev/modem`, to point to the real device name, `/dev/cuaaN`. This allows you to abstract the actual device name should you ever need to move the modem to a different port. It can become quite cumbersome when you need to fix a bunch of files in `/etc` and `.kernrc` files all over the system!

Note: `/dev/cuaa0` is COM1, `cuaa1` is COM2, etc.

Make sure you have the following in your kernel configuration file:

```
pseudo-device    sl        1
```

It is included in the `GENERIC` kernel, so this should not be a problem unless you have deleted it.

18.6.1.1 Things You Have to Do Only Once

1. Add your home machine, the gateway and nameservers to your `/etc/hosts` file. Mine looks like this:

```
127.0.0.1          localhost loghost
136.152.64.181    water.CS.Example.EDU water.CS water
136.152.64.1      inr-3.CS.Example.EDU inr-3 slip-gateway
128.32.136.9      ns1.Example.EDU ns1
128.32.136.12     ns2.Example.EDU ns2
```

2. Make sure you have `hosts` before `bind` in your `/etc/host.conf`.
3. Edit the `/etc/rc.conf` file.

1. Set your hostname by editing the line that says:

```
hostname="myname.my.domain"
```

Your machine's full Internet hostname should be placed here.

2. Add `s10` to the list of network interfaces by changing the line that says:

```
network_interfaces="lo0"
```

to:

```
network_interfaces="lo0 s10"
```

3. Set the startup flags of `s10` by adding a line:

```
ifconfig_s10="inet ${hostname} slip-gateway netmask 0xffffffff00 up"
```

4. Designate the default router by changing the line:

```
defaultrouter="NO"
```

to:

```
defaultrouter="slip-gateway"
```

4. Make a file `/etc/resolv.conf` which contains:

```
domain CS.Example.EDU
nameserver 128.32.136.9
nameserver 128.32.136.12
```

As you can see, these set up the nameserver hosts. Of course, the actual domain names and addresses depend on your environment.

5. Set the password for `root` and `toor` (and any other accounts that do not have a password).
6. Reboot your machine and make sure it comes up with the correct hostname.

18.6.1.2 Making a SLIP Connection

1. Dial up, type `slip` at the prompt, enter your machine name and password. What is required to be entered depends on your environment. If you use `kermit`, you can try a script like this:

```
# kermit setup
set modem hayes
set line /dev/modem
```

```

set speed 115200
set parity none
set flow rts/cts
set terminal bytesize 8
set file type binary
# The next macro will dial up and login
define slip dial 643-9600, input 10 =>, if failure stop, -
output slip\x0d, input 10 Username:, if failure stop, -
output silvia\x0d, input 10 Password:, if failure stop, -
output ***\x0d, echo \x0aCONNECTED\x0a

```

Of course, you have to change the hostname and password to fit yours. After doing so, you can just type `slip` from the kermit prompt to connect.

Note: Leaving your password in plain text anywhere in the filesystem is generally a *bad* idea. Do it at your own risk.

2. Leave the kermit there (you can suspend it by **Ctrl-z**) and as `root`, type:

```
# slattach -h -c -s 115200 /dev/modem
```

If you are able to ping hosts on the other side of the router, you are connected! If it does not work, you might want to try `-a` instead of `-c` as an argument to `slattach`.

18.6.1.3 How to Shutdown the Connection

Do the following:

```
# kill -INT `cat /var/run/slattach.modem.pid`
```

to kill `slattach`. Keep in mind you must be `root` to do the above. Then go back to kermit (by running `fg` if you suspended it) and exit from it (`q`).

The `slattach` manual page says you have to use `ifconfig s10 down` to mark the interface down, but this does not seem to make any difference for me. (`ifconfig s10` reports the same thing.)

Some times, your modem might refuse to drop the carrier (mine often does). In that case, simply start kermit and quit it again. It usually goes out on the second try.

18.6.1.4 Troubleshooting

If it does not work, feel free to ask me. The things that people tripped over so far:

- Not using `-c` or `-a` in `slattach` (This should not be fatal, but some users have reported that this solves their problems.)
- Using `s10` instead of `sl0` (might be hard to see the difference on some fonts).
- Try `ifconfig s10` to see your interface status. For example, you might get:

```
# ifconfig s10
s10: flags=10<POINTOPOINT>
```

```
inet 136.152.64.181 --> 136.152.64.1 netmask fffffff0
```

- If you get no route to host messages from ping, there may be a problem with your routing table. You can use the `netstat -r` command to display the current routes :

```
# netstat -r
Routing tables
Destination      Gateway          Flags           Refs      Use  IfaceMTU    Rtt      Netmasks:

(root node)
(root node)

Route Tree for Protocol Family inet:
(root node) =>
default          inr-3.Example.EDU  UG              8    224515  sl0 -        -
localhost.Exampl localhost.Example. UH              5     42127  lo0 -        0.438
inr-3.Example.ED water.CS.Example.E UH              1         0    sl0 -        -
water.CS.Example localhost.Example. UGH            34  47641234  lo0 -        0.438
(root node)
```

The preceding examples are from a relatively busy system. The numbers on your system will vary depending on network activity.

18.6.2 Setting Up a SLIP Server

This document provides suggestions for setting up SLIP Server services on a DragonFly system, which typically means configuring your system to automatically startup connections upon login for remote SLIP clients.

18.6.2.1 Prerequisites

This section is very technical in nature, so background knowledge is required. It is assumed that you are familiar with the TCP/IP network protocol, and in particular, network and node addressing, network address masks, subnetting, routing, and routing protocols, such as RIP. Configuring SLIP services on a dial-up server requires a knowledge of these concepts, and if you are not familiar with them, please read a copy of either Craig Hunt's *TCP/IP Network Administration* published by O'Reilly & Associates, Inc. (ISBN Number 0-937175-82-X), or Douglas Comer's books on the TCP/IP protocol.

It is further assumed that you have already set up your modem(s) and configured the appropriate system files to allow logins through your modems. If you have not prepared your system for this yet, please see the tutorial for configuring dialup services. You may also want to check the manual pages for `sio(4)` for information on the serial port device driver and `ttys(5)`, `gettytab(5)`, `getty(8)`, & `init(8)` for information relevant to configuring the system to accept logins on modems, and perhaps `stty(1)` for information on setting serial port parameters (such as `cllocal` for directly-connected serial interfaces).

18.6.2.2 Quick Overview

In its typical configuration, using DragonFly as a SLIP server works as follows: a SLIP user dials up your DragonFly SLIP Server system and logs in with a special SLIP login ID that uses `/usr/sbin/sliplogin` as the special user's shell. The `sliplogin` program browses the file `/etc/sliphome/slip.hosts` to find a matching line for the

special user, and if it finds a match, connects the serial line to an available SLIP interface and then runs the shell script `/etc/sliphome/slip.login` to configure the SLIP interface.

18.6.2.2.1 An Example of a SLIP Server Login

For example, if a SLIP user ID were `Shelmerg`, `Shelmerg`'s entry in `/etc/master.passwd` would look something like this:

```
Shelmerg:password:1964:89::0:0:Guy Helmer - SLIP:/usr/users/Shelmerg:/usr/sbin/sliplogin
```

When `Shelmerg` logs in, `sliplogin` will search `/etc/sliphome/slip.hosts` for a line that had a matching user ID; for example, there may be a line in `/etc/sliphome/slip.hosts` that reads:

```
Shelmerg          dc-slip sl-helmer          0xfffffc00  autocomp
```

`sliplogin` will find that matching line, hook the serial line into the next available SLIP interface, and then execute `/etc/sliphome/slip.login` like this:

```
/etc/sliphome/slip.login 0 19200 Shelmerg dc-slip sl-helmer 0xfffffc00 autocomp
```

If all goes well, `/etc/sliphome/slip.login` will issue an `ifconfig` for the SLIP interface to which `sliplogin` attached itself (slip interface 0, in the above example, which was the first parameter in the list given to `slip.login`) to set the local IP address (`dc-slip`), remote IP address (`sl-helmer`), network mask for the SLIP interface (`0xfffffc00`), and any additional flags (`autocomp`). If something goes wrong, `sliplogin` usually logs good informational messages via the daemon `syslog` facility, which usually logs to `/var/log/messages` (see the manual pages for `syslogd(8)` and `syslog.conf(5)` and perhaps check `/etc/syslog.conf` to see to what `syslogd` is logging and where it is logging to).

OK, enough of the examples — let us dive into setting up the system.

18.6.2.3 Kernel Configuration

DragonFly's default kernels usually come with two SLIP interfaces defined (`s10` and `s11`); you can use `netstat -i` to see whether these interfaces are defined in your kernel.

Sample output from `netstat -i`:

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
ed0	1500	<Link>	0.0.c0.2c.5f.4a	291311	0	174209	0	133
ed0	1500	138.247.224	ivory	291311	0	174209	0	133
lo0	65535	<Link>		79	0	79	0	0
lo0	65535	loop	localhost	79	0	79	0	0
s10*	296	<Link>		0	0	0	0	0
s11*	296	<Link>		0	0	0	0	0

The `s10` and `s11` interfaces shown from `netstat -i` indicate that there are two SLIP interfaces built into the kernel. (The asterisks after the `s10` and `s11` indicate that the interfaces are “down”.)

However, DragonFly's default kernel does not come configured to forward packets (by default, your DragonFly machine will not act as a router) due to Internet RFC requirements for Internet hosts (see RFCs 1009 [Requirements for Internet Gateways], 1122 [Requirements for Internet Hosts — Communication Layers], and perhaps 1127 [A

Perspective on the Host Requirements RFCs]). If you want your DragonFly SLIP Server to act as a router, you will have to edit the `/etc/rc.conf` file and change the setting of the `gateway_enable` variable to `YES`.

You will then need to reboot for the new settings to take effect.

You will notice that near the end of the default kernel configuration file (`/sys/i386/conf/GENERIC`) is a line that reads:

```
pseudo-device sl 2
```

This is the line that defines the number of SLIP devices available in the kernel; the number at the end of the line is the maximum number of SLIP connections that may be operating simultaneously.

Please refer to Chapter 9 on Configuring the DragonFly Kernel for help in reconfiguring your kernel.

18.6.2.4 Sliplogin Configuration

As mentioned earlier, there are three files in the `/etc/sliphome` directory that are part of the configuration for `/usr/sbin/sliplogin` (see `sliplogin(8)` for the actual manual page for `sliplogin`): `slip.hosts`, which defines the SLIP users and their associated IP addresses; `slip.login`, which usually just configures the SLIP interface; and (optionally) `slip.logout`, which undoes `slip.login`'s effects when the serial connection is terminated.

18.6.2.4.1 `slip.hosts` Configuration

`/etc/sliphome/slip.hosts` contains lines which have at least four items separated by whitespace:

- SLIP user's login ID
- Local address (local to the SLIP server) of the SLIP link
- Remote address of the SLIP link
- Network mask

The local and remote addresses may be host names (resolved to IP addresses by `/etc/hosts` or by the domain name service, depending on your specifications in `/etc/host.conf`), and the network mask may be a name that can be resolved by a lookup into `/etc/networks`. On a sample system, `/etc/sliphome/slip.hosts` looks like this:

```
#
# login local-addr      remote-addr      mask              opt1    opt2
#                               (normal,compress,noicmp)
#
Shelmerg dc-slip        sl-helmerg       0xfffffc00       autocomp
```

At the end of the line is one or more of the options.

- `normal` — no header compression
- `compress` — compress headers
- `autocomp` — compress headers if the remote end allows it
- `noicmp` — disable ICMP packets (so any “ping” packets will be dropped instead of using up your bandwidth)

Your choice of local and remote addresses for your SLIP links depends on whether you are going to dedicate a TCP/IP subnet or if you are going to use “proxy ARP” on your SLIP server (it is not “true” proxy ARP, but that is the

terminology used in this section to describe it). If you are not sure which method to select or how to assign IP addresses, please refer to the TCP/IP books referenced in the SLIP Prerequisites (Section 18.6.2.1) and/or consult your IP network manager.

If you are going to use a separate subnet for your SLIP clients, you will need to allocate the subnet number out of your assigned IP network number and assign each of your SLIP client's IP numbers out of that subnet. Then, you will probably need to configure a static route to the SLIP subnet via your SLIP server on your nearest IP router.

Otherwise, if you will use the "proxy ARP" method, you will need to assign your SLIP client's IP addresses out of your SLIP server's Ethernet subnet, and you will also need to adjust your `/etc/sliphome/slip.login` and `/etc/sliphome/slip.logout` scripts to use `arp(8)` to manage the proxy-ARP entries in the SLIP server's ARP table.

18.6.2.4.2 `slip.login` Configuration

The typical `/etc/sliphome/slip.login` file looks like this:

```
#!/bin/sh -
#
#      @(#)slip.login  5.1 (Berkeley) 7/1/90

#
# generic login file for a slip line.  sliplogin invokes this with
# the parameters:
#      1      2      3      4      5      6      7-n
#  slipunit ttyspeed loginname local-addr remote-addr mask opt-args
#
/sbin/ifconfig sl$1 inet $4 $5 netmask $6
```

This `slip.login` file merely runs `ifconfig` for the appropriate SLIP interface with the local and remote addresses and network mask of the SLIP interface.

If you have decided to use the "proxy ARP" method (instead of using a separate subnet for your SLIP clients), your `/etc/sliphome/slip.login` file will need to look something like this:

```
#!/bin/sh -
#
#      @(#)slip.login  5.1 (Berkeley) 7/1/90

#
# generic login file for a slip line.  sliplogin invokes this with
# the parameters:
#      1      2      3      4      5      6      7-n
#  slipunit ttyspeed loginname local-addr remote-addr mask opt-args
#
/sbin/ifconfig sl$1 inet $4 $5 netmask $6
# Answer ARP requests for the SLIP client with our Ethernet addr
/usr/sbin/arp -s $5 00:11:22:33:44:55 pub
```

The additional line in this `slip.login`, `arp -s $5 00:11:22:33:44:55 pub`, creates an ARP entry in the SLIP server's ARP table. This ARP entry causes the SLIP server to respond with the SLIP server's Ethernet MAC address whenever another IP node on the Ethernet asks to speak to the SLIP client's IP address.

When using the example above, be sure to replace the Ethernet MAC address (00:11:22:33:44:55) with the MAC address of your system's Ethernet card, or your "proxy ARP" will definitely not work! You can discover your SLIP server's Ethernet MAC address by looking at the results of running `netstat -i`; the second line of the output should look something like:

```
ed0    1500    <Link>0.2.c1.28.5f.4a          191923 0    129457    0    116
```

This indicates that this particular system's Ethernet MAC address is 00:02:c1:28:5f:4a — the periods in the Ethernet MAC address given by `netstat -i` must be changed to colons and leading zeros should be added to each single-digit hexadecimal number to convert the address into the form that `arp(8)` desires; see the manual page on `arp(8)` for complete information on usage.

Note: When you create `/etc/sliphome/slip.login` and `/etc/sliphome/slip.logout`, the "execute" bit (`chmod 755 /etc/sliphome/slip.login /etc/sliphome/slip.logout`) must be set, or `sliplogin` will be unable to execute it.

18.6.2.4.3 `slip.logout` Configuration

`/etc/sliphome/slip.logout` is not strictly needed (unless you are implementing "proxy ARP"), but if you decide to create it, this is an example of a basic `slip.logout` script:

```
#!/bin/sh -
#
#      slip.logout

#
# logout file for a slip line.  sliplogin invokes this with
# the parameters:
#      1      2      3      4      5      6      7-n
#  slipunit ttyspeed loginname local-addr remote-addr mask opt-args
#
/sbin/ifconfig sl$1 down
```

If you are using "proxy ARP", you will want to have `/etc/sliphome/slip.logout` remove the ARP entry for the SLIP client:

```
#!/bin/sh -
#
#      @(#)slip.logout

#
# logout file for a slip line.  sliplogin invokes this with
# the parameters:
#      1      2      3      4      5      6      7-n
#  slipunit ttyspeed loginname local-addr remote-addr mask opt-args
#
/sbin/ifconfig sl$1 down
# Quit answering ARP requests for the SLIP client
/usr/sbin/arp -d $5
```

The `arp -d $5` removes the ARP entry that the “proxy ARP” `slip.login` added when the SLIP client logged in. It bears repeating: make sure `/etc/sliphome/slip.logout` has the execute bit set after you create it (ie, `chmod 755 /etc/sliphome/slip.logout`).

18.6.2.5 Routing Considerations

If you are not using the “proxy ARP” method for routing packets between your SLIP clients and the rest of your network (and perhaps the Internet), you will probably have to add static routes to your closest default router(s) to route your SLIP client subnet via your SLIP server.

18.6.2.5.1 Static Routes

Adding static routes to your nearest default routers can be troublesome (or impossible if you do not have authority to do so...). If you have a multiple-router network in your organization, some routers, such as those made by Cisco and Proteon, may not only need to be configured with the static route to the SLIP subnet, but also need to be told which static routes to tell other routers about, so some expertise and troubleshooting/tweaking may be necessary to get static-route-based routing to work.

18.6.2.5.2 Running **GateD**[®]

Note: **GateD**[®] is proprietary software now and will not be available as source code to the public anymore (more info on the GateD (<http://www.gated.org/>) website). This section only exists to ensure backwards compatibility for those that are still using an older version.

An alternative to the headaches of static routes is to install **GateD** on your DragonFly SLIP server and configure it to use the appropriate routing protocols (RIP/OSPF/BGP/EGP) to tell other routers about your SLIP subnet. You’ll need to write a `/etc/gated.conf` file to configure your gated; here is a sample, similar to what the author used on a FreeBSD SLIP server:

```
#
# gated configuration file for dc.dsu.edu; for gated version 3.5alpha5
# Only broadcast RIP information for xxx.xxx.yy out the ed Ethernet interface
#
#
# tracing options
#
traceoptions "/var/tmp/gated.output" replace size 100k files 2 general ;

rip yes {
    interface sl noripout noripin ;
    interface ed ripin ripout version 1 ;
    traceoptions route ;
} ;

#
# Turn on a bunch of tracing info for the interface to the kernel:
kernel {
```

```

    traceoptions remnants request routes info interface ;
} ;

#
# Propagate the route to xxx.xxx.yy out the Ethernet interface via RIP
#

export proto rip interface ed {
    proto direct {
        xxx.xxx.yy mask 255.255.252.0 metric 1; # SLIP connections
    } ;
} ;

#
# Accept routes from RIP via ed Ethernet interfaces

import proto rip interface ed {
    all ;
} ;

```

The above sample `gated.conf` file broadcasts routing information regarding the SLIP subnet `xxx.xxx.yy` via RIP onto the Ethernet; if you are using a different Ethernet driver than the `ed` driver, you will need to change the references to the `ed` interface appropriately. This sample file also sets up tracing to `/var/tmp/gated.output` for debugging **GateD**'s activity; you can certainly turn off the tracing options if **GateD** works OK for you. You will need to change the `xxx.xxx.yy`'s into the network address of your own SLIP subnet (be sure to change the net mask in the `proto direct` clause as well).

Once you have installed and configured **GateD** on your system, you will need to tell the DragonFly startup scripts to run **GateD** in place of **routed**. The easiest way to accomplish this is to set the `router` and `router_flags` variables in `/etc/rc.conf`. Please see the manual page for **GateD** for information on command-line parameters.

Chapter 19 Advanced Networking

19.1 Synopsis

This chapter will cover some of the more frequently used network services on UNIX systems. We will cover how to define, set up, test and maintain all of the network services that DragonFly utilizes. In addition, there have been example configuration files included throughout this chapter for you to benefit from.

After reading this chapter, you will know:

- The basics of gateways and routes.
- How to set up IEEE 802.11 and Bluetooth® devices.
- How to make DragonFly act as a bridge.
- How to set up a network filesystem.
- How to set up network booting on a diskless machine.
- How to set up a network information server for sharing user accounts.
- How to set up automatic network settings using DHCP.
- How to set up a domain name server.
- How to synchronize the time and date, and set up a time server, with the NTP protocol.
- How to set up network address translation.
- How to manage the **inetd** daemon.
- How to connect two computers via PLIP.
- How to set up IPv6 on a DragonFly machine.

Before reading this chapter, you should:

- Understand the basics of the `/etc/rc` scripts.
- Be familiar with basic network terminology.

19.2 Gateways and Routes

Contributed by Coranth Gryphon.

For one machine to be able to find another over a network, there must be a mechanism in place to describe how to get from one to the other. This is called *routing*. A “route” is a defined pair of addresses: a “destination” and a “gateway”. The pair indicates that if you are trying to get to this *destination*, communicate through this *gateway*. There are three types of destinations: individual hosts, subnets, and “default”. The “default route” is used if none of the other routes apply. We will talk a little bit more about default routes later on. There are also three types of gateways: individual hosts, interfaces (also called “links”), and Ethernet hardware addresses (MAC addresses).

19.2.1 An Example

To illustrate different aspects of routing, we will use the following example from `netstat`:

```
% netstat -r
Routing tables

Destination      Gateway          Flags    Refs     Use      Netif  Expire
default          outside-gw      UGSc     37       418     ppp0
localhost        localhost       UH        0        181     lo0
test0            0:e0:b5:36:cf:4f UHLW     5       63288   ed0    77
10.20.30.255     link#1         UHLW     1        2421
example.com      link#1         UC        0         0
host1            0:e0:a8:37:8:1e UHLW     3       4601   lo0
host2            0:e0:a8:37:8:1e UHLW     0         5     lo0 =>
host2.example.com link#1         UC        0         0
224              link#1         UC        0         0
```

The first two lines specify the default route (which we will cover in the next section) and the `localhost` route.

The interface (`Netif` column) that this routing table specifies to use for `localhost` is `lo0`, also known as the loopback device. This says to keep all traffic for this destination internal, rather than sending it out over the LAN, since it will only end up back where it started.

The next thing that stands out are the addresses beginning with `0:e0:`. These are Ethernet hardware addresses, which are also known as MAC addresses. DragonFly will automatically identify any hosts (`test0` in the example) on the local Ethernet and add a route for that host, directly to it over the Ethernet interface, `ed0`. There is also a timeout (`Expire` column) associated with this type of route, which is used if we fail to hear from the host in a specific amount of time. When this happens, the route to this host will be automatically deleted. These hosts are identified using a mechanism known as RIP (Routing Information Protocol), which figures out routes to local hosts based upon a shortest path determination.

DragonFly will also add subnet routes for the local subnet (`10.20.30.255` is the broadcast address for the subnet `10.20.30`, and `example.com` is the domain name associated with that subnet). The designation `link#1` refers to the first Ethernet card in the machine. You will notice no additional interface is specified for those.

Both of these groups (local network hosts and local subnets) have their routes automatically configured by a daemon called **routed**. If this is not run, then only routes which are statically defined (i.e. entered explicitly) will exist.

The `host1` line refers to our host, which it knows by Ethernet address. Since we are the sending host, DragonFly knows to use the loopback interface (`lo0`) rather than sending it out over the Ethernet interface.

The two `host2` lines are an example of what happens when we use an `ifconfig(8)` alias (see the section on Ethernet for reasons why we would do this). The `=>` symbol after the `lo0` interface says that not only are we using the loopback (since this address also refers to the local host), but specifically it is an alias. Such routes only show up on the host that supports the alias; all other hosts on the local network will simply have a `link#1` line for such routes.

The final line (destination subnet `224`) deals with multicasting, which will be covered in another section.

Finally, various attributes of each route can be seen in the `Flags` column. Below is a short table of some of these flags and their meanings:

U	Up: The route is active.
H	Host: The route destination is a single host.

G	Gateway: Send anything for this destination on to this remote system, which will figure out from there where to send it.
S	Static: This route was configured manually, not automatically generated by the system.
C	Clone: Generates a new route based upon this route for machines we connect to. This type of route is normally used for local networks.
W	WasCloned: Indicated a route that was auto-configured based upon a local area network (Clone) route.
L	Link: Route involves references to Ethernet hardware.


19.2.2 Default Routes

When the local system needs to make a connection to a remote host, it checks the routing table to determine if a known path exists. If the remote host falls into a subnet that we know how to reach (Cloned routes), then the system checks to see if it can connect along that interface.

If all known paths fail, the system has one last option: the “default” route. This route is a special type of gateway route (usually the only one present in the system), and is always marked with a `c` in the flags field. For hosts on a local area network, this gateway is set to whatever machine has a direct connection to the outside world (whether via PPP link, DSL, cable modem, T1, or another network interface).

If you are configuring the default route for a machine which itself is functioning as the gateway to the outside world, then the default route will be the gateway machine at your Internet Service Provider’s (ISP) site.

Let us look at an example of default routes. This is a common configuration:



The hosts `Local1` and `Local2` are at your site. `Local1` is connected to an ISP via a dial up PPP connection. This PPP server computer is connected through a local area network to another gateway computer through an external interface to the ISP’s Internet feed.

The default routes for each of your machines will be:

Host	Default Gateway	Interface
<code>Local2</code>	<code>Local1</code>	Ethernet
<code>Local1</code>	<code>T1-GW</code>	PPP

A common question is “Why (or how) would we set the `T1-GW` to be the default gateway for `Local1`, rather than the ISP server it is connected to?”.

Remember, since the PPP interface is using an address on the ISP’s local network for your side of the connection, routes for any other machines on the ISP’s local network will be automatically generated. Hence, you will already know how to reach the `T1-GW` machine, so there is no need for the intermediate step of sending traffic to the ISP

server.

It is common to use the address `x.x.x.1` as the gateway address for your local network. So (using the same example), if your local class-C address space was `10.20.30` and your ISP was using `10.9.9` then the default routes would be:

Host	Default Route
Local2 (10.20.30.2)	Local1 (10.20.30.1)
Local1 (10.20.30.1, 10.9.9.30)	T1-GW (10.9.9.1)

You can easily define the default route via the `/etc/rc.conf` file. In our example, on the `Local2` machine, we added the following line in `/etc/rc.conf`:

```
defaultrouter="10.20.30.1"
```

It is also possible to do it directly from the command line with the `route(8)` command:

```
# route add default 10.20.30.1
```

For more informations on manual manipulation of network routing tables, consult `route(8)` manual page.

19.2.3 Dual Homed Hosts

There is one other type of configuration that we should cover, and that is a host that sits on two different networks. Technically, any machine functioning as a gateway (in the example above, using a PPP connection) counts as a dual-homed host. But the term is really only used to refer to a machine that sits on two local-area networks.

In one case, the machine has two Ethernet cards, each having an address on the separate subnets. Alternately, the machine may only have one Ethernet card, and be using `ifconfig(8)` aliasing. The former is used if two physically separate Ethernet networks are in use, the latter if there is one physical network segment, but two logically separate subnets.

Either way, routing tables are set up so that each subnet knows that this machine is the defined gateway (inbound route) to the other subnet. This configuration, with the machine acting as a router between the two subnets, is often used when we need to implement packet filtering or firewall security in either or both directions.

If you want this machine to actually forward packets between the two interfaces, you need to tell DragonFly to enable this ability. See the next section for more details on how to do this.

19.2.4 Building a Router

A network router is simply a system that forwards packets from one interface to another. This is not enabled by default in DragonFly. You can enable this feature by changing the following variable to `YES` in `rc.conf(5)`:

```
gateway_enable=YES          # Set to YES if this host will be a gateway
```

This option will set the `sysctl(8)` variable `net.inet.ip.forwarding` to `1`. If you should need to stop routing temporarily, you can reset this to `0` temporarily.

Your new router will need routes to know where to send the traffic. If your network is simple enough you can use static routes. DragonFly also comes with the standard BSD routing daemon `routed(8)`, which speaks RIP (both

version 1 and version 2) and IRDP. Support for BGP v4, OSPF v2, and other sophisticated routing protocols is available with the `net/zebra` package. Commercial products such as **GateD** are also available for more complex network routing solutions.

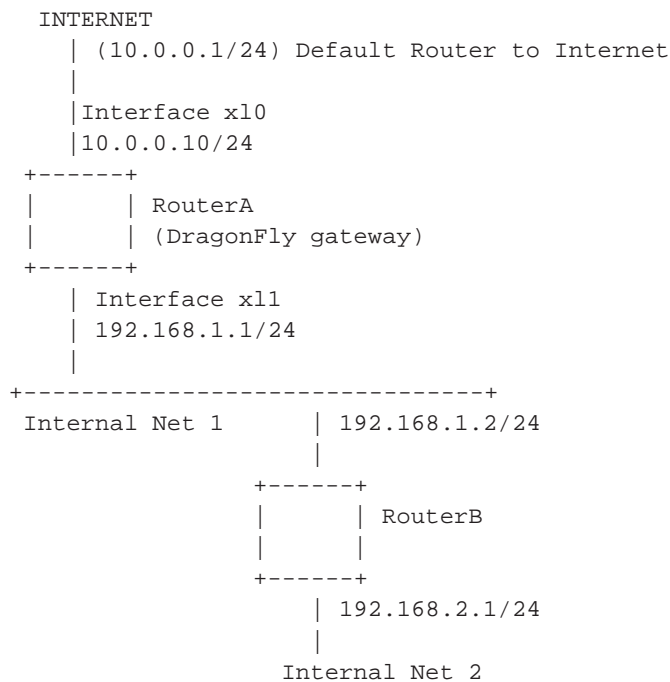
Even when DragonFly is configured in this way, it does not completely comply with the Internet standard requirements for routers. It comes close enough for ordinary use, however.

19.2.5 Setting Up Static Routes

Contributed by Al Hoang.

19.2.5.1 Manual Configuration

Let us assume we have a network as follows:



In this scenario, RouterA is our DragonFly machine that is acting as a router to the rest of the Internet. It has a default route set to `10.0.0.1` which allows it to connect with the outside world. We will assume that RouterB is already configured properly and knows how to get wherever it needs to go. (This is simple in this picture. Just add a default route on RouterB using `192.168.1.1` as the gateway.)

If we look at the routing table for RouterA we would see something like the following:

```

% netstat -nr
Routing tables

Internet:
Destination      Gateway          Flags    Refs      Use  Netif  Expire
default          10.0.0.1        UGS      0         49378  x10
127.0.0.1       127.0.0.1      UH       0          6     lo0

```



```

10.0.0/24          link#1          UC          0          0          x10
192.168.1/24      link#2          UC          0          0          x11

```

With the current routing table RouterA will not be able to reach our Internal Net 2. It does not have a route for 192.168.2.0/24. One way to alleviate this is to manually add the route. The following command would add the Internal Net 2 network to RouterA's routing table using 192.168.1.2 as the next hop:

```
# route add -net 192.168.2.0/24 192.168.1.2
```

Now RouterA can reach any hosts on the 192.168.2.0/24 network.

19.2.5.2 Persistent Configuration

The above example is perfect for configuring a static route on a running system. However, one problem is that the routing information will not persist if you reboot your DragonFly machine. The way to handle the addition of a static route is to put it in your `/etc/rc.conf` file:

```
# Add Internal Net 2 as a static route
static_routes="internalnet2"
route_internalnet2="-net 192.168.2.0/24 192.168.1.2"
```

The `static_routes` configuration variable is a list of strings separated by a space. Each string references to a route name. In our above example we only have one string in `static_routes`. This string is `internalnet2`. We then add a configuration variable called `route_internalnet2` where we put all of the configuration parameters we would give to the `route(8)` command. For our example above we would have used the command:

```
# route add -net 192.168.2.0/24 192.168.1.2
```

so we need `"-net 192.168.2.0/24 192.168.1.2"`.

As said above, we can have more than one string in `static_routes`. This allows us to create multiple static routes. The following lines shows an example of adding static routes for the 192.168.0.0/24 and 192.168.1.0/24 networks on an imaginary router:

```
static_routes="net1 net2"
route_net1="-net 192.168.0.0/24 192.168.0.1"
route_net2="-net 192.168.1.0/24 192.168.1.1"
```

19.2.6 Routing Propagation

We have already talked about how we define our routes to the outside world, but not about how the outside world finds us.

We already know that routing tables can be set up so that all traffic for a particular address space (in our examples, a class-C subnet) can be sent to a particular host on that network, which will forward the packets inbound.

When you get an address space assigned to your site, your service provider will set up their routing tables so that all traffic for your subnet will be sent down your PPP link to your site. But how do sites across the country know to send to your ISP?

There is a system (much like the distributed DNS information) that keeps track of all assigned address-spaces, and defines their point of connection to the Internet Backbone. The “Backbone” are the main trunk lines that carry Internet traffic across the country, and around the world. Each backbone machine has a copy of a master set of tables, which direct traffic for a particular network to a specific backbone carrier, and from there down the chain of service providers until it reaches your network.

It is the task of your service provider to advertise to the backbone sites that they are the point of connection (and thus the path inward) for your site. This is known as route propagation.

19.2.7 Troubleshooting

Sometimes, there is a problem with routing propagation, and some sites are unable to connect to you. Perhaps the most useful command for trying to figure out where routing is breaking down is the `traceroute(8)` command. It is equally useful if you cannot seem to make a connection to a remote machine (i.e. `ping(8)` fails).

The `traceroute(8)` command is run with the name of the remote host you are trying to connect to. It will show the gateway hosts along the path of the attempt, eventually either reaching the target host, or terminating because of a lack of connection.

For more information, see the manual page for `traceroute(8)`.

19.2.8 Multicast Routing

DragonFly supports both multicast applications and multicast routing natively. Multicast applications do not require any special configuration of DragonFly; applications will generally run out of the box. Multicast routing requires that support be compiled into the kernel:

```
options MROUTING
```

In addition, the multicast routing daemon, `mROUTED(8)` must be configured to set up tunnels and DVMRP via `/etc/mROUTED.conf`. More details on multicast configuration may be found in the manual page for `mROUTED(8)`.

19.3 Wireless Networking

Written by Eric Anderson.

19.3.1 Introduction

It can be very useful to be able to use a computer without the annoyance of having a network cable attached at all times. DragonFly can be used as a wireless client, and even as a wireless “access point”.

19.3.2 Wireless Modes of Operation

There are two different ways to configure 802.11 wireless devices: BSS and IBSS.

19.3.2.1 BSS Mode

BSS mode is the mode that typically is used. BSS mode is also called infrastructure mode. In this mode, a number of wireless access points are connected to a wired network. Each wireless network has its own name. This name is called the SSID of the network.

Wireless clients connect to these wireless access points. The IEEE 802.11 standard defines the protocol that wireless networks use to connect. A wireless client can be tied to a specific network, when a SSID is set. A wireless client can also attach to any network by not explicitly setting a SSID.

19.3.2.2 IBSS Mode

IBSS mode, also called ad-hoc mode, is designed for point to point connections. There are actually two types of ad-hoc mode. One is IBSS mode, also called ad-hoc or IEEE ad-hoc mode. This mode is defined by the IEEE 802.11 standards. The second is called demo ad-hoc mode or Lucent ad-hoc mode (and sometimes, confusingly, ad-hoc mode). This is the old, pre-802.11 ad-hoc mode and should only be used for legacy installations. We will not cover either of the ad-hoc modes further.

19.3.3 Infrastructure Mode

19.3.3.1 Access Points

Access points are wireless networking devices that allow one or more wireless clients to use the device as a central hub. When using an access point, all clients communicate through the access point. Multiple access points are often used to cover a complete area such as a house, business, or park with a wireless network.

Access points typically have multiple network connections: the wireless card, and one or more wired Ethernet adapters for connection to the rest of the network.

Access points can either be purchased prebuilt, or you can build your own with DragonFly and a supported wireless card. Several vendors make wireless access points and wireless cards with various features.

19.3.3.2 Building a DragonFly Access Point

19.3.3.2.1 Requirements

In order to set up a wireless access point with DragonFly, you need to have a compatible wireless card. Currently, only cards with the Prism chipset are supported. You will also need a wired network card that is supported by DragonFly (this should not be difficult to find, DragonFly supports a lot of different devices). For this guide, we will assume you want to bridge(4) all traffic between the wireless device and the network attached to the wired network card.

The hostap functionality that DragonFly uses to implement the access point works best with certain versions of firmware. Prism 2 cards should use firmware version 1.3.4 or newer. Prism 2.5 and Prism 3 cards should use firmware 1.4.9. Older versions of the firmware may or may not function correctly. At this time, the only way to update cards is with Windows firmware update utilities available from your card's manufacturer.

19.3.3.2.2 Setting It Up

First, make sure your system can see the wireless card:

```
# ifconfig -a
wi0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    inet6 fe80::202:2dff:fe2d:c938%wi0 prefixlen 64 scopeid 0x7
    inet 0.0.0.0 netmask 0xff000000 broadcast 255.255.255.255
    ether 00:09:2d:2d:c9:50
    media: IEEE 802.11 Wireless Ethernet autoselect (DS/2Mbps)
    status: no carrier
    ssid ""
    stationname "DragonFly Wireless node"
    channel 10 authmode OPEN powersavemode OFF powersavesleep 100
    wepmode OFF weptxkey 1
```

Do not worry about the details now, just make sure it shows you something to indicate you have a wireless card installed. If you have trouble seeing the wireless interface, and you are using a PC Card, you may want to check out `pccardc(8)` and `pccardd(8)` manual pages for more information.

Next, you will need to load a module in order to get the bridging part of DragonFly ready for the access point. To load the `bridge(4)` module, simply run the following command:

```
# kldload bridge
```

It should not have produced any errors when loading the module. If it did, you may need to compile the `bridge(4)` code into your kernel. The Bridging section of this handbook should be able to help you accomplish that task.

Now that you have the bridging stuff done, we need to tell the DragonFly kernel which interfaces to bridge together. We do that by using `sysctl(8)`:

```
# sysctl net.link.ether.bridge=1
# sysctl net.link.ether.bridge_cfg="wi0,x10"
# sysctl net.inet.ip.forwarding=1
```

Now it is time for the wireless card setup. The following command will set the card into an access point:

```
# ifconfig wi0 ssid my_net channel 11 media DS/11Mbps mediaopt hostap up stationname "DragonFly AP"
```

The `ifconfig(8)` line brings the `wi0` interface up, sets its SSID to `my_net`, and sets the station name to `DragonFly AP`. The `media DS/11Mbps` sets the card into 11Mbps mode and is needed for any `mediaopt` to take effect. The `mediaopt hostap` option places the interface into access point mode. The `channel 11` option sets the 802.11b channel to use. The `wicontrol(8)` manual page has valid channel options for your regulatory domain.

Now you should have a complete functioning access point up and running. You are encouraged to read `wicontrol(8)`, `ifconfig(8)`, and `wi(4)` for further information.

It is also suggested that you read the section on encryption that follows.

19.3.3.2.3 Status Information

Once the access point is configured and operational, operators will want to see the clients that are associated with the access point. At any time, the operator may type:

```
# wicontrol -l
l station:
00:09:b7:7b:9d:16  asid=04c0, flags=3<ASSOC,AUTH>, caps=1<ESS>, rates=f<1M,2M,5.5M,11M>, sig=38
```

This shows that there is one station associated, along with its parameters. The signal indicated should be used as a relative indication of strength only. Its translation to dBm or other units varies between different firmware revisions.

19.3.3.3 Clients

A wireless client is a system that accesses an access point or another client directly.

Typically, wireless clients only have one network device, the wireless networking card.

There are a few different ways to configure a wireless client. These are based on the different wireless modes, generally BSS (infrastructure mode, which requires an access point), and IBSS (ad-hoc, or peer-to-peer mode). In our example, we will use the most popular of the two, BSS mode, to talk to an access point.

19.3.3.3.1 Requirements

There is only one real requirement for setting up DragonFly as a wireless client. You will need a wireless card that is supported by DragonFly.

19.3.3.3.2 Setting Up a Wireless DragonFly Client

You will need to know a few things about the wireless network you are joining before you start. In this example, we are joining a network that has a name of *my_net*, and encryption turned off.

Note: In this example, we are not using encryption, which is a dangerous situation. In the next section, you will learn how to turn on encryption, why it is important to do so, and why some encryption technologies still do not completely protect you.

Make sure your card is recognized by DragonFly:

```
# ifconfig -a
wi0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    inet6 fe80::202:2dff:fe2d:c938%wi0 prefixlen 64 scopeid 0x7
    inet 0.0.0.0 netmask 0xff000000 broadcast 255.255.255.255
    ether 00:09:2d:2d:c9:50
    media: IEEE 802.11 Wireless Ethernet autoselect (DS/2Mbps)
    status: no carrier
    ssid ""
    stationname "DragonFly Wireless node"
    channel 10 authmode OPEN powersavemode OFF powersavesleep 100
    wepmode OFF weptxkey 1
```

Now, we can set the card to the correct settings for our network:

```
# ifconfig wi0 inet 192.168.0.20 netmask 255.255.255.0 ssid my_net
```

Replace `192.168.0.20` and `255.255.255.0` with a valid IP address and netmask on your wired network. Remember, our access point is bridging the data between the wireless network, and the wired network, so it will appear to the other devices on your network that you are on the wired network just as they are.

Once you have done that, you should be able to ping hosts on the wired network just as if you were connected using a standard wired connection.

If you are experiencing problems with your wireless connection, check to make sure that you are associated (connected) to the access point:

```
# ifconfig wi0
```

should return some information, and you should see:

```
status: associated
```

If it does not show `associated`, then you may be out of range of the access point, have encryption on, or possibly have a configuration problem.

19.3.3.4 Encryption

Encryption on a wireless network is important because you no longer have the ability to keep the network contained in a well protected area. Your wireless data will be broadcast across your entire neighborhood, so anyone who cares to read it can. This is where encryption comes in. By encrypting the data that is sent over the airwaves, you make it much more difficult for any interested party to grab your data right out of the air.

The two most common ways to encrypt the data between your client and the access point are WEP, and ipsec(4).

19.3.3.4.1 WEP

WEP is an abbreviation for Wired Equivalency Protocol. WEP is an attempt to make wireless networks as safe and secure as a wired network. Unfortunately, it has been cracked, and is fairly trivial to break. This also means it is not something to rely on when it comes to encrypting sensitive data.

It is better than nothing, so use the following to turn on WEP on your new DragonFly access point:

```
# ifconfig wi0 inet up ssid my_net wepmode on wepkey 0x1234567890 media DS/11Mbps mediaopt hostap
```

And you can turn on WEP on a client with this command:

```
# ifconfig wi0 inet 192.168.0.20 netmask 255.255.255.0 ssid my_net wepmode on wepkey 0x1234567890
```

Note that you should replace the `0x1234567890` with a more unique key.

19.3.3.4.2 IPsec

ipsec(4) is a much more robust and powerful tool for encrypting data across a network. This is definitely the preferred way to encrypt data over a wireless network. You can read more about ipsec(4) security and how to implement it in the IPsec section of this handbook.

19.3.3.5 Tools

There are a small number of tools available for use in debugging and setting up your wireless network, and here we will attempt to describe some of them and what they do.

19.3.3.5.1 The *wicontrol*, *ancontrol* and *raycontrol* Utilities

These are the tools you can use to control how your wireless card behaves on the wireless network. In the examples above, we have chosen to use `wicontrol(8)`, since our wireless card is a `wi0` interface. If you had a Cisco wireless device, it would come up as `an0`, and therefore you would use `ancontrol(8)`.

19.3.3.5.2 The *ifconfig* Command

The `ifconfig(8)` command can be used to do many of the same options as `wicontrol(8)`, however it does lack a few options. Check `ifconfig(8)` for command line parameters and options.

19.3.3.6 Supported Cards

19.3.3.6.1 Access Points

The only cards that are currently supported for BSS (as an access point) mode are devices based on the Prism 2, 2.5, or 3 chipsets. For a complete list, look at `wi(4)`.

19.3.3.6.2 Clients

Almost all 802.11b wireless cards are currently supported under DragonFly. Most cards based on Prism, Spectrum24, Hermes, Aironet, and Raylink will work as a wireless network card in IBSS (ad-hoc, peer-to-peer, and BSS) mode.

19.4 Bluetooth

Written by Pav Lucistnik.

19.4.1 Introduction

Bluetooth is a wireless technology for creating personal networks operating in the 2.4 GHz unlicensed band, with a range of 10 meters. Networks are usually formed ad-hoc from portable devices such as cellular phones, handhelds and laptops. Unlike the other popular wireless technology, Wi-Fi, Bluetooth offers higher level service profiles, e.g. FTP-like file servers, file pushing, voice transport, serial line emulation, and more.

The Bluetooth stack in DragonFly is implemented using the Netgraph framework (see `netgraph(4)`). A broad variety of Bluetooth USB dongles is supported by the `ng_ubt(4)` driver. The Broadcom BCM2033 chip based Bluetooth devices are supported via the `ubtbcmfw(4)` and `ng_ubt(4)` drivers. The 3Com Bluetooth PC Card 3CRWB60-A is supported by the `ng_bt3c(4)` driver. Serial and UART based Bluetooth devices are supported via `sio(4)`, `ng_h4(4)` and

hcseriald(8). This chapter describes the use of the USB Bluetooth dongle. Bluetooth support is available in DragonFly 5.0 and newer systems.

19.4.2 Plugging in the Device

By default Bluetooth device drivers are available as kernel modules. Before attaching a device, you will need to load the driver into the kernel.

```
# kldload ng_ubt
```

If the Bluetooth device is present in the system during system startup, load the module from `/boot/loader.conf`.

```
ng_ubt_load="YES"
```

Plug in your USB dongle. The output similar to the following will appear on the console (or in syslog).

```
ubt0: vendor 0x0a12 product 0x0001, rev 1.10/5.25, addr 2
ubt0: Interface 0 endpoints: interrupt=0x81, bulk-in=0x82, bulk-out=0x2
ubt0: Interface 1 (alt.config 5) endpoints: isoc-in=0x83, isoc-out=0x3,
      wMaxPacketSize=49, nframes=6, buffer size=294
```

Copy `/usr/share/examples/netgraph/bluetooth/rc.bluetooth` into some convenient place, like `/etc/rc.bluetooth`. This script is used to start and stop the Bluetooth stack. It is a good idea to stop the stack before unplugging the device, but it is not (usually) fatal. When starting the stack, you will receive output similar to the following:

```
# /etc/rc.bluetooth start ubt0
BD_ADDR: 00:02:72:00:d4:1a
Features: 0xff 0xff 0xf 00 00 00 00 00
<3-Slot> <5-Slot> <Encryption> <Slot offset>
<Timing accuracy> <Switch> <Hold mode> <Sniff mode>
<Park mode> <RSSI> <Channel quality> <SCO link>
<HV2 packets> <HV3 packets> <u-law log> <A-law log> <CVSD>
<Paging scheme> <Power control> <Transparent SCO data>
Max. ACL packet size: 192 bytes
Number of ACL packets: 8
Max. SCO packet size: 64 bytes
Number of SCO packets: 8
```

19.4.3 Host Controller Interface (HCI)

Host Controller Interface (HCI) provides a command interface to the baseband controller and link manager, and access to hardware status and control registers. This interface provides a uniform method of accessing the Bluetooth baseband capabilities. HCI layer on the Host exchanges data and commands with the HCI firmware on the Bluetooth hardware. The Host Controller Transport Layer (i.e. physical bus) driver provides both HCI layers with the ability to exchange information with each other.

A single Netgraph node of type `hci` is created for a single Bluetooth device. The HCI node is normally connected to the Bluetooth device driver node (downstream) and the L2CAP node (upstream). All HCI operations must be

performed on the HCI node and not on the device driver node. Default name for the HCI node is “devicehci”. For more details refer to the `ng_hci(4)` man page.

One of the most common tasks is discovery of Bluetooth devices in RF proximity. This operation is called *inquiry*. Inquiry and other HCI related operations are done with the `hccontrol(8)` utility. The example below shows how to find out which Bluetooth devices are in range. You should receive the list of devices in a few seconds. Note that a remote device will only answer the inquiry if it put into *discoverable* mode.

```
% hccontrol -n ubt0hci inquiry
Inquiry result, num_responses=1
Inquiry result #0
    BD_ADDR: 00:80:37:29:19:a4
    Page Scan Rep. Mode: 0x1
    Page Scan Period Mode: 00
    Page Scan Mode: 00
    Class: 52:02:04
    Clock offset: 0x78ef
Inquiry complete. Status: No error [00]
```

`BD_ADDR` is unique address of a Bluetooth device, similar to MAC addresses of a network card. This address is needed for further communication with a device. It is possible to assign human readable name to a `BD_ADDR`. The `/etc/bluetooth/hosts` file contains information regarding the known Bluetooth hosts. The following example shows how to obtain human readable name that was assigned to the remote device.

```
% hccontrol -n ubt0hci remote_name_request 00:80:37:29:19:a4
BD_ADDR: 00:80:37:29:19:a4
Name: Pav's T39
```

If you perform an inquiry on a remote Bluetooth device, it will find your computer as “your.host.name (ubt0)”. The name assigned to the local device can be changed at any time.

The Bluetooth system provides a point-to-point connection (only two Bluetooth units involved), or a point-to-multipoint connection. In the point-to-multipoint connection the connection is shared among several Bluetooth devices. The following example shows how to obtain the list of active baseband connections for the local device.

```
% hccontrol -n ubt0hci read_connection_list
Remote BD_ADDR    Handle Type Mode Role Encrypt Pending Queue State
00:80:37:29:19:a4    41  ACL    0 MAST  NONE      0      0 OPEN
```

A *connection handle* is useful when termination of the baseband connection is required. Note, that it is normally not required to do it by hand. The stack will automatically terminate inactive baseband connections.

```
# hccontrol -n ubt0hci disconnect 41
Connection handle: 41
Reason: Connection terminated by local host [0x16]
```

Refer to `hccontrol help` for a complete listing of available HCI commands. Most of the HCI commands do not require superuser privileges.

19.4.4 Logical Link Control and Adaptation Protocol (L2CAP)

Logical Link Control and Adaptation Protocol (L2CAP) provides connection-oriented and connectionless data services to upper layer protocols with protocol multiplexing capability and segmentation and reassembly operation. L2CAP permits higher level protocols and applications to transmit and receive L2CAP data packets up to 64 kilobytes in length.

L2CAP is based around the concept of *channels*. Channel is a logical connection on top of baseband connection. Each channel is bound to a single protocol in a many-to-one fashion. Multiple channels can be bound to the same protocol, but a channel cannot be bound to multiple protocols. Each L2CAP packet received on a channel is directed to the appropriate higher level protocol. Multiple channels can share the same baseband connection.

A single Netgraph node of type *l2cap* is created for a single Bluetooth device. The L2CAP node is normally connected to the Bluetooth HCI node (downstream) and Bluetooth sockets nodes (upstream). Default name for the L2CAP node is “devicel2cap”. For more details refer to the `ng_l2cap(4)` man page.

A useful command is `l2ping(8)`, which can be used to ping other devices. Some Bluetooth implementations might not return all of the data sent to them, so *0 bytes* in the following example is normal.

```
# l2ping -a 00:80:37:29:19:a4
0 bytes from 0:80:37:29:19:a4 seq_no=0 time=48.633 ms result=0
0 bytes from 0:80:37:29:19:a4 seq_no=1 time=37.551 ms result=0
0 bytes from 0:80:37:29:19:a4 seq_no=2 time=28.324 ms result=0
0 bytes from 0:80:37:29:19:a4 seq_no=3 time=46.150 ms result=0
```

The `l2control(8)` utility is used to perform various operations on L2CAP nodes. This example shows how to obtain the list of logical connections (channels) and the list of baseband connections for the local device.

```
% l2control -a 00:02:72:00:d4:1a read_channel_list
L2CAP channels:
Remote BD_ADDR      SCID/ DCID   PSM  IMTU/ OMTU State
00:07:e0:00:0b:ca   66/   64     3   132/  672 OPEN
% l2control -a 00:02:72:00:d4:1a read_connection_list
L2CAP connections:
Remote BD_ADDR      Handle Flags Pending State
00:07:e0:00:0b:ca   41  O           0  OPEN
```

Another diagnostic tool is `btsockstat(1)`. It does a job similar to as `netstat(1)` does, but for Bluetooth network-related data structures. The example below shows the same logical connection as `l2control(8)` above.

```
% btsockstat
Active L2CAP sockets
PCB      Recv-Q Send-Q Local address/PSM      Foreign address  CID   State
c2afe900  0      0  00:02:72:00:d4:1a/3   00:07:e0:00:0b:ca 66    OPEN
Active RFCOMM sessions
L2PCB    PCB      Flag MTU   Out-Q DLCs State
c2afe900 c2b53380 1    127    0    Yes  OPEN
Active RFCOMM sockets
PCB      Recv-Q Send-Q Local address      Foreign address  Chan DLCI State
c2e8bc80  0      250  00:02:72:00:d4:1a  00:07:e0:00:0b:ca 3     6    OPEN
```

19.4.5 RFCOMM Protocol

The RFCOMM protocol provides emulation of serial ports over the L2CAP protocol. The protocol is based on the ETSI standard TS 07.10. RFCOMM is a simple transport protocol, with additional provisions for emulating the 9 circuits of RS-232 (EIA/TIA-232-E) serial ports. The RFCOMM protocol supports up to 60 simultaneous connections (RFCOMM channels) between two Bluetooth devices.

For the purposes of RFCOMM, a complete communication path involves two applications running on different devices (the communication endpoints) with a communication segment between them. RFCOMM is intended to cover applications that make use of the serial ports of the devices in which they reside. The communication segment is a Bluetooth link from one device to another (direct connect).

RFCOMM is only concerned with the connection between the devices in the direct connect case, or between the device and a modem in the network case. RFCOMM can support other configurations, such as modules that communicate via Bluetooth wireless technology on one side and provide a wired interface on the other side.

In DragonFly the RFCOMM protocol is implemented at the Bluetooth sockets layer.

19.4.6 Pairing of Devices

By default, Bluetooth communication is not authenticated, and any device can talk to any other device. A Bluetooth device (for example, cellular phone) may choose to require authentication to provide a particular service (for example, Dial-Up service). Bluetooth authentication is normally done with *PIN codes*. A PIN code is an ASCII string up to 16 characters in length. User is required to enter the same PIN code on both devices. Once user has entered the PIN code, both devices will generate a *link key*. After that the link key can be stored either in the devices themselves or in a persistent storage. Next time both devices will use previously generated link key. The described above procedure is called *pairing*. Note that if the link key is lost by any device then pairing must be repeated.

The `hcsecd(8)` daemon is responsible for handling of all Bluetooth authentication requests. The default configuration file is `/etc/bluetooth/hcsecd.conf`. An example section for a cellular phone with the PIN code arbitrarily set to "1234" is shown below.

```
device {
    bdaddr 00:80:37:29:19:a4;
    name   "Pav's T39";
    key    nokey;
    pin    "1234";
}
```

There is no limitation on PIN codes (except length). Some devices (for example Bluetooth headsets) may have a fixed PIN code built in. The `-d` switch forces the `hcsecd(8)` daemon to stay in the foreground, so it is easy to see what is happening. Set the remote device to receive pairing and initiate the Bluetooth connection to the remote device. The remote device should say that pairing was accepted, and request the PIN code. Enter the same PIN code as you have in `hcsecd.conf`. Now your PC and the remote device are paired. Alternatively, you can initiate pairing on the remote device. Below in the sample `hcsecd` output.

```
hcsecd[16484]: Got Link_Key_Request event from 'ubt0hci', remote bdaddr 0:80:37:29:19:a4
hcsecd[16484]: Found matching entry, remote bdaddr 0:80:37:29:19:a4, name 'Pav's T39', link key d
hcsecd[16484]: Sending Link_Key_Negative_Reply to 'ubt0hci' for remote bdaddr 0:80:37:29:19:a4
hcsecd[16484]: Got PIN_Code_Request event from 'ubt0hci', remote bdaddr 0:80:37:29:19:a4
hcsecd[16484]: Found matching entry, remote bdaddr 0:80:37:29:19:a4, name 'Pav's T39', PIN code e
hcsecd[16484]: Sending PIN_Code_Reply to 'ubt0hci' for remote bdaddr 0:80:37:29:19:a4
```

19.4.7 Service Discovery Protocol (SDP)

The Service Discovery Protocol (SDP) provides the means for client applications to discover the existence of services provided by server applications as well as the attributes of those services. The attributes of a service include the type or class of service offered and the mechanism or protocol information needed to utilize the service.

SDP involves communication between a SDP server and a SDP client. The server maintains a list of service records that describe the characteristics of services associated with the server. Each service record contains information about a single service. A client may retrieve information from a service record maintained by the SDP server by issuing a SDP request. If the client, or an application associated with the client, decides to use a service, it must open a separate connection to the service provider in order to utilize the service. SDP provides a mechanism for discovering services and their attributes, but it does not provide a mechanism for utilizing those services.

Normally, a SDP client searches for services based on some desired characteristics of the services. However, there are times when it is desirable to discover which types of services are described by an SDP server's service records without any a priori information about the services. This process of looking for any offered services is called *browsing*.

The Bluetooth SDP server `sdpd(8)` and command line client `sdpcontrol(8)` are included in the standard DragonFly installation. The following example shows how to perform a SDP browse query.

```
% sdpcontrol -a 00:01:03:fc:6e:ec browse
Record Handle: 00000000
Service Class ID List:
    Service Discovery Server (0x1000)
Protocol Descriptor List:
    L2CAP (0x0100)
        Protocol specific parameter #1: u/int/uuid16 1
        Protocol specific parameter #2: u/int/uuid16 1

Record Handle: 0x00000001
Service Class ID List:
    Browse Group Descriptor (0x1001)

Record Handle: 0x00000002
Service Class ID List:
    LAN Access Using PPP (0x1102)
Protocol Descriptor List:
    L2CAP (0x0100)
    RFCOMM (0x0003)
        Protocol specific parameter #1: u/int8/bool 1
Bluetooth Profile Descriptor List:
    LAN Access Using PPP (0x1102) ver. 1.0
```

... and so on. Note that each service has a list of attributes (RFCOMM channel for example). Depending on the service you might need to make a note of some of the attributes. Some Bluetooth implementations do not support service browsing and may return an empty list. In this case it is possible to search for the specific service. The example below shows how to search for the OBEX Object Push (OPUSH) service.

```
% sdpcontrol -a 00:01:03:fc:6e:ec search OPUSH
```

Offering services on DragonFly to Bluetooth clients is done with the `sdpd(8)` server.

```
# sdpd
```

The local server application that wants to provide Bluetooth service to the remote clients will register service with the local SDP daemon. The example of such application is `rfcomm_pppd(8)`. Once started it will register Bluetooth LAN service with the local SDP daemon.

The list of services registered with the local SDP server can be obtained by issuing SDP browse query via local control channel.

```
# sdpcontrol -l browse
```

19.4.8 Dial-Up Networking (DUN) and Network Access with PPP (LAN) Profiles

The Dial-Up Networking (DUN) profile is mostly used with modems and cellular phones. The scenarios covered by this profile are the following:

- use of a cellular phone or modem by a computer as a wireless modem for connecting to a dial-up internet access server, or using other dial-up services;
- use of a cellular phone or modem by a computer to receive data calls.

Network Access with PPP (LAN) profile can be used in the following situations:

- LAN access for a single Bluetooth device;
- LAN access for multiple Bluetooth devices;
- PC to PC (using PPP networking over serial cable emulation).

In DragonFly both profiles are implemented with `ppp(8)` and `rfcomm_pppd(8)` - a wrapper that converts RFCOMM Bluetooth connection into something PPP can operate with. Before any profile can be used, a new PPP label in the `/etc/ppp/ppp.conf` must be created. Consult `rfcomm_pppd(8)` manual page for examples.

In the following example `rfcomm_pppd(8)` will be used to open RFCOMM connection to remote device with `BD_ADDR 00:80:37:29:19:a4` on DUN RFCOMM channel. The actual RFCOMM channel number will be obtained from the remote device via SDP. It is possible to specify RFCOMM channel by hand, and in this case `rfcomm_pppd(8)` will not perform SDP query. Use `sdpcontrol(8)` to find out RFCOMM channel on the remote device.

```
# rfcomm_pppd -a 00:80:37:29:19:a4 -c -C dun -l rfcomm-dialup
```

In order to provide Network Access with PPP (LAN) service the `sdpd(8)` server must be running. A new entry for LAN clients must be created in the `/etc/ppp/ppp.conf` file. Consult `rfcomm_pppd(8)` manual page for examples. Finally, start RFCOMM PPP server on valid RFCOMM channel number. The RFCOMM PPP server will automatically register Bluetooth LAN service with the local SDP daemon. The example below shows how to start RFCOMM PPP server.

```
# rfcomm_pppd -s -C 7 -l rfcomm-server
```

19.4.9 OBEX Object Push (OPUSH) Profile

OBEX is a widely used protocol for simple file transfers between mobile devices. Its main use is in infrared communication, where it is used for generic file transfers between notebooks or Palm handhelds, and for sending business cards or calendar entries between cellular phones and other devices with PIM applications.

The OBEX server and client are implemented as a third-party package **obexapp**, which is available as `comms/obexapp` port.

OBEX client is used to push and/or pull objects from the OBEX server. An object can, for example, be a business card or an appointment. The OBEX client can obtain RFCOMM channel number from the remote device via SDP. This can be done by specifying service name instead of RFCOMM channel number. Supported service names are: IrMC, FTRN and OPUSH. It is possible to specify RFCOMM channel as a number. Below is an example of an OBEX session, where device information object is pulled from the cellular phone, and a new object (business card) is pushed into the phone's directory.

```
% obexapp -a 00:80:37:29:19:a4 -C IrMC
obex> get
get: remote file> telecom/devinfo.txt
get: local file> devinfo-t39.txt
Success, response: OK, Success (0x20)
obex> put
put: local file> new.vcf
put: remote file> new.vcf
Success, response: OK, Success (0x20)
obex> di
Success, response: OK, Success (0x20)
```

In order to provide OBEX Object Push service, `sdpd(8)` server must be running. A root folder, where all incoming objects will be stored, must be created. The default path to the root folder is `/var/spool/obex`. Finally, start OBEX server on valid RFCOMM channel number. The OBEX server will automatically register OBEX Object Push service with the local SDP daemon. The example below shows how to start OBEX server.

```
# obexapp -s -C 10
```

19.4.10 Serial Port (SP) Profile

The Serial Port (SP) profile allows Bluetooth device to perform RS232 (or similar) serial cable emulation. The scenario covered by this profile deals with legacy applications using Bluetooth as a cable replacement, through a virtual serial port abstraction.

The `rfcomm_sppd(1)` utility implements the Serial Port profile. Pseudo tty is used as a virtual serial port abstraction. The example below shows how to connect to a remote device Serial Port service. Note that you do not have to specify RFCOMM channel - `rfcomm_sppd(1)` can obtain it from the remote device via SDP. If you would like to override this, specify RFCOMM channel in the command line.

```
# rfcomm_sppd -a 00:07:E0:00:0B:CA -t /dev/tty6
rfcomm_sppd[94692]: Starting on /dev/tty6...
```

Once connected pseudo tty can be used as serial port.

```
# cu -l tty6
```

19.4.11 Troubleshooting

19.4.11.1 A remote device cannot connect

Some older Bluetooth devices do not support role switching. By default, when DragonFly is accepting a new connection, it tries to perform role switch and become a master. Devices, which do not support this will not be able to connect. Note the role switching is performed when a new connection is being established, so it is not possible to ask the remote device if it does support role switching. There is a HCI option to disable role switching on the local side.

```
# hccontrol -n ubt0hci write_node_role_switch 0
```

19.4.11.2 Something is going wrong, can I see what exactly is happening?

Yes, you can. Use the **hcidump-1.5** third-party package that can be downloaded from here (http://www.geocities.com/m_evmenkin/). The **hcidump** utility is similar to `tcpdump(1)`. It can be used to display the content of the Bluetooth packets on the terminal and to dump the Bluetooth packets to a file.

19.5 Bridging

Written by Steve Peterson.

19.5.1 Introduction

It is sometimes useful to divide one physical network (such as an Ethernet segment) into two separate network segments without having to create IP subnets and use a router to connect the segments together. A device that connects two networks together in this fashion is called a “bridge”. A DragonFly system with two network interface cards can act as a bridge.

The bridge works by learning the MAC layer addresses (Ethernet addresses) of the devices on each of its network interfaces. It forwards traffic between two networks only when its source and destination are on different networks.

In many respects, a bridge is like an Ethernet switch with very few ports.

19.5.2 Situations Where Bridging Is Appropriate

There are two common situations in which a bridge is used today.

19.5.2.1 High Traffic on a Segment

Situation one is where your physical network segment is overloaded with traffic, but you do not want for whatever reason to subnet the network and interconnect the subnets with a router.

Let us consider an example of a newspaper where the Editorial and Production departments are on the same subnetwork. The Editorial users all use server A for file service, and the Production users are on server B. An Ethernet network is used to connect all users together, and high loads on the network are slowing things down.

If the Editorial users could be segregated on one network segment and the Production users on another, the two network segments could be connected with a bridge. Only the network traffic destined for interfaces on the “other” side of the bridge would be sent to the other network, reducing congestion on each network segment.

19.5.2.2 Filtering/Traffic Shaping Firewall

The second common situation is where firewall functionality is needed without network address translation (NAT).

An example is a small company that is connected via DSL or ISDN to their ISP. They have a 13 globally-accessible IP addresses from their ISP and have 10 PCs on their network. In this situation, using a router-based firewall is difficult because of subnetting issues.

A bridge-based firewall can be configured and dropped into the path just downstream of their DSL/ISDN router without any IP numbering issues.

19.5.3 Configuring a Bridge

19.5.3.1 Network Interface Card Selection

A bridge requires at least two network cards to function. Unfortunately, not all network interface cards support bridging. Read `bridge(4)` for details on the cards that are supported.

Install and test the two network cards before continuing.

19.5.3.2 Kernel Configuration Changes

To enable kernel support for bridging, add the:

```
options BRIDGE
```

statement to your kernel configuration file, and rebuild your kernel.

19.5.3.3 Firewall Support

If you are planning to use the bridge as a firewall, you will need to add the `IPFIREWALL` option as well. Read Section 10.7 for general information on configuring the bridge as a firewall.

If you need to allow non-IP packets (such as ARP) to flow through the bridge, there is a firewall option that must be set. This option is `IPFIREWALL_DEFAULT_TO_ACCEPT`. Note that this changes the default rule for the firewall to accept any packet. Make sure you know how this changes the meaning of your ruleset before you set it.

19.5.3.4 Traffic Shaping Support

If you want to use the bridge as a traffic shaper, you will need to add the `DUMMYNET` option to your kernel configuration. Read `dummynet(4)` for further information.

19.5.4 Enabling the Bridge

Add the line:

```
net.link.ether.bridge=1
```

to `/etc/sysctl.conf` to enable the bridge at runtime, and the line:

```
net.link.ether.bridge_cfg=if1,if2
```

to enable bridging on the specified interfaces (replace `if1` and `if2` with the names of your two network interfaces). If you want the bridged packets to be filtered by `ipfw(8)`, you should add:

```
net.link.ether.bridge_ipfw=1
```

as well.

19.5.5 Other Information

If you want to be able to `telnet(1)` into the bridge from the network, it is correct to assign one of the network cards an IP address. The consensus is that assigning both cards an address is a bad idea.

If you have multiple bridges on your network, there cannot be more than one path between any two workstations. Technically, this means that there is no support for spanning tree link management.

A bridge can add latency to your `ping(8)` times, especially for traffic from one segment to another.

19.6 NFS

Reorganized and enhanced by Tom Rhodes. Written by Bill Swingle.

Among the many different filesystems that DragonFly supports is the Network File System, also known as NFS. NFS allows a system to share directories and files with others over a network. By using NFS, users and programs can access files on remote systems almost as if they were local files.

Some of the most notable benefits that NFS can provide are:

- Local workstations use less disk space because commonly used data can be stored on a single machine and still remain accessible to others over the network.
- There is no need for users to have separate home directories on every network machine. Home directories could be set up on the NFS server and made available throughout the network.
- Storage devices such as floppy disks, CDROM drives, and ZIP drives can be used by other machines on the network. This may reduce the number of removable media drives throughout the network.

19.6.1 How NFS Works

NFS consists of at least two main parts: a server and one or more clients. The client remotely accesses the data that is stored on the server machine. In order for this to function properly a few processes have to be configured and running:

The server has to be running the following daemons:

Daemon	Description
nfsd	The NFS daemon which services requests from the NFS clients.
mountd	The NFS mount daemon which carries out the requests that nfsd(8) passes on to it.
portmap	The portmapper daemon allows NFS clients to discover which port the NFS server is using.

The client can also run a daemon, known as **nfsiod**. The **nfsiod** daemon services the requests from the NFS server. This is optional, and improves performance, but is not required for normal and correct operation. See the `nfsiod(8)` manual page for more information.

19.6.2 Configuring NFS

NFS configuration is a relatively straightforward process. The processes that need to be running can all start at boot time with a few modifications to your `/etc/rc.conf` file.

On the NFS server, make sure that the following options are configured in the `/etc/rc.conf` file:

```
portmap_enable="YES"
nfs_server_enable="YES"
mountd_flags="-r"
```

`mountd` runs automatically whenever the NFS server is enabled.

On the client, make sure this option is present in `/etc/rc.conf`:

```
nfs_client_enable="YES"
```

The `/etc/exports` file specifies which filesystems NFS should export (sometimes referred to as “share”). Each line in `/etc/exports` specifies a filesystem to be exported and which machines have access to that filesystem. Along with what machines have access to that filesystem, access options may also be specified. There are many such options that can be used in this file but only a few will be mentioned here. You can easily discover other options by reading over the `exports(5)` manual page.

Here are a few example `/etc/exports` entries:

The following examples give an idea of how to export filesystems, although the settings may be different depending on your environment and network configuration. For instance, to export the `/cdrom` directory to three example machines that have the same domain name as the server (hence the lack of a domain name for each) or have entries in your `/etc/hosts` file. The `-ro` flag makes the exported filesystem read-only. With this flag, the remote system will not be able to write any changes to the exported filesystem.

```
/cdrom -ro host1 host2 host3
```

The following line exports `/home` to three hosts by IP address. This is a useful setup if you have a private network without a DNS server configured. Optionally the `/etc/hosts` file could be configured for internal hostnames; please review `hosts(5)` for more information. The `-alldirs` flag allows the subdirectories to be mount points. In

other words, it will not mount the subdirectories but permit the client to mount only the directories that are required or needed.

```
/home -alldirs 10.0.0.2 10.0.0.3 10.0.0.4
```

The following line exports `/a` so that two clients from different domains may access the filesystem. The `-maproot=root` flag allows the `root` user on the remote system to write data on the exported filesystem as `root`. If the `-maproot=root` flag is not specified, then even if a user has `root` access on the remote system, they will not be able to modify files on the exported filesystem.

```
/a -maproot=root host.example.com box.example.org
```

In order for a client to access an exported filesystem, the client must have permission to do so. Make sure the client is listed in your `/etc/exports` file.

In `/etc/exports`, each line represents the export information for one filesystem to one host. A remote host can only be specified once per filesystem, and may only have one default entry. For example, assume that `/usr` is a single filesystem. The following `/etc/exports` would be invalid:

```
/usr/src client
/usr/ports client
```

One filesystem, `/usr`, has two lines specifying exports to the same host, `client`. The correct format for this situation is:

```
/usr/src /usr/ports client
```

The properties of one filesystem exported to a given host must all occur on one line. Lines without a client specified are treated as a single host. This limits how you can export filesystems, but for most people this is not an issue.

The following is an example of a valid export list, where `/usr` and `/exports` are local filesystems:

```
# Export src and ports to client01 and client02, but only
# client01 has root privileges on it
/usr/src /usr/ports -maproot=root client01
/usr/src /usr/ports client02
# The client machines have root and can mount anywhere
# on /exports. Anyone in the world can mount /exports/obj read-only
/exports -alldirs -maproot=root client01 client02
/exports/obj -ro
```

You must restart `mountd` whenever you modify `/etc/exports` so the changes can take effect. This can be accomplished by sending the HUP signal to the `mountd` process:

```
# kill -HUP `cat /var/run/mountd.pid`
```

Alternatively, a reboot will make DragonFly set everything up properly. A reboot is not necessary though. Executing the following commands as `root` should start everything up.

On the NFS server:

```
# portmap
# nfsd -u -t -n 4
# mountd -r
```

On the NFS client:

```
# nfsiod -n 4
```

Now everything should be ready to actually mount a remote file system. In these examples the server's name will be `server` and the client's name will be `client`. If you only want to temporarily mount a remote filesystem or would rather test the configuration, just execute a command like this as `root` on the client:

```
# mount server:/home /mnt
```

This will mount the `/home` directory on the server at `/mnt` on the client. If everything is set up correctly you should be able to enter `/mnt` on the client and see all the files that are on the server.

If you want to automatically mount a remote filesystem each time the computer boots, add the filesystem to the `/etc/fstab` file. Here is an example:

```
server:/home /mnt nfs rw 0 0
```

The `fstab(5)` manual page lists all the available options.

19.6.3 Practical Uses

NFS has many practical uses. Some of the more common ones are listed below:

- Set several machines to share a CDROM or other media among them. This is cheaper and often a more convenient method to install software on multiple machines.
- On large networks, it might be more convenient to configure a central NFS server in which to store all the user home directories. These home directories can then be exported to the network so that users would always have the same home directory, regardless of which workstation they log in to.
- Several machines could have a common `/usr/ports/distfiles` directory. That way, when you need to install a port on several machines, you can quickly access the source without downloading it on each machine.

19.6.4 Automatic Mounts with amd

Contributed by Wylie Stilwell. Rewritten by Chern Lee.

`amd(8)` (the automatic mounter daemon) automatically mounts a remote filesystem whenever a file or directory within that filesystem is accessed. Filesystems that are inactive for a period of time will also be automatically unmounted by `amd`. Using `amd` provides a simple alternative to permanent mounts, as permanent mounts are usually listed in `/etc/fstab`.

`amd` operates by attaching itself as an NFS server to the `/host` and `/net` directories. When a file is accessed within one of these directories, `amd` looks up the corresponding remote mount and automatically mounts it. `/net` is used to mount an exported filesystem from an IP address, while `/host` is used to mount an export from a remote hostname.

An access to a file within `/host/foobar/usr` would tell `amd` to attempt to mount the `/usr` export on the host `foobar`.

Example 19-1. Mounting an Export with amd

You can view the available mounts of a remote host with the `showmount` command. For example, to view the mounts of a host named `foobar`, you can use:

```
% showmount -e foobar
Exports list on foobar:
/usr                               10.10.10.0
/a                                 10.10.10.0
% cd /host/foobar/usr
```

As seen in the example, the `showmount` shows `/usr` as an export. When changing directories to `/host/foobar/usr`, **amd** attempts to resolve the hostname `foobar` and automatically mount the desired export.

amd can be started by the startup scripts by placing the following lines in `/etc/rc.conf`:

```
amd_enable="YES"
```

Additionally, custom flags can be passed to **amd** from the `amd_flags` option. By default, `amd_flags` is set to:

```
amd_flags="-a /.amd_mnt -l syslog /host /etc/amd.map /net /etc/amd.map"
```

The `/etc/amd.map` file defines the default options that exports are mounted with. The `/etc/amd.conf` file defines some of the more advanced features of **amd**.

Consult the `amd(8)` and `amd.conf(5)` manual pages for more information.

19.6.5 Problems Integrating with Other Systems

Contributed by John Lind.

Certain Ethernet adapters for ISA PC systems have limitations which can lead to serious network problems, particularly with NFS. This difficulty is not specific to DragonFly, but DragonFly systems are affected by it.

The problem nearly always occurs when (DragonFly) PC systems are networked with high-performance workstations, such as those made by Silicon Graphics, Inc., and Sun Microsystems, Inc. The NFS mount will work fine, and some operations may succeed, but suddenly the server will seem to become unresponsive to the client, even though requests to and from other systems continue to be processed. This happens to the client system, whether the client is the DragonFly system or the workstation. On many systems, there is no way to shut down the client gracefully once this problem has manifested itself. The only solution is often to reset the client, because the NFS situation cannot be resolved.

Though the “correct” solution is to get a higher performance and capacity Ethernet adapter for the DragonFly system, there is a simple workaround that will allow satisfactory operation. If the DragonFly system is the *server*, include the option `-w=1024` on the mount from the client. If the DragonFly system is the *client*, then mount the NFS filesystem with the option `-r=1024`. These options may be specified using the fourth field of the `fstab` entry on the client for automatic mounts, or by using the `-o` parameter of the `mount(8)` command for manual mounts.

It should be noted that there is a different problem, sometimes mistaken for this one, when the NFS servers and clients are on different networks. If that is the case, make *certain* that your routers are routing the necessary UDP information, or you will not get anywhere, no matter what else you are doing.

In the following examples, `fastws` is the host (interface) name of a high-performance workstation, and `freebox` is the host (interface) name of a DragonFly system with a lower-performance Ethernet adapter. Also, `/sharedfs` will

be the exported NFS filesystem (see `exports(5)`), and `/project` will be the mount point on the client for the exported filesystem. In all cases, note that additional options, such as `hard` or `soft` and `bg` may be desirable in your application.

Examples for the DragonFly system (`freebox`) as the client in `/etc/fstab` on `freebox`:

```
fastws:/sharedfs /project nfs rw,-r=1024 0 0
```

As a manual mount command on `freebox`:

```
# mount -t nfs -o -r=1024 fastws:/sharedfs /project
```

Examples for the DragonFly system as the server in `/etc/fstab` on `fastws`:

```
freebox:/sharedfs /project nfs rw,-w=1024 0 0
```

As a manual mount command on `fastws`:

```
# mount -t nfs -o -w=1024 freebox:/sharedfs /project
```

Nearly any 16-bit Ethernet adapter will allow operation without the above restrictions on the read or write size.

For anyone who cares, here is what happens when the failure occurs, which also explains why it is unrecoverable. NFS typically works with a “block” size of 8 k (though it may do fragments of smaller sizes). Since the maximum Ethernet packet is around 1500 bytes, the NFS “block” gets split into multiple Ethernet packets, even though it is still a single unit to the upper-level code, and must be received, assembled, and *acknowledged* as a unit. The high-performance workstations can pump out the packets which comprise the NFS unit one right after the other, just as close together as the standard allows. On the smaller, lower capacity cards, the later packets overrun the earlier packets of the same unit before they can be transferred to the host and the unit as a whole cannot be reconstructed or acknowledged. As a result, the workstation will time out and try again, but it will try again with the entire 8 K unit, and the process will be repeated, ad infinitum.

By keeping the unit size below the Ethernet packet size limitation, we ensure that any complete Ethernet packet received can be acknowledged individually, avoiding the deadlock situation.

Overruns may still occur when a high-performance workstation is slamming data out to a PC system, but with the better cards, such overruns are not guaranteed on NFS “units”. When an overrun occurs, the units affected will be retransmitted, and there will be a fair chance that they will be received, assembled, and acknowledged.

19.7 Diskless Operation

Updated by Jean-François Dockès. Reorganized and enhanced by Alex Dupre.

A DragonFly machine can boot over the network and operate without a local disk, using filesystems mounted from an NFS server. No system modification is necessary, beyond standard configuration files. Such a system is relatively easy to set up because all the necessary elements are readily available:

- There are at least two possible methods to load the kernel over the network:
 - PXE: The Intel Preboot Execution Environment system is a form of smart boot ROM built into some networking cards or motherboards. See `pxeboot(8)` for more details.

- The **etherboot** port (`net/etherboot`) produces ROM-able code to boot kernels over the network. The code can be either burnt into a boot PROM on a network card, or loaded from a local floppy (or hard) disk drive, or from a running MS-DOS system. Many network cards are supported.
- A sample script (`/usr/share/examples/diskless/clone_root`) eases the creation and maintenance of the workstation's root filesystem on the server. The script will probably require a little customization but it will get you started very quickly.
- Standard system startup files exist in `/etc` to detect and support a diskless system startup.
- Swapping, if needed, can be done either to an NFS file or to a local disk.

There are many ways to set up diskless workstations. Many elements are involved, and most can be customized to suit local taste. The following will describe variations on the setup of a complete system, emphasizing simplicity and compatibility with the standard DragonFly startup scripts. The system described has the following characteristics:

- The diskless workstations use a shared read-only `root` filesystem, and a shared read-only `/usr`.
The `root` filesystem is a copy of a standard DragonFly root (typically the server's), with some configuration files overridden by ones specific to diskless operation or, possibly, to the workstation they belong to.
The parts of the `root` which have to be writable are overlaid with `mfs(8)` filesystems. Any changes will be lost when the system reboots.
- The kernel is transferred and loaded either with **etherboot** or PXE as some situations may mandate the use of either method.

Caution: As described, this system is insecure. It should live in a protected area of a network, and be untrusted by other hosts.

19.7.1 Background Information

Setting up diskless workstations is both relatively straightforward and prone to errors. These are sometimes difficult to diagnose for a number of reasons. For example:

- Compile time options may determine different behaviours at runtime.
- Error messages are often cryptic or totally absent.

In this context, having some knowledge of the background mechanisms involved is very useful to solve the problems that may arise.

Several operations need to be performed for a successful bootstrap:

- The machine needs to obtain initial parameters such as its IP address, executable filename, server name, root path. This is done using the DHCP or BOOTP protocols. DHCP is a compatible extension of BOOTP, and uses the same port numbers and basic packet format.

It is possible to configure a system to use only BOOTP. The `bootpd(8)` server program is included in the base DragonFly system.

However, DHCP has a number of advantages over BOOTP (nicer configuration files, possibility of using PXE, plus many others not directly related to diskless operation), and we will describe mainly a DHCP configuration, with equivalent examples using bootpd(8) when possible. The sample configuration will use the **ISC DHCP** software package (release 3.0.1.r12 was installed on the test server).

- The machine needs to transfer one or several programs to local memory. Either TFTP or NFS are used. The choice between TFTP and NFS is a compile time option in several places. A common source of error is to specify filenames for the wrong protocol: TFTP typically transfers all files from a single directory on the server, and would expect filenames relative to this directory. NFS needs absolute file paths.
- The possible intermediate bootstrap programs and the kernel need to be initialized and executed. There are several important variations in this area:
 - PXE will load pxeboot(8), which is a modified version of the DragonFly third stage loader. The loader(8) will obtain most parameters necessary to system startup, and leave them in the kernel environment before transferring control. It is possible to use a `GENERIC` kernel in this case.
 - **etherboot**, will directly load the kernel, with less preparation. You will need to build a kernel with specific options.
- Finally, the machine needs to access its filesystems. NFS is used in all cases.

See also `diskless(8)` manual page.

19.7.2 Setup Instructions

19.7.2.1 Configuration Using ISC DHCP

The **ISC DHCP** server can answer both BOOTP and DHCP requests.

ISC DHCP needs a configuration file to run, (normally named `/usr/local/etc/dhcpd.conf`). Here follows a commented example, where host `margaux` uses **etherboot** and host `corbieres` uses PXE:

```
default-lease-time 600;
max-lease-time 7200;
authoritative;

option domain-name "example.com";
option domain-name-servers 192.168.4.1;
option routers 192.168.4.1;

subnet 192.168.4.0 netmask 255.255.255.0 {
    use-host-decl-names on; ❶
    option subnet-mask 255.255.255.0;
    option broadcast-address 192.168.4.255;

    host margaux {
        hardware ethernet 01:23:45:67:89:ab;
        fixed-address margaux.example.com;
        next-server 192.168.4.4; ❷
        filename "/data/misc/kernel.diskless"; ❸
        option root-path "192.168.4.4:/data/misc/diskless"; ❹
    }
}
```



```

}
host corbieres {
    hardware ethernet 00:02:b3:27:62:df;
    fixed-address corbieres.example.com;
    next-server 192.168.4.4;
    filename "pxeboot";
    option root-path "192.168.4.4:/data/misc/diskless";
}
}

```

- ❶ This option tells **dhcpd** to send the value in the host declarations as the hostname for the diskless host. An alternate way would be to add an option `host-name margaux` inside the host declarations.
- ❷ The `next-server` directive designates the TFTP or NFS server to use for loading loader or kernel file (the default is to use the same host as the DHCP server).
- ❸ The `filename` directive defines the file that **etherboot** or PXE will load for the next execution step. It must be specified according to the transfer method used. **etherboot** can be compiled to use NFS or TFTP. The DragonFly port configures NFS by default. PXE uses TFTP, which is why a relative filename is used here (this may depend on the TFTP server configuration, but would be fairly typical). Also, PXE loads `pxeboot`, not the kernel. There are other interesting possibilities, like loading `pxeboot` from a DragonFly CD-ROM `/boot` directory (as `pxeboot(8)` can load a `GENERIC` kernel, this makes it possible to use PXE to boot from a remote CD-ROM).
- ❹ The `root-path` option defines the path to the root filesystem, in usual NFS notation. When using PXE, it is possible to leave off the host's IP as long as you do not enable the kernel option `BOOTP`. The NFS server will then be the same as the TFTP one.

19.7.2.2 Configuration Using BOOTP

Here follows an equivalent **bootpd** configuration (reduced to one client). This would be found in `/etc/bootptab`.

Please note that **etherboot** must be compiled with the non-default option `NO_DHCP_SUPPORT` in order to use `BOOTP`, and that PXE *needs* DHCP. The only obvious advantage of **bootpd** is that it exists in the base system.

```

.def100:\
    :hn:ht=1:sa=192.168.4.4:vm=rfc1048:\
    :sm=255.255.255.0:\
    :ds=192.168.4.1:\
    :gw=192.168.4.1:\
    :hd="/tftpboot":\
    :bf="/kernel.diskless":\
    :rp="192.168.4.4:/data/misc/diskless":

margaux:ha=0123456789ab:tc=.def100

```

19.7.2.3 Booting with PXE

By default, the pxeboot(8) loader loads the kernel via NFS. It can be compiled to use TFTP instead by specifying the `LOADER_TFTP_SUPPORT` option in `/etc/make.conf`. See the comments in `/etc/defaults/make.conf` (or `/usr/share/examples/etc/make.conf` for 5.X systems) for instructions.

There are two other undocumented `make.conf` options which may be useful for setting up a serial console diskless machine: `BOOT_PXEldr_PROBE_KEYBOARD`, and `BOOT_PXEldr_ALWAYS_SERIAL`.

To use PXE when the machine starts, you will usually need to select the `Boot from network` option in your BIOS setup, or type a function key during the PC initialization.

19.7.2.4 Configuring the TFTP and NFS Servers

If you are using PXE or **etherboot** configured to use TFTP, you need to enable **tftpd** on the file server:

1. Create a directory from which **tftpd** will serve the files, e.g. `/tftpboot`.
2. Add this line to your `/etc/inetd.conf`:

```
tftp dgram udp wait root /usr/libexec/tftpd tftpd -l -s /tftpboot
```

Note: It appears that at least some PXE versions want the TCP version of TFTP. In this case, add a second line, replacing `dgram udp` with `stream tcp`.

3. Tell **inetd** to reread its configuration file:

```
# kill -HUP `cat /var/run/inetd.pid`
```

You can place the `tftpboot` directory anywhere on the server. Make sure that the location is set in both `inetd.conf` and `dhcpd.conf`.

In all cases, you also need to enable NFS and export the appropriate filesystem on the NFS server.

1. Add this to `/etc/rc.conf`:

```
nfs_server_enable="YES"
```

2. Export the filesystem where the diskless root directory is located by adding the following to `/etc/exports` (adjust the volume mount point and replace *margaux corbieres* with the names of the diskless workstations):

```
/data/misc -alldirs -ro margaux corbieres
```

3. Tell **mountd** to reread its configuration file. If you actually needed to enable NFS in `/etc/rc.conf` at the first step, you probably want to reboot instead.

```
# kill -HUP `cat /var/run/mountd.pid`
```

19.7.2.5 Building a Diskless Kernel

If using **etherboot**, you need to create a kernel configuration file for the diskless client with the following options (in addition to the usual ones):

```
options          BOOTP          # Use BOOTP to obtain IP address/hostname
```

```
options      BOOTP_NFSROOT # NFS mount root filesystem using BOOTP info
```

You may also want to use `BOOTP_NFSV3`, `BOOT_COMPAT` and `BOOTP_WIRED_TO` (refer to `LINT`).

These option names are historical and slightly misleading as they actually enable indifferent use of DHCP and BOOTP inside the kernel (it is also possible to force strict BOOTP or DHCP use).

Build the kernel (see Chapter 9), and copy it to the place specified in `dhcpcd.conf`.

Note: When using PXE, building a kernel with the above options is not strictly necessary (though suggested). Enabling them will cause more DHCP requests to be issued during kernel startup, with a small risk of inconsistency between the new values and those retrieved by `pxeboot(8)` in some special cases. The advantage of using them is that the host name will be set as a side effect. Otherwise you will need to set the host name by another method, for example in a client-specific `rc.conf` file.

19.7.2.6 Preparing the Root Filesystem

You need to create a root filesystem for the diskless workstations, in the location listed as `root-path` in `dhcpcd.conf`. The following sections describe two ways to do it.

19.7.2.6.1 Using the `clone_root` Script

This is the quickest way to create a root filesystem. This shell script is located at `/usr/share/examples/diskless/clone_root` and needs customization, at least to adjust the place where the filesystem will be created (the `DEST` variable).

Refer to the comments at the top of the script for instructions. They explain how the base filesystem is built, and how files may be selectively overridden by versions specific to diskless operation, to a subnetwork, or to an individual workstation. They also give examples for the diskless `/etc/fstab` and `/etc/rc.conf` files.

The `README` files in `/usr/share/examples/diskless` contain a lot of interesting background information, but, together with the other examples in the `diskless` directory, they actually document a configuration method which is distinct from the one used by `clone_root` and the system startup scripts in `/etc`, which is a little confusing. Use them for reference only, except if you prefer the method that they describe, in which case you will need customized `rc` scripts.

19.7.2.6.2 Using the Standard `make world` Procedure

This method will install a complete virgin system (not only the root filesystem) into `DESTDIR`. All you have to do is simply execute the following script:

```
#!/bin/sh
export DESTDIR=/data/misc/diskless
mkdir -p ${DESTDIR}
cd /usr/src; make world && make kernel
cd /usr/src/etc; make distribution
```

Once done, you may need to customize your `/etc/rc.conf` and `/etc/fstab` placed into `DESTDIR` according to your needs.

19.7.2.7 Configuring Swap

If needed, a swap file located on the server can be accessed via NFS.

19.7.2.7.1 NFS Swap with DragonFly

The swap file location and size can be specified with BOOTP/DHCP DragonFly-specific options 128 and 129. Examples of configuration files for **ISC DHCP 3.0** or **bootpd** follow:

1. Add the following lines to `dhcpd.conf`:

```
# Global section
option swap-path code 128 = string;
option swap-size code 129 = integer 32;

host margaux {
    ... # Standard lines, see above
    option swap-path "192.168.4.4:/netswapvolume/netswap";
    option swap-size 64000;
}
```

`swap-path` is the path to a directory where swap files will be located. Each file will be named `swap.client-ip`.

Older versions of **dhcpd** used a syntax of `option option-128 "...`, which is no longer supported.

`/etc/bootptab` would use the following syntax instead:

```
T128="192.168.4.4:/netswapvolume/netswap":T129=0000fa00
```

Note: In `/etc/bootptab`, the swap size must be expressed in hexadecimal format.

2. On the NFS swap file server, create the swap file(s)

```
# mkdir /netswapvolume/netswap
# cd /netswapvolume/netswap
# dd if=/dev/zero bs=1024 count=64000 of=swap.192.168.4.6
# chmod 0600 swap.192.168.4.6
```

`192.168.4.6` is the IP address for the diskless client.

3. On the NFS swap file server, add the following line to `/etc/exports`:

```
/netswapvolume -maproot=0:10 -alldirs margaux corbieres
```

Then tell **mountd** to reread the exports file, as above.

19.7.2.8 Miscellaneous Issues

19.7.2.8.1 Running with a Read-only /usr

If the diskless workstation is configured to run X, you will have to adjust the xdm configuration file, which puts the error log on /usr by default.

19.7.2.8.2 Using a Non-DragonFly Server

When the server for the root filesystem is not running DragonFly, you will have to create the root filesystem on a DragonFly machine, then copy it to its destination, using `tar` or `cpio`.

In this situation, there are sometimes problems with the special files in /dev, due to differing major/minor integer sizes. A solution to this problem is to export a directory from the non-DragonFly server, mount this directory onto a DragonFly machine, and run `MAKEDEV` on the DragonFly machine to create the correct device entries.

19.8 ISDN

A good resource for information on ISDN technology and hardware is Dan Kegel's ISDN Page (<http://www.alumni.caltech.edu/~dank/isdn/>).

A quick simple road map to ISDN follows:

- If you live in Europe you might want to investigate the ISDN card section.
- If you are planning to use ISDN primarily to connect to the Internet with an Internet Provider on a dial-up non-dedicated basis, you might look into Terminal Adapters. This will give you the most flexibility, with the fewest problems, if you change providers.
- If you are connecting two LANs together, or connecting to the Internet with a dedicated ISDN connection, you might consider the stand alone router/bridge option.

Cost is a significant factor in determining what solution you will choose. The following options are listed from least expensive to most expensive.

19.8.1 ISDN Cards

Contributed by Hellmuth Michaelis.

DragonFly's ISDN implementation supports only the DSS1/Q.931 (or Euro-ISDN) standard using passive cards. Some active cards are supported where the firmware also supports other signaling protocols; this also includes the first supported Primary Rate (PRI) ISDN card.

The `isdn4bsd` software allows you to connect to other ISDN routers using either IP over raw HDLC or by using synchronous PPP: either by using kernel PPP with `isppp`, a modified `sppp(4)` driver, or by using userland `ppp(8)`. By using userland `ppp(8)`, channel bonding of two or more ISDN B-channels is possible. A telephone answering machine application is also available as well as many utilities such as a software 300 Baud modem.

Some growing number of PC ISDN cards are supported under DragonFly and the reports show that it is successfully used all over Europe and in many other parts of the world.

The passive ISDN cards supported are mostly the ones with the Infineon (formerly Siemens) ISAC/HSCX/IPAC ISDN chipsets, but also ISDN cards with chips from Cologne Chip (ISA bus only), PCI cards with Winbond W6692 chips, some cards with the Tiger300/320/ISAC chipset combinations and some vendor specific chipset based cards such as the AVM Fritz!Card PCI V.1.0 and the AVM Fritz!Card PnP.

Currently the active supported ISDN cards are the AVM B1 (ISA and PCI) BRI cards and the AVM T1 PCI PRI cards.

For documentation on **isdn4bsd**, have a look at `/usr/share/examples/isdn/` directory on your DragonFly system or at the homepage of **isdn4bsd** (<http://www.freebsd-support.de/i4b/>) which also has pointers to hints, erratas and much more documentation such as the **isdn4bsd** handbook (<http://people.FreeBSD.org/~hm/>).

For questions regarding the installation, configuration and troubleshooting **isdn4bsd**, a **freebsd-isdn** (<http://lists.FreeBSD.org/mailman/listinfo/freebsd-isdn>) mailing list is available.

19.8.2 ISDN Terminal Adapters

Terminal adapters (TA), are to ISDN what modems are to regular phone lines.

Most TA's use the standard Hayes modem AT command set, and can be used as a drop in replacement for a modem.

A TA will operate basically the same as a modem except connection and throughput speeds will be much faster than your old modem. You will need to configure PPP exactly the same as for a modem setup. Make sure you set your serial speed as high as possible.

The main advantage of using a TA to connect to an Internet Provider is that you can do Dynamic PPP. As IP address space becomes more and more scarce, most providers are not willing to provide you with a static IP anymore. Most stand-alone routers are not able to accommodate dynamic IP allocation.

TA's completely rely on the PPP daemon that you are running for their features and stability of connection. This allows you to upgrade easily from using a modem to ISDN on a DragonFly machine, if you already have PPP set up. However, at the same time any problems you experienced with the PPP program and are going to persist.

If you want maximum stability, use the kernel PPP option, not the userland PPP.

The following TA's are known to work with DragonFly:

- Motorola BitSurfer and Bitsurfer Pro
- Adtran

Most other TA's will probably work as well, TA vendors try to make sure their product can accept most of the standard modem AT command set.

The real problem with external TA's is that, like modems, you need a good serial card in your computer.

You should read the DragonFly Serial Hardware ([../articles/serial-uart/index.html](http://www.freebsd.org/doc/en_US/8.0/articles/serial-uart/index.html)) tutorial for a detailed understanding of serial devices, and the differences between asynchronous and synchronous serial ports.

A TA running off a standard PC serial port (asynchronous) limits you to 115.2 Kbs, even though you have a 128 Kbs connection. To fully utilize the 128 Kbs that ISDN is capable of, you must move the TA to a synchronous serial card.

Do not be fooled into buying an internal TA and thinking you have avoided the synchronous/asynchronous issue. Internal TA's simply have a standard PC serial port chip built into them. All this will do is save you having to buy another serial cable and find another empty electrical socket.

A synchronous card with a TA is at least as fast as a stand-alone router, and with a simple 386 DragonFly box driving it, probably more flexible.

The choice of synchronous card/TA v.s. stand-alone router is largely a religious issue. There has been some discussion of this in the mailing lists. We suggest you search the archives ([../..../search/index.html](#)) for the complete discussion.

19.8.3 Stand-alone ISDN Bridges/Routers

ISDN bridges or routers are not at all specific to DragonFly or any other operating system. For a more complete description of routing and bridging technology, please refer to a networking reference book.

In the context of this section, the terms router and bridge will be used interchangeably.

As the cost of low end ISDN routers/bridges comes down, it will likely become a more and more popular choice. An ISDN router is a small box that plugs directly into your local Ethernet network, and manages its own connection to the other bridge/router. It has built in software to communicate via PPP and other popular protocols.

A router will allow you much faster throughput than a standard TA, since it will be using a full synchronous ISDN connection.


The main problem with ISDN routers and bridges is that interoperability between manufacturers can still be a problem. If you are planning to connect to an Internet provider, you should discuss your needs with them.

If you are planning to connect two LAN segments together, such as your home LAN to the office LAN, this is the simplest lowest maintenance solution. Since you are buying the equipment for both sides of the connection you can be assured that the link will work.

For example to connect a home computer or branch office network to a head office network the following setup could be used:

Example 19-2. Branch Office or Home Network


Network uses a bus based topology with 10 base 2 Ethernet ("thinnet"). Connect router to network cable with AUI/10BT transceiver, if necessary.



If your home/branch office is only one computer you can use a twisted pair crossover cable to connect to the stand-alone router directly.

Example 19-3. Head Office or Other LAN

Network uses a star topology with 10 base T Ethernet (“Twisted Pair”).



One large advantage of most routers/bridges is that they allow you to have two *separate independent* PPP connections to two separate sites at the *same* time. This is not supported on most TA's, except for specific (usually expensive) models that have two serial ports. Do not confuse this with channel bonding, MPP, etc.

This can be a very useful feature if, for example, you have a dedicated ISDN connection at your office and would like to tap into it, but do not want to get another ISDN line at work. A router at the office location can manage a dedicated B channel connection (64 Kbps) to the Internet and use the other B channel for a separate data connection.

The second B channel can be used for dial-in, dial-out or dynamically bonding (MPP, etc.) with the first B channel for more bandwidth.

An Ethernet bridge will also allow you to transmit more than just IP traffic. You can also send IPX/SPX or whatever other protocols you use.

19.9 NIS/YP

Written by Bill Swingle. Enhanced by Eric Ogren and Udo Erdelhoff.

19.9.1 What Is It?

NIS, which stands for Network Information Services, was developed by Sun Microsystems to centralize administration of UNIX (originally SunOS) systems. It has now essentially become an industry standard; all major UNIX like systems (Solaris, HP-UX, AIX®, Linux, NetBSD, OpenBSD, FreeBSD, DragonFly, etc.) support NIS.

NIS was formerly known as Yellow Pages, but because of trademark issues, Sun changed the name. The old term (and yp) is still often seen and used.

It is a RPC-based client/server system that allows a group of machines within an NIS domain to share a common set of configuration files. This permits a system administrator to set up NIS client systems with only minimal configuration data and add, remove or modify configuration data from a single location.

It is similar to the Windows NT® domain system; although the internal implementation of the two are not at all similar, the basic functionality can be compared.

19.9.2 Terms/Processes You Should Know

There are several terms and several important user processes that you will come across when attempting to implement NIS on DragonFly, whether you are trying to create an NIS server or act as an NIS client:

Term	Description
NIS domainname	An NIS master server and all of its clients (including its slave servers) have a NIS domainname. Similar to an Windows NT domain name, the NIS domainname does not have anything to do with DNS.
portmap	Must be running in order to enable RPC (Remote Procedure Call, a network protocol used by NIS). If portmap is not running, it will be impossible to run an NIS server, or to act as an NIS client.
ypbind	“Binds” an NIS client to its NIS server. It will take the NIS domainname from the system, and using RPC, connect to the server. ypbind is the core of client-server communication in an NIS environment; if ypbind dies on a client machine, it will not be able to access the NIS server.

Term	Description
ypserv	Should only be running on NIS servers; this is the NIS server process itself. If ypserv(8) dies, then the server will no longer be able to respond to NIS requests (hopefully, there is a slave server to take over for it). There are some implementations of NIS (but not the DragonFly one), that do not try to reconnect to another server if the server it used before dies. Often, the only thing that helps in this case is to restart the server process (or even the whole server) or the <code>ypbind</code> process on the client.
rpc.yppasswdd	Another process that should only be running on NIS master servers; this is a daemon that will allow NIS clients to change their NIS passwords. If this daemon is not running, users will have to login to the NIS master server and change their passwords there.

19.9.3 How Does It Work?

There are three types of hosts in an NIS environment: master servers, slave servers, and clients. Servers act as a central repository for host configuration information. Master servers hold the authoritative copy of this information, while slave servers mirror this information for redundancy. Clients rely on the servers to provide this information to them.

Information in many files can be shared in this manner. The `master.passwd`, `group`, and `hosts` files are commonly shared via NIS. Whenever a process on a client needs information that would normally be found in these files locally, it makes a query to the NIS server that it is bound to instead.

19.9.3.1 Machine Types

- A *NIS master server*. This server, analogous to a Windows NT primary domain controller, maintains the files used by all of the NIS clients. The `passwd`, `group`, and other various files used by the NIS clients live on the master server.

Note: It is possible for one machine to be an NIS master server for more than one NIS domain. However, this will not be covered in this introduction, which assumes a relatively small-scale NIS environment.

- *NIS slave servers*. Similar to the Windows NT backup domain controllers, NIS slave servers maintain copies of the NIS master's data files. NIS slave servers provide the redundancy, which is needed in important environments. They also help to balance the load of the master server: NIS Clients always attach to the NIS server whose response they get first, and this includes slave-server-replies.
- *NIS clients*. NIS clients, like most Windows NT workstations, authenticate against the NIS server (or the Windows NT domain controller in the Windows NT Workstation case) to log on.

19.9.4 Using NIS/YP

This section will deal with setting up a sample NIS environment.

19.9.4.1 Planning

Let us assume that you are the administrator of a small university lab. This lab, which consists of 15 DragonFly machines, currently has no centralized point of administration; each machine has its own `/etc/passwd` and `/etc/master.passwd`. These files are kept in sync with each other only through manual intervention; currently, when you add a user to the lab, you must run `adduser` on all 15 machines. Clearly, this has to change, so you have decided to convert the lab to use NIS, using two of the machines as servers.

Therefore, the configuration of the lab now looks something like:

Machine name	IP address	Machine role
ellington	10.0.0.2	NIS master
coltrane	10.0.0.3	NIS slave
basie	10.0.0.4	Faculty workstation
bird	10.0.0.5	Client machine
cli[1-11]	10.0.0.[6-17]	Other client machines

If you are setting up a NIS scheme for the first time, it is a good idea to think through how you want to go about it. No matter what the size of your network, there are a few decisions that need to be made.

19.9.4.1.1 Choosing a NIS Domain Name

This might not be the “domainname” that you are used to. It is more accurately called the “NIS domainname”. When a client broadcasts its requests for info, it includes the name of the NIS domain that it is part of. This is how multiple servers on one network can tell which server should answer which request. Think of the NIS domainname as the name for a group of hosts that are related in some way.

Some organizations choose to use their Internet domainname for their NIS domainname. This is not recommended as it can cause confusion when trying to debug network problems. The NIS domainname should be unique within your network and it is helpful if it describes the group of machines it represents. For example, the Art department at Acme Inc. might be in the “acme-art” NIS domain. For this example, assume you have chosen the name *test-domain*.

However, some operating systems (notably SunOS) use their NIS domain name as their Internet domain name. If one or more machines on your network have this restriction, you *must* use the Internet domain name as your NIS domain name.

19.9.4.1.2 Physical Server Requirements

There are several things to keep in mind when choosing a machine to use as a NIS server. One of the unfortunate things about NIS is the level of dependency the clients have on the server. If a client cannot contact the server for its NIS domain, very often the machine becomes unusable. The lack of user and group information causes most systems to temporarily freeze up. With this in mind you should make sure to choose a machine that will not be prone to being rebooted regularly, or one that might be used for development. The NIS server should ideally be a stand alone machine whose sole purpose in life is to be an NIS server. If you have a network that is not very heavily used, it is

acceptable to put the NIS server on a machine running other services, just keep in mind that if the NIS server becomes unavailable, it will affect *all* of your NIS clients adversely.

19.9.4.2 NIS Servers

The canonical copies of all NIS information are stored on a single machine called the NIS master server. The databases used to store the information are called NIS maps. In DragonFly, these maps are stored in `/var/yp/[domainname]` where `[domainname]` is the name of the NIS domain being served. A single NIS server can support several domains at once, therefore it is possible to have several such directories, one for each supported domain. Each domain will have its own independent set of maps.

NIS master and slave servers handle all NIS requests with the `yplib` daemon. `yplib` is responsible for receiving incoming requests from NIS clients, translating the requested domain and map name to a path to the corresponding database file and transmitting data from the database back to the client.

19.9.4.2.1 Setting Up a NIS Master Server

Setting up a master NIS server can be relatively straight forward, depending on your needs. DragonFly comes with support for NIS out-of-the-box. All you need is to add the following lines to `/etc/rc.conf`, and DragonFly will do the rest for you.

1.

```
nisdomainname="test-domain"
```

This line will set the NIS domainname to `test-domain` upon network setup (e.g. after reboot).

2.

```
nis_server_enable="YES"
```

This will tell DragonFly to start up the NIS server processes when the networking is next brought up.

3.

```
nis_yppasswdd_enable="YES"
```

This will enable the `rpc.yppasswdd` daemon which, as mentioned above, will allow users to change their NIS password from a client machine.

Note: Depending on your NIS setup, you may need to add further entries. See the section about NIS servers that are also NIS clients, below, for details.

Now, all you have to do is to run the command `/etc/netstart` as superuser. It will set up everything for you, using the values you defined in `/etc/rc.conf`.

19.9.4.2.2 Initializing the NIS Maps

The *NIS maps* are database files, that are kept in the `/var/yp` directory. They are generated from configuration files in the `/etc` directory of the NIS master, with one exception: the `/etc/master.passwd` file. This is for a good reason; you do not want to propagate passwords to your `root` and other administrative accounts to all the servers in the NIS domain. Therefore, before we initialize the NIS maps, you should:

```
# cp /etc/master.passwd /var/yp/master.passwd
# cd /var/yp
# vi master.passwd
```

You should remove all entries regarding system accounts (`bin`, `tty`, `kmem`, `games`, etc), as well as any accounts that you do not want to be propagated to the NIS clients (for example `root` and any other UID 0 (superuser) accounts).

Note: Make sure the `/var/yp/master.passwd` is neither group nor world readable (mode 600)! Use the `chmod` command, if appropriate.

When you have finished, it is time to initialize the NIS maps! DragonFly includes a script named `ypinit` to do this for you (see its manual page for more information). Note that this script is available on most UNIX Operating Systems, but not on all. On Digital UNIX/Compaq Tru64 UNIX it is called `ypsetup`. Because we are generating maps for an NIS master, we are going to pass the `-m` option to `ypinit`. To generate the NIS maps, assuming you already performed the steps above, run:

```
ellington# ypinit -m test-domain
Server Type: MASTER Domain: test-domain
Creating an YP server will require that you answer a few questions.
Questions will all be asked at the beginning of the procedure.
Do you want this procedure to quit on non-fatal errors? [y/n: n] n
Ok, please remember to go back and redo manually whatever fails.
If you don't, something might not work.
At this point, we have to construct a list of this domains YP servers.
rod.darktech.org is already known as master server.
Please continue to add any slave servers, one per line. When you are
done with the list, type a <control D>.
master server   : ellington
next host to add: coltrane
next host to add: ^D
The current list of NIS servers looks like this:
ellington
coltrane
Is this correct? [y/n: y] y
```

```
[..output from map generation..]
```

```
NIS Map update completed.
ellington has been setup as an YP master server without any errors.
```

`ypinit` should have created `/var/yp/Makefile` from `/var/yp/Makefile.dist`. When created, this file assumes that you are operating in a single server NIS environment with only DragonFly machines. Since `test-domain` has a slave server as well, you must edit `/var/yp/Makefile`:

```
ellington# vi /var/yp/Makefile
```

You should comment out the line that says

```
NOPUSH = "True"
```

(if it is not commented out already).

19.9.4.2.3 Setting up a NIS Slave Server

Setting up an NIS slave server is even more simple than setting up the master. Log on to the slave server and edit the file `/etc/rc.conf` as you did before. The only difference is that we now must use the `-s` option when running `ypinit`. The `-s` option requires the name of the NIS master be passed to it as well, so our command line looks like:

```
coltrane# ypinit -s ellington test-domain
```

```
Server Type: SLAVE Domain: test-domain Master: ellington
```

Creating an YP server will require that you answer a few questions. Questions will all be asked at the beginning of the procedure.

```
Do you want this procedure to quit on non-fatal errors? [y/n: n]  n
```

Ok, please remember to go back and redo manually whatever fails. If you don't, something might not work.

There will be no further questions. The remainder of the procedure should take a few minutes, to copy the databases from ellington.

```
Transferring netgroup...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring netgroup.byuser...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring netgroup.byhost...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring master.passwd.byuid...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring passwd.byuid...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring passwd.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring group.bygid...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring group.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring services.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring rpc.bynumber...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring rpc.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring protocols.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring master.passwd.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring networks.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring networks.byaddr...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring netid.byname...
```

```
ypxfr: Exiting: Map successfully transferred
```

```
Transferring hosts.byaddr...
```

```
ypxfr: Exiting: Map successfully transferred
```

```

Transferring protocols.bynumber...
ypxfr: Exiting: Map successfully transferred
Transferring ypservers...
ypxfr: Exiting: Map successfully transferred
Transferring hosts.byname...
ypxfr: Exiting: Map successfully transferred

```

coltrane has been setup as an YP slave server without any errors.
Don't forget to update map ypservers on ellington.

You should now have a directory called `/var/yp/test-domain`. Copies of the NIS master server's maps should be in this directory. You will need to make sure that these stay updated. The following `/etc/crontab` entries on your slave servers should do the job:

```

20      *      *      *      *      root    /usr/libexec/ypxfr passwd.byname
21      *      *      *      *      root    /usr/libexec/ypxfr passwd.byuid

```

These two lines force the slave to sync its maps with the maps on the master server. Although these entries are not mandatory, since the master server attempts to ensure any changes to its NIS maps are communicated to its slaves and because password information is vital to systems depending on the server, it is a good idea to force the updates. This is more important on busy networks where map updates might not always complete.

Now, run the command `/etc/netstart` on the slave server as well, which again starts the NIS server.

19.9.4.3 NIS Clients

An NIS client establishes what is called a binding to a particular NIS server using the `ypbind` daemon. `ypbind` checks the system's default domain (as set by the `domainname` command), and begins broadcasting RPC requests on the local network. These requests specify the name of the domain for which `ypbind` is attempting to establish a binding. If a server that has been configured to serve the requested domain receives one of the broadcasts, it will respond to `ypbind`, which will record the server's address. If there are several servers available (a master and several slaves, for example), `ypbind` will use the address of the first one to respond. From that point on, the client system will direct all of its NIS requests to that server. `ypbind` will occasionally "ping" the server to make sure it is still up and running. If it fails to receive a reply to one of its pings within a reasonable amount of time, `ypbind` will mark the domain as unbound and begin broadcasting again in the hopes of locating another server.

19.9.4.3.1 Setting Up a NIS Client

Setting up a DragonFly machine to be a NIS client is fairly straightforward.

1. Edit the file `/etc/rc.conf` and add the following lines in order to set the NIS domainname and start `ypbind` upon network startup:

```

nisdomainname="test-domain"
nis_client_enable="YES"

```

2. To import all possible password entries from the NIS server, remove all user accounts from your `/etc/master.passwd` file and use `vipw` to add the following line to the end of the file:

```

+:::~::~:

```

Note: This line will afford anyone with a valid account in the NIS server's password maps an account. There are many ways to configure your NIS client by changing this line. See the `netgroups` section below for more information. For more detailed reading see O'Reilly's book on `Managing NFS and NIS`.

Note: You should keep at least one local account (i.e. not imported via NIS) in your `/etc/master.passwd` and this account should also be a member of the group `wheel`. If there is something wrong with NIS, this account can be used to log in remotely, become root, and fix things.

3. To import all possible group entries from the NIS server, add this line to your `/etc/group` file:

```
+:*:*:
```

After completing these steps, you should be able to run `ypcat passwd` and see the NIS server's `passwd` map.

19.9.5 NIS Security

In general, any remote user can issue an RPC to `ypserv(8)` and retrieve the contents of your NIS maps, provided the remote user knows your domainname. To prevent such unauthorized transactions, `ypserv(8)` supports a feature called `securenets` which can be used to restrict access to a given set of hosts. At startup, `ypserv(8)` will attempt to load the `securenets` information from a file called `/var/yp/securenets`.

Note: This path varies depending on the path specified with the `-p` option. This file contains entries that consist of a network specification and a network mask separated by white space. Lines starting with `#` are considered to be comments. A sample `securenets` file might look like this:

```
# allow connections from local host -- mandatory
127.0.0.1      255.255.255.255
# allow connections from any host
# on the 192.168.128.0 network
192.168.128.0 255.255.255.0
# allow connections from any host
# between 10.0.0.0 to 10.0.15.255
# this includes the machines in the testlab
10.0.0.0      255.255.240.0
```

If `ypserv(8)` receives a request from an address that matches one of these rules, it will process the request normally. If the address fails to match a rule, the request will be ignored and a warning message will be logged. If the `/var/yp/securenets` file does not exist, `ypserv` will allow connections from any host.

The `ypserv` program also has support for Wietse Venema's `tcpwrapper` package. This allows the administrator to use the `tcpwrapper` configuration files for access control instead of `/var/yp/securenets`.

Note: While both of these access control mechanisms provide some security, they, like the privileged port test, are vulnerable to "IP spoofing" attacks. All NIS-related traffic should be blocked at your firewall.

Servers using `/var/yp/securenets` may fail to serve legitimate NIS clients with archaic TCP/IP implementations. Some of these implementations set all host bits to zero when doing broadcasts and/or fail to observe the subnet mask when calculating the broadcast address. While some of these problems can be fixed by changing the client configuration, other problems may force the retirement of the client systems in question or the abandonment of `/var/yp/securenets`.

Using `/var/yp/securenets` on a server with such an archaic implementation of TCP/IP is a really bad idea and will lead to loss of NIS functionality for large parts of your network.

The use of the **tcpwrapper** package increases the latency of your NIS server. The additional delay may be long enough to cause timeouts in client programs, especially in busy networks or with slow NIS servers. If one or more of your client systems suffers from these symptoms, you should convert the client systems in question into NIS slave servers and force them to bind to themselves.

19.9.6 Barring Some Users from Logging On

In our lab, there is a machine `basie` that is supposed to be a faculty only workstation. We do not want to take this machine out of the NIS domain, yet the `passwd` file on the master NIS server contains accounts for both faculty and students. What can we do?

There is a way to bar specific users from logging on to a machine, even if they are present in the NIS database. To do this, all you must do is add `-username` to the end of the `/etc/master.passwd` file on the client machine, where `username` is the username of the user you wish to bar from logging in. This should preferably be done using `vipw`, since `vipw` will sanity check your changes to `/etc/master.passwd`, as well as automatically rebuild the password database when you finish editing. For example, if we wanted to bar user `bill` from logging on to `basie` we would:

```
basie# vipw
[add -bill to the end, exit]
vipw: rebuilding the database...
vipw: done

basie# cat /etc/master.passwd

root:[password]:0:0::0:0:The super-user:/root:/bin/csh
toor:[password]:0:0::0:0:The other super-user:/root:/bin/sh
daemon:*:1:1::0:0:Owner of many system processes:/root:/sbin/nologin
operator:*:2:5::0:0:System &:/sbin/nologin
bin:*:3:7::0:0:Binaries Commands and Source,,:/sbin/nologin
tty:*:4:65533::0:0:Tty Sandbox:/sbin/nologin
kmem:*:5:65533::0:0:KMem Sandbox:/sbin/nologin
games:*:7:13::0:0:Games pseudo-user:/usr/games:/sbin/nologin
news:*:8:8::0:0:News Subsystem:/sbin/nologin
man:*:9:9::0:0:Mister Man Pages:/usr/share/man:/sbin/nologin
bind:*:53:53::0:0:Bind Sandbox:/sbin/nologin
uucp:*:66:66::0:0:UUCP pseudo-user:/var/spool/uucppublic:/usr/libexec/uucp/uucico
xten:*:67:67::0:0:X-10 daemon:/usr/local/xten:/sbin/nologin
pop:*:68:6::0:0:Post Office Owner:/nonexistent:/sbin/nologin
nobody:*:65534:65534::0:0:Unprivileged user:/nonexistent:/sbin/nologin
+:::~::~:
-bill
```

basie#

19.9.7 Using Netgroups

Contributed by Udo Erdelhoff.

The method shown in the previous section works reasonably well if you need special rules for a very small number of users and/or machines. On larger networks, you *will* forget to bar some users from logging onto sensitive machines, or you may even have to modify each machine separately, thus losing the main benefit of NIS, *centralized* administration.

The NIS developers' solution for this problem is called *netgroups*. Their purpose and semantics can be compared to the normal groups used by UNIX file systems. The main differences are the lack of a numeric id and the ability to define a netgroup by including both user accounts and other netgroups.

Netgroups were developed to handle large, complex networks with hundreds of users and machines. On one hand, this is a Good Thing if you are forced to deal with such a situation. On the other hand, this complexity makes it almost impossible to explain netgroups with really simple examples. The example used in the remainder of this section demonstrates this problem.

Let us assume that your successful introduction of NIS in your laboratory caught your superiors' interest. Your next job is to extend your NIS domain to cover some of the other machines on campus. The two tables contain the names of the new users and new machines as well as brief descriptions of them.

User Name(s)	Description
alpha, beta	Normal employees of the IT department
charlie, delta	The new apprentices of the IT department
echo, foxtrott, golf, ...	Ordinary employees
able, baker, ...	The current interns

Machine Name(s)	Description
war, death, famine, pollution	Your most important servers. Only the IT employees are allowed to log onto these machines.
pride, greed, envy, wrath, lust, sloth	Less important servers. All members of the IT department are allowed to login onto these machines.
one, two, three, four, ...	Ordinary workstations. Only the <i>real</i> employees are allowed to use these machines.
trashcan	A very old machine without any critical data. Even the intern is allowed to use this box.

If you tried to implement these restrictions by separately blocking each user, you would have to add one *-user* line to each system's `passwd` for each user who is not allowed to login onto that system. If you forget just one entry, you could be in trouble. It may be feasible to do this correctly during the initial setup, however you *will* eventually forget to add the lines for new users during day-to-day operations. After all, Murphy was an optimist.

Handling this situation with netgroups offers several advantages. Each user need not be handled separately; you assign a user to one or more netgroups and allow or forbid logins for all members of the netgroup. If you add a new

machine, you will only have to define login restrictions for netgroups. If a new user is added, you will only have to add the user to one or more netgroups. Those changes are independent of each other; no more “for each combination of user and machine do...” If your NIS setup is planned carefully, you will only have to modify exactly one central configuration file to grant or deny access to machines.

The first step is the initialization of the NIS map netgroup. DragonFly’s ypinit(8) does not create this map by default, but its NIS implementation will support it once it has been created. To create an empty map, simply type

```
ellington# vi /var/yp/netgroup
```

and start adding content. For our example, we need at least four netgroups: IT employees, IT apprentices, normal employees and interns.

```
IT_EMP ( ,alpha,test-domain) ( ,beta,test-domain)
IT_APP ( ,charlie,test-domain) ( ,delta,test-domain)
USERS ( ,echo,test-domain) ( ,foxtrott,test-domain) \
      ( ,golf,test-domain)
INTERNS ( ,able,test-domain) ( ,baker,test-domain)
```

IT_EMP, IT_APP etc. are the names of the netgroups. Each bracketed group adds one or more user accounts to it. The three fields inside a group are:

1. The name of the host(s) where the following items are valid. If you do not specify a hostname, the entry is valid on all hosts. If you do specify a hostname, you will enter a realm of darkness, horror and utter confusion.
2. The name of the account that belongs to this netgroup.
3. The NIS domain for the account. You can import accounts from other NIS domains into your netgroup if you are one of the unlucky fellows with more than one NIS domain.

Each of these fields can contain wildcards. See netgroup(5) for details.

Note: Netgroup names longer than 8 characters should not be used, especially if you have machines running other operating systems within your NIS domain. The names are case sensitive; using capital letters for your netgroup names is an easy way to distinguish between user, machine and netgroup names.

Some NIS clients (other than DragonFly) cannot handle netgroups with a large number of entries. For example, some older versions of SunOS start to cause trouble if a netgroup contains more than 15 *entries*. You can circumvent this limit by creating several sub-netgroups with 15 users or less and a real netgroup that consists of the sub-netgroups:

```
BIGGRP1 ( ,joe1,domain) ( ,joe2,domain) ( ,joe3,domain) [...]
BIGGRP2 ( ,joe16,domain) ( ,joe17,domain) [...]
BIGGRP3 ( ,joe31,domain) ( ,joe32,domain)
BIGGROUP BIGGRP1 BIGGRP2 BIGGRP3
```

You can repeat this process if you need more than 225 users within a single netgroup.

Activating and distributing your new NIS map is easy:

```
ellington# cd /var/yp
ellington# make
```

This will generate the three NIS maps `netgroup`, `netgroup.byhost` and `netgroup.byuser`. Use `ypcat(1)` to check if your new NIS maps are available:

```
ellington% ypcat -k netgroup
ellington% ypcat -k netgroup.byhost
ellington% ypcat -k netgroup.byuser
```

The output of the first command should resemble the contents of `/var/yp/netgroup`. The second command will not produce output if you have not specified host-specific netgroups. The third command can be used to get the list of netgroups for a user.

The client setup is quite simple. To configure the server `war`, you only have to start `vipw(8)` and replace the line

```
+:::~::~:
```

with

```
+@IT_EMP:~::~:
```

Now, only the data for the users defined in the netgroup `IT_EMP` is imported into `war`'s password database and only these users are allowed to login.

Unfortunately, this limitation also applies to the `~` function of the shell and all routines converting between user names and numerical user IDs. In other words, `cd ~user` will not work, `ls -l` will show the numerical id instead of the username and `find . -user joe -print` will fail with `No such user`. To fix this, you will have to import all user entries *without allowing them to login onto your servers*.

This can be achieved by adding another line to `/etc/master.passwd`. This line should contain:

```
+:::~::~:/sbin/nologin, meaning "Import all entries but replace the shell with /sbin/nologin in the
imported entries". You can replace any field in the passwd entry by placing a default value in your
/etc/master.passwd.
```

Warning: Make sure that the line `+:::~::~:/sbin/nologin` is placed after `+@IT_EMP:~::~:`. Otherwise, all user accounts imported from NIS will have `/sbin/nologin` as their login shell.

After this change, you will only have to change one NIS map if a new employee joins the IT department. You could use a similar approach for the less important servers by replacing the old `+:::~::~:` in their local version of `/etc/master.passwd` with something like this:

```
+@IT_EMP:~::~:
+@IT_APP:~::~:
+:::~::~:/sbin/nologin
```

The corresponding lines for the normal workstations could be:

```
+@IT_EMP:~::~:
+@USERS:~::~:
+:::~::~:/sbin/nologin
```

And everything would be fine until there is a policy change a few weeks later: The IT department starts hiring interns. The IT interns are allowed to use the normal workstations and the less important servers; and the IT

apprentices are allowed to login onto the main servers. You add a new netgroup `IT_INTERN`, add the new IT interns to this netgroup and start to change the config on each and every machine... As the old saying goes: “Errors in centralized planning lead to global mess”.

NIS’ ability to create netgroups from other netgroups can be used to prevent situations like these. One possibility is the creation of role-based netgroups. For example, you could create a netgroup called `BIGSRV` to define the login restrictions for the important servers, another netgroup called `SMALLSRV` for the less important servers and a third netgroup called `USERBOX` for the normal workstations. Each of these netgroups contains the netgroups that are allowed to login onto these machines. The new entries for your NIS map netgroup should look like this:

```
BIGSRV    IT_EMP  IT_APP
SMALLSRV  IT_EMP  IT_APP  ITINTERN
USERBOX   IT_EMP  ITINTERN  USERS
```

This method of defining login restrictions works reasonably well if you can define groups of machines with identical restrictions. Unfortunately, this is the exception and not the rule. Most of the time, you will need the ability to define login restrictions on a per-machine basis.

Machine-specific netgroup definitions are the other possibility to deal with the policy change outlined above. In this scenario, the `/etc/master.passwd` of each box contains two lines starting with “+”. The first of them adds a netgroup with the accounts allowed to login onto this machine, the second one adds all other accounts with `/sbin/nologin` as shell. It is a good idea to use the ALL-CAPS version of the machine name as the name of the netgroup. In other words, the lines should look like this:

```
+@BOXNAME:::::::::
+:::::::::/sbin/nologin
```

Once you have completed this task for all your machines, you will not have to modify the local versions of `/etc/master.passwd` ever again. All further changes can be handled by modifying the NIS map. Here is an example of a possible netgroup map for this scenario with some additional goodies.

```
# Define groups of users first
IT_EMP    (,alpha,test-domain)    (,beta,test-domain)
IT_APP    (,charlie,test-domain)  (,delta,test-domain)
DEPT1     (,echo,test-domain)     (,foxtrott,test-domain)
DEPT2     (,golf,test-domain)     (,hotel,test-domain)
DEPT3     (,india,test-domain)    (,juliet,test-domain)
ITINTERN  (,kilo,test-domain)     (,lima,test-domain)
D_INTERNS (,able,test-domain)     (,baker,test-domain)
#
# Now, define some groups based on roles
USERS     DEPT1    DEPT2    DEPT3
BIGSRV    IT_EMP  IT_APP
SMALLSRV  IT_EMP  IT_APP  ITINTERN
USERBOX   IT_EMP  ITINTERN  USERS
#
# And a groups for a special tasks
# Allow echo and golf to access our anti-virus-machine
SECURITY  IT_EMP  (,echo,test-domain) (,golf,test-domain)
#
# machine-based netgroups
# Our main servers
WAR       BIGSRV
```

```

FAMINE    BIGSRV
# User india needs access to this server
POLLUTION BIGSRV (,india,test-domain)
#
# This one is really important and needs more access restrictions
DEATH     IT_EMP
#
# The anti-virus-machine mentioned above
ONE       SECURITY
#
# Restrict a machine to a single user
TWO       (,hotel,test-domain)
# [...more groups to follow]

```

If you are using some kind of database to manage your user accounts, you should be able to create the first part of the map with your database's report tools. This way, new users will automatically have access to the boxes.

One last word of caution: It may not always be advisable to use machine-based netgroups. If you are deploying a couple of dozen or even hundreds of identical machines for student labs, you should use role-based netgroups instead of machine-based netgroups to keep the size of the NIS map within reasonable limits.

19.9.8 Important Things to Remember

There are still a couple of things that you will need to do differently now that you are in an NIS environment.

- Every time you wish to add a user to the lab, you must add it to the master NIS server *only*, and *you must remember to rebuild the NIS maps*. If you forget to do this, the new user will not be able to login anywhere except on the NIS master. For example, if we needed to add a new user “jsmith” to the lab, we would:

```

# pw useradd jsmith
# cd /var/yp
# make test-domain

```

You could also run `adduser jsmith` instead of `pw useradd jsmith`.

- *Keep the administration accounts out of the NIS maps*. You do not want to be propagating administrative accounts and passwords to machines that will have users that should not have access to those accounts.
- *Keep the NIS master and slave secure, and minimize their downtime*. If somebody either hacks or simply turns off these machines, they have effectively rendered many people without the ability to login to the lab.

This is the chief weakness of any centralized administration system. If you do not protect your NIS servers, you will have a lot of angry users!

19.9.9 NIS v1 Compatibility

DragonFly's `ypserv` has some support for serving NIS v1 clients. DragonFly's NIS implementation only uses the NIS v2 protocol, however other implementations include support for the v1 protocol for backwards compatibility with older systems. The `ybind` daemons supplied with these systems will try to establish a binding to an NIS v1 server even though they may never actually need it (and they may persist in broadcasting in search of one even after they receive a response from a v2 server). Note that while support for normal client calls is provided, this version of `ypserv` does not handle v1 map transfer requests; consequently, it cannot be used as a master or slave in conjunction

with older NIS servers that only support the v1 protocol. Fortunately, there probably are not any such servers still in use today.

19.9.10 NIS Servers That Are Also NIS Clients

Care must be taken when running `ypserv` in a multi-server domain where the server machines are also NIS clients. It is generally a good idea to force the servers to bind to themselves rather than allowing them to broadcast bind requests and possibly become bound to each other. Strange failure modes can result if one server goes down and others are dependent upon it. Eventually all the clients will time out and attempt to bind to other servers, but the delay involved can be considerable and the failure mode is still present since the servers might bind to each other all over again.

You can force a host to bind to a particular server by running `yplibind` with the `-s` flag. If you do not want to do this manually each time you reboot your NIS server, you can add the following lines to your `/etc/rc.conf`:

```
nis_client_enable="YES" # run client stuff as well
nis_client_flags="-S NIS domain,server"
```

See `yplibind(8)` for further information.

19.9.11 Password Formats

One of the most common issues that people run into when trying to implement NIS is password format compatibility. If your NIS server is using DES encrypted passwords, it will only support clients that are also using DES. For example, if you have Solaris NIS clients in your network, then you will almost certainly need to use DES encrypted passwords.

To check which format your servers and clients are using, look at `/etc/login.conf`. If the host is configured to use DES encrypted passwords, then the `default` class will contain an entry like this:

```
default:\
:passwd_format=des:\
:copyright=/etc/COPYRIGHT:\
[Further entries elided]
```

Other possible values for the `passwd_format` capability include `blf` and `md5` (for Blowfish and MD5 encrypted passwords, respectively).

If you have made changes to `/etc/login.conf`, you will also need to rebuild the login capability database, which is achieved by running the following command as `root`:

```
# cap_mkdb /etc/login.conf
```

Note: The format of passwords already in `/etc/master.passwd` will not be updated until a user changes their password for the first time *after* the login capability database is rebuilt.

Next, in order to ensure that passwords are encrypted with the format that you have chosen, you should also check that the `crypt_default` in `/etc/auth.conf` gives precedence to your chosen password format. To do this, place

the format that you have chosen first in the list. For example, when using DES encrypted passwords, the entry would be:

```
crypt_default = des blf md5
```

Having followed the above steps on each of the DragonFly based NIS servers and clients, you can be sure that they all agree on which password format is used within your network. If you have trouble authenticating on an NIS client, this is a pretty good place to start looking for possible problems. Remember: if you want to deploy an NIS server for a heterogenous network, you will probably have to use DES on all systems because it is the lowest common standard.

19.10 DHCP

Written by Greg Sutter.

19.10.1 What Is DHCP?

DHCP, the Dynamic Host Configuration Protocol, describes the means by which a system can connect to a network and obtain the necessary information for communication upon that network. DragonFly uses the ISC (Internet Software Consortium) DHCP implementation, so all implementation-specific information here is for use with the ISC distribution.

19.10.2 What This Section Covers

This section describes both the client-side and server-side components of the ISC DHCP system. The client-side program, `dhclient`, and the server, come integrated within DragonFly. The `dhclient(8)`, `dhcp-options(5)`, and `dhclient.conf(5)` manual pages, in addition to the references below, are useful resources.

19.10.3 How It Works

When `dhclient`, the DHCP client, is executed on the client machine, it begins broadcasting requests for configuration information. By default, these requests are on UDP port 68. The server replies on UDP 67, giving the client an IP address and other relevant network information such as netmask, router, and DNS servers. All of this information comes in the form of a DHCP “lease” and is only valid for a certain time (configured by the DHCP server maintainer). In this manner, stale IP addresses for clients no longer connected to the network can be automatically reclaimed.

DHCP clients can obtain a great deal of information from the server. An exhaustive list may be found in `dhcp-options(5)`.

19.10.4 DragonFly Integration

DragonFly fully integrates the ISC DHCP client, `dhclient`. DHCP client support is provided within both the installer and the base system, obviating the need for detailed knowledge of network configurations on any network that runs a DHCP server.

There are two things you must do to have your system use DHCP upon startup:

- Make sure that the `bpf` device is compiled into your kernel. To do this, add `pseudo-device bpf` to your kernel configuration file, and rebuild the kernel. For more information about building kernels, see Chapter 9.

The `bpf` device is already part of the `GENERIC` kernel that is supplied with DragonFly, so if you do not have a custom kernel, you should not need to create one in order to get DHCP working.

Note: For those who are particularly security conscious, you should be warned that `bpf` is also the device that allows packet sniffers to work correctly (although they still have to be run as `root`). `bpf` is required to use DHCP, but if you are very sensitive about security, you probably should not add `bpf` to your kernel in the expectation that at some point in the future you will be using DHCP.

- Edit your `/etc/rc.conf` to include the following:

```
ifconfig_fxp0="DHCP"
```

Note: Be sure to replace `fxp0` with the designation for the interface that you wish to dynamically configure, as described in Section 6.8.

If you are using a different location for `dhclient`, or if you wish to pass additional flags to `dhclient`, also include the following (editing as necessary):

```
dhcp_program="/sbin/dhclient"
dhcp_flags=""
```

The DHCP server, **dhcpcd**, is included as part of the `net/isc-dhcp3-server` port in the ports collection. This port contains the ISC DHCP server and documentation.

19.10.5 Files

- `/etc/dhclient.conf`

`dhclient` requires a configuration file, `/etc/dhclient.conf`. Typically the file contains only comments, the defaults being reasonably sane. This configuration file is described by the `dhclient.conf(5)` manual page.

- `/sbin/dhclient`

`dhclient` is statically linked and resides in `/sbin`. The `dhclient(8)` manual page gives more information about `dhclient`.

- `/sbin/dhclient-script`

`dhclient-script` is the DragonFly-specific DHCP client configuration script. It is described in `dhclient-script(8)`, but should not need any user modification to function properly.

- `/var/db/dhclient.leases`

The DHCP client keeps a database of valid leases in this file, which is written as a log. `dhclient.leases(5)` gives a slightly longer description.

19.10.6 Further Reading

The DHCP protocol is fully described in RFC 2131 (<http://www.freesoft.org/CIE/RFC/2131/>). An informational resource has also been set up at [dhcp.org](http://www.dhcp.org) (<http://www.dhcp.org/>).

19.10.7 Installing and Configuring a DHCP Server

19.10.7.1 What This Section Covers

This section provides information on how to configure a DragonFly system to act as a DHCP server using the ISC (Internet Software Consortium) implementation of the DHCP suite.

19.10.7.2 DHCP Server Installation

In order to configure your DragonFly system as a DHCP server, you will need to ensure that the `bpf(4)` device is compiled into your kernel. To do this, add `pseudo-device bpf` to your kernel configuration file, and rebuild the kernel. For more information about building kernels, see Chapter 9.

The `bpf` device is already part of the `GENERIC` kernel that is supplied with DragonFly, so you do not need to create a custom kernel in order to get DHCP working.

Note: Those who are particularly security conscious should note that `bpf` is also the device that allows packet sniffers to work correctly (although such programs still need privileged access). `bpf` is required to use DHCP, but if you are very sensitive about security, you probably should not include `bpf` in your kernel purely because you expect to use DHCP at some point in the future.

The next thing that you will need to do is edit the sample `dhcpd.conf` which was installed by the `net/isc-dhcp3-server` port. By default, this will be `/usr/local/etc/dhcpd.conf.sample`, and you should copy this to `/usr/local/etc/dhcpd.conf` before proceeding to make changes.

19.10.7.3 Configuring the DHCP Server

`dhcpd.conf` is comprised of declarations regarding subnets and hosts, and is perhaps most easily explained using an example :

```
option domain-name "example.com";❶
option domain-name-servers 192.168.4.100;❷
option subnet-mask 255.255.255.0;❸

default-lease-time 3600;❹
max-lease-time 86400;❺
ddns-update-style none;❻

subnet 192.168.4.0 netmask 255.255.255.0 {
    range 192.168.4.129 192.168.4.254;❼
    option routers 192.168.4.1;❽
}
```

```
host mailhost {
    hardware ethernet 02:03:04:05:06:07;❹
    fixed-address mailhost.example.com;❿❶
}
```

- ❶ This option specifies the domain that will be provided to clients as the default search domain. See `resolv.conf(5)` for more information on what this means.
- ❷ This option specifies a comma separated list of DNS servers that the client should use.
- ❸ The netmask that will be provided to clients.
- ❹ A client may request a specific length of time that a lease will be valid. Otherwise the server will assign a lease with this expiry value (in seconds).
- ❺ This is the maximum length of time that the server will lease for. Should a client request a longer lease, a lease will be issued, although it will only be valid for `max-lease-time` seconds.
- ❻ This option specifies whether the DHCP server should attempt to update DNS when a lease is accepted or released. In the ISC implementation, this option is *required*.
- ❼ This denotes which IP addresses should be used in the pool reserved for allocating to clients. IP addresses between, and including, the ones stated are handed out to clients.
- ❽ Declares the default gateway that will be provided to clients.
- ❾ The hardware MAC address of a host (so that the DHCP server can recognize a host when it makes a request).
- ❿❶ Specifies that the host should always be given the same IP address. Note that using a hostname is correct here, since the DHCP server will resolve the hostname itself before returning the lease information.

Once you have finished writing your `dhcpd.conf`, you can proceed to start the server by issuing the following command:

```
# /usr/local/etc/rc.d/isc-dhcpd.sh start
```

Should you need to make changes to the configuration of your server in the future, it is important to note that sending a `SIGHUP` signal to **dhcpd** does *not* result in the configuration being reloaded, as it does with most daemons. You will need to send a `SIGTERM` signal to stop the process, and then restart it using the command above.

19.10.7.4 Files

- `/usr/sbin/dhcpd`

dhcpd is statically linked and resides in `/usr/local/sbin`. The `dhcpd(8)` manual page installed with the port gives more information about **dhcpd**.

- `/etc/dhcpd.conf`

dhcpd requires a configuration file, `/usr/local/etc/dhcpd.conf` before it will start providing service to clients. This file needs to contain all the information that should be provided to clients that are being serviced, along with information regarding the operation of the server. This configuration file is described by the `dhcpd.conf(5)` manual page installed by the port.

- `/var/db/dhcpd.leases`

The DHCP server keeps a database of leases it has issued in this file, which is written as a log. The manual page `dhcpd.leases(5)`, installed by the port gives a slightly longer description.

- `/usr/sbin/dhcrelay`

dhcrelay is used in advanced environments where one DHCP server forwards a request from a client to another DHCP server on a separate network.

19.11 DNS

Contributed by Chern Lee.

19.11.1 Overview

DragonFly utilizes, by default, a version of BIND (Berkeley Internet Name Domain), which is the most common implementation of the DNS protocol. DNS is the protocol through which names are mapped to IP addresses, and vice versa. For example, a query for `www.dragonflybsd.org` will receive a reply with the IP address of The DragonFly Project's web server, whereas, a query for `ftp.dragonflybsd.org` will return the IP address of the corresponding FTP machine. Likewise, the opposite can happen. A query for an IP address can resolve its hostname. It is not necessary to run a name server to perform DNS lookups on a system.

DNS is coordinated across the Internet through a somewhat complex system of authoritative root name servers, and other smaller-scale name servers who host and cache individual domain information.

This document refers to BIND 9.x.

RFC1034 and RFC1035 dictate the DNS protocol.

Currently, BIND is maintained by the Internet Software Consortium (www.isc.org) (<http://www.isc.org/>).

19.11.2 Terminology

To understand this document, some terms related to DNS must be understood.

Term	Definition
Forward DNS	Mapping of hostnames to IP addresses
Origin	Refers to the domain covered in a particular zone file
named , BIND, name server	Common names for the BIND name server package within DragonFly
Resolver	A system process through which a machine queries a name server for zone information
Reverse DNS	The opposite of forward DNS; mapping of IP addresses to hostnames
Root zone	The beginning of the Internet zone hierarchy. All zones fall under the root zone, similar to how all files in a file system fall under the root directory.

Term	Definition
Zone	An individual domain, subdomain, or portion of the DNS administered by the same authority

Examples of zones:

- `.` is the root zone
- `org.` is a zone under the root zone
- `example.org` is a zone under the `org.` zone
- `foo.example.org.` is a subdomain, a zone under the `example.org.` zone
- `1.2.3.in-addr.arpa` is a zone referencing all IP addresses which fall under the `3.2.1.*` IP space.

As one can see, the more specific part of a hostname appears to its left. For example, `example.org.` is more specific than `org.`, as `org.` is more specific than the root zone. The layout of each part of a hostname is much like a filesystem: the `/dev` directory falls within the root, and so on.

19.11.3 Reasons to Run a Name Server

Name servers usually come in two forms: an authoritative name server, and a caching name server.

An authoritative name server is needed when:

- one wants to serve DNS information to the world, replying authoritatively to queries.
- a domain, such as `example.org`, is registered and IP addresses need to be assigned to hostnames under it.
- an IP address block requires reverse DNS entries (IP to hostname).
- a backup name server, called a slave, must reply to queries when the primary is down or inaccessible.

A caching name server is needed when:

- a local DNS server may cache and respond more quickly than querying an outside name server.
- a reduction in overall network traffic is desired (DNS traffic has been measured to account for 5% or more of total Internet traffic).

When one queries for `www.dragonflybsd.org`, the resolver usually queries the uplink ISP's name server, and retrieves the reply. With a local, caching DNS server, the query only has to be made once to the outside world by the caching DNS server. Every additional query will not have to look to the outside of the local network, since the information is cached locally.

19.11.4 How It Works

In DragonFly, the BIND daemon is called **named** for obvious reasons.

File	Description
named	the BIND daemon
<code>ndc</code>	name daemon control program

File	Description
<code>/etc/namedb</code>	directory where BIND zone information resides
<code>/etc/namedb/named.conf</code>	daemon configuration file

Zone files are usually contained within the `/etc/namedb` directory, and contain the DNS zone information served by the name server.

19.11.5 Starting BIND

Since BIND is installed by default, configuring it all is relatively simple.

To ensure the named daemon is started at boot, put the following modifications in `/etc/rc.conf`:

```
named_enable="YES"
```

To start the daemon manually (after configuring it)

```
# ndc start
```

19.11.6 Configuration Files

19.11.6.1 Using `make-localhost`

Be sure to:

```
# cd /etc/namedb
# sh make-localhost
```

to properly create the local reverse DNS zone file in `/etc/namedb/localhost.rev`.

19.11.6.2 `/etc/namedb/named.conf`

```
// $DragonFlyBSD$
//
// Refer to the named(8) manual page for details.  If you are ever going
// to setup a primary server, make sure you've understood the hairy
// details of how DNS is working.  Even with simple mistakes, you can
// break connectivity for affected parties, or cause huge amount of
// useless Internet traffic.

options {
    directory "/etc/namedb";

// In addition to the "forwarders" clause, you can force your name
// server to never initiate queries of its own, but always ask its
// forwarders only, by enabling the following line:
//
//     forward only;
```

```
// If you've got a DNS server around at your upstream provider, enter
// its IP address here, and enable the line below. This will make you
// benefit from its cache, thus reduce overall DNS traffic in the
// Internet.
/*
    forwarders {
        127.0.0.1;
    };
*/
```

Just as the comment says, to benefit from an uplink's cache, `forwarders` can be enabled here. Under normal circumstances, a name server will recursively query the Internet looking at certain name servers until it finds the answer it is looking for. Having this enabled will have it query the uplink's name server (or name server provided) first, taking advantage of its cache. If the uplink name server in question is a heavily trafficked, fast name server, enabling this may be worthwhile.

Warning: `127.0.0.1` will *not* work here. Change this IP address to a name server at your uplink.

```
/*
 * If there is a firewall between you and name servers you want
 * to talk to, you might need to uncomment the query-source
 * directive below. Previous versions of BIND always asked
 * questions using port 53, but BIND 8.1 uses an unprivileged
 * port by default.
 */
// query-source address * port 53;

/*
 * If running in a sandbox, you may have to specify a different
 * location for the dumpfile.
 */
// dump-file "s/named_dump.db";
};

// Note: the following will be supported in a future release.
/*
host { any; } {
    topology {
        127.0.0.0/8;
    };
};
*/

// Setting up secondaries is way easier and the rough picture for this
// is explained below.
//
// If you enable a local name server, don't forget to enter 127.0.0.1
// into your /etc/resolv.conf so this server will be queried first.
// Also, make sure to enable it in /etc/rc.conf.

zone "." {
```



```
zone "0.168.192.in-addr.arpa" {
    type slave;
    file "s/0.168.192.in-addr.arpa.bak";
    masters {
        192.168.1.1;
    };
};
*/
```

In `named.conf`, these are examples of slave entries for a forward and reverse zone.

For each new zone served, a new zone entry must be added to `named.conf`

For example, the simplest zone entry for `example.org` can look like:

```
zone "example.org" {
    type master;
    file "example.org";
};
```

The zone is a master, as indicated by the `type` statement, holding its zone information in `/etc/namedb/example.org` indicated by the `file` statement.

```
zone "example.org" {
    type slave;
    file "example.org";
};
```

In the slave case, the zone information is transferred from the master name server for the particular zone, and saved in the file specified. If and when the master server dies or is unreachable, the slave name server will have the transferred zone information and will be able to serve it.

19.11.6.3 Zone Files

An example master zone file for `example.org` (existing within `/etc/namedb/example.org`) is as follows:

```
$TTL 3600

example.org. IN SOA ns1.example.org. admin.example.org. (
                    5           ; Serial
                    10800        ; Refresh
                    3600         ; Retry
                    604800       ; Expire
                    86400 )      ; Minimum TTL

; DNS Servers
@           IN NS      ns1.example.org.
@           IN NS      ns2.example.org.

; Machine Names
localhost   IN A       127.0.0.1
ns1         IN A       3.2.1.2
ns2         IN A       3.2.1.3
```

```

mail           IN A      3.2.1.10
@             IN A      3.2.1.30

; Aliases
www           IN CNAME  @

; MX Record
@             IN MX    10      mail.example.org.

```

Note that every hostname ending in a “.” is an exact hostname, whereas everything without a trailing “.” is referenced to the origin. For example, `www` is translated into `www + origin`. In our fictitious zone file, our origin is `example.org.`, so `www` would translate to `www.example.org.`

The format of a zone file follows:

```
recordname     IN recordtype  value
```

The most commonly used DNS records:

SOA

start of zone authority

NS

an authoritative name server

A

A host address

CNAME

the canonical name for an alias

MX

mail exchanger

PTR

a domain name pointer (used in reverse DNS)

```

example.org. IN SOA ns1.example.org. admin.example.org. (
                    5                ; Serial
                    10800             ; Refresh after 3 hours
                    3600              ; Retry after 1 hour
                    604800            ; Expire after 1 week
                    86400 )          ; Minimum TTL of 1 day

```

`example.org.`

the domain name, also the origin for this zone file.

```
ns1.example.org.
```

the primary/authoritative name server for this zone

```
admin.example.org.
```

the responsible person for this zone, email address with @ replaced. (<admin@example.org> becomes admin.example.org)

```
5
```

the serial number of the file. this must be incremented each time the zone file is modified. Nowadays, many admins prefer a `yyymmddrr` format for the serial number. 2001041002 would mean last modified 04/10/2001, the latter 02 being the second time the zone file has been modified this day. The serial number is important as it alerts slave name servers for a zone when it is updated.

```
@          IN NS          ns1.example.org.
```

This is an NS entry. Every name server that is going to reply authoritatively for the zone must have one of these entries. The @ as seen here could have been `example.org`. The @ translates to the origin.

```
localhost      IN A      127.0.0.1
ns1             IN A      3.2.1.2
ns2            IN A      3.2.1.3
mail           IN A      3.2.1.10
@              IN A      3.2.1.30
```

The A record indicates machine names. As seen above, `ns1.example.org` would resolve to `3.2.1.2`. Again, the origin symbol, @, is used here, thus meaning `example.org` would resolve to `3.2.1.30`.

```
www           IN CNAME  @
```

The canonical name record is usually used for giving aliases to a machine. In the example, `www` is aliased to the machine addressed to the origin, or `example.org` (`3.2.1.30`). CNAMEs can be used to provide alias hostnames, or round robin one hostname among multiple machines.

```
@           IN MX     10      mail.example.org.
```

The MX record indicates which mail servers are responsible for handling incoming mail for the zone. `mail.example.org` is the hostname of the mail server, and 10 being the priority of that mail server.

One can have several mail servers, with priorities of 3, 2, 1. A mail server attempting to deliver to `example.org` would first try the highest priority MX, then the second highest, etc, until the mail can be properly delivered.

For `in-addr.arpa` zone files (reverse DNS), the same format is used, except with PTR entries instead of A or CNAME.

```
$TTL 3600
```

```
1.2.3.in-addr.arpa. IN SOA ns1.example.org. admin.example.org. (
                        5                ; Serial
                        10800             ; Refresh
                        3600              ; Retry
                        604800            ; Expire
                        3600 )            ; Minimum
```

```

@      IN NS    ns1.example.org.
@      IN NS    ns2.example.org.

2      IN PTR   ns1.example.org.
3      IN PTR   ns2.example.org.
10     IN PTR   mail.example.org.
30     IN PTR   example.org.

```

This file gives the proper IP address to hostname mappings of our above fictitious domain.

19.11.7 Caching Name Server

A caching name server is a name server that is not authoritative for any zones. It simply asks queries of its own, and remembers them for later use. To set one up, just configure the name server as usual, omitting any inclusions of zones.

19.11.8 Running named in a Sandbox

For added security you may want to run `named(8)` as an unprivileged user, and configure it to `chroot(8)` into a sandbox directory. This makes everything outside of the sandbox inaccessible to the `named` daemon. Should `named` be compromised, this will help to reduce the damage that can be caused. By default, DragonFly has a user and a group called `bind`, intended for this use.

Note: Various people would recommend that instead of configuring `named` to `chroot`, you should run `named` inside a `jail(8)`. This section does not attempt to cover this situation.

Since `named` will not be able to access anything outside of the sandbox (such as shared libraries, log sockets, and so on), there are a number of steps that need to be followed in order to allow `named` to function correctly. In the following checklist, it is assumed that the path to the sandbox is `/etc/namedb` and that you have made no prior modifications to the contents of this directory. Perform the following steps as `root`.

- Create all directories that `named` expects to see:

```

# cd /etc/namedb
# mkdir -p bin dev etc var/tmp var/run master slave
# chown bind:bind slave var/*❶

```

- ❶ `named` only needs write access to these directories, so that is all we give it.

- Rearrange and create basic zone and configuration files:

```

# cp /etc/localtime etc❶
# mv named.conf etc && ln -sf etc/named.conf
# mv named.root master
# sh make-localhost && mv localhost.rev localhost-v6.rev master
# cat > master/named.localhost
$ORIGIN localhost.

```

```
$TTL 6h
@ IN SOA localhost. postmaster.localhost. (
1 ; serial
3600 ; refresh
1800 ; retry
604800 ; expiration
3600 ) ; minimum
IN NS localhost.
IN A 127.0.0.1
^D
```

❶ This allows **named** to log the correct time to syslogd(8)

- Use cp(1) to copy named-xfer in /usr/libexec into your sandbox.
- Make a dev/null that **named** can see and write to:

```
# cd /etc/namedb/dev && mknod null c 2 2
# chmod 666 null
```

- Symlink /var/run/ndc to /etc/namedb/var/run/ndc:

```
# ln -sf /etc/namedb/var/run/ndc /var/run/ndc
```

Note: This simply avoids having to specify the `-c` option to ndc(8) every time you run it. Since the contents of /var/run are deleted on boot, if this is something that you find useful you may wish to add this command to root's crontab, making use of the `@reboot` option. See crontab(5) for more information regarding this.

- Configure syslogd(8) to create an extra log socket that **named** can write to. To do this, add `-l /etc/namedb/dev/log` to the `syslogd_flags` variable in `/etc/rc.conf`.
- Arrange to have **named** start and chroot itself to the sandbox by adding the following to `/etc/rc.conf`:

```
named_enable="YES"
named_flags="-u bind -g bind -t /etc/namedb /etc/named.conf"
```

Note: Note that the configuration file `/etc/named.conf` is denoted by a full pathname *relative to the sandbox*, i.e. in the line above, the file referred to is actually `/etc/namedb/etc/named.conf`.

The next step is to edit `/etc/namedb/etc/named.conf` so that **named** knows which zones to load and where to find them on the disk. There follows a commented example (anything not specifically commented here is no different from the setup for a DNS server not running in a sandbox):

```
options {
    directory "/";❶
    named-xfer "/bin/named-xfer";❷
    version ""; // Don't reveal BIND version
    query-source address * port 53;
};
```


After completing the steps above, either reboot your server or restart `syslogd(8)` and start `named(8)`, making sure to use the new options specified in `syslogd_flags` and `named_flags`. You should now be running a sandboxed copy of **named!**

19.11.9 Security

Although BIND is the most common implementation of DNS, there is always the issue of security. Possible and exploitable security holes are sometimes found.

It is a good idea to subscribe to CERT (<http://www.cert.org/>) and `freebsd-security-notifications` ([../handbook/eresources.html#ERESOURCES-MAIL](mailto:..handbook/eresources.html#ERESOURCES-MAIL)) to stay up to date with the current Internet and FreeBSD security issues.

Tip: If a problem arises, keeping sources up to date and having a fresh build of `named` would not hurt.

19.11.10 Further Reading

BIND/named manual pages: `ndc(8)` `named(8)` `named.conf(5)`

- Official ISC Bind Page (<http://www.isc.org/products/BIND/>)
- BIND FAQ (<http://www.nominum.com/getOpenSourceResource.php?id=6>)
- O'Reilly DNS and BIND 4th Edition (<http://www.oreilly.com/catalog/dns4/>)
- RFC1034 - Domain Names - Concepts and Facilities (<ftp://ftp.isi.edu/in-notes/rfc1034.txt>)
- RFC1035 - Domain Names - Implementation and Specification (<ftp://ftp.isi.edu/in-notes/rfc1035.txt>)

19.12 NTP

Contributed by Tom Hukins.

19.12.1 Overview

Over time, a computer's clock is prone to drift. As time passes, the computer's clock becomes less accurate. NTP (Network Time Protocol) is one way to ensure your clock is right.

Many Internet services rely on, or greatly benefit from, computers' clocks being accurate. For example, a Web server may receive requests to send a file if it has modified since a certain time. Services such as `cron(8)` run commands at a given time. If the clock is inaccurate, these commands may not run when expected.

DragonFly ships with the `ntpd(8)` NTP server which can be used to query other NTP servers to set the clock on your machine or provide time services to others.

19.12.2 Choosing Appropriate NTP Servers

In order to synchronize your clock, you will need to find one or more NTP servers to use. Your network administrator or ISP may have set up an NTP server for this purpose—check their documentation to see if this is the case. There is a list of publicly accessible NTP servers (<http://www.eecis.udel.edu/~mills/ntp/servers.html>) which you can use to find an NTP server near to you. Make sure you are aware of the policy for any servers you choose, and ask for permission if required.

Choosing several unconnected NTP servers is a good idea in case one of the servers you are using becomes unreachable or its clock is unreliable. `ntpd(8)` uses the responses it receives from other servers intelligently—it will favor unreliable servers less than reliable ones.

19.12.3 Configuring Your Machine

19.12.3.1 Basic Configuration

If you only wish to synchronize your clock when the machine boots up, you can use `ntpdate(8)`. This may be appropriate for some desktop machines which are frequently rebooted and only require infrequent synchronization, but most machines should run `ntpd(8)`.

Using `ntpdate(8)` at boot time is also a good idea for machines that run `ntpd(8)`. The `ntpd(8)` program changes the clock gradually, whereas `ntpdate(8)` sets the clock, no matter how great the difference between a machine's current clock setting and the correct time.

To enable `ntpdate(8)` at boot time, add `ntpdate_enable="YES"` to `/etc/rc.conf`. You will also need to specify all servers you wish to synchronize with and any flags to be passed to `ntpdate(8)` in `ntpdate_flags`.

19.12.3.2 General Configuration

NTP is configured by the `/etc/ntp.conf` file in the format described in `ntp.conf(5)`. Here is a simple example:

```
server ntplocal.example.com prefer
server timeserver.example.org
server ntp2a.example.net

driftfile /var/db/ntp.drift
```

The `server` option specifies which servers are to be used, with one server listed on each line. If a server is specified with the `prefer` argument, as with `ntplocal.example.com`, that server is preferred over other servers. A response from a preferred server will be discarded if it differs significantly from other servers' responses, otherwise it will be used without any consideration to other responses. The `prefer` argument is normally used for NTP servers that are known to be highly accurate, such as those with special time monitoring hardware.

The `driftfile` option specifies which file is used to store the system clock's frequency offset. The `ntpd(8)` program uses this to automatically compensate for the clock's natural drift, allowing it to maintain a reasonably correct setting even if it is cut off from all external time sources for a period of time.

The `driftfile` option specifies which file is used to store information about previous responses from the NTP servers you are using. This file contains internal information for NTP. It should not be modified by any other process.

19.12.3.3 Controlling Access to Your Server

By default, your NTP server will be accessible to all hosts on the Internet. The `restrict` option in `/etc/ntp.conf` allows you to control which machines can access your server.

If you want to deny all machines from accessing your NTP server, add the following line to `/etc/ntp.conf`:

```
restrict default ignore
```

If you only want to allow machines within your own network to synchronize their clocks with your server, but ensure they are not allowed to configure the server or used as peers to synchronize against, add

```
restrict 192.168.1.0 mask 255.255.255.0 notrust nomodify notrap
```

instead, where `192.168.1.0` is an IP address on your network and `255.255.255.0` is your network's netmask.

`/etc/ntp.conf` can contain multiple `restrict` options. For more details, see the `Access Control Support` subsection of `ntp.conf(5)`.

19.12.4 Running the NTP Server

To ensure the NTP server is started at boot time, add the line `xntpd_enable="YES"` to `/etc/rc.conf`. If you wish to pass additional flags to `ntpd(8)`, edit the `xntpd_flags` parameter in `/etc/rc.conf`.

To start the server without rebooting your machine, run `ntpd` being sure to specify any additional parameters from `xntpd_flags` in `/etc/rc.conf`. For example:

```
# ntpd -p /var/run/ntpd.pid
```

19.12.5 Using ntpd with a Temporary Internet Connection

The `ntpd(8)` program does not need a permanent connection to the Internet to function properly. However, if you have a temporary connection that is configured to dial out on demand, it is a good idea to prevent NTP traffic from triggering a dial out or keeping the connection alive. If you are using user PPP, you can use `filter` directives in `/etc/ppp/ppp.conf`. For example:

```
set filter dial 0 deny udp src eq 123
# Prevent NTP traffic from initiating dial out
set filter dial 1 permit 0 0
set filter alive 0 deny udp src eq 123
# Prevent incoming NTP traffic from keeping the connection open
set filter alive 1 deny udp dst eq 123
# Prevent outgoing NTP traffic from keeping the connection open
set filter alive 2 permit 0/0 0/0
```

For more details see the `PACKET FILTERING` section in `ppp(8)` and the examples in `/usr/share/examples/ppp/`.

Note: Some Internet access providers block low-numbered ports, preventing NTP from functioning since replies never reach your machine.

19.12.6 Further Information

Documentation for the NTP server can be found in `/usr/share/doc/ntp/` in HTML format.

19.13 Network Address Translation

Contributed by Chern Lee.

19.13.1 Overview

DragonFly's Network Address Translation daemon, commonly known as `natd(8)` is a daemon that accepts incoming raw IP packets, changes the source to the local machine and re-injects these packets back into the outgoing IP packet stream. `natd(8)` does this by changing the source IP address and port such that when data is received back, it is able to determine the original location of the data and forward it back to its original requester.


The most common use of NAT is to perform what is commonly known as Internet Connection Sharing.

19.13.2 Setup

Due to the diminishing IP space in IPv4, and the increased number of users on high-speed consumer lines such as cable or DSL, people are increasingly in need of an Internet Connection Sharing solution. The ability to connect several computers online through one connection and IP address makes `natd(8)` a reasonable choice.

Most commonly, a user has a machine connected to a cable or DSL line with one IP address and wishes to use this one connected computer to provide Internet access to several more over a LAN.

To do this, the DragonFly machine on the Internet must act as a gateway. This gateway machine must have two NICs—one for connecting to the Internet router, the other connecting to a LAN. All the machines on the LAN are connected through a hub or switch.



A setup like this is commonly used to share an Internet connection. One of the LAN machines is connected to the Internet. The rest of the machines access the Internet through that “gateway” machine.

19.13.3 Configuration

The following options must be in the kernel configuration file:

```
options IPFIREWALL
options IPDIVERT
```

Additionally, at choice, the following may also be suitable:

```
options IPFIREWALL_DEFAULT_TO_ACCEPT
options IPFIREWALL_VERBOSE
```

The following must be in `/etc/rc.conf`:

```
gateway_enable="YES"
firewall_enable="YES"
firewall_type="OPEN"
natd_enable="YES"
natd_interface="fxp0"
natd_flags=""
```

`gateway_enable="YES"`

Sets up the machine to act as a gateway. Running `sysctl net.inet.ip.forwarding=1` would have the same effect.

`firewall_enable="YES"`

Enables the firewall rules in `/etc/rc.firewall` at boot.

`firewall_type="OPEN"`

This specifies a predefined firewall ruleset that allows anything in. See `/etc/rc.firewall` for additional types.

`natd_interface="fxp0"`

Indicates which interface to forward packets through (the interface connected to the Internet).

`natd_flags=""`

Any additional configuration options passed to `natd(8)` on boot.

Having the previous options defined in `/etc/rc.conf` would run `natd -interface fxp0` at boot. This can also be run manually.

Note: It is also possible to use a configuration file for `natd(8)` when there are too many options to pass. In this case, the configuration file must be defined by adding the following line to `/etc/rc.conf`:

```
natd_flags="-f /etc/natd.conf"
```

The `/etc/natd.conf` file will contain a list of configuration options, one per line. For example the next section case would use the following file:

```
redirect_port tcp 192.168.0.2:6667 6667
redirect_port tcp 192.168.0.3:80 80
```

For more information about the configuration file, consult the `natd(8)` manual page about the `-f` option.

Each machine and interface behind the LAN should be assigned IP address numbers in the private network space as defined by RFC 1918 (<ftp://ftp.isi.edu/in-notes/rfc1918.txt>) and have a default gateway of the **natd** machine's internal IP address.

For example, client A and B behind the LAN have IP addresses of 192.168.0.2 and 192.168.0.3, while the natd machine's LAN interface has an IP address of 192.168.0.1. Client A and B's default gateway must be set to that of the **natd** machine, 192.168.0.1. The **natd** machine's external, or Internet interface does not require any special modification for natd(8) to work.

19.13.4 Port Redirection

The drawback with natd(8) is that the LAN clients are not accessible from the Internet. Clients on the LAN can make outgoing connections to the world but cannot receive incoming ones. This presents a problem if trying to run Internet services on one of the LAN client machines. A simple way around this is to redirect selected Internet ports on the **natd** machine to a LAN client.

For example, an IRC server runs on client A, and a web server runs on client B. For this to work properly, connections received on ports 6667 (IRC) and 80 (web) must be redirected to the respective machines.

The `-redirect_port` must be passed to natd(8) with the proper options. The syntax is as follows:

```
-redirect_port proto targetIP:targetPORT[-targetPORT]
                [aliasIP:]aliasPORT[-aliasPORT]
                [remoteIP[:remotePORT[-remotePORT]]]
```

In the above example, the argument should be:

```
-redirect_port tcp 192.168.0.2:6667 6667
-redirect_port tcp 192.168.0.3:80 80
```

This will redirect the proper *tcp* ports to the LAN client machines.

The `-redirect_port` argument can be used to indicate port ranges over individual ports. For example, `tcp 192.168.0.2:2000-3000 2000-3000` would redirect all connections received on ports 2000 to 3000 to ports 2000 to 3000 on client A.

These options can be used when directly running natd(8), placed within the `natd_flags=""` option in `/etc/rc.conf`, or passed via a configuration file.

For further configuration options, consult natd(8)

19.13.5 Address Redirection

Address redirection is useful if several IP addresses are available, yet they must be on one machine. With this, natd(8) can assign each LAN client its own external IP address. natd(8) then rewrites outgoing packets from the LAN clients with the proper external IP address and redirects all traffic incoming on that particular IP address back to the specific LAN client. This is also known as static NAT. For example, the IP addresses 128.1.1.1, 128.1.1.2, and 128.1.1.3 belong to the **natd** gateway machine. 128.1.1.1 can be used as the **natd** gateway machine's external IP address, while 128.1.1.2 and 128.1.1.3 are forwarded back to LAN clients A and B.

The `-redirect_address` syntax is as follows:

```
-redirect_address localIP publicIP
```

localIP	The internal IP address of the LAN client.
publicIP	The external IP address corresponding to the LAN client.

In the example, this argument would read:

```
-redirect_address 192.168.0.2 128.1.1.2
-redirect_address 192.168.0.3 128.1.1.3
```

Like `-redirect_port`, these arguments are also placed within the `natd_flags=""` option of `/etc/rc.conf`, or passed via a configuration file. With address redirection, there is no need for port redirection since all data received on a particular IP address is redirected.

The external IP addresses on the **natd** machine must be active and aliased to the external interface. Look at `rc.conf(5)` to do so.

19.14 The `inetd` “Super-Server”

Contributed by Chern Lee.

19.14.1 Overview

`inetd(8)` is referred to as the “Internet Super-Server” because it manages connections for several daemons. Programs that provide network service are commonly known as daemons. **inetd** serves as a managing server for other daemons. When a connection is received by **inetd**, it determines which daemon the connection is destined for, spawns the particular daemon and delegates the socket to it. Running one instance of **inetd** reduces the overall system load as compared to running each daemon individually in stand-alone mode.

Primarily, **inetd** is used to spawn other daemons, but several trivial protocols are handled directly, such as **chargen**, **auth**, and **daytime**.

This section will cover the basics in configuring **inetd** through its command-line options and its configuration file, `/etc/inetd.conf`.

19.14.2 Settings

inetd is initialized through the `/etc/rc.conf` system. The `inetd_enable` option is set to `NO` by default. Placing:

```
inetd_enable="YES"
```

or

```
inetd_enable="NO"
```

into `/etc/rc.conf` can enable or disable **inetd** starting at boot time.

Additionally, different command-line options can be passed to **inetd** via the `inetd_flags` option.

19.14.3 Command-Line Options

inetd synopsis:

```
inetd [-d] [-l] [-w] [-W] [-c maximum] [-C rate] [-a address | hostname] [-p filename]
[-R rate] [configuration file]
```

-d

Turn on debugging.

-l

Turn on logging of successful connections.

-w

Turn on TCP Wrapping for external services (on by default).

-W

Turn on TCP Wrapping for internal services which are built into **inetd** (on by default).

-c maximum

Specify the default maximum number of simultaneous invocations of each service; the default is unlimited. May be overridden on a per-service basis with the `max-child` parameter.

-C rate

Specify the default maximum number of times a service can be invoked from a single IP address in one minute; the default is unlimited. May be overridden on a per-service basis with the `max-connections-per-ip-per-minute` parameter.

-R rate

Specify the maximum number of times a service can be invoked in one minute; the default is 256. A rate of 0 allows an unlimited number of invocations.

-a

Specify one specific IP address to bind to. Alternatively, a hostname can be specified, in which case the IPv4 or IPv6 address which corresponds to that hostname is used. Usually a hostname is specified when **inetd** is run inside a `jail(8)`, in which case the hostname corresponds to the `jail(8)` environment.

When hostname specification is used and both IPv4 and IPv6 bindings are desired, one entry with the appropriate protocol type for each binding is required for each service in `/etc/inetd.conf`. For example, a TCP-based service would need two entries, one using `tcp4` for the protocol and the other using `tcp6`.

-P

Specify an alternate file in which to store the process ID.

These options can be passed to **inetd** using the `inetd_flags` option in `/etc/rc.conf`. By default, `inetd_flags` is set to `-wW`, which turns on TCP wrapping for **inetd**'s internal and external services. For novice users, these parameters usually do not need to be modified or even entered in `/etc/rc.conf`.

Note: An external service is a daemon outside of **inetd**, which is invoked when a connection is received for it. On the other hand, an internal service is one that **inetd** has the facility of offering within itself.

19.14.4 inetd.conf

Configuration of **inetd** is controlled through the `/etc/inetd.conf` file.

When a modification is made to `/etc/inetd.conf`, **inetd** can be forced to re-read its configuration file by sending a HangUP signal to the **inetd** process as shown:

Example 19-4. Sending inetd a HangUP Signal

```
# kill -HUP `cat /var/run/inetd.pid`
```

Each line of the configuration file specifies an individual daemon. Comments in the file are preceded by a “#”. The format of `/etc/inetd.conf` is as follows:

```
service-name
socket-type
protocol
{wait|nowait}[/max-child[/max-connections-per-ip-per-minute]]
user[:group][[/login-class]]
server-program
server-program-arguments
```

An example entry for the **ftpd** daemon using IPv4:

```
ftp      stream  tcp      nowait  root    /usr/libexec/ftpd      ftpd -l
```

service-name

This is the service name of the particular daemon. It must correspond to a service listed in `/etc/services`.

This determines which port **inetd** must listen to. If a new service is being created, it must be placed in `/etc/services` first.

socket-type

Either `stream`, `dgram`, `raw`, or `seqpacket`. `stream` must be used for connection-based, TCP daemons, while `dgram` is used for daemons utilizing the UDP transport protocol.

protocol

One of the following:

Protocol	Explanation
tcp, tcp4	TCP IPv4
udp, udp4	UDP IPv4
tcp6	TCP IPv6
udp6	UDP IPv6

Protocol	Explanation
tcp46	Both TCP IPv4 and v6
udp46	Both UDP IPv4 and v6

{wait|nowait}[/max-child[/max-connections-per-ip-per-minute]]

`wait|nowait` indicates whether the daemon invoked from **inetd** is able to handle its own socket or not. `dgram` socket types must use the `wait` option, while stream socket daemons, which are usually multi-threaded, should use `nowait`. `wait` usually hands off multiple sockets to a single daemon, while `nowait` spawns a child daemon for each new socket.

The maximum number of child daemons **inetd** may spawn can be set using the `max-child` option. If a limit of ten instances of a particular daemon is needed, a `/10` would be placed after `nowait`.

In addition to `max-child`, another option limiting the maximum connections from a single place to a particular daemon can be enabled. `max-connections-per-ip-per-minute` does just this. A value of ten here would limit any particular IP address connecting to a particular service to ten attempts per minute. This is useful to prevent intentional or unintentional resource consumption and Denial of Service (DoS) attacks to a machine.

In this field, `wait` or `nowait` is mandatory. `max-child` and `max-connections-per-ip-per-minute` are optional.

A stream-type multi-threaded daemon without any `max-child` or `max-connections-per-ip-per-minute` limits would simply be: `nowait`.

The same daemon with a maximum limit of ten daemons would read: `nowait/10`.

Additionally, the same setup with a limit of twenty connections per IP address per minute and a maximum total limit of ten child daemons would read: `nowait/10/20`.

These options are all utilized by the default settings of the **fingerd** daemon, as seen here:

```
finger stream tcp      nowait/3/10 nobody /usr/libexec/fingerd fingerd -s
```

user

This is the username that the particular daemon should run as. Most commonly, daemons run as the `root` user. For security purposes, it is common to find some servers running as the `daemon` user, or the least privileged `nobody` user.

server-program

The full path of the daemon to be executed when a connection is received. If the daemon is a service provided by **inetd** internally, then `internal` should be used.

server-program-arguments

This works in conjunction with `server-program` by specifying the arguments, starting with `argv[0]`, passed to the daemon on invocation. If `mydaemon -d` is the command line, `mydaemon -d` would be the value of `server-program-arguments`. Again, if the daemon is an internal service, use `internal` here.

19.14.5 Security

Depending on the security profile chosen at install, many of **inetd**'s daemons may be enabled by default. If there is no apparent need for a particular daemon, disable it! Place a “#” in front of the daemon in question, and send a hangup signal to **inetd**. Some daemons, such as **fingerd**, may not be desired at all because they provide an attacker with too much information.

Some daemons are not security-conscious and have long, or non-existent timeouts for connection attempts. This allows an attacker to slowly send connections to a particular daemon, thus saturating available resources. It may be a good idea to place `max-connections-per-ip-per-minute` and `max-child` limitations on certain daemons.

By default, TCP wrapping is turned on. Consult the `hosts_access(5)` manual page for more information on placing TCP restrictions on various **inetd** invoked daemons.

19.14.6 Miscellaneous

daytime, **time**, **echo**, **discard**, **chargen**, and **auth** are all internally provided services of **inetd**.

The **auth** service provides identity (**ident**, **identd**) network services, and is configurable to a certain degree.

Consult the `inetd(8)` manual page for more in-depth information.

19.15 Parallel Line IP (PLIP)

PLIP lets us run TCP/IP between parallel ports. It is useful on machines without network cards, or to install on laptops. In this section, we will discuss:

- Creating a parallel (laplink) cable.
- Connecting two computers with PLIP.

19.15.1 Creating a Parallel Cable

You can purchase a parallel cable at most computer supply stores. If you cannot do that, or you just want to know how it is done, the following table shows how to make one out of a normal parallel printer cable.

Table 19-1. Wiring a Parallel Cable for Networking

A-name	A-End	B-End	Descr.	Post/Bit
DATA0 -ERROR	2 15	15 2	Data	0/0x01 1/0x08
DATA1 +SLCT	3 13	13 3	Data	0/0x02 1/0x10
DATA2 +PE	4 12	12 4	Data	0/0x04 1/0x20
DATA3 -ACK	5 10	10 5	Strobe	0/0x08 1/0x40

A-name	A-End	B-End	Descr.	Post/Bit
DATA4	6	11	Data	0/0x10
BUSY	11	6		1/0x80
GND	18-25	18-25	GND	-

19.15.2 Setting Up PLIP

First, you have to get a laplink cable. Then, confirm that both computers have a kernel with lpt(4) driver support:

```
# grep lp /var/run/dmesg.boot
lpt0: <Printer> on ppbus0
lpt0: Interrupt-driven port
```

The parallel port must be an interrupt driven port. You should have a line similar to the following in your kernel configuration file:

```
device ppc0 at isa? irq 7
```

Then check if the kernel configuration file has a `device plip` line or if the `plip.ko` kernel module is loaded. In both cases the parallel networking interface should appear when you directly use the `ifconfig(8)` command.

```
# ifconfig lp0
lp0: flags=8810<POINTOPOINT,SIMPLEX,MULTICAST> mtu 1500
```

Plug in the laplink cable into the parallel interface on both computers.

Configure the network interface parameters on both sites as `root`. For example, if you want connect the host `host1` with `host2`:

```

                host1 <-----> host2
IP Address    10.0.0.1      10.0.0.2
```

Configure the interface on `host1` by doing:

```
# ifconfig lp0 10.0.0.1 10.0.0.2
```

Configure the interface on `host2` by doing:

```
# ifconfig lp0 10.0.0.2 10.0.0.1
```

You now should have a working connection. Please read the manual pages `lp(4)` and `lpt(4)` for more details.

You should also add both hosts to `/etc/hosts`:

```
127.0.0.1          localhost.my.domain localhost
10.0.0.1          host1.my.domain host1
10.0.0.2          host2.my.domain
```

To confirm the connection works, go to each host and ping the other. For example, on `host1`:

```
# ifconfig lp0
lp0: flags=8851<UP,POINTOPOINT,RUNNING,SIMPLEX,MULTICAST> mtu 1500
```

```

    inet 10.0.0.1 --> 10.0.0.2 netmask 0xff000000
# netstat -r
Routing tables

Internet:
Destination          Gateway             Flags      Refs      Use      Netif Expire
host2                 host1              UH         0         0         lp0
# ping -c 4 host2
PING host2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=255 time=2.774 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=255 time=2.530 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=255 time=2.556 ms
64 bytes from 10.0.0.2: icmp_seq=3 ttl=255 time=2.714 ms

--- host2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.530/2.643/2.774/0.103 ms

```

19.16 IPv6

Originally Written by Aaron Kaplan. Restructured and Added by Tom Rhodes.

IPv6 (also known as IPng “IP next generation”) is the new version of the well known IP protocol (also known as IPv4). Like the other current *BSD systems, DragonFly includes the KAME IPv6 reference implementation. So your DragonFly system comes with all you will need to experiment with IPv6. This section focuses on getting IPv6 configured and running.

In the early 1990s, people became aware of the rapidly diminishing address space of IPv4. Given the expansion rate of the Internet there were two major concerns:

- Running out of addresses. Today this is not so much of a concern anymore since private address spaces (10.0.0.0/8, 192.168.0.0/24, etc.) and Network Address Translation (NAT) are being employed.
- Router table entries were getting too large. This is still a concern today.

IPv6 deals with these and many other issues:

- 128 bit address space. In other words theoretically there are 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses available. This means there are approximately $6.67 * 10^{27}$ IPv6 addresses per square meter on our planet.
- Routers will only store network aggregation addresses in their routing tables thus reducing the average space of a routing table to 8192 entries.

There are also lots of other useful features of IPv6 such as:

- Address autoconfiguration (RFC2462 (<http://www.ietf.org/rfc/rfc2462.txt>))
- Anycast addresses (“one-out-of many”)
- Mandatory multicast addresses

- IPsec (IP security)
- Simplified header structure
- Mobile IP
- IPv4-to-IPv6 transition mechanisms

For more information see:

- IPv6 overview at playground.sun.com (<http://playground.sun.com/pub/ipng/html/ipng-main.html>)
- KAME.net (<http://www.kame.net>)
- 6bone.net (<http://www.6bone.net>)

19.16.1 Background on IPv6 Addresses

There are different types of IPv6 addresses: Unicast, Anycast and Multicast.

Unicast addresses are the well known addresses. A packet sent to a unicast address arrives exactly at the interface belonging to the address.

Anycast addresses are syntactically indistinguishable from unicast addresses but they address a group of interfaces. The packet destined for an anycast address will arrive at the nearest (in router metric) interface. Anycast addresses may only be used by routers.

Multicast addresses identify a group of interfaces. A packet destined for a multicast address will arrive at all interfaces belonging to the multicast group.

Note: The IPv4 broadcast address (usually `xxx.xxx.xxx.255`) is expressed by multicast addresses in IPv6.

Table 19-2. Reserved IPv6 addresses

IPv6 address	Prefixlength (Bits)	Description	Notes
::	128 bits	unspecified	cf. 0.0.0.0 in IPv4
:::1	128 bits	loopback address	cf. 127.0.0.1 in IPv4
::00:xx:xx:xx:xx	96 bits	embedded IPv4	The lower 32 bits are the IPv4 address. Also called "IPv4 compatible IPv6 address"
::ff:xx:xx:xx:xx	96 bits	IPv4 mapped IPv6 address	The lower 32 bits are the IPv4 address. For hosts which do not support IPv6.
fe80:: - feb::	10 bits	link-local	cf. loopback address in IPv4
fec0:: - fef::	10 bits	site-local	
ff::	8 bits	multicast	

IPv6 address	Prefixlength (Bits)	Description	Notes
001 (base 2)	3 bits	global unicast	All global unicast addresses are assigned from this pool. The first 3 bits are "001".

19.16.2 Reading IPv6 Addresses

The canonical form is represented as: `x:x:x:x:x:x:x`, each "x" being a 16 Bit hex value. For example `FEBC:A574:382B:23C1:AA49:4592:4EFE:9982`

Often an address will have long substrings of all zeros therefore each such substring can be abbreviated by "::". For example `fe80::1` corresponds to the canonical form `fe80:0000:0000:0000:0000:0000:0000:0001`.

A third form is to write the last 32 Bit part in the well known (decimal) IPv4 style with dots "." as separators. For example `2002::10.0.0.1` corresponds to the (hexadecimal) canonical representation `2002:0000:0000:0000:0000:0000:0a00:0001` which in turn is equivalent to writing `2002::a00:1`.

By now the reader should be able to understand the following:

```
# ifconfig
r10: flags=8943<UP,BROADCAST,RUNNING,PROMISC,SIMPLEX,MULTICAST> mtu 1500
    inet 10.0.0.10 netmask 0xffffffff broadcast 10.0.0.255
    inet6 fe80::200:21ff:fe03:8e1%r10 prefixlen 64 scopeid 0x1
    ether 00:00:21:03:08:e1
    media: Ethernet autoselect (100baseTX )
    status: active
```

`fe80::200:21ff:fe03:8e1%r10` is an auto configured link-local address. It includes the scrambled Ethernet MAC as part of the auto configuration.

For further information on the structure of IPv6 addresses see RFC3513 (<http://www.ietf.org/rfc/rfc3513.txt>).

19.16.3 Getting Connected

Currently there are four ways to connect to other IPv6 hosts and networks:

- Join the experimental 6bone
- Getting an IPv6 network from your upstream provider. Talk to your Internet provider for instructions.
- Tunnel via 6-to-4
- Use the `net/freenet6` port if you are on a dial-up connection.

Here we will talk on how to connect to the 6bone since it currently seems to be the most popular way.

First take a look at the 6bone (<http://www.6bone.net/>) site and find a 6bone connection nearest to you. Write to the responsible person and with a little bit of luck you will be given instructions on how to set up your connection. Usually this involves setting up a GRE (gif) tunnel.

Here is a typical example on setting up a gif(4) tunnel:

```
# ifconfig gif0 create
# ifconfig gif0
gif0: flags=8010<POINTOPOINT,MULTICAST> mtu 1280
# ifconfig gif0 tunnel MY_IPv4_ADDR HIS_IPv4_ADDR
# ifconfig gif0 inet6 alias MY_ASSIGNED_IPv6_TUNNEL_ENDPOINT_ADDR
```

Replace the capitalized words by the information you received from the upstream 6bone node.

This establishes the tunnel. Check if the tunnel is working by ping6(8) 'ing ff02::1%gif0. You should receive two ping replies.

Note: In case you are intrigued by the address ff02:1%gif0, this is a multicast address. %gif0 states that the multicast address at network interface gif0 is to be used. Since we ping a multicast address the other endpoint of the tunnel should reply as well.

By now setting up a route to your 6bone uplink should be rather straightforward:

```
# route add -inet6 default -interface gif0
# ping6 -n MY_UPLINK

# traceroute6 www.jp.FreeBSD.org
(3ffe:505:2008:1:2a0:24ff:fe57:e561) from 3ffe:8060:100::40:2, 30 hops max, 12 byte packets
 1  atnet-meta6  14.147 ms  15.499 ms  24.319 ms
 2  6bone-gw2-ATNET-NT.ipv6.tilab.com  103.408 ms  95.072 ms *
 3  3ffe:1831:0:ffff::4  138.645 ms  134.437 ms  144.257 ms
 4  3ffe:1810:0:6:290:27ff:fe79:7677  282.975 ms  278.666 ms  292.811 ms
 5  3ffe:1800:0:ff00::4  400.131 ms  396.324 ms  394.769 ms
 6  3ffe:1800:0:3:290:27ff:fe14:cdee  394.712 ms  397.19 ms  394.102 ms
```

This output will differ from machine to machine. By now you should be able to reach the IPv6 site www.kame.net (<http://www.kame.net>) and see the dancing tortoise — that is if you have a IPv6 enabled browser such as www/mozilla.

19.16.4 DNS in the IPv6 World

There are two new types of DNS records for IPv6:

- AAAA records,
- A6 records

Using AAAA records is straightforward. Assign your hostname to the new IPv6 address you just got by adding:

```
MYHOSTNAME          AAAA      MYIPv6ADDR
```

To your primary zone DNS file. In case you do not serve your own DNS zones ask your DNS provider. Current versions of **bind** (version 8.3 and 9) support AAAA records.

Chapter 20 Electronic Mail

Original work by Bill Lloyd. Rewritten by Jim Mock.

20.1 Synopsis

“Electronic Mail”, better known as email, is one of the most widely used forms of communication today. This chapter provides a basic introduction to running a mail server on DragonFly, as well as an introduction to sending and receiving email using DragonFly; however, it is not a complete reference and in fact many important considerations are omitted. For more complete coverage of the subject, the reader is referred to the many excellent books listed in Appendix B.

After reading this chapter, you will know:

- What software components are involved in sending and receiving electronic mail.
- Where basic **sendmail** configuration files are located in DragonFly.
- The difference between remote and local mailboxes.
- How to block spammers from illegally using your mail server as a relay.
- How to install and configure an alternate Mail Transfer Agent on your system, replacing **sendmail**.
- How to troubleshoot common mail server problems.
- How to use SMTP with UUCP.
- How to set up the system to send mail only.
- How to use mail with a dialup connection.
- How to configure SMTP Authentication for added security.
- How to install and use a Mail User Agent, such as **mutt** to send and receive email.
- How to download your mail from a remote POP or IMAP server.
- How to automatically apply filters and rules to incoming email.

Before reading this chapter, you should:

- Properly set up your network connection (Chapter 19).
- Properly set up the DNS information for your mail host (Chapter 19).
- Know how to install additional third-party software (Chapter 4).

20.2 Using Electronic Mail

There are five major parts involved in an email exchange. They are: the user program, the server daemon, DNS, a remote or local mailbox, and of course, the mailhost itself.

20.2.1 The User Program

This includes command line programs such as **mutt**, **pine**, **elm**, and **mail**, and GUI programs such as **balsa**, **xfmail** to name a few, and something more “sophisticated” like a WWW browser. These programs simply pass off the email transactions to the local “mailhost”, either by calling one of the server daemons available, or delivering it over TCP.

20.2.2 Mailhost Server Daemon

DragonFly ships with **sendmail** by default, but also support numerous other mail server daemons, just some of which include:

- **exim**;
- **postfix**;
- **qmail**.

The server daemon usually has two functions—it is responsible for receiving incoming mail as well as delivering outgoing mail. It is *not* responsible for the collection of mail using protocols such as POP or IMAP to read your email, nor does it allow connecting to local `mbox` or Maildir mailboxes. You may require an additional daemon for that.

Warning: Older versions of **sendmail** have some serious security issues which may result in an attacker gaining local and/or remote access to your machine. Make sure that you are running a current version to avoid these problems. Optionally, install an alternative MTA from the DragonFly Pkgsrc Collection.

20.2.3 Email and DNS

The Domain Name System (DNS) and its daemon `named` play a large role in the delivery of email. In order to deliver mail from your site to another, the server daemon will look up the remote site in the DNS to determine the host that will receive mail for the destination. This process also occurs when mail is sent from a remote host to your mail server.

DNS is responsible for mapping hostnames to IP addresses, as well as for storing information specific to mail delivery, known as MX records. The MX (Mail eXchanger) record specifies which host, or hosts, will receive mail for a particular domain. If you do not have an MX record for your hostname or domain, the mail will be delivered directly to your host provided you have an A record pointing your hostname to your IP address.

You may view the MX records for any domain by using the `host(1)` command, as seen in the example below:

```
% host -t mx DragonflyBSD.org
DragonflyBSD.org mail is handled (pri=10) by crater.dragonflybsd.org
```

20.2.4 Receiving Mail

Receiving mail for your domain is done by the mail host. It will collect all mail sent to your domain and store it either in `mbox` (the default method for storing mail) or Maildir format, depending on your configuration. Once mail has been stored, it may either be read locally using applications such as `mail(1)` or **mutt**, or remotely accessed and

collected using protocols such as POP or IMAP. This means that should you only wish to read mail locally, you are not required to install a POP or IMAP server.

20.2.4.1 Accessing remote mailboxes using POP and IMAP

In order to access mailboxes remotely, you are required to have access to a POP or IMAP server. These protocols allow users to connect to their mailboxes from remote locations with ease. Though both POP and IMAP allow users to remotely access mailboxes, IMAP offers many advantages, some of which are:

- IMAP can store messages on a remote server as well as fetch them.
- IMAP supports concurrent updates.
- IMAP can be extremely useful over low-speed links as it allows users to fetch the structure of messages without downloading them; it can also perform tasks such as searching on the server in order to minimize data transfer between clients and servers.

In order to install a POP or IMAP server, the following steps should be performed:

1. Choose an IMAP or POP server that best suits your needs. The following POP and IMAP servers are well known and serve as some good examples:
 - **qpopper**;
 - **teapop**;
 - **imap-uw**;
 - **courier-imap**;
2. Install the POP or IMAP daemon of your choosing from the ports collection.
3. Where required, modify `/etc/inetd.conf` to load the POP or IMAP server.

Warning: It should be noted that both POP and IMAP transmit information, including username and password credentials in clear-text. This means that if you wish to secure the transmission of information across these protocols, you should consider tunneling sessions over `ssh(1)`. Tunneling sessions is described in Section 10.10.7.

20.2.4.2 Accessing local mailboxes

Mailboxes may be accessed locally by directly utilizing MUAs on the server on which the mailbox resides. This can be done using applications such as **mutt** or `mail(1)`.

20.2.5 The Mail Host

The mail host is the name given to a server that is responsible for delivering and receiving mail for your host, and possibly your network.

20.3 sendmail Configuration

Contributed by Christopher Shumway.

sendmail(8) is the default Mail Transfer Agent (MTA) in DragonFly. **sendmail**'s job is to accept mail from Mail User Agents (MUA) and deliver it to the appropriate mailer as defined by its configuration file. **sendmail** can also accept network connections and deliver mail to local mailboxes or deliver it to another program.

sendmail uses the following configuration files:

Filename	Function
/etc/mail/access	sendmail access database file
/etc/mail/aliases	Mailbox aliases
/etc/mail/local-host-names	Lists of hosts sendmail accepts mail for
/etc/mail/mailer.conf	Mailer program configuration
/etc/mail/mailertable	Mailer delivery table
/etc/mail/sendmail.cf	sendmail master configuration file
/etc/mail/virtusertable	Virtual users and domain tables

20.3.1 /etc/mail/access

The access database defines what host(s) or IP addresses have access to the local mail server and what kind of access they have. Hosts can be listed as OK, REJECT, RELAY or simply passed to **sendmail**'s error handling routine with a given mailer error. Hosts that are listed as OK, which is the default, are allowed to send mail to this host as long as the mail's final destination is the local machine. Hosts that are listed as REJECT are rejected for all mail connections. Hosts that have the RELAY option for their hostname are allowed to send mail for any destination through this mail server.

Example 20-1. Configuring the sendmail Access Database

```
cyberspammer.com          550 We don't accept mail from spammers
FREE.STEALTH.MAILER@     550 We don't accept mail from spammers
another.source.of.spam   REJECT
okay.cyberspammer.com    OK
128.32                   RELAY
```

In this example we have five entries. Mail senders that match the left hand side of the table are affected by the action on the right side of the table. The first two examples give an error code to **sendmail**'s error handling routine. The message is printed to the remote host when a mail matches the left hand side of the table. The next entry rejects mail from a specific host on the Internet, `another.source.of.spam`. The next entry accepts mail connections from a host `okay.cyberspammer.com`, which is more exact than the `cyberspammer.com` line above. More specific matches override less exact matches. The last entry allows relaying of electronic mail from hosts with an IP address that begins with `128.32`. These hosts would be able to send mail through this mail server that are destined for other mail servers.

When this file is updated, you need to run `make in /etc/mail/` to update the database.

20.3.2 /etc/mail/aliases

The aliases database contains a list of virtual mailboxes that are expanded to other user(s), files, programs or other aliases. Here are a few examples that can be used in /etc/mail/aliases:

Example 20-2. Mail Aliases

```
root: localuser
ftp-bugs: joe,eric,paul
bit.bucket: /dev/null
procmail: "|/usr/local/bin/procmail"
```

The file format is simple; the mailbox name on the left side of the colon is expanded to the target(s) on the right. The first example simply expands the mailbox `root` to the mailbox `localuser`, which is then looked up again in the aliases database. If no match is found, then the message is delivered to the local user `localuser`. The next example shows a mail list. Mail to the mailbox `ftp-bugs` is expanded to the three local mailboxes `joe`, `eric`, and `paul`. Note that a remote mailbox could be specified as `user@example.com`. The next example shows writing mail to a file, in this case `/dev/null`. The last example shows sending mail to a program, in this case the mail message is written to the standard input of `/usr/local/bin/procmail` through a UNIX pipe.

When this file is updated, you need to run `make` in `/etc/mail/` to update the database.

20.3.3 /etc/mail/local-host-names

This is a list of hostnames `sendmail(8)` is to accept as the local host name. Place any domains or hosts that **sendmail** is to be receiving mail for. For example, if this mail server was to accept mail for the domain `example.com` and the host `mail.example.com`, its `local-host-names` might look something like this:

```
example.com
mail.example.com
```

When this file is updated, `sendmail(8)` needs to be restarted to read the changes.

20.3.4 /etc/mail/sendmail.cf

sendmail's master configuration file, `sendmail.cf` controls the overall behavior of **sendmail**, including everything from rewriting e-mail addresses to printing rejection messages to remote mail servers. Naturally, with such a diverse role, this configuration file is quite complex and its details are a bit out of the scope of this section. Fortunately, this file rarely needs to be changed for standard mail servers.

The master **sendmail** configuration file can be built from `m4(1)` macros that define the features and behavior of **sendmail**. Please see `/usr/src/contrib/sendmail/cf/README` for some of the details.

When changes to this file are made, **sendmail** needs to be restarted for the changes to take effect.

20.3.5 /etc/mail/virtusertable

The `virtusertable` maps mail addresses for virtual domains and mailboxes to real mailboxes. These mailboxes can be local, remote, aliases defined in `/etc/mail/aliases` or files.

Example 20-3. Example Virtual Domain Mail Map

```

root@example.com          root
postmaster@example.com   postmaster@noc.example.net
@example.com              joe

```

In the above example, we have a mapping for a domain `example.com`. This file is processed in a first match order down the file. The first item maps `root@example.com` to the local mailbox `root`. The next entry maps `postmaster@example.com` to the mailbox `postmaster` on the host `noc.example.net`. Finally, if nothing from `example.com` has matched so far, it will match the last mapping, which matches every other mail message addressed to someone at `example.com`. This will be mapped to the local mailbox `joe`.

20.4 Changing Your Mail Transfer Agent

Written by Andrew Boothman. Information taken from e-mails written by Gregory Neil Shapiro.

As already mentioned, DragonFly comes with **sendmail** already installed as your MTA (Mail Transfer Agent). Therefore by default it is in charge of your outgoing and incoming mail.

However, for a variety of reasons, some system administrators want to change their system's MTA. These reasons range from simply wanting to try out another MTA to needing a specific feature or package which relies on another mailer. Fortunately, whatever the reason, DragonFly makes it easy to make the change.

20.4.1 Install a New MTA

You have a wide choice of MTAs available. A good starting point is the `pkgsrc` collection or where you will be able to find many. Of course you are free to use any MTA you want from any location, as long as you can make it run under DragonFly.

Start by installing your new MTA. Once it is installed it gives you a chance to decide if it really fulfills your needs, and also gives you the opportunity to configure your new software before getting it to take over from **sendmail**.

When doing this, you should be sure that installing the new software will not attempt to overwrite system binaries such as `/usr/bin/sendmail`. Otherwise, your new mail software has essentially been put into service before you have configured it.

Please refer to your chosen MTA's documentation for information on how to configure the software you have chosen.

20.4.2 Disable sendmail

In order to completely disable **sendmail** you must use

```
sendmail_enable="NONE"
```

```
in /etc/rc.conf.
```

Warning: If you disable **sendmail**'s outgoing mail service in this way, it is important that you replace it with a fully working alternative mail delivery system. If you choose not to, system functions such as `periodic(8)` will be unable to deliver their results by e-mail as they would normally expect to. Many parts of your system may expect to have

a functional **sendmail**-compatible system. If applications continue to use **sendmail**'s binaries to try to send e-mail after you have disabled them, mail could go into an inactive **sendmail** queue, and never be delivered.

If you only want to disable **sendmail**'s incoming mail service, you should set

```
sendmail_enable="NO"
```

in `/etc/rc.conf`. More information on **sendmail**'s startup options is available from the `rc.sendmail(8)` manual page.

20.4.3 Running Your New MTA on Boot

You may have a choice of two methods for running your new MTA on boot, again depending on what version of DragonFly you are running.

With later versions of DragonFly, you can use the above method or you can set

```
mta_start_script="filename"
```

in `/etc/rc.conf`, where *filename* is the name of some script that you want executed at boot to start your MTA.

20.4.4 Replacing sendmail as the System's Default Mailer

The program **sendmail** is so ubiquitous as standard software on UNIX systems that some software just assumes it is already installed and configured. For this reason, many alternative MTA's provide their own compatible implementations of the **sendmail** command-line interface; this facilitates using them as "drop-in" replacements for **sendmail**.

Therefore, if you are using an alternative mailer, you will need to make sure that software trying to execute standard **sendmail** binaries such as `/usr/bin/sendmail` actually executes your chosen mailer instead. Fortunately, DragonFly provides a system called `mailwrapper(8)` that does this job for you.

When **sendmail** is operating as installed, you will find something like the following in `/etc/mail/mailler.conf`:

```
sendmail /usr/libexec/sendmail/sendmail
send-mail /usr/libexec/sendmail/sendmail
mailq /usr/libexec/sendmail/sendmail
newaliases /usr/libexec/sendmail/sendmail
hoststat /usr/libexec/sendmail/sendmail
purgestat /usr/libexec/sendmail/sendmail
```

This means that when any of these common commands (such as `sendmail` itself) are run, the system actually invokes a copy of `mailwrapper` named `sendmail`, which checks `mailler.conf` and executes `/usr/libexec/sendmail/sendmail` instead. This system makes it easy to change what binaries are actually executed when these default `sendmail` functions are invoked.

Therefore if you wanted `/usr/local/supermailer/bin/sendmail-compat` to be run instead of **sendmail**, you could change `/etc/mail/mailler.conf` to read:

```
sendmail /usr/local/supermailer/bin/sendmail-compat
send-mail /usr/local/supermailer/bin/sendmail-compat
```

```
mailq /usr/local/supermailer/bin/mailq-compat
newaliases /usr/local/supermailer/bin/newaliases-compat
hoststat /usr/local/supermailer/bin/hoststat-compat
purgestat /usr/local/supermailer/bin/purgestat-compat
```

20.4.5 Finishing

Once you have everything configured the way you want it, you should either kill the **sendmail** processes that you no longer need and start the processes belonging to your new software, or simply reboot. Rebooting will also give you the opportunity to ensure that you have correctly configured your system to start your new MTA automatically on boot.

20.5 Troubleshooting

1. Why do I have to use the FQDN for hosts on my site?

You will probably find that the host is actually in a different domain; for example, if you are in `foo.bar.edu` and you wish to reach a host called `mumble` in the `bar.edu` domain, you will have to refer to it by the fully-qualified domain name, `mumble.bar.edu`, instead of just `mumble`.

Traditionally, this was allowed by BSD BIND resolvers. However the current version of **BIND** that ships with DragonFly no longer provides default abbreviations for non-fully qualified domain names other than the domain you are in. So an unqualified host `mumble` must either be found as `mumble.foo.bar.edu`, or it will be searched for in the root domain.

This is different from the previous behavior, where the search continued across `mumble.bar.edu`, and `mumble.edu`. Have a look at RFC 1535 for why this was considered bad practice, or even a security hole.

As a good workaround, you can place the line:

```
search foo.bar.edu bar.edu
```

instead of the previous:

```
domain foo.bar.edu
```

into your `/etc/resolv.conf`. However, make sure that the search order does not go beyond the “boundary between local and public administration”, as RFC 1535 calls it.

2. **sendmail** says mail loops back to myself

This is answered in the **sendmail** FAQ as follows:

I’m getting these error messages:

```
553 MX list for domain.net points back to relay.domain.net
554 <user@domain.net>... Local configuration error
```

How can I solve this problem?

You have asked mail to the domain (e.g., domain.net) to be forwarded to a specific host (in this case, relay.domain.net) by using an MX record, but the relay machine does not recognize itself as domain.net. Add domain.net to /etc/mail/local-host-names [known as /etc/sendmail.cw prior to version 8.10] (if you are using FEATURE(use_cw_file)) or add "Cw domain.net" to /etc/mail/sendmail.cf.

The **sendmail** FAQ can be found at <http://www.sendmail.org/faq/> and is recommended reading if you want to do any "tweaking" of your mail setup.

3. How can I run a mail server on a dial-up PPP host?

You want to connect a DragonFly box on a LAN to the Internet. The DragonFly box will be a mail gateway for the LAN. The PPP connection is non-dedicated.

There are at least two ways to do this. One way is to use UUCP.

Another way is to get a full-time Internet server to provide secondary MX services for your domain. For example, if your company's domain is example.com and your Internet service provider has set example.net up to provide secondary MX services to your domain:

```
example.com.      MX      10      example.com.
                  MX      20      example.net.
```

Only one host should be specified as the final recipient (add Cw example.com in /etc/mail/sendmail.cf on example.com).

When the sending sendmail is trying to deliver the mail it will try to connect to you (example.com) over the modem link. It will most likely time out because you are not online. The program **sendmail** will automatically deliver it to the secondary MX site, i.e. your Internet provider (example.net). The secondary MX site will then periodically try to connect to your host and deliver the mail to the primary MX host (example.com).

You might want to use something like this as a login script:

```
#!/bin/sh
# Put me in /usr/local/bin/pppmyisp
( sleep 60 ; /usr/sbin/sendmail -q ) &
/usr/sbin/ppp -direct pppmyisp
```

If you are going to create a separate login script for a user you could use `sendmail -qRexample.com` instead in the script above. This will force all mail in your queue for example.com to be processed immediately.

A further refinement of the situation is as follows:

Message stolen from the FreeBSD Internet service provider's mailing list (<http://lists.FreeBSD.org/mailman/listinfo/freebsd-isp>).

```
> we provide the secondary MX for a customer. The customer connects to
> our services several times a day automatically to get the mails to
> his primary MX (We do not call his site when a mail for his domains
> arrived). Our sendmail sends the mailqueue every 30 minutes. At the
> moment he has to stay 30 minutes online to be sure that all mail is
```

```
> gone to the primary MX.
>
> Is there a command that would initiate sendmail to send all the mails
> now? The user has not root-privileges on our machine of course.
```

In the "privacy flags" section of `sendmail.cf`, there is a definition `Opgoaway,restrictqrun`

Remove `restrictqrun` to allow non-root users to start the queue processing. You might also like to rearrange the MXs. We are the 1st MX for our customers like this, and we have defined:

```
# If we are the best MX for a host, try directly instead of generating
# local config error.
OwTrue
```

That way a remote site will deliver straight to you, without trying the customer connection. You then send to your customer. Only works for "hosts", so you need to get your customer to name their mail machine "customer.com" as well as "hostname.customer.com" in the DNS. Just put an A record in the DNS for "customer.com".

4. Why do I keep getting Relaying Denied errors when sending mail from other hosts?

In default DragonFly installations, **sendmail** is configured to only send mail from the host it is running on. For example, if a POP server is available, then users will be able to check mail from school, work, or other remote locations but they still will not be able to send outgoing emails from outside locations. Typically, a few moments after the attempt, an email will be sent from **MAILER-DAEMON** with a 5.7 Relaying Denied error message.

There are several ways to get around this. The most straightforward solution is to put your ISP's address in a relay-domains file at `/etc/mail/relay-domains`. A quick way to do this would be:

```
# echo "your.isp.example.com" > /etc/mail/relay-domains
```

After creating or editing this file you must restart **sendmail**. This works great if you are a server administrator and do not wish to send mail locally, or would like to use a point and click client/system on another machine or even another ISP. It is also very useful if you only have one or two email accounts set up. If there is a large number of addresses to add, you can simply open this file in your favorite text editor and then add the domains, one per line:

```
your.isp.example.com
other.isp.example.net
users-isp.example.org
www.example.org
```

Now any mail sent through your system, by any host in this list (provided the user has an account on your system), will succeed. This is a very nice way to allow users to send mail from your system remotely without allowing people to send SPAM through your system.

20.6 Advanced Topics

The following section covers more involved topics such as mail configuration and setting up mail for your entire domain.

20.6.1 Basic Configuration

Out of the box, you should be able to send email to external hosts as long as you have set up `/etc/resolv.conf` or are running your own name server. If you would like to have mail for your host delivered to the MTA (e.g., **sendmail**) on your own DragonFly host, there are two methods:

- Run your own name server and have your own domain. For example, `dragonflybsd.org`
- Get mail delivered directly to your host. This is done by delivering mail directly to the current DNS name for your machine. For example, `example.dragonflybsd.org`.

Regardless of which of the above you choose, in order to have mail delivered directly to your host, it must have a permanent static IP address (not a dynamic address, as with most PPP dial-up configurations). If you are behind a firewall, it must pass SMTP traffic on to you. If you want to receive mail directly at your host, you need to be sure of either of two things:

- Make sure that the (lowest-numbered) MX record in your DNS points to your host's IP address.
- Make sure there is no MX entry in your DNS for your host.

Either of the above will allow you to receive mail directly at your host.

Try this:

```
# hostname
example.dragonflybsd.org
# host example.dragonflybsd.org
example.dragonflybsd.org has address 204.216.27.XX
```

If that is what you see, mail directly to `<yourlogin@example.dragonflybsd.org>` should work without problems (assuming **sendmail** is running correctly on `example.dragonflybsd.org`).

If instead you see something like this:

```
# host example.dragonflybsd.org
example.dragonflybsd.org has address 204.216.27.XX
example.dragonflybsd.org mail is handled (pri=10) by hub.dragonflybsd.org
```

All mail sent to your host (`example.dragonflybsd.org`) will end up being collected on `hub` under the same username instead of being sent directly to your host.

The above information is handled by your DNS server. The DNS record that carries mail routing information is the *Mail eXchange* entry. If no MX record exists, mail will be delivered directly to the host by way of its IP address.

The MX entry for `freefall.FreeBSD.org` at one time looked like this:

```
freefall MX 30 mail.crl.net
freefall MX 40 agora.rdrop.com
freefall MX 10 freefall.FreeBSD.org
freefall MX 20 who.cdrom.com
```

As you can see, `freefall` had many MX entries. The lowest MX number is the host that receives mail directly if available; if it is not accessible for some reason, the others (sometimes called “backup MXes”) accept messages temporarily, and pass it along when a lower-numbered host becomes available, eventually to the lowest-numbered host.

Alternate MX sites should have separate Internet connections from your own in order to be most useful. Your ISP or another friendly site should have no problem providing this service for you.

20.6.2 Mail for Your Domain

In order to set up a “mailhost” (a.k.a. mail server) you need to have any mail sent to various workstations directed to it. Basically, you want to “claim” any mail for any hostname in your domain (in this case `*.dragonflybsd.org`) and divert it to your mail server so your users can receive their mail on the master mail server.

To make life easiest, a user account with the same *username* should exist on both machines. Use `adduser(8)` to do this.

The mailhost you will be using must be the designated mail exchanger for each workstation on the network. This is done in your DNS configuration like so:

```
example.dragonflybsd.org A 204.216.27.XX ; Workstation
MX 10 hub.dragonflybsd.org ; Mailhost
```

This will redirect mail for the workstation to the mailhost no matter where the A record points. The mail is sent to the MX host.

You cannot do this yourself unless you are running a DNS server. If you are not, or cannot run your own DNS server, talk to your ISP or whoever provides your DNS.

If you are doing virtual email hosting, the following information will come in handy. For this example, we will assume you have a customer with his own domain, in this case `customer1.org`, and you want all the mail for `customer1.org` sent to your mailhost, `mail.myhost.com`. The entry in your DNS should look like this:

```
customer1.org MX 10 mail.myhost.com
```

You do *not* need an A record for `customer1.org` if you only want to handle email for that domain.

Note: Be aware that pinging `customer1.org` will not work unless an A record exists for it.

The last thing that you must do is tell **sendmail** on your mailhost what domains and/or hostnames it should be accepting mail for. There are a few different ways this can be done. Either of the following will work:

- Add the hosts to your `/etc/mail/local-host-names` file if you are using the `FEATURE(use_cw_file)`.
- Add a `Cyour.host.com` line to your `/etc/mail/sendmail.cf`.

20.7 SMTP with UUCP

The **sendmail** configuration that ships with DragonFly is designed for sites that connect directly to the Internet. Sites that wish to exchange their mail via UUCP must install another **sendmail** configuration file.

Tweaking `/etc/mail/sendmail.cf` manually is an advanced topic. **sendmail** version 8 generates config files via m4(1) preprocessing, where the actual configuration occurs on a higher abstraction level. The m4(1) configuration files can be found under `/usr/src/usr.sbin/sendmail/cf`.

If you did not install your system with full sources, the **sendmail** configuration set has been broken out into a separate source distribution tarball. Assuming you have your DragonFly source code CDROM mounted, do:

```
# cd /cdrom/src
# cat scontrib.?? | tar xzf - -C /usr/src/contrib/sendmail
```

This extracts to only a few hundred kilobytes. The file `README` in the `cf` directory can serve as a basic introduction to m4(1) configuration.

The best way to support UUCP delivery is to use the `mailertable` feature. This creates a database that **sendmail** can use to make routing decisions.

First, you have to create your `.mc` file. The directory `/usr/src/usr.sbin/sendmail/cf/cf` contains a few examples. Assuming you have named your file `foo.mc`, all you need to do in order to convert it into a valid `sendmail.cf` is:

```
# cd /usr/src/usr.sbin/sendmail/cf/cf
# make foo.cf
# cp foo.cf /etc/mail/sendmail.cf
```

A typical `.mc` file might look like:

```
VERSIONID('Your version number') OSTYPE(bsd4.4)

FEATURE(accept_unresolvable_domains)
FEATURE(nocanonify)
FEATURE(mailertable, 'hash -o /etc/mail/mailertable')

define('UUCP_RELAY', your.uucp.relay)
define('UUCP_MAX_SIZE', 200000)
define('confDONT_PROBE_INTERFACES')

MAILER(local)
MAILER(smtp)
MAILER(uucp)

Cw    your.alias.host.name
Cw    youruucpnodename.UUCP
```

The lines containing `accept_unresolvable_domains`, `nocanonify`, and `confDONT_PROBE_INTERFACES` features will prevent any usage of the DNS during mail delivery. The `UUCP_RELAY` clause is needed to support UUCP delivery. Simply put an Internet hostname there that is able to handle `.UUCP` pseudo-domain addresses; most likely, you will enter the mail relay of your ISP there.

Once you have this, you need an `/etc/mail/mailertable` file. If you have only one link to the outside that is used for all your mails, the following file will suffice:

```
#
# makemap hash /etc/mail/mailertable.db < /etc/mail/mailertable
.
      uucp-dom:your.uucp.relay
```

A more complex example might look like this:

```
#
# makemap hash /etc/mail/mailertable.db < /etc/mail/mailertable
#
horus.interface-business.de    uucp-dom:horus
.interface-business.de        uucp-dom:if-bus
interface-business.de         uucp-dom:if-bus
.heep.sax.de                   smtp8:%1
horus.UUCP                     uucp-dom:horus
if-bus.UUCP                    uucp-dom:if-bus
.                               uucp-dom:
```

The first three lines handle special cases where domain-addressed mail should not be sent out to the default route, but instead to some UUCP neighbor in order to “shortcut” the delivery path. The next line handles mail to the local Ethernet domain that can be delivered using SMTP. Finally, the UUCP neighbors are mentioned in the .UUCP pseudo-domain notation, to allow for a *uucp-neighbor !recipient* override of the default rules. The last line is always a single dot, matching everything else, with UUCP delivery to a UUCP neighbor that serves as your universal mail gateway to the world. All of the node names behind the `uucp-dom:` keyword must be valid UUCP neighbors, as you can verify using the command `uname`.

As a reminder that this file needs to be converted into a DBM database file before use. The command line to accomplish this is best placed as a comment at the top of the `mailertable` file. You always have to execute this command each time you change your `mailertable` file.

Final hint: if you are uncertain whether some particular mail routing would work, remember the `-bt` option to **sendmail**. It starts **sendmail** in *address test mode*; simply enter `3,0`, followed by the address you wish to test for the mail routing. The last line tells you the used internal mail agent, the destination host this agent will be called with, and the (possibly translated) address. Leave this mode by typing **Ctrl+D**.

```
% sendmail -bt
ADDRESS TEST MODE (ruleset 3 NOT automatically invoked)
Enter <ruleset> <address>
> 3,0 foo@example.com
canonify          input: foo @ example . com
...
parse            returns: $# uucp-dom $@ your.uucp.relay $: foo < @ example . com . >
> ^D
```

20.8 Setting up to send only

Contributed by Bill Moran.

There are many instances where you may only want to send mail through a relay. Some examples are:

- Your computer is a desktop machine, but you want to use programs from the command line that send mail. To do so, you should use your ISP’s mail relay.

- The computer is a server that does not handle mail locally, but needs to pass off all mail to a relay for processing.

Just about any MTA is capable of filling this particular niche. Unfortunately, it can be very difficult to properly configure a full-featured MTA just to handle offloading mail. Programs such as **sendmail** and **postfix** are largely overkill for this use.

Additionally, if you are using a typical Internet access service, your agreement may forbid you from running a “mail server”.

The easiest way to fulfill those needs is to install the `mail/ssmtp` port. Execute the following commands as `root`:

```
# cd /usr/ports/mail/ssmtp
# make install replace clean
```

Once installed, `mail/ssmtp` can be configured with a four-line file located at `/usr/local/etc/ssmtp/ssmtp.conf`:

```
root=yourrealemail@example.com
mailhub=mail.example.com
rewriteDomain=example.com
hostname=_HOSTNAME_
```

Make sure you use your real email address for `root`. Enter your ISP’s outgoing mail relay in place of `mail.example.com` (some ISPs call this the “outgoing mail server” or “SMTP server”).

Make sure you disable **sendmail** by setting `sendmail_enable="NONE"` in `/etc/rc.conf`.

`mail/ssmtp` has some other options available. See the example configuration file in `/usr/local/etc/ssmtp` or the manual page of **ssmtp** for some examples and more information.

Setting up **ssmtp** in this manner will allow any software on your computer that needs to send mail to function properly, while not violating your ISP’s usage policy or allowing your computer to be hijacked for spamming.

20.9 Using Mail with a Dialup Connection

If you have a static IP address, you should not need to adjust anything from the defaults. Set your host name to your assigned Internet name and **sendmail** will do the rest.

If you have a dynamically assigned IP number and use a dialup PPP connection to the Internet, you will probably have a mailbox on your ISP’s mail server. Let’s assume your ISP’s domain is `example.net`, and that your user name is `user`, you have called your machine `bsd.home`, and your ISP has told you that you may use `relay.example.net` as a mail relay.

In order to retrieve mail from your mailbox, you must install a retrieval agent. The **fetchmail** utility is a good choice as it supports many different protocols. This program is available as a package or from the ports collection (`mail/fetchmail`). Usually, your ISP will provide POP. If you are using user PPP, you can automatically fetch your mail when an Internet connection is established with the following entry in `/etc/ppp/ppp.linkup`:

```
MYADDR:
!bg su user -c fetchmail
```

If you are using **sendmail** (as shown below) to deliver mail to non-local accounts, you probably want to have **sendmail** process your mailqueue as soon as your Internet connection is established. To do this, put this command after the `fetchmail` command in `/etc/ppp/ppp.linkup`:

```
!bg su user -c "sendmail -q"
```

Assume that you have an account for `user` on `bsd.home`. In the home directory of `user` on `bsd.home`, create a `.fetchmailrc` file:

```
poll example.net protocol pop3 fetchall pass MySecret
```

This file should not be readable by anyone except `user` as it contains the password `MySecret`.

In order to send mail with the correct `from:` header, you must tell **sendmail** to use `user@example.net` rather than `user@bsd.home`. You may also wish to tell **sendmail** to send all mail via `relay.example.net`, allowing quicker mail transmission.

The following `.mc` file should suffice:

```
VERSIONID('bsd.home.mc version 1.0')
OSTYPE(bsd4.4)dnl
FEATURE(nouucp)dnl
MAILER(local)dnl
MAILER(smtp)dnl
Cwlocalhost
Cwbsd.home
MASQUERADE_AS('example.net')dnl
FEATURE(allmasquerade)dnl
FEATURE(masquerade_envelope)dnl
FEATURE(nocanonify)dnl
FEATURE(nodns)dnl
define('SMART_HOST', 'relay.example.net')
Dmbsd.home
define('confDOMAIN_NAME', 'bsd.home')dnl
define('confDELIVERY_MODE', 'deferred')dnl
```

Refer to the previous section for details of how to turn this `.mc` file into a `sendmail.cf` file. Also, do not forget to restart **sendmail** after updating `sendmail.cf`.

20.10 SMTP Authentication

Written by James Gorham.

Having SMTP Authentication in place on your mail server has a number of benefits. SMTP Authentication can add another layer of security to **sendmail**, and has the benefit of giving mobile users who switch hosts the ability to use the same mail server without the need to reconfigure their mail client settings each time.

1. Install `security/cyrus-sasl` from the ports. You can find this port in `security/cyrus-sasl`. `security/cyrus-sasl` has a number of compile time options to choose from and, for the method we will be using here, make sure to select the `pwcheck` option.
2. After installing `security/cyrus-sasl`, edit `/usr/local/lib/sasl/Sendmail.conf` (or create it if it does not exist) and add the following line:

```
pwcheck_method: passwd
```

This method will enable **sendmail** to authenticate against your DragonFly `passwd` database. This saves the trouble of creating a new set of usernames and passwords for each user that needs to use SMTP authentication, and keeps the login and mail password the same.

3. Now edit `/etc/make.conf` and add the following lines:

```
SENDMAIL_CFLAGS=-I/usr/local/include/sasl -DSASL
SENDMAIL_LDFLAGS=-L/usr/local/lib
SENDMAIL_LDADD=-lsasl
```

These lines will give **sendmail** the proper configuration options for linking to `cyrus-sasl` at compile time. Make sure that `cyrus-sasl` has been installed before recompiling **sendmail**.

4. Recompile **sendmail** by executing the following commands:

```
# cd /usr/src/usr.sbin/sendmail
# make cleandir
# make obj
# make
# make install
```

The compile of **sendmail** should not have any problems if `/usr/src` has not been changed extensively and the shared libraries it needs are available.

5. After **sendmail** has been compiled and reinstalled, edit your `/etc/mail/freebsd.mc` file (or whichever file you use as your `.mc` file. Many administrators choose to use the output from `hostname(1)` as the `.mc` file for uniqueness). Add these lines to it:

```
dn1 set SASL options
TRUST_AUTH_MECH('GSSAPI DIGEST-MD5 CRAM-MD5 LOGIN')dn1
define('confAUTH_MECHANISMS', 'GSSAPI DIGEST-MD5 CRAM-MD5 LOGIN')dn1
define('confDEF_AUTH_INFO', '/etc/mail/auth-info')dn1
```

These options configure the different methods available to **sendmail** for authenticating users. If you would like to use a method other than **pwcheck**, please see the included documentation.

6. Finally, run `make(1)` while in `/etc/mail`. That will run your new `.mc` file and create a `.cf` file named `freebsd.cf` (or whatever name you have used for your `.mc` file). Then use the command `make install restart`, which will copy the file to `sendmail.cf`, and will properly restart **sendmail**. For more information about this process, you should refer to `/etc/mail/Makefile`.

If all has gone correctly, you should be able to enter your login information into the mail client and send a test message. For further investigation, set the `LogLevel` of **sendmail** to 13 and watch `/var/log/maillog` for any errors.

You may wish to add the following lines to `/etc/rc.conf` so this service will be available after every system boot:

```
sasl_pwcheck_enable="YES"
sasl_pwcheck_program="/usr/local/sbin/pwcheck"
```

This will ensure the initialization of `SMTP_AUTH` upon system boot.

For more information, please see the **sendmail** page regarding SMTP authentication (<http://www.sendmail.org/~ca/email/auth.html>).

20.11 Mail User Agents

Contributed by Marc Silver.

A Mail User Agent (MUA) is an application that is used to send and receive email. Furthermore, as email “evolves” and becomes more complex, MUA’s are becoming increasingly powerful in the way they interact with email; this gives users increased functionality and flexibility. DragonFly contains support for numerous mail user agents, all of which can be easily installed using the `pkgsrc` collection. Users may choose between graphical email clients such as **evolution** or **balsa**, console based clients such as **mutt**, **pine** or `mail`, or the web interfaces used by some large organizations.

20.11.1 mail

`mail(1)` is the default Mail User Agent (MUA) in DragonFly. It is a console based MUA that offers all the basic functionality required to send and receive text-based email, though it is limited in interaction abilities with attachments and can only support local mailboxes.

Although `mail` does not natively support interaction with POP or IMAP servers, these mailboxes may be downloaded to a local `mbox` file using an application such as **fetchmail**, which will be discussed later in this chapter (Section 20.12).

In order to send and receive email, simply invoke the `mail` command as per the following example:

```
% mail
```

The contents of the user mailbox in `/var/mail` are automatically read by the `mail` utility. Should the mailbox be empty, the utility exits with a message indicating that no mails could be found. Once the mailbox has been read, the application interface is started, and a list of messages will be displayed. Messages are automatically numbered, as can be seen in the following example:

```
Mail version 8.1 6/6/93.  Type ? for help.
"/var/mail/marcs": 3 messages 3 new
>N  1 root@localhost      Mon Mar  8 14:05  14/510  "test"
   N  2 root@localhost      Mon Mar  8 14:05  14/509  "user account"
   N  3 root@localhost      Mon Mar  8 14:05  14/509  "sample"
```

Messages can now be read by using the `t mail` command, suffixed by the message number that should be displayed. In this example, we will read the first email:

```
& t 1
Message 1:
From root@localhost  Mon Mar  8 14:05:52 2004
X-Original-To: marcs@localhost
Delivered-To: marcs@localhost
To: marcs@localhost
Subject: test
Date: Mon,  8 Mar 2004 14:05:52 +0200 (SAST)
From: root@localhost (Charlie Root)
```

This is a test message, please reply if you receive it.

As can be seen in the example above, the `t` key will cause the message to be displayed with full headers. To display the list of messages again, the `h` key should be used.

If the email requires a response, you may use `mail` to reply, by using either the **R** or **r** `mail` keys. The **R** key instructs `mail` to reply only to the sender of the email, while **r** replies not only to the sender, but also to other recipients of the message. You may also suffix these commands with the mail number which you would like make a reply to. Once this has been done, the response should be entered, and the end of the message should be marked by a single `.` on a new line. An example can be seen below:

```
& R 1
To: root@localhost
Subject: Re: test
```

Thank you, I did get your email.

```
.
EOT
```

In order to send new email, the **m** key should be used, followed by the recipient email address. Multiple recipients may also be specified by separating each address with the `,` delimiter. The subject of the message may then be entered, followed by the message contents. The end of the message should be specified by putting a single `.` on a new line.

```
& mail root@localhost
Subject: I mastered mail
```

Now I can send and receive email using mail ... :)

```
.
EOT
```

While inside the `mail` utility, the `?` command may be used to display help at any time, the `mail(1)` manual page should also be consulted for more help with `mail`.

Note: As previously mentioned, the `mail(1)` command was not originally designed to handle attachments, and thus deals with them very poorly. Newer MUAs such as **mutt** handle attachments in a much more intelligent way. But should you still wish to use the `mail` command, the `converters/mpack` port may be of considerable use.

20.11.2 mutt

mutt is a small yet very powerful Mail User Agent, with excellent features, just some of which include:

- The ability to thread messages;
- PGP support for digital signing and encryption of email;
- MIME Support;
- Maildir Support;
- Highly customizable.

All of these features help to make **mutt** one of the most advanced mail user agents available. See <http://www.mutt.org> for more information on **mutt**.

The stable version of **mutt** may be installed using the `mail/mutt` port, while the current development version may be installed via the `mail/mutt-devel` port. After the port has been installed, **mutt** can be started by issuing the following command:

```
% mutt
```

mutt will automatically read the contents of the user mailbox in `/var/mail` and display the contents if applicable. If no mails are found in the user mailbox, then **mutt** will wait for commands from the user. The example below shows **mutt** displaying a list of messages:

```
q:Quit d:Del u:Undel s:Save n:Mail r:Reply g:Group ?:Help
1 N Mar 09 Super-User ( 1) test
2 N Mar 09 Super-User ( 1) user account
3 N Mar 09 Super-User ( 1) sample

--*Mutt: /var/mail/marcs [Msgs:3 New:3 1.6K]--(date/date)----- (all)-----
```

In order to read an email, simply select it using the cursor keys, and press the **Enter** key. An example of **mutt** displaying email can be seen below:

```
i:Exit -:PrevPg <Space>:NextPg v:View Attachm. d:Del r:Reply j:Next ?:Help
X-Original-To: marcs@localhost
Delivered-To: marcs@localhost
To: marcs@localhost
Subject: test
Date: Tue, 9 Mar 2004 10:28:36 +0200 (SAST)
From: Super-User <root@localhost>

This is a test message, please reply if you receive it.

--N - 1/1: Super-User test -- (all)
```

As with the `mail(1)` command, **mutt** allows users to reply only to the sender of the message as well as to all recipients. To reply only to the sender of the email, use the `r` keyboard shortcut. To send a group reply, which will be sent to the original sender as well as all the message recipients, use the `g` shortcut.

Note: **mutt** makes use of the `vi(1)` command as an editor for creating and replying to emails. This may be customized by the user by creating or editing their own `.muttrc` file in their home directory and setting the

editor variable.

In order to compose a new mail message, press **m**. After a valid subject has been given, **mutt** will start `vi(1)` and the mail can be written. Once the contents of the mail are complete, save and quit from `vi` and **mutt** will resume, displaying a summary screen of the mail that is to be delivered. In order to send the mail, press **y**. An example of the summary screen can be seen below:

```

y:Send q:Abort t:To c:CC s:Subj a:Attach file d:Descrip ?:Help
  From: Marc Silver <marcs@localhost>
  To: Super-User <root@localhost>
  Cc:
  Bcc:
  Subject: Re: test
  Reply-To:
  Fcc:
  Security: Clear

-- Attachments
- I 1 /tmp/mutt-bsd-c0hobscQ [text/plain, 7bit, us-ascii, 1.1K]

-----
Mutt: Compose [Approx. msg size: 1.1K Atts: 1]

```

mutt also contains extensive help, which can be accessed from most of the menus by pressing the **?** key. The top line also displays the keyboard shortcuts where appropriate.

20.11.3 pine

pine is aimed at a beginner user, but also includes some advanced features.

Warning: The **pine** software has had several remote vulnerabilities discovered in the past, which allowed remote attackers to execute arbitrary code as users on the local system, by the action of sending a specially-prepared email. All such *known* problems have been fixed, but the **pine** code is written in a very insecure style and the DragonFly Security Officer believes there are likely to be other undiscovered vulnerabilities. You install **pine** at your own risk.

The current version of **pine** may be installed using the `mail/pine4` port. Once the port has installed, **pine** can be started by issuing the following command:

```
% pine
```

The first time that **pine** is run it displays a greeting page with a brief introduction, as well as a request from the **pine** development team to send an anonymous email message allowing them to judge how many users are using their client. To send this anonymous message, press **Enter**, or alternatively press **E** to exit the greeting without sending an anonymous message. An example of the greeting page can be seen below:

```

PINE 4.58  GREETING TEXT                                     No Messages
<<<This message will appear only once>>>

Welcome to Pine ... a Program for Internet News and Email

We hope you will explore Pine's many capabilities. From the Main Menu,
select Setup/Config to see many of the options available to you. Also
note that all screens have context-sensitive help text available.

SPECIAL REQUEST: This software is made available world-wide as a public
service of the University of Washington in Seattle. In order to justify
continuing development, it is helpful to have an idea of how many people
are using Pine. Are you willing to be counted as a Pine user? Pressing
Return will send an anonymous (meaning, your real email address will not
be revealed) message to the Pine development team at the University of
Washington for purposes of tallying.

Pine is a trademark of the University of Washington.

[ALL of greeting text]
? Help      [E] Exit this greeting      [P] PreVPage  [Z] Print
[Ret] [Be Counted!]                [SpC] NextPage

```

Users are then presented with the main menu, which can be easily navigated using the cursor keys. This main menu provides shortcuts for the composing new mails, browsing of mail directories, and even the administration of address book entries. Below the main menu, relevant keyboard shortcuts to perform functions specific to the task at hand are shown.

The default directory opened by **pine** is the `inbox`. To view the message index, press **I**, or select the MESSAGE INDEX option as seen below:

```

PINE 4.58  MAIN MENU                                       Folder: INBOX  3 Messages
?  HELP          - Get help using Pine
C  COMPOSE MESSAGE - Compose and send a message
I  MESSAGE INDEX - View messages in current folder
L  FOLDER LIST   - Select a folder to view
A  ADDRESS BOOK  - Update address book
S  SETUP         - Configure Pine Options
Q  QUIT          - Leave the Pine program

Copyright 1989-2003. PINE is a trademark of the University of Washington.

? Help      [P] PreVPcmd      [R] ReVPNotes
[O] OTHER CMDS [Z] [Index]  [N] NextCmd      [K] KBlock

```

The message index shows messages in the current directory, and can be navigated by using the cursor keys. Highlighted messages can be read by pressing the **Enter** key.

```

PINE 4.58 MESSAGE INDEX Folder: INBOX Message 1 of 3 ANS
A 1 Mar 9 Super-User (471) test
A 2 Mar 9 Super-User (479) user account
A 3 Mar 9 Super-User (473) sample

? Help < FldrList P PrevMsg - PrePage D Delete R Reply
0 OTHER CMDS > |ViewMsg| N NextMsg Spc NextPage U Undelete F Forward

```

In the screenshot below, a sample message is displayed by **pine**. Keyboard shortcuts are displayed as a reference at the bottom of the screen. An example of one of these shortcuts is the **r** key, which tells the MUA to reply to the current message being displayed.

```

PINE 4.58 MESSAGE TEXT Folder: INBOX Message 1 of 3 ALL ANS
Date: Tue, 9 Mar 2004 10:28:36 +0200 (SAST)
From: Super-User <root@localhost>
To: marcs@localhost
Subject: test

This is a test message, please reply if you receive it.

[ALL of message]
? Help < MsgIndex P PrevMsg - PrePage D Delete R Reply
0 OTHER CMDS > |ViewAtch| N NextMsg Spc NextPage U Undelete F Forward

```

Replying to an email in **pine** is done using the **pico** editor, which is installed by default with **pine**. The **pico** utility makes it easy to navigate around the message and is slightly more forgiving on novice users than **vi(1)** or **mail(1)**. Once the reply is complete, the message can be sent by pressing **Ctrl+X**. The **pine** application will ask for confirmation.

```

PINE 4.58  COMPOSE MESSAGE REPLY  Folder: INBOX  3 Messages
To      : Super-User <root@localhost>
Cc      :
Atchmnt:
Subject : Re: test
----- Message Text -----
I did recieve your message...

^G Get Help  ^X Send      ^R Read File ^V Prev Pg  ^K Cut Text  ^O Postpone
^C Cancel    ^J Justify   ^U Where is  ^U Next Pg  ^U UnCut Text ^T To Spell

```

The **pine** application can be customized using the **SETUP** option from the main menu. Consult <http://www.washington.edu/pine/> for more information.

20.12 Using fetchmail

Contributed by Marc Silver.

fetchmail is a full-featured IMAP and POP client which allows users to automatically download mail from remote IMAP and POP servers and save it into local mailboxes; there it can be accessed more easily. **fetchmail** can be installed using the `mail/fetchmail` port, and offers various features, some of which include:

- Support of POP3, APOP, KPOP, IMAP, ETRN and ODMR protocols.
- Ability to forward mail using SMTP, which allows filtering, forwarding, and aliasing to function normally.
- May be run in daemon mode to check periodically for new messages.
- Can retrieve multiple mailboxes and forward them based on configuration, to different local users.

While it is outside the scope of this document to explain all of **fetchmail**'s features, some basic features will be explained. The **fetchmail** utility requires a configuration file known as `.fetchmailrc`, in order to run correctly. This file includes server information as well as login credentials. Due to the sensitive nature of the contents of this file, it is advisable to make it readable only by the owner, with the following command:

```
% chmod 600 .fetchmailrc
```

The following `.fetchmailrc` serves as an example for downloading a single user mailbox using POP. It tells **fetchmail** to connect to `example.com` using a username of `joesoap` and a password of `XXX`. This example assumes that the user `joesoap` is also a user on the local system.

```
poll example.com protocol pop3 username "joesoap" password "XXX"
```

The next example connects to multiple POP and IMAP servers and redirects to different local usernames where applicable:

```
poll example.com proto pop3:
user "joesoap", with password "XXX", is "jsoap" here;
user "andrea", with password "XXXX";
poll example2.net proto imap:
user "john", with password "XXXXX", is "myth" here;
```

The **fetchmail** utility can be run in daemon mode by running it with the `-d` flag, followed by the interval (in seconds) that **fetchmail** should poll servers listed in the `.fetchmailrc` file. The following example would cause **fetchmail** to poll every 60 seconds:

```
% fetchmail -d 60
```

More information on **fetchmail** can be found at <http://www.catb.org/~esr/fetchmail/>.

20.13 Using procmail

Contributed by Marc Silver.

The **procmail** utility is an incredibly powerful application used to filter incoming mail. It allows users to define “rules” which can be matched to incoming mails to perform specific functions or to reroute mail to alternative mailboxes and/or email addresses. **procmail** can be installed using the `mail/procmail` port. Once installed, it can be directly integrated into most MTAs; consult your MTA documentation for more information. Alternatively, **procmail** can be integrated by adding the following line to a `.forward` in the home directory of the user utilizing **procmail** features:

```
"|exec /usr/local/bin/procmail || exit 75"
```

The following section will display some basic **procmail** rules, as well as brief descriptions on what they do. These rules, and others must be inserted into a `.procmailrc` file, which must reside in the user’s home directory.

The majority of these rules can also be found in the `procmailex(5)` manual page.

Forward all mail from `user@example.com` to an external address of `goodmail@example2.com`:

```
:0
* ^From.*user@example.com
! goodmail@example2.com
```

Forward all mails shorter than 1000 bytes to an external address of `goodmail@example2.com`:

```
:0
* < 1000
! goodmail@example2.com
```

Send all mail sent to `alternate@example.com` into a mailbox called `alternate`:

```
:0
* ^TOalternate@example.com
alternate
```

Send all mail with a subject of “Spam” to `/dev/null`:

```
:0
^Subject: .*Spam
/dev/null
```

A useful recipe that parses incoming dragonflybsd.org mailing lists and places each list in its own mailbox:

```
:0
* ^List-Post: <mailto:\[^@]+
{
LISTNAME=${MATCH}
:0
* LISTNAME??^\[^-]+
DragonFly-${MATCH}
}
```


Chapter 21 Updating DragonFly

Written by Justin Sherrill.

21.1 Initial Setup

Updates to the DragonFly source code is performed using **cvsup**. **cvsup** compares your local system source or ports files to a remote repository, and downloads any changes. Only the differences in the files are downloaded, saving on bandwidth and time.

cvsup exists as a port (`devel/cvsup`) and traditionally had to be installed separately on FreeBSD. With DragonFly, the binary is installed as part of the base system.

21.2 Configuration

cvsup is guided by a configuration file that describes what files to update, and the source from which to update them.

Here is a basic DragonFly cvsup configuration file:

```
*default host=cvsup.dragonflybsd.org
*default base=/usr
*default prefix=/usr
*default release=cvs
*default release=cvs tag=.
*default delete use-rel-suffix
*default compress
```

```
cvs-src
```

Alternately, the file `/usr/share/examples/cvsup/DragonFly-src-supfile` can be used as-is to update system source.

Run `cvsup` using `/usr/share/examples/cvsup/DragonFly-src-supfile` as an argument or with a separate file containing the above example text. Your system source files will be updated.

21.3 Preparing to Update

If you want to create a custom kernel, see Chapter 9. This is not needed to update your system unless there is a specific feature that needs to be added to the kernel.

Check recent mail traffic on DragonFly Kernel mailing list (<http://leaf.dragonflybsd.org/mailarchive/>) and the file `/usr/src/UPDATING`. Any recent problems or changes should be described there. `/usr/src/UPDATING` also contains abbreviated build instructions in case these directions are not available.

21.4 Updating the System

Updating the system is a relatively simple process. As root, in `/usr/src`:

```
% make buildworld
% make buildkernel KERNCONF=GENERIC
% make installkernel KERNCONF=GENERIC
% make installworld
% make upgrade
(reboot)
```

An explanation of each step follows.

- **make buildworld** : This command rebuilds all userland programs. This is the most time-consuming step.
- **make buildkernel KERNCONF=GENERIC** : This builds the kernel using the config file specified by `KERNCONF`. If you've created a different kernel configuration file as detailed in Chapter 9, use that instead of `GENERIC`. If `KERNCONF` isn't specified, the `GENERIC` configuration file (installed by default) is used.
- **make installkernel KERNCONF=GENERIC** : This installs the kernel using the config file specified by `KERNCONF`. The value of `KERNCONF` must match what was specified in the **make buildkernel** command, so that files that match this configuration can be installed correctly. As with **make buildkernel**, `KERNCONF` will be set to `GENERIC` if not otherwise specified.
- **make installworld** : This copies all the files built in the `buildworld` step (i.e. everything that is not the kernel) to the proper places in the filesystem.
- **make upgrade** : This cleans out any files made unnecessary by this upgrade.
- **(reboot)** : Reboot the computer to load the new kernel and use the new files installed as part of this process.

If your computer fails to reboot, check the Section 9.6 section of the handbook.

Chapter 22 Linux Binary Compatibility

Restructured and parts updated by Jim Mock. Originally contributed by Brian N. Handy and Rich Murphey.

22.1 Synopsis

DragonFly provides binary compatibility with several other UNIX like operating systems, including Linux. At this point, you may be asking yourself why exactly, does DragonFly need to be able to run Linux binaries? The answer to that question is quite simple. Many companies and developers develop only for Linux, since it is the latest “hot thing” in the computing world. That leaves the rest of us DragonFly users bugging these same companies and developers to put out native DragonFly versions of their applications. The problem is, that most of these companies do not really realize how many people would use their product if there were DragonFly versions too, and most continue to only develop for Linux. So what is a DragonFly user to do? This is where the Linux binary compatibility of DragonFly comes into play.

In a nutshell, the compatibility allows DragonFly users to run about 90% of all Linux applications without modification. This includes applications such as **StarOffice**, the Linux version of **Netscape**, **Adobe® Acrobat®**, **RealPlayer® 5 and 7**, **VMware™**, **Oracle**, **WordPerfect®**, **Doom**, **Quake**, and more. It is also reported that in some situations, Linux binaries perform better on DragonFly than they do under Linux.

There are, however, some Linux-specific operating system features that are not supported under DragonFly. Linux binaries will not work on DragonFly if they overly use the Linux `/proc` file system (which is different from DragonFly’s `/proc` file system), or i386 specific calls, such as enabling virtual 8086 mode.

After reading this chapter, you will know:

- How to enable Linux binary compatibility on your system.
- How to install additional Linux shared libraries.
- How to install Linux applications on your DragonFly system.
- The implementation details of Linux compatibility in DragonFly.

Before reading this chapter, you should:

- Know how to install additional third-party software (Chapter 4).

22.2 Installation

Linux binary compatibility is not turned on by default. The easiest way to enable this functionality is to load the `linux` KLD object (“Kernel Loadable Device”). You can load this module by simply typing `linux` at the command prompt.

If you would like Linux compatibility to always be enabled, then you should add the following line to `/etc/rc.conf`:

```
linux_enable="YES"
```

The `kldstat(8)` command can be used to verify that the KLD is loaded:

```
% kldstat
Id Refs Address      Size      Name
  1     2 0xc0100000 16bdb8   kernel
  7     1 0xc24db000 d000     linux.ko
```

If for some reason you do not want to or cannot load the KLD, then you may statically link Linux binary compatibility into the kernel by adding options `COMPAT_LINUX` to your kernel configuration file. Then install your new kernel as described in Chapter 9.

22.2.1 Installing Linux Runtime Libraries

This can be done one of two ways, either by using the suse package, or by installing them manually.

22.2.1.1 Installing Using suse Package

This is by far the easiest method to use when installing the runtime libraries. It is just like installing any other package from the `pkgsrc` collection (`/usr/pkgsrc/`). Simply do the following:

```
# cd /usr/pkgsrc/meta-pkgs/suse9
# make install distclean
# ln -s /usr/pkg/emul /compat
```

You should now have working Linux binary compatibility. Some programs may complain about incorrect minor versions of the system libraries. In general, however, this does not seem to be a problem.

Note: There may be multiple versions of the `meta-pkgs/suse` package available, corresponding to different versions of various Linux distributions. You should install the package most closely resembling the requirements of the Linux applications you would like to install.

22.2.1.2 Installing Libraries Manually

If you do not have the “`pkgsrc`” collection installed, you can install the libraries by hand instead. You will need the Linux shared libraries that the program depends on and the runtime linker. Also, you will need to create a “shadow root” directory, `/compat/linux`, for Linux libraries on your DragonFly system. Any shared libraries opened by Linux programs run under DragonFly will look in this tree first. So, if a Linux program loads, for example, `/lib/libc.so`, DragonFly will first try to open `/compat/linux/lib/libc.so`, and if that does not exist, it will then try `/lib/libc.so`. Shared libraries should be installed in the shadow tree `/compat/linux/lib` rather than the paths that the Linux `ld.so` reports.

Generally, you will need to look for the shared libraries that Linux binaries depend on only the first few times that you install a Linux program on your DragonFly system. After a while, you will have a sufficient set of Linux shared libraries on your system to be able to run newly imported Linux binaries without any extra work.

22.2.1.3 How to Install Additional Shared Libraries

What if you install the `suse` package and your application still complains about missing shared libraries? How do you know which shared libraries Linux binaries need, and where to get them? Basically, there are 2 possibilities (when following these instructions you will need to be `root` on your DragonFly system).

If you have access to a Linux system, see what shared libraries the application needs, and copy them to your DragonFly system. Look at the following example:

Let us assume you used FTP to get the Linux binary of **Doom**, and put it on a Linux system you have access to. You then can check which shared libraries it needs by running `ldd linuxdoom`, like so:

```
% ldd linuxdoom
libXt.so.3 (DLL Jump 3.1) => /usr/X11/lib/libXt.so.3.1.0
libX11.so.3 (DLL Jump 3.1) => /usr/X11/lib/libX11.so.3.1.0
libc.so.4 (DLL Jump 4.5p126) => /lib/libc.so.4.6.29
```

You would need to get all the files from the last column, and put them under `/compat/linux`, with the names in the first column as symbolic links pointing to them. This means you eventually have these files on your DragonFly system:

```
/compat/linux/usr/X11/lib/libXt.so.3.1.0
/compat/linux/usr/X11/lib/libXt.so.3 -> libXt.so.3.1.0
/compat/linux/usr/X11/lib/libX11.so.3.1.0
/compat/linux/usr/X11/lib/libX11.so.3 -> libX11.so.3.1.0
/compat/linux/lib/libc.so.4.6.29
/compat/linux/lib/libc.so.4 -> libc.so.4.6.29
```

Note: Note that if you already have a Linux shared library with a matching major revision number to the first column of the `ldd` output, you will not need to copy the file named in the last column to your system, the one you already have should work. It is advisable to copy the shared library anyway if it is a newer version, though. You can remove the old one, as long as you make the symbolic link point to the new one. So, if you have these libraries on your system:

```
/compat/linux/lib/libc.so.4.6.27
/compat/linux/lib/libc.so.4 -> libc.so.4.6.27
```

and you find a new binary that claims to require a later version according to the output of `ldd`:

```
libc.so.4 (DLL Jump 4.5p126) -> libc.so.4.6.29
```

If it is only one or two versions out of date in the trailing digit then do not worry about copying `/lib/libc.so.4.6.29` too, because the program should work fine with the slightly older version. However, if you like, you can decide to replace the `libc.so` anyway, and that should leave you with:

```
/compat/linux/lib/libc.so.4.6.29
/compat/linux/lib/libc.so.4 -> libc.so.4.6.29
```

Note: The symbolic link mechanism is *only* needed for Linux binaries. The DragonFly runtime linker takes care of looking for matching major revision numbers itself and you do not need to worry about it.

22.2.2 Installing Linux ELF Binaries

ELF binaries sometimes require an extra step of “branding”. If you attempt to run an unbranded ELF binary, you will get an error message like the following:

```
% ./my-linux-elf-binary
ELF binary type not known
Abort
```

To help the DragonFly kernel distinguish between a DragonFly ELF binary from a Linux binary, use the `brandelf(1)` utility.

```
% brandelf -t Linux my-linux-elf-binary
```

The GNU toolchain now places the appropriate branding information into ELF binaries automatically, so this step should only be needed for older, `libc5` binaries.

22.2.3 Configuring the Hostname Resolver

If DNS does not work or you get this message:

```
resolv+: "bind" is an invalid keyword resolv+:
"hosts" is an invalid keyword
```

You will need to configure a `/compat/linux/etc/host.conf` file containing:

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

The order here specifies that `/etc/hosts` is searched first and DNS is searched second. When `/compat/linux/etc/host.conf` is not installed, Linux applications find DragonFly’s `/etc/host.conf` and complain about the incompatible DragonFly syntax. You should remove `bind` if you have not configured a name server using the `/etc/resolv.conf` file.

22.3 Installing Mathematica®

Updated for Mathematica 4.X by Murray Stokely. Merged with work by Bojan Bistrovic.

This document describes the process of installing the Linux version of **Mathematica 4.X** onto a DragonFly system.

Warning: This description applies to FreeBSD, for which it was originally written. This may or may not apply to DragonFly at this point; while FreeBSD 4.x features usually translate over to DragonFly well, your mileage may vary.

The Linux version of **Mathematica** runs perfectly under DragonFly however the binaries shipped by Wolfram need to be branded so that DragonFly knows to use the Linux ABI to execute them.

The Linux version of **Mathematica** or **Mathematica for Students** can be ordered directly from Wolfram at <http://www.wolfram.com/>.

22.3.1 Branding the Linux Binaries

The Linux binaries are located in the `Unix` directory of the **Mathematica** CDROM distributed by Wolfram. You need to copy this directory tree to your local hard drive so that you can brand the Linux binaries with `brandelf(1)` before running the installer:

```
# mount /cdrom
# cp -rp /cdrom/Unix/ /localdir/
# brandelf -t Linux /localdir/Files/SystemFiles/Kernel/Binaries/Linux/*
# brandelf -t Linux /localdir/Files/SystemFiles/FrontEnd/Binaries/Linux/*
# brandelf -t Linux /localdir/Files/SystemFiles/Installation/Binaries/Linux/*
# brandelf -t Linux /localdir/Files/SystemFiles/Graphics/Binaries/Linux/*
# brandelf -t Linux /localdir/Files/SystemFiles/Converters/Binaries/Linux/*
# brandelf -t Linux /localdir/Files/SystemFiles/LicenseManager/Binaries/Linux/mathlm
# cd /localdir/Installers/Linux/
# ./MathInstaller
```

Alternatively, you can simply set the default ELF brand to Linux for all unbranded binaries with the command:

```
# sysctl kern.fallback_elf_brand=3
```

This will make DragonFly assume that unbranded ELF binaries use the Linux ABI and so you should be able to run the installer straight from the CDROM.

22.3.2 Obtaining Your Mathematica Password

Before you can run **Mathematica** you will have to obtain a password from Wolfram that corresponds to your “machine ID”.

Once you have installed the Linux compatibility runtime libraries and unpacked **Mathematica** you can obtain the “machine ID” by running the program `mathinfo` in the installation directory. This machine ID is based solely on the MAC address of your first Ethernet card.

```
# cd /localdir/Files/SystemFiles/Installation/Binaries/Linux
# mathinfo
disco.example.com 7115-70839-20412
```

When you register with Wolfram, either by email, phone or fax, you will give them the “machine ID” and they will respond with a corresponding password consisting of groups of numbers. You can then enter this information when you attempt to run **Mathematica** for the first time exactly as you would for any other **Mathematica** platform.

22.3.3 Running the Mathematica Frontend over a Network

Mathematica uses some special fonts to display characters not present in any of the standard font sets (integrals, sums, Greek letters, etc.). The X protocol requires these fonts to be installed *locally*. This means you will have to copy these fonts from the CDROM or from a host with **Mathematica** installed to your local machine. These fonts are normally stored in `/cdrom/Unix/Files/SystemFiles/Fonts` on the CDROM, or `/usr/local/mathematica/SystemFiles/Fonts` on your hard drive. The actual fonts are in the subdirectories `Type1` and `X`. There are several ways to use them, as described below.

The first way is to copy them into one of the existing font directories in `/usr/X11R6/lib/X11/fonts`. This will require editing the `fonts.dir` file, adding the font names to it, and changing the number of fonts on the first line. Alternatively, you should also just be able to run `mkfontdir(1)` in the directory you have copied them to.

The second way to do this is to copy the directories to `/usr/X11R6/lib/X11/fonts`:

```
# cd /usr/X11R6/lib/X11/fonts
# mkdir X
# mkdir MathType1
# cd /cdrom/Unix/Files/SystemFiles/Fonts
# cp X/* /usr/X11R6/lib/X11/fonts/X
# cp Type1/* /usr/X11R6/lib/X11/fonts/MathType1
# cd /usr/X11R6/lib/X11/fonts/X
# mkfontdir
# cd ../MathType1
# mkfontdir
```

Now add the new font directories to your font path:

```
# xset fp+ /usr/X11R6/lib/X11/fonts/X
# xset fp+ /usr/X11R6/lib/X11/fonts/MathType1
# xset fp rehash
```

If you are using the **XFree86** server, you can have these font directories loaded automatically by adding them to your `XF86Config` file.

If you *do not* already have a directory called `/usr/X11R6/lib/X11/fonts/Type1`, you can change the name of the `MathType1` directory in the example above to `Type1`.

22.4 Installing Maple™

Contributed by Aaron Kaplan. Thanks to Robert Getschmann.

Maple™ is a commercial mathematics program similar to **Mathematica**. You must purchase this software from <http://www.maplesoft.com/> and then register there for a license file. To install this software on DragonFly, please follow these simple steps.

Warning: This description applies to FreeBSD, for which it was originally written. This may or may not apply to DragonFly at this point; while FreeBSD 4.x features usually translate over to DragonFly well, your mileage may vary.

1. Execute the `INSTALL` shell script from the product distribution. Choose the “RedHat” option when prompted by the installation program. A typical installation directory might be `/usr/local/maple`.
2. If you have not done so, order a license for **Maple** from Maple Waterloo Software (<http://register.maplesoft.com/>) and copy it to `/usr/local/maple/license/license.dat`.
3. Install the **FLEXlm** license manager by running the `INSTALL_LIC` install shell script that comes with **Maple**. Specify the primary hostname for your machine for the license server.

4. Patch the `/usr/local/maple/bin/maple.system.type` file with the following:

```

----- snip -----
*** maple.system.type.orig      Sun Jul  8 16:35:33 2001
--- maple.system.type      Sun Jul  8 16:35:51 2001
*****
*** 72,77 ****
--- 72,78 ----
        # the IBM RS/6000 AIX case
        MAPLE_BIN="bin.IBM_RISC_UNIX"
        ;;
+   "DragonFly" |\
    "Linux")
        # the Linux/x86 case
        # We have two Linux implementations, one for Red Hat and
----- snip end of patch -----

```

Please note that after the "DragonFly" |\ no other whitespace should be present.

This patch instructs **Maple** to recognize "DragonFly" as a type of Linux system. The `bin/maple` shell script calls the `bin/maple.system.type` shell script which in turn calls `uname -a` to find out the operating system name. Depending on the OS name it will find out which binaries to use.

5. Start the license server.

The following script, installed as `/usr/local/etc/rc.d/lmgrd.sh` is a convenient way to start up `lmgrd`:

```

----- snip -----

#!/bin/sh
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin:/usr/X11R6/bin
PATH=${PATH}:/usr/local/maple/bin:/usr/local/maple/FLEXlm/UNIX/LINUX
export PATH

LICENSE_FILE=/usr/local/maple/license/license.dat
LOG=/var/log/lmgrd.log

case "$1" in
start)
lmgrd -c ${LICENSE_FILE} 2>> ${LOG} 1>&2
echo -n " lmgrd"
;;
stop)
lmgrd -c ${LICENSE_FILE} -x lmdown 2>> ${LOG} 1>&2
;;
*)
echo "Usage: `basename $0` {start|stop}" 1>&2
exit 64
;;
esac

exit 0
----- snip -----

```

6. Test-start **Maple**:

```
% cd /usr/local/maple/bin
```

```
% ./xmaple
```

You should be up and running. Make sure to write Maplesoft to let them know you would like a native DragonFly version!

22.4.1 Common Pitfalls

- The **FLEXlm** license manager can be a difficult tool to work with. Additional documentation on the subject can be found at <http://www.globetrotter.com/>.
- `lmgrd` is known to be very picky about the license file and to core dump if there are any problems. A correct license file should look like this:

```
# =====
# License File for UNIX Installations ("Pointer File")
# =====
SERVER chillig ANY
#USE_SERVER
VENDOR maplelmg

FEATURE Maple maplelmg 2000.0831 permanent 1 XXXXXXXXXXXX \
    PLATFORMS=i86_r ISSUER="Waterloo Maple Inc." \
    ISSUED=11-may-2000 NOTICE=" Technische Universitat Wien" \
    SN=XXXXXXXXXX
```

Note: Serial number and key 'X'ed out. `chillig` is a hostname.

Editing the license file works as long as you do not touch the "FEATURE" line (which is protected by the license key).

22.5 Installing MATLAB®

Contributed by Dan Pelleg.

This document describes the process of installing the Linux version of **MATLAB® version 6.5** onto a DragonFly system. It works quite well, with the exception of the **Java Virtual Machine™** (see Section 22.5.3).

The Linux version of **MATLAB** can be ordered directly from The MathWorks at <http://www.mathworks.com>. Make sure you also get the license file or instructions how to create it. While you are there, let them know you would like a native DragonFly version of their software.

22.5.1 Installing MATLAB

To install **MATLAB**, do the following:

1. Insert the installation CD and mount it. Become `root`, as recommended by the installation script. To start the installation script type:

```
# /compat/linux/bin/sh /cdrom/install
```

Tip: The installer is graphical. If you get errors about not being able to open a display, type `setenv HOME ~USER`, where `USER` is the user you did a `su(1)` as.

2. When asked for the **MATLAB** root directory, type: `/compat/linux/usr/local/matlab`.

Tip: For easier typing on the rest of the installation process, type this at your shell prompt: `set MATLAB=/compat/linux/usr/local/matlab`

3. Edit the license file as instructed when obtaining the **MATLAB** license.

Tip: You can prepare this file in advance using your favorite editor, and copy it to `$MATLAB/license.dat` before the installer asks you to edit it.

4. Complete the installation process.

At this point your **MATLAB** installation is complete. The following steps apply “glue” to connect it to your DragonFly system.

22.5.2 License Manager Startup

1. Create symlinks for the license manager scripts:

```
# ln -s $MATLAB/etc/lmboot /usr/local/etc/lmboot_TMW
# ln -s $MATLAB/etc/lmdown /usr/local/etc/lmdown_TMW
```

2. Create a startup file at `/usr/local/etc/rc.d/flexlm.sh`. The example below is a modified version of the distributed `$MATLAB/etc/rc.lm.glnx86`. The changes are file locations, and startup of the license manager under Linux emulation.

```
#!/bin/sh
case "$1" in
  start)
    if [ -f /usr/local/etc/lmboot_TMW ]; then
      /compat/linux/bin/sh /usr/local/etc/lmboot_TMW -u username && echo 'MATLAB_lmgrd'
    fi
    ;;
  stop)
    if [ -f /usr/local/etc/lmdown_TMW ]; then
      /compat/linux/bin/sh /usr/local/etc/lmdown_TMW > /dev/null 2>&1
    fi
    ;;
  *)
    echo "Usage: $0 {start|stop}"
    exit 1
  ;;
```

```
esac

exit 0
```

Important: The file must be made executable:

```
# chmod +x /usr/local/etc/rc.d/flexlm.sh
```

You must also replace *username* above with the name of a valid user on your system (and not *root*).

3. Start the license manager with the command:

```
# /usr/local/etc/rc.d/flexlm.sh start
```

22.5.3 Linking the Java Runtime Environment

Change the **Java** Runtime Environment (JRE) link to one working under DragonFly:

```
# cd $MATLAB/sys/java/jre/glnx86/
# unlink jre; ln -s ./jre1.1.8 ./jre
```

22.5.4 Creating a MATLAB Startup Script

1. Place the following startup script in `/usr/local/bin/matlab`:

```
#!/bin/sh
/compat/linux/bin/sh /compat/linux/usr/local/matlab/bin/matlab "$@"
```

2. Then type the command `chmod +x /usr/local/bin/matlab`.

Tip: Depending on your version of `emulators/linux_base`, you may run into errors when running this script. To avoid that, edit the file `/compat/linux/usr/local/matlab/bin/matlab`, and change the line that says:

```
if [ `expr "$lsCMD" : '.*->.*'` -ne 0 ]; then
```

(in version 13.0.1 it is on line 410) to this line:

```
if test -L $newbase; then
```

22.5.5 Creating a MATLAB Shutdown Script

The following is needed to solve a problem with MATLAB not exiting correctly.

1. Create a file `$MATLAB/toolbox/local/finish.m`, and in it put the single line:

```
! $MATLAB/bin/finish.sh
```

Note: The `$MATLAB` is literal.

Tip: In the same directory, you will find the files `finishsav.m` and `finishdlg.m`, which let you save your workspace before quitting. If you use either of them, insert the line above immediately after the `save` command.

2. Create a file `$MATLAB/bin/finish.sh`, which will contain the following:

```
#!/usr/compat/linux/bin/sh
(sleep 5; killall -1 matlab_helper) &
exit 0
```

3. Make the file executable:

```
# chmod +x $MATLAB/bin/finish.sh
```

22.5.6 Using MATLAB

At this point you are ready to type `matlab` and start using it.

22.6 Installing Oracle®

Contributed by Marcel Moolenaar.

22.6.1 Preface

This document describes the process of installing **Oracle 8.0.5** and **Oracle 8.0.5.1 Enterprise Edition** for Linux onto a DragonFly machine.

Warning: This description applies to FreeBSD, for which it was originally written. This may or may not apply to DragonFly at this point; while FreeBSD 4.x features usually translate over to DragonFly well, your mileage may vary.

22.6.2 Installing the Linux Environment

Make sure you have both `emulators/linux_base` and `devel/linux_devtools` from the ports collection installed. If you run into difficulties with these ports, you may have to use the packages or older versions available in the ports collection.

If you want to run the intelligent agent, you will also need to install the Red Hat Tcl package: `tcl-8.0.3-20.i386.rpm`. The general command for installing packages with the official **RPM** port (`archivers/rpm`) is:

```
# rpm -i --ignoreos --root /compat/linux --dbpath /var/lib/rpm package
```

Installation of the *package* should not generate any errors.

22.6.3 Creating the Oracle Environment

Before you can install **Oracle**, you need to set up a proper environment. This document only describes what to do *especially* to run **Oracle** for Linux on DragonFly, not what has been described in the **Oracle** installation guide.

22.6.3.1 Kernel Tuning

As described in the **Oracle** installation guide, you need to set the maximum size of shared memory. Do not use SHMMAX under DragonFly. SHMMAX is merely calculated out of SHMMAXPGS and PGSIZE. Therefore define SHMMAXPGS. All other options can be used as described in the guide. For example:

```
options SHMMAXPGS=10000
options SHMMNI=100
options SHMSEG=10
options SEMMNS=200
options SEMMNI=70
options SEMMSL=61
```

Set these options to suit your intended use of **Oracle**.

Also, make sure you have the following options in your kernel configuration file:

```
options SYSVSHM #SysV shared memory
options SYSVSEM #SysV semaphores
options SYSVMSG #SysV interprocess communication
```

22.6.3.2 Oracle Account

Create an `oracle` account just as you would create any other account. The `oracle` account is special only that you need to give it a Linux shell. Add `/compat/linux/bin/bash` to `/etc/shells` and set the shell for the `oracle` account to `/compat/linux/bin/bash`.

22.6.3.3 Environment

Besides the normal **Oracle** variables, such as `ORACLE_HOME` and `ORACLE_SID` you must set the following environment variables:

Variable	Value
<code>LD_LIBRARY_PATH</code>	<code>\$ORACLE_HOME/lib</code>
<code>CLASSPATH</code>	<code>\$ORACLE_HOME/jdbc/lib/classes111.zip</code>
<code>PATH</code>	<code>/compat/linux/bin /compat/linux/sbin /compat/linux/usr/bin /compat/linux/usr/sbin /bin /sbin /usr/bin /usr/sbin /usr/local/bin \$ORACLE_HOME/bin</code>

It is advised to set all the environment variables in `.profile`. A complete example is:

```

ORACLE_BASE=/oracle; export ORACLE_BASE
ORACLE_HOME=/oracle; export ORACLE_HOME
LD_LIBRARY_PATH=$ORACLE_HOME/lib
export LD_LIBRARY_PATH
ORACLE_SID=ORCL; export ORACLE_SID
ORACLE_TERM=386x; export ORACLE_TERM
CLASSPATH=$ORACLE_HOME/jdbc/lib/classes111.zip
export CLASSPATH
PATH=/compat/linux/bin:/compat/linux/sbin:/compat/linux/usr/bin
PATH=$PATH:/compat/linux/usr/sbin:/bin:/sbin:/usr/bin:/usr/sbin
PATH=$PATH:/usr/local/bin:$ORACLE_HOME/bin
export PATH

```

22.6.4 Installing Oracle

Due to a slight inconsistency in the Linux emulator, you need to create a directory named `.oracle` in `/var/tmp` before you start the installer. Either make it world writable or let it be owned by the `oracle` user. You should be able to install **Oracle** without any problems. If you have problems, check your **Oracle** distribution and/or configuration first! After you have installed **Oracle**, apply the patches described in the next two subsections.

A frequent problem is that the TCP protocol adapter is not installed right. As a consequence, you cannot start any TCP listeners. The following actions help solve this problem:

```

# cd $ORACLE_HOME/network/lib
# make -f ins_network.mk ntcontab.o
# cd $ORACLE_HOME/lib
# ar r libnetwork.a ntcontab.o
# cd $ORACLE_HOME/network/lib
# make -f ins_network.mk install

```

Do not forget to run `root.sh` again!

22.6.4.1 Patching root.sh

When installing **Oracle**, some actions, which need to be performed as `root`, are recorded in a shell script called `root.sh`. This script is written in the `oraInst` directory. Apply the following patch to `root.sh`, to have it use to proper location of `chown` or alternatively run the script under a Linux native shell.

```

*** oraInst/root.sh.orig Tue Oct 6 21:57:33 1998
--- oraInst/root.sh Mon Dec 28 15:58:53 1998
*****
*** 31,37 ****
# This is the default value for CHOWN
# It will redefined later in this script for those ports
# which have it conditionally defined in ss_install.h
! CHOWN=/bin/chown
#
# Define variables to be used in this script
--- 31,37 ----
# This is the default value for CHOWN
# It will redefined later in this script for those ports

```

```
# which have it conditionally defined in ss_install.h
! CHOWN=/usr/sbin/chown
#
# Define variables to be used in this script
```

When you do not install **Oracle** from CD, you can patch the source for `root.sh`. It is called `rthd.sh` and is located in the `orainst` directory in the source tree.

22.6.4.2 Patching `genclntsh`

The script `genclntsh` is used to create a single shared client library. It is used when building the demos. Apply the following patch to comment out the definition of `PATH`:

```
*** bin/genclntsh.orig Wed Sep 30 07:37:19 1998
--- bin/genclntsh Tue Dec 22 15:36:49 1998
*****
*** 32,38 ****
#
# Explicit path to ensure that we're using the correct commands
#PATH=/usr/bin:/usr/ccs/bin export PATH
! PATH=/usr/local/bin:/bin:/usr/bin:/usr/X11R6/bin export PATH
#
# each product MUST provide a $PRODUCT/admin/shrept.lst
--- 32,38 ----
#
# Explicit path to ensure that we're using the correct commands
#PATH=/usr/bin:/usr/ccs/bin export PATH
! #PATH=/usr/local/bin:/bin:/usr/bin:/usr/X11R6/bin export PATH
#
# each product MUST provide a $PRODUCT/admin/shrept.lst
```

22.6.5 Running Oracle

When you have followed the instructions, you should be able to run **Oracle** as if it was run on Linux itself.

22.7 Installing SAP® R/3®

Contributed by Holger Kipp. Original version converted to SGML by Valentino Vaschetto.

Installations of **SAP** Systems using DragonFly will not be supported by the SAP support team — they only offer support for certified platforms.

Warning: This description applies to FreeBSD, for which it was originally written. This may or may not apply to DragonFly at this point; while FreeBSD 4.x features usually translate over to DragonFly well, your mileage may vary.

22.7.1 Preface

This document describes a possible way of installing a **SAP R/3 System** with **Oracle Database** for Linux onto a DragonFly machine, including the installation of DragonFly and **Oracle**. Two different configurations will be described:

- **SAP R/3 4.6B (IDES)** with **Oracle 8.0.5** on FreeBSD 4.3-STABLE
- **SAP R/3 4.6C** with **Oracle 8.1.7** on FreeBSD 4.5-STABLE

Even though this document tries to describe all important steps in a greater detail, it is not intended as a replacement for the **Oracle** and **SAP R/3** installation guides.

Please see the documentation that comes with the **SAP R/3** Linux edition for **SAP** and **Oracle** specific questions, as well as resources from **Oracle** and **SAP OSS**.

22.7.2 Software

The following CD-ROMs have been used for **SAP** installations:

22.7.2.1 SAP R/3 4.6B, Oracle 8.0.5

Name	Number	Description
KERNEL	51009113	SAP Kernel Oracle / Installation / AIX, Linux, Solaris
RDBMS	51007558	Oracle / RDBMS 8.0.5.X / Linux
EXPORT1	51010208	IDES / DB-Export / Disc 1 of 6
EXPORT2	51010209	IDES / DB-Export / Disc 2 of 6
EXPORT3	51010210	IDES / DB-Export / Disc 3 of 6
EXPORT4	51010211	IDES / DB-Export / Disc 4 of 6
EXPORT5	51010212	IDES / DB-Export / Disc 5 of 6
EXPORT6	51010213	IDES / DB-Export / Disc 6 of 6

Additionally, we used the **Oracle 8 Server** (Pre-production version 8.0.5 for Linux, Kernel Version 2.0.33) CD which is not really necessary, and FreeBSD 4.3-STABLE (it was only a few days past 4.3 RELEASE).

22.7.2.2 SAP R/3 4.6C SR2, Oracle 8.1.7

Name	Number	Description
KERNEL	51014004	SAP Kernel Oracle / SAP Kernel Version 4.6D / DEC, Linux
RDBMS	51012930	Oracle 8.1.7/ RDBMS / Linux
EXPORT1	51013953	Release 4.6C SR2 / Export / Disc 1 of 4

Name	Number	Description
EXPORT1	51013953	Release 4.6C SR2 / Export / Disc 2 of 4
EXPORT1	51013953	Release 4.6C SR2 / Export / Disc 3 of 4
EXPORT1	51013953	Release 4.6C SR2 / Export / Disc 4 of 4
LANG1	51013954	Release 4.6C SR2 / Language / DE, EN, FR / Disc 1 of 3

Depending on the languages you would like to install, additional language CDs might be necessary. Here we are just using DE and EN, so the first language CD is the only one needed. As a little note, the numbers for all four EXPORT CDs are identical. All three language CDs also have the same number (this is different from the 4.6B IDES release CD numbering). At the time of writing this installation is running on FreeBSD 4.5-STABLE (20.03.2002).

22.7.3 SAP Notes

The following notes should be read before installing **SAP R/3** and proved to be useful during installation:

22.7.3.1 SAP R/3 4.6B, Oracle 8.0.5

Number	Title
0171356	SAP Software on Linux: Essential Comments
0201147	INST: 4.6C R/3 Inst. on UNIX - Oracle
0373203	Update / Migration Oracle 8.0.5 --> 8.0.6/8.1.6 LINUX
0072984	Release of Digital UNIX 4.0B for Oracle
0130581	R3SETUP step DIPGNTAB terminates
0144978	Your system has not been installed correctly
0162266	Questions and tips for R3SETUP on Windows NT / W2K

22.7.3.2 SAP R/3 4.6C, Oracle 8.1.7

Number	Title
0015023	Initializing table TCPDB (RSXP0004) (EBCDIC)
0045619	R/3 with several languages or typefaces
0171356	SAP Software on Linux: Essential Comments
0195603	RedHat 6.1 Enterprise version: Known problems
0212876	The new archiving tool SAPCAR
0300900	Linux: Released DELL Hardware
0377187	RedHat 6.2: important remarks

Number	Title
0387074	INST: R/3 4.6C SR2 Installation on UNIX
0387077	INST: R/3 4.6C SR2 Inst. on UNIX - Oracle
0387078	SAP Software on UNIX: OS Dependencies 4.6C SR2

22.7.4 Hardware Requirements

The following equipment is sufficient for the installation of a **SAP R/3 System**. For production use, a more exact sizing is of course needed:

Component	4.6B	4.6C
Processor	2 x 800MHz Pentium III	2 x 800MHz Pentium III
Memory	1GB ECC	2GB ECC
Hard Disk Space	50-60GB (IDES)	50-60GB (IDES)

For use in production, Xeon Processors with large cache, high-speed disk access (SCSI, RAID hardware controller), USV and ECC-RAM is recommended. The large amount of hard disk space is due to the preconfigured IDES System, which creates 27 GB of database files during installation. This space is also sufficient for initial production systems and application data.

22.7.4.1 SAP R/3 4.6B, Oracle 8.0.5

The following off-the-shelf hardware was used: a dual processor board with two 800 MHz Pentium III processors, Adaptec 29160 Ultra160 SCSI adapter (for accessing a 40/80 GB DLT tape drive and CDROM), Mylex AcceleRAID™ (2 channels, firmware 6.00-1-00 with 32 MB RAM). To the Mylex RAID controller are attached two 17 GB hard disks (mirrored) and four 36 GB hard disks (RAID level 5).

22.7.4.2 SAP R/3 4.6C, Oracle 8.1.7

For this installation a Dell™ PowerEdge™ 2500 was used: a dual processor board with two 1000 MHz Pentium III processors (256 kB Cache), 2 GB PC133 ECC SDRAM, PERC/3 DC PCI RAID Controller with 128 MB, and an EIDE DVD-ROM drive. To the RAID controller are attached two 18 GB hard disks (mirrored) and four 36 GB hard disks (RAID level 5).

22.7.5 Installation of FreeBSD

First you have to install FreeBSD.

22.7.5.1 Disk Layout

To keep it simple, the same disk layout both for the **SAP R/3 46B** and **SAP R/3 46C SR2** installation was used. Only the device names changed, as the installations were on different hardware (`/dev/da` and `/dev/amr` respectively, so if using an AMI MegaRAID, one will see `/dev/amr0s1a` instead of `/dev/da0s1a`):

File system	Size (1k-blocks)	Size (GB)	Mounted on
/dev/da0s1a	1.016.303	1	/
/dev/da0s1b		6	swap
/dev/da0s1e	2.032.623	2	/var
/dev/da0s1f	8.205.339	8	/usr
/dev/da1s1e	45.734.361	45	/compat/linux/oracle
/dev/da1s1f	2.032.623	2	/compat/linux/sapmnt
/dev/da1s1g	2.032.623	2	/compat/linux/usr/sap

Configure and initialize the two logical drives with the Mylex or PERC/3 RAID software beforehand. The software can be started during the BIOS boot phase.

Please note that this disk layout differs slightly from the SAP recommendations, as SAP suggests mounting the **Oracle** subdirectories (and some others) separately — we decided to just create them as real subdirectories for simplicity.

22.7.5.2 make world and a New Kernel

Download the latest -STABLE sources. Rebuild world and your custom kernel after configuring your kernel configuration file. Here you should also include the kernel parameters which are required for both **SAP R/3** and **Oracle**.

22.7.6 Installing the Linux Environment

22.7.6.1 Installing the Linux Base System

First the linux_base port needs to be installed (as root):

```
# cd /usr/ports/emulators/linux_base
# make install distclean
```

22.7.6.2 Installing Linux Development Environment

The Linux development environment is needed, if you want to install **Oracle** on FreeBSD according to the Section 22.6:

```
# cd /usr/ports/devel/linux_devtools
# make install distclean
```

The Linux development environment has only been installed for the **SAP R/3 46B IDES** installation. It is not needed, if the **Oracle DB** is not relinked on the FreeBSD system. This is the case if you are using the **Oracle** tarball from a Linux system.

22.7.6.3 Installing the Necessary RPMs

To start the R3SETUP program, PAM support is needed. During the first **SAP** Installation on FreeBSD 4.3-STABLE we tried to install PAM with all the required packages and finally forced the installation of the PAM package, which worked. For **SAP R/3 4.6C SR2** we directly forced the installation of the PAM RPM, which also works, so it seems the dependent packages are not needed:

```
# rpm -i --ignoreos --nodeps --root /compat/linux --dbpath /var/lib/rpm \
pam-0.68-7.i386.rpm
```

For **Oracle 8.0.5** to run the intelligent agent, we also had to install the RedHat Tcl package `tcl-8.0.5-30.i386.rpm` (otherwise the relinking during **Oracle** installation will not work). There are some other issues regarding relinking of **Oracle**, but that is a **Oracle** Linux issue, not DragonFly specific.

22.7.6.4 Some Additional Hints

It might also be a good idea to add `linprocfs` to `/etc/fstab`, for more informations, see the `linprocfs(5)` manual page. Another parameter to set is `kern.fallback_elf_brand=3` which is done in the file `/etc/sysctl.conf`.

22.7.7 Creating the SAP R/3 Environment

22.7.7.1 Creating the Necessary File Systems and Mountpoints

For a simple installation, it is sufficient to create the following file systems:

mount point	size in GB
<code>/compat/linux/oracle</code>	45 GB
<code>/compat/linux/sapmnt</code>	2 GB
<code>/compat/linux/usr/sap</code>	2 GB

It is also necessary to create some links. Otherwise the **SAP** Installer will complain, as it is checking the created links:

```
# ln -s /compat/linux/oracle /oracle
# ln -s /compat/linux/sapmnt /sapmnt
# ln -s /compat/linux/usr/sap /usr/sap
```

Possible error message during installation (here with System *PRD* and the **SAP R/3 4.6C SR2** installation):

```
INFO 2002-03-19 16:45:36 R3LINKS_IND_IND SyLinkCreate:200
    Checking existence of symbolic link /usr/sap/PRD/SYS/exe/dbg to
    /sapmnt/PRD/exe. Creating if it does not exist...

WARNING 2002-03-19 16:45:36 R3LINKS_IND_IND SyLinkCreate:400
    Link /usr/sap/PRD/SYS/exe/dbg exists but it points to file
    /compat/linux/sapmnt/PRD/exe instead of /sapmnt/PRD/exe. The
    program cannot go on as long as this link exists at this
    location. Move the link to another location.
```

```
ERROR 2002-03-19 16:45:36 R3LINKS_IND_IND Ins_SetupLinks:0
  can not setup link '/usr/sap/PRD/SYS/exe/dbg' with content
  '/sapmnt/PRD/exe'
```

22.7.7.2 Creating Users and Directories

SAP R/3 needs two users and three groups. The user names depend on the **SAP** system ID (SID) which consists of three letters. Some of these SIDs are reserved by **SAP** (for example **SAP** and **NIX**. For a complete list please see the **SAP** documentation). For the IDES installation we used **IDS**, for the 4.6C SR2 installation **PRD**, as that system is intended for production use. We have therefore the following groups (group IDs might differ, these are just the values we used with our installation):

group ID	group name	description
100	dba	Data Base Administrator
101	sapsys	SAP System
102	oper	Data Base Operator

For a default **Oracle** installation, only group **dba** is used. As **oper** group, one also uses group **dba** (see **Oracle** and **SAP** documentation for further information).

We also need the following users:

user ID	user name	generic name	group	additional groups	description
1000	idsadm/prdadm	sidadm	sapsys	oper	SAP Administrator
1002	oraids/oraprd	orasid	dba	oper	Oracle Administrator

Adding the users with `adduser(8)` requires the following (please note shell and home directory) entries for “SAP Administrator”:

```
Name: sidadm
Password: *****
Fullname: SAP Administrator SID
Uid: 1000
Gid: 101 (sapsys)
Class:
Groups: sapsys dba
HOME: /home/sidadm
Shell: bash (/compat/linux/bin/bash)
```

and for “Oracle Administrator”:

```
Name: orasid
Password: *****
Fullname: Oracle Administrator SID
Uid: 1002
```

```
Gid: 100 (dba)
Class:
Groups: dba
HOME: /oracle/sid
Shell: bash (/compat/linux/bin/bash)
```

This should also include group `oper` in case you are using both groups `dba` and `oper`.

22.7.7.3 Creating Directories

These directories are usually created as separate file systems. This depends entirely on your requirements. We choose to create them as simple directories, as they are all located on the same RAID 5 anyway:

First we will set owners and rights of some directories (as user `root`):

```
# chmod 775 /oracle
# chmod 777 /sapmnt
# chown root:dba /oracle
# chown sidadm:sapsys /compat/linux/usr/sap
# chmod 775 /compat/linux/usr/sap
```

Second we will create directories as user `orasid`. These will all be subdirectories of `/oracle/SID`:

```
# su - orasid
# cd /oracle/SID
# mkdir mirrlogA mirrlogB origlogA origlogB
# mkdir sapdata1 sapdata2 sapdata3 sapdata4 sapdata5 sapdata6
# mkdir saparch sapreorg
# exit
```

For the **Oracle 8.1.7** installation some additional directories are needed:

```
# su - orasid
# cd /oracle
# mkdir 805_32
# mkdir client stage
# mkdir client/80x_32
# mkdir stage/817_32
# cd /oracle/SID
# mkdir 817_32
```

Note: The directory `client/80x_32` is used with exactly this name. Do not replace the `x` with some number or anything.

In the third step we create directories as user `sidadm`:

```
# su - sidadm
# cd /usr/sap
# mkdir SID
# mkdir trans
# exit
```

22.7.7.4 Entries in `/etc/services`

SAP R/3 requires some entries in file `/etc/services`, which will not be set correctly during installation under FreeBSD. Please add the following entries (you need at least those entries corresponding to the instance number — in this case, 00. It will do no harm adding all entries from 00 to 99 for `dp`, `gw`, `sp` and `ms`). If you are going to use a **SAProuter** or need to access **SAP OSS**, you also need 99, as port 3299 is usually used for the **SAProuter** process on the target system:

```
sapdp00    3200/tcp # SAP Dispatcher.          3200 + Instance-Number
sapgw00    3300/tcp # SAP Gateway.                3300 + Instance-Number
sapsp00    3400/tcp #                               3400 + Instance-Number
sapms00    3500/tcp #                               3500 + Instance-Number
sapmsSID   3600/tcp # SAP Message Server.     3600 + Instance-Number
sapgw00s   4800/tcp # SAP Secure Gateway          4800 + Instance-Number
```

22.7.7.5 Necessary Locales

SAP requires at least two locales that are not part of the default RedHat installation. **SAP** offers the required RPMs as download from their FTP server (which is only accessible if you are a customer with OSS access). See note 0171356 for a list of RPMs you need.

It is also possible to just create appropriate links (for example from `de_DE` and `en_US`), but we would not recommend this for a production system (so far it worked with the IDES system without any problems, though). The following locales are needed:

```
de_DE.ISO-8859-1
en_US.ISO-8859-1
```

Create the links like this:

```
# cd /compat/linux/usr/share/locale
# ln -s de_DE de_DE.ISO-8859-1
# ln -s en_US en_US.ISO-8859-1
```

If they are not present, there will be some problems during the installation. If these are then subsequently ignored (by setting the `STATUS` of the offending steps to `OK` in file `CENTRDB.R3S`), it will be impossible to log onto the **SAP** system without some additional effort.

22.7.7.6 Kernel Tuning

SAP R/3 systems need a lot of resources. We therefore added the following parameters to the kernel configuration file:

```
# Set these for memory pigs (SAP and Oracle):
options MAXDSIZ="(1024*1024*1024)"
options DFLDSIZ="(1024*1024*1024)"
# System V options needed.
options SYSVSHM #SYSV-style shared memory
options SHMMAXPGS=262144 #max amount of shared mem. pages
#options SHMMAXPGS=393216 #use this for the 46C inst.parameters
options SHMMNI=256 #max number of shared memory ident if.
```



```

options SHMSEG=100 #max shared mem.segs per process
options SYSVMSG #SYSV-style message queues
options MSGSEG=32767 #max num. of mes.segments in system
options MSGSSZ=32 #size of msg-seg. MUST be power of 2
options MSGMNB=65535 #max char. per message queue
options MSGTQL=2046 #max amount of msgs in system
options SYSVSEM #SYSV-style semaphores
options SEMMNU=256 #number of semaphore UNDO structures
options SEMMNS=1024 #number of semaphores in system
options SEMMNI=520 #number of semaphore identifiers
options SEMUME=100      #number of UNDO keys

```

The minimum values are specified in the documentation that comes from SAP. As there is no description for Linux, see the HP-UX section (32-bit) for further information. As the system for the 4.6C SR2 installation has more main memory, the shared segments can be larger both for **SAP** and **Oracle**, therefore choose a larger number of shared memory pages.

Note: With the default installation of FreeBSD 4.5 on i386, leave `MAXDSIZ` and `DFLDSIZ` at 1 GB maximum. Otherwise, strange errors like `ORA-27102: out of memory` and `Linux Error: 12: Cannot allocate memory` might happen.

22.7.8 Installing SAP R/3

22.7.8.1 Preparing SAP CDRoms

There are many CDRoms to mount and unmount during the installation. Assuming you have enough CDRom drives, you can just mount them all. We decided to copy the CDRoms contents to corresponding directories:

```
/oracle/SID/sapreorg/cd-name
```

where `cd-name` was one of `KERNEL`, `RDBMS`, `EXPORT1`, `EXPORT2`, `EXPORT3`, `EXPORT4`, `EXPORT5` and `EXPORT6` for the 4.6B/IDES installation, and `KERNEL`, `RDBMS`, `DISK1`, `DISK2`, `DISK3`, `DISK4` and `LANG` for the 4.6C SR2 installation. All the filenames on the mounted CDs should be in capital letters, otherwise use the `-g` option for mounting. So use the following commands:

```

# mount_cd9660 -g /dev/cd0a /mnt
# cp -R /mnt/* /oracle/SID/sapreorg/cd-name
# umount /mnt

```

22.7.8.2 Running the Installation Script

First you have to prepare an `install` directory:

```

# cd /oracle/SID/sapreorg
# mkdir install
# cd install

```

Then the installation script is started, which will copy nearly all the relevant files into the `install` directory:

```
# /oracle/SID/sapreorg/KERNEL/UNIX/INSTTOOL.SH
```

The IDES installation (4.6B) comes with a fully customized SAP R/3 demonstration system, so there are six instead of just three EXPORT CDs. At this point the installation template `CENTRDB.R3S` is for installing a standard central instance (**R/3** and database), not the IDES central instance, so one needs to copy the corresponding `CENTRDB.R3S` from the `EXPORT1` directory, otherwise `R3SETUP` will only ask for three EXPORT CDs.

The newer **SAP 4.6C SR2** release comes with four EXPORT CDs. The parameter file that controls the installation steps is `CENTRAL.R3S`. Contrary to earlier releases there are no separate installation templates for a central instance with or without database. **SAP** is using a separate template for database installation. To restart the installation later it is however sufficient to restart with the original file.

During and after installation, **SAP** requires `hostname` to return the computer name only, not the fully qualified domain name. So either set the `hostname` accordingly, or set an alias with `alias hostname='hostname -s'` for both `orasid` and `sidadm` (and for `root` at least during installation steps performed as `root`). It is also possible to adjust the installed `.profile` and `.login` files of both users that are installed during **SAP** installation.

22.7.8.3 Start R3SETUP 4.6B

Make sure `LD_LIBRARY_PATH` is set correctly:

```
# export LD_LIBRARY_PATH=/oracle/IDS/lib:/sapmnt/IDS/exe:/oracle/805_32/lib
```

Start `R3SETUP` as `root` from installation directory:

```
# cd /oracle/IDS/sapreorg/install
# ./R3SETUP -f CENTRDB.R3S
```

The script then asks some questions (defaults in brackets, followed by actual input):

Question	Default	Input
Enter SAP System ID	[C11]	IDSEnter
Enter SAP Instance Number	[00]	Enter
Enter SAPMOUNT Directory	[/sapmnt]	Enter
Enter name of SAP central host	[troubadix.domain.de]	Enter
Enter name of SAP db host	[troubadix]	Enter
Select character set	[1] (WE8DEC)	Enter
Enter Oracle server version (1) Oracle 8.0.5, (2) Oracle 8.0.6, (3) Oracle 8.1.5, (4) Oracle 8.1.6		1Enter
Extract Oracle Client archive	[1] (Yes, extract)	Enter
Enter path to KERNEL CD	[/sapcd]	/oracle/IDS/sapreorg/KERNEL
Enter path to RDBMS CD	[/sapcd]	/oracle/IDS/sapreorg/RDBMS
Enter path to EXPORT1 CD	[/sapcd]	/oracle/IDS/sapreorg/EXPORT1
Directory to copy EXPORT1 CD	[/oracle/IDS/sapreorg/CD4_DIR]	Enter
Enter path to EXPORT2 CD	[/sapcd]	/oracle/IDS/sapreorg/EXPORT2

Question	Default	Input
Directory to copy EXPORT2 CD	[/oracle/IDS/sapreorg/CD5_DIR]	Enter
Enter path to EXPORT3 CD	[/sapcd]	/oracle/IDS/sapreorg/EXPORT3
Directory to copy EXPORT3 CD	[/oracle/IDS/sapreorg/CD6_DIR]	Enter
Enter path to EXPORT4 CD	[/sapcd]	/oracle/IDS/sapreorg/EXPORT4
Directory to copy EXPORT4 CD	[/oracle/IDS/sapreorg/CD7_DIR]	Enter
Enter path to EXPORT5 CD	[/sapcd]	/oracle/IDS/sapreorg/EXPORT5
Directory to copy EXPORT5 CD	[/oracle/IDS/sapreorg/CD8_DIR]	Enter
Enter path to EXPORT6 CD	[/sapcd]	/oracle/IDS/sapreorg/EXPORT6
Directory to copy EXPORT6 CD	[/oracle/IDS/sapreorg/CD9_DIR]	Enter
Enter amount of RAM for SAP + DB		850 Enter (in Megabytes)
Service Entry Message Server	[3600]	Enter
Enter Group-ID of sapsys	[101]	Enter
Enter Group-ID of oper	[102]	Enter
Enter Group-ID of dba	[100]	Enter
Enter User-ID of <i>sidadm</i>	[1000]	Enter
Enter User-ID of <i>orasid</i>	[1002]	Enter
Number of parallel procs	[2]	Enter

If you had not copied the CDs to the different locations, then the **SAP** installer cannot find the CD needed (identified by the LABEL.ASC file on the CD) and would then ask you to insert and mount the CD and confirm or enter the mount path.

The CENTRDB.R3S might not be error free. In our case, it requested EXPORT4 CD again but indicated the correct key (6_LOCATION, then 7_LOCATION etc.), so one can just continue with entering the correct values.

Apart from some problems mentioned below, everything should go straight through up to the point where the Oracle database software needs to be installed.

22.7.8.4 Start R3SETUP 4.6C SR2

Make sure LD_LIBRARY_PATH is set correctly. This is a different value from the 4.6B installation with **Oracle 8.0.5**:

```
# export LD_LIBRARY_PATH=/sapmnt/PRD/exe:/oracle/PRD/817_32/lib
```

Start R3SETUP as user `root` from installation directory:

```
# cd /oracle/PRD/sapreorg/install
```

```
# ./R3SETUP -f CENTRAL.R3S
```

The script then asks some questions (defaults in brackets, followed by actual input):

Question	Default	Input
Enter SAP System ID	[C11]	PRD Enter
Enter SAP Instance Number	[00]	Enter
Enter SAPMOUNT Directory	[/sapmnt]	Enter

Question	Default	Input
Enter name of SAP central host	[majestix]	Enter
Enter Database System ID	[PRD]	PRD Enter
Enter name of SAP db host	[majestix]	Enter
Select character set	[1] (WE8DEC)	Enter
Enter Oracle server version (2) Oracle 8.1.7		2 Enter
Extract Oracle Client archive	[1] (Yes, extract)	Enter
Enter path to KERNEL CD	[/sapcd]	/oracle/PRD/sapreorg/KERNEL
Enter amount of RAM for SAP + DB	2044	1800 Enter (in Megabytes)
Service Entry Message Server	[3600]	Enter
Enter Group-ID of sapsys	[100]	Enter
Enter Group-ID of oper	[101]	Enter
Enter Group-ID of dba	[102]	Enter
Enter User-ID of oraprd	[1002]	Enter
Enter User-ID of prdadm	[1000]	Enter
LDAP support		3 Enter (no support)
Installation step completed	[1] (continue)	Enter
Choose installation service	[1] (DB inst,file)	Enter

So far, creation of users gives an error during installation in phases OSUSERDBSID_IND_ORA (for creating user *orasid*) and OSUSERSIDADM_IND_ORA (creating user *sidadm*).

Apart from some problems mentioned below, everything should go straight through up to the point where the Oracle database software needs to be installed.

22.7.9 Installing Oracle 8.0.5

Please see the corresponding SAP Notes and Oracle Readmes regarding Linux and **Oracle DB** for possible problems. Most if not all problems stem from incompatible libraries.

For more information on installing **Oracle**, refer to the Installing Oracle chapter.

22.7.9.1 Installing the Oracle 8.0.5 with `orainst`

If **Oracle 8.0.5** is to be used, some additional libraries are needed for successfully relinking, as **Oracle 8.0.5** was linked with an old glibc (RedHat 6.0), but RedHat 6.1 already uses a new glibc. So you have to install the following additional packages to ensure that linking will work:

```
compat-libs-5.2-2.i386.rpm
compat-glibc-5.2-2.0.7.2.i386.rpm
compat-egcs-5.2-1.0.3a.1.i386.rpm
compat-egcs-c++-5.2-1.0.3a.1.i386.rpm
compat-binutils-5.2-2.9.1.0.23.1.i386.rpm
```

See the corresponding SAP Notes or Oracle Readmes for further information. If this is no option (at the time of installation we did not have enough time to check this), one could use the original binaries, or use the relinked binaries from an original RedHat system.

For compiling the intelligent agent, the RedHat Tcl package must be installed. If you cannot get `tcl-8.0.3-20.i386.rpm`, a newer one like `tcl-8.0.5-30.i386.rpm` for RedHat 6.1 should also do.

Apart from relinking, the installation is straightforward:

```
# su - oraids
# export TERM=xterm
# export ORACLE_TERM=xterm
# export ORACLE_HOME=/oracle/IDS
# cd /ORACLE_HOME/orainst_sap
# ./orainst
```

Confirm all screens with **Enter** until the software is installed, except that one has to deselect the *Oracle On-Line Text Viewer*, as this is not currently available for Linux. **Oracle** then wants to relink with `i386-glibc20-linux-gcc` instead of the available `gcc`, `egcs` or `i386-redhat-linux-gcc`.

Due to time constrains we decided to use the binaries from an **Oracle 8.0.5 PreProduction** release, after the first attempt at getting the version from the RDBMS CD working, failed, and finding and accessing the correct RPMs was a nightmare at that time.

22.7.9.2 Installing the Oracle 8.0.5 Pre-production Release for Linux (Kernel 2.0.33)

This installation is quite easy. Mount the CD, start the installer. It will then ask for the location of the Oracle home directory, and copy all binaries there. We did not delete the remains of our previous RDBMS installation tries, though.

Afterwards, **Oracle** Database could be started with no problems.

22.7.10 Installing the Oracle 8.1.7 Linux Tarball

Take the tarball `oracle81732.tgz` you produced from the installation directory on a Linux system and untar it to `/oracle/SID/817_32/`.

22.7.11 Continue with SAP R/3 Installation

First check the environment settings of users `idsamd` (`sidadm`) and `oraids` (`orasid`). They should now both have the files `.profile`, `.login` and `.cshrc` which are all using `hostname`. In case the system's hostname is the fully qualified name, you need to change `hostname` to `hostname -s` within all three files.

22.7.11.1 Database Load

Afterwards, `R3SETUP` can either be restarted or continued (depending on whether `exit` was chosen or not). `R3SETUP` then creates the tablespaces and loads the data (for 46B IDES, from `EXPORT1` to `EXPORT6`, for 46C from `DISK1` to `DISK4`) with `R3load` into the database.

When the database load is finished (might take a few hours), some passwords are requested. For test installations, one can use the well known default passwords (use different ones if security is an issue!):

Question	Input
Enter Password for sapr3	sapEnter
Confirm Password for sapr3	sapEnter
Enter Password for sys	change_on_installEnter
Confirm Password for sys	change_on_installEnter
Enter Password for system	managerEnter
Confirm Password for system	managerEnter

At this point We had a few problems with `diagntab` during the 4.6B installation.

22.7.11.2 Listener

Start the **Oracle** Listener as user `orasid` as follows:

```
% umask 0; lsnrctl start
```

Otherwise you might get the error ORA-12546 as the sockets will not have the correct permissions. See SAP Note 072984.

22.7.11.3 Updating MNLS Tables

If you plan to import non-Latin-1 languages into the **SAP** system, you have to update the Multi National Language Support tables. This is described in the SAP OSS Notes 15023 and 45619. Otherwise, you can skip this question during **SAP** installation.

Note: If you do not need MNLS, it is still necessary to check the table TCPDB and initializing it if this has not been done. See SAP note 0015023 and 0045619 for further information.

22.7.12 Post-installation Steps

22.7.12.1 Request SAP R/3 License Key

You have to request your **SAP R/3** License Key. This is needed, as the temporary license that was installed during installation is only valid for four weeks. First get the hardware key. Log on as user `idsadm` and call `saplicense`:

```
# /sapmnt/IDS/exe/saplicense -get
```

Calling `saplicense` without parameters gives a list of options. Upon receiving the license key, it can be installed using:

```
# /sapmnt/IDS/exe/saplicense -install
```

You are then required to enter the following values:

```
SAP SYSTEM ID   = SID, 3 chars
CUSTOMER KEY    = hardware key, 11 chars
INSTALLATION NO = installation, 10 digits
EXPIRATION DATE = yyyyymmdd, usually "99991231"
LICENSE KEY     = license key, 24 chars
```

22.7.12.2 Creating Users

Create a user within client 000 (for some tasks required to be done within client 000, but with a user different from users `sap*` and `ddic`). As a user name, We usually choose `wartung` (or `service` in English). Profiles required are `sap_new` and `sap_all`. For additional safety the passwords of default users within all clients should be changed (this includes users `sap*` and `ddic`).

22.7.12.3 Configure Transport System, Profile, Operation Modes, Etc.

Within client 000, user different from `ddic` and `sap*`, do at least the following:

Task	Transaction
Configure Transport System, e.g. as <i>Stand-Alone Transport Domain Entity</i>	STMS
Create / Edit Profile for System	RZ10
Maintain Operation Modes and Instances	RZ04

These and all the other post-installation steps are thoroughly described in **SAP** installation guides.

22.7.12.4 Edit `initsid.sap` (`initIDS.sap`)

The file `/oracle/IDS/dbs/initIDS.sap` contains the **SAP** backup profile. Here the size of the tape to be used, type of compression and so on need to be defined. To get this running with `sapdba / brbackup`, we changed the following values:

```
compress = hardware
archive_function = copy_delete_save
cpio_flags = "-ov --format=newc --block-size=128 --quiet"
cpio_in_flags = "-iuv --block-size=128 --quiet"
tape_size = 38000M
tape_address = /dev/nsa0
tape_address_rew = /dev/sa0
```

Explanations:

`compress`: The tape we use is a HP DLT1 which does hardware compression.

`archive_function`: This defines the default behavior for saving Oracle archive logs: new logfiles are saved to tape, already saved logfiles are saved again and are then deleted. This prevents lots of trouble if you need to recover the database, and one of the archive-tapes has gone bad.

`cpio_flags`: Default is to use `-B` which sets block size to 5120 Bytes. For DLT Tapes, HP recommends at least 32 K block size, so we used `--block-size=128` for 64 K. `--format=newc` is needed because we have inode numbers greater than 65535. The last option `--quiet` is needed as otherwise `brbackup` complains as soon as `cpio` outputs the numbers of blocks saved.

`cpio_in_flags`: Flags needed for loading data back from tape. Format is recognized automatically.

`tape_size`: This usually gives the raw storage capability of the tape. For security reason (we use hardware compression), the value is slightly lower than the actual value.

`tape_address`: The non-rewindable device to be used with `cpio`.

`tape_address_rew`: The rewindable device to be used with `cpio`.

22.7.12.5 Configuration Issues after Installation

The following SAP parameters should be tuned after installation (examples for IDES 46B, 1 GB memory):

Name	Value
ztta/roll_extension	250000000
abap/heap_area_dia	300000000
abap/heap_area_nondia	400000000
em/initial_size_MB	256
em/blocksize_kB	1024
ipc/shm_psize_40	70000000

SAP Note 0013026:

Name	Value
ztta/dynpro_area	2500000

SAP Note 0157246:

Name	Value
rdisp/ROLL_MAXFS	16000
rdisp/PG_MAXFS	30000

Note: With the above parameters, on a system with 1 gigabyte of memory, one may find memory consumption similar to:

```
Mem: 547M Active, 305M Inact, 109M Wired, 40M Cache, 112M Buf, 3492K Free
```


22.7.13 Problems during Installation

22.7.13.1 Restart R3SETUP after Fixing a Problem

R3SETUP stops if it encounters an error. If you have looked at the corresponding logfiles and fixed the error, you have to start R3SETUP again, usually selecting REPEAT as option for the last step R3SETUP complained about.

To restart R3SETUP, just start it with the corresponding R3S file:

```
# ./R3SETUP -f CENTRDB.R3S
```

for 4.6B, or with

```
# ./R3SETUP -f CENTRAL.R3S
```

for 4.6C, no matter whether the error occurred with CENTRAL.R3S or DATABASE.R3S.

Note: At some stages, R3SETUP assumes that both database and **SAP** processes are up and running (as those were steps it already completed). Should errors occur and for example the database could not be started, you have to start both database and **SAP** by hand after you fixed the errors and before starting R3SETUP again.

Do not forget to also start the **Oracle** listener again (as `orasid` with `umask 0; lsnrctl start`) if it was also stopped (for example due to a necessary reboot of the system).

22.7.13.2 OSUSERSIDADM_IND_ORA during R3SETUP

If R3SETUP complains at this stage, edit the template file R3SETUP used at that time (CENTRDB.R3S (4.6B) or either CENTRAL.R3S or DATABASE.R3S (4.6C)). Locate [OSUSERSIDADM_IND_ORA] or search for the only STATUS=ERROR entry and edit the following values:

```
HOME=/home/sidadm (was empty)
STATUS=OK (had status ERROR)
```

Then you can restart R3SETUP again.

22.7.13.3 OSUSERDBSID_IND_ORA during R3SETUP

Possibly R3SETUP also complains at this stage. The error here is similar to the one in phase OSUSERSIDADM_IND_ORA. Just edit the template file R3SETUP used at that time (CENTRDB.R3S (4.6B) or either CENTRAL.R3S or DATABASE.R3S (4.6C)). Locate [OSUSERDBSID_IND_ORA] or search for the only STATUS=ERROR entry and edit the following value in that section:

```
STATUS=OK
```

Then restart R3SETUP.

22.7.13.4 oraview.vrf FILE NOT FOUND during Oracle Installation

You have not deselected *Oracle On-Line Text Viewer* before starting the installation. This is marked for installation even though this option is currently not available for Linux. Deselect this product inside the **Oracle** installation menu and restart installation.

22.7.13.5 TEXTENV_INVALID during R3SETUP, RFC or SAPgui Start

If this error is encountered, the correct locale is missing. SAP Note 0171356 lists the necessary RPMs that need be installed (e.g. `saplocales-1.0-3`, `saposcheck-1.0-1` for RedHat 6.1). In case you ignored all the related errors and set the corresponding STATUS from ERROR to OK (in `CENTRDB.R3S`) every time `R3SETUP` complained and just restarted `R3SETUP`, the **SAP** system will not be properly configured and you will then not be able to connect to the system with a **SAPgui**, even though the system can be started. Trying to connect with the old Linux **SAPgui** gave the following messages:

```
Sat May 5 14:23:14 2001
*** ERROR => no valid userarea given [trgmsggo. 0401]
Sat May 5 14:23:22 2001
*** ERROR => ERROR NR 24 occurred [trgmsggi. 0410]
*** ERROR => Error when generating text environment. [trgmsggi. 0435]
*** ERROR => function failed [trgmsggi. 0447]
*** ERROR => no socket operation allowed [trxio.c 3363]
Speicherzugriffsfehler
```

This behavior is due to **SAP R/3** being unable to correctly assign a locale and also not being properly configured itself (missing entries in some database tables). To be able to connect to **SAP**, add the following entries to file `DEFAULT.PFL` (see Note 0043288):

```
abap/set_etct_env_at_new_mode = 0
install/collate/active = 0
rscp/TCP0B = TCP0B
```

Restart the **SAP** system. Now you can connect to the system, even though country-specific language settings might not work as expected. After correcting country settings (and providing the correct locales), these entries can be removed from `DEFAULT.PFL` and the **SAP** system can be restarted.

22.7.13.6 ORA-00001

This error only happened with **Oracle 8.1.7** on FreeBSD 4.5. The reason was that the **Oracle** database could not initialize itself properly and crashed, leaving semaphores and shared memory on the system. The next try to start the database then returned ORA-00001.

Find them with `ipcs -a` and remove them with `ipcrm`.

22.7.13.7 ORA-00445 (Background Process PMON Did Not Start)

This error happened with **Oracle 8.1.7**. This error is reported if the database is started with the usual `startsap` script (for example `startsap_majestix_00`) as user `prdadm`.

A possible workaround is to start the database as user `oraprd` instead with `svrmgr1`:

```
% svrmgrl
SVRMGR> connect internal;
SVRMGR> startup;
SVRMGR> exit
```

22.7.13.8 ORA-12546 (Start Listener with Correct Permissions)

Start the **Oracle** listener as user `oraids` with the following commands:

```
# umask 0; lsnrctl start
```

Otherwise you might get ORA-12546 as the sockets will not have the correct permissions. See SAP Note 0072984.

22.7.13.9 ORA-27102 (Out of Memory)

This error happened whilst trying to use values for `MAXDSIZ` and `DFLDSIZ` greater than 1 GB (1024x1024x1024). Additionally, we got `Linux Error 12: Cannot allocate memory`.

22.7.13.10 [DIPGNTAB_IND_IND] during R3SETUP

In general, see SAP Note 0130581 (R3SETUP step DIPGNTAB terminates). During the IDES-specific installation, for some reasons the installation process was not using the proper **SAP** system name “IDS”, but the empty string “ ” instead. This led to some minor problems with accessing directories, as the paths are generated dynamically using `SID` (in this case IDS). So instead of accessing:

```
/usr/sap/IDS/SYS/...
/usr/sap/IDS/DVMGS00
```

the following paths were used:

```
/usr/sap//SYS/...
/usr/sap/D00
```

To continue with the installation, we created a link and an additional directory:

```
# pwd
/compat/linux/usr/sap
# ls -l
total 4
drwxr-xr-x 3 idsadm sapsys 512 May 5 11:20 D00
drwxr-x--x 5 idsadm sapsys 512 May 5 11:35 IDS
lrwxr-xr-x 1 root sapsys 7 May 5 11:35 SYS -> IDS/SYS
drwxrwxr-x 2 idsadm sapsys 512 May 5 13:00 tmp
drwxrwxr-x 11 idsadm sapsys 512 May 4 14:20 trans
```

We also found SAP Notes (0029227 and 0008401) describing this behavior. We did not encounter any of these problems with the **SAP 4.6C** installation.

22.7.13.11 [RFCRSWBOINI_IND_IND] during R3SETUP

During installation of **SAP 4.6C**, this error was just the result of another error happening earlier during installation. In this case, you have to look through the corresponding logfiles and correct the real problem.

If after looking through the logfiles this error is indeed the correct one (check the SAP Notes), you can set `STATUS` of the offending step from `ERROR` to `OK` (file `CENTRDB.R3S`) and restart `R3SETUP`. After installation, you have to execute the report `RSWBOINS` from transaction `SE38`. See SAP Note 0162266 for additional information about phase `RFCRSWBOINI` and `RFCRADDBDIF`.

22.7.13.12 [RFCRADDBDIF_IND_IND] during R3SETUP

Here the same restrictions apply: make sure by looking through the logfiles, that this error is not caused by some previous problems.

If you can confirm that SAP Note 0162266 applies, just set `STATUS` of the offending step from `ERROR` to `OK` (file `CENTRDB.R3S`) and restart `R3SETUP`. After installation, you have to execute the report `RADDBDIF` from transaction `SE38`.

22.7.13.13 sigaction sig31: File size limit exceeded

This error occurred during start of **SAP** processes `disp+work`. If starting **SAP** with the `startsap` script, subprocesses are then started which detach and do the dirty work of starting all other **SAP** processes. As a result, the script itself will not notice if something goes wrong.

To check whether the **SAP** processes did start properly, have a look at the process status with `ps ax | grep SID`, which will give you a list of all **Oracle** and **SAP** processes. If it looks like some processes are missing or if you cannot connect to the **SAP** system, look at the corresponding logfiles which can be found at `/usr/sap/SID/DVEBMGSnr/work/`. The files to look at are `dev_ms` and `dev_disp`.

Signal 31 happens here if the amount of shared memory used by **Oracle** and **SAP** exceed the one defined within the kernel configuration file and could be resolved by using a larger value:

```
# larger value for 46C production systems:
options SHMMAXPGS=393216
# smaller value sufficient for 46B:
#options SHMMAXPGS=262144
```

22.7.13.14 Start of saposcol Failed

There are some problems with the program `saposcol` (version 4.6D). The **SAP** system is using `saposcol` to collect data about the system performance. This program is not needed to use the **SAP** system, so this problem can be considered a minor one. The older versions (4.6B) does work, but does not collect all the data (many calls will just return 0, for example for CPU usage).

22.8 Advanced Topics

If you are curious as to how the Linux binary compatibility works, this is the section you want to read. Most of what follows is based heavily on an email written to FreeBSD chat mailing list (<http://lists.FreeBSD.org/mailman/listinfo/freebsd-chat>) by Terry Lambert <tlambert@primenet.com> (Message ID: <199906020108.SAA07001@usr09.primenet.com>).

Warning: This description applies to FreeBSD, for which it was originally written. This may or may not apply to DragonFly at this point; while FreeBSD 4.x features usually translate over to DragonFly well, your mileage may vary.

22.8.1 How Does It Work?

DragonFly has an abstraction called an “execution class loader”. This is a wedge into the `execve(2)` system call.

What happens is that DragonFly has a list of loaders, instead of a single loader with a fallback to the `#!` loader for running any shell interpreters or shell scripts.

Historically, the only loader on the UNIX platform examined the magic number (generally the first 4 or 8 bytes of the file) to see if it was a binary known to the system, and if so, invoked the binary loader.

If it was not the binary type for the system, the `execve(2)` call returned a failure, and the shell attempted to start executing it as shell commands.

The assumption was a default of “whatever the current shell is”.

Later, a hack was made for `sh(1)` to examine the first two characters, and if they were `:\n`, then it invoked the `csh(1)` shell instead (we believe SCO first made this hack).

What DragonFly does now is go through a list of loaders, with a generic `#!` loader that knows about interpreters as the characters which follow to the next whitespace next to last, followed by a fallback to `/bin/sh`.

For the Linux ABI support, DragonFly sees the magic number as an ELF binary (it makes no distinction between FreeBSD, Solaris, Linux, or any other OS which has an ELF image type, at this point).

The ELF loader looks for a specialized *brand*, which is a comment section in the ELF image, and which is not present on SVR4/Solaris ELF binaries.

For Linux binaries to function, they must be *branded* as type `Linux` from `brandelf(1)`:

```
# brandelf -t Linux file
```

When this is done, the ELF loader will see the `Linux` brand on the file.

When the ELF loader sees the `Linux` brand, the loader replaces a pointer in the `proc` structure. All system calls are indexed through this pointer (in a traditional UNIX system, this would be the `sysent[]` structure array, containing the system calls). In addition, the process is flagged for special handling of the trap vector for the signal trampoline code, and several other (minor) fix-ups that are handled by the Linux kernel module.

The Linux system call vector contains, among other things, a list of `sysent[]` entries whose addresses reside in the kernel module.

When a system call is called by the Linux binary, the trap code dereferences the system call function pointer off the `proc` structure, and gets the Linux, not the DragonFly, system call entry points.

In addition, the Linux mode dynamically *reroots* lookups; this is, in effect, what the `union` option to file system mounts (*not* the `unionfs` file system type!) does. First, an attempt is made to lookup the file in the `/compat/linux/original-path` directory, *then* only if that fails, the lookup is done in the `/original-path` directory. This makes sure that binaries that require other binaries can run (e.g., the Linux toolchain can all run under Linux ABI support). It also means that the Linux binaries can load and execute DragonFly binaries, if there are no corresponding Linux binaries present, and that you could place a `uname(1)` command in the `/compat/linux` directory tree to ensure that the Linux binaries could not tell they were not running on Linux.

In effect, there is a Linux kernel in the DragonFly kernel; the various underlying functions that implement all of the services provided by the kernel are identical to both the DragonFly system call table entries, and the Linux system call table entries: file system operations, virtual memory operations, signal delivery, System V IPC, etc. . . . The only difference is that DragonFly binaries get the DragonFly *glue* functions, and Linux binaries get the Linux *glue* functions (most older OS's only had their own *glue* functions: addresses of functions in a static global `sysent[]` structure array, instead of addresses of functions dereferenced off a dynamically initialized pointer in the `proc` structure of the process making the call).

Which one is the native DragonFly ABI? It does not matter. Basically the only difference is that (currently; this could easily be changed in a future release, and probably will be after this) the DragonFly *glue* functions are statically linked into the kernel, and the Linux *glue* functions can be statically linked, or they can be accessed via a kernel module.

Yeah, but is this really emulation? No. It is an ABI implementation, not an emulation. There is no emulator (or simulator, to cut off the next question) involved.

So why is it sometimes called “Linux emulation”? To make it hard to sell DragonFly! Really, it is because the historical implementation was done at a time when there was really no word other than that to describe what was going on; saying that FreeBSD¹ ran Linux binaries was not true, if you did not compile the code in or load a module, and there needed to be a word to describe what was being loaded—hence “the Linux emulator”.

Notes

1. FreeBSD's original Linux compatibility code was committed in June 1995. It fulfilled milestone number one: running DOOM.

III. Appendices

Appendix A. Obtaining DragonFly

A.1 CDRom and DVD Publishers

A.1.1 Retail Products

DragonFly is available as a purchasable CD:

- Crescent Anchor
WWW: <http://www.crescentanchor.com/>

A.1.2 CD and DVD Sets

DragonFly BSD CD and DVD sets are available from online retailers:

- Daemon News Mall
560 South State Street, Suite A2
Orem, UT 84058
USA
Phone: +1 800 407-5170
Fax: +1 1 801 765-0877
Email: <sales@bsdmail.com>
WWW: <http://www.bsdmail.com/>
- BSD-Systems
Email: <sales@bsd-systems.co.uk>
WWW: <http://www.bsd-systems.co.uk/>

A.2 FTP Sites

The official sources for DragonFly are available via anonymous FTP from a worldwide set of mirror sites. The site <ftp://ftp.dragonflybsd.org/> is well connected and allows a large number of connections to it, but you are probably better off finding a “closer” mirror site (especially if you decide to set up some sort of mirror site).

The DragonFly mirror sites list (<http://www.dragonflybsd.org/main/download.cgi>) is the best, most up-to-date source.

Additionally, DragonFly is available via anonymous FTP from the following mirror sites. If you choose to obtain DragonFly via anonymous FTP, please try to use a site near you. The mirror sites listed as “Primary Mirror Sites” typically have the entire DragonFly archive (all the currently available versions for each of the architectures) but you will probably have faster download times from a site that is in your country or region. The regional sites carry the most recent versions for the most popular architecture(s) but might not carry the entire DragonFly archive. All sites provide access via anonymous FTP but some sites also provide access via other methods. The access methods available for each site are provided in parentheses after the hostname.

Central Servers, Primary Mirror Sites, Australia, USA.

(as of 2005/06/27 23:37:47 UTC)

Central Servers

- <ftp://ftp.DragonFlyBSD.org/> (ftp)

Primary Mirror Sites

In case of problems, please contact <docs@DragonFlyBSD.org>.

- <ftp://chlamydia.fs.ei.tum.de/pub/DragonFly/> (ftp)
- <ftp://ftp.allbsd.org/pub/DragonFly/> (ftp)
- <ftp://ftp.esat.net/mirrors/chlamydia.fs.ei.tum.de/pub/DragonFly/> (ftp)
- <ftp://ftp.fortunaty.net/DragonFly/> (ftp)
- <ftp://ftp.univie.ac.at/systems/DragonFly/> (ftp)
- <ftp://cvsup.math.uic.edu/dragonflybsd/> (ftp)
- <ftp://ftp.starkast.net/pub/DragonFly/> (ftp)
- <ftp://mirror.bgp4.net/pub/DragonFly/> (ftp)
- <ftp://ftp.csie.chu.edu.tw/> (ftp)
- <ftp://dragonflybsd.cs.pu.edu.tw/DragonFLYBSD> (ftp)
- <ftp://mirror.isp.net.au/pub/DragonFly/> (ftp)
- <ftp://ftp.theshell.com/pub/DragonFly/> (ftp)
- <ftp://mirror.macomnet.net/pub/DragonFlyBSD/> (ftp)
- <ftp://ftp.tu-clausthal.de/pub/DragonFly/> (ftp)

Australia

In case of problems, please contact <docs@DragonFlyBSD.org>.

- <ftp://mirror.isp.net.au/pub/DragonFly/> (ftp / [http \(http://mirror.isp.net.au/ftp/pub/DragonFly/\)](http://mirror.isp.net.au/ftp/pub/DragonFly/) / rsync)

USA

In case of problems, please contact <docs@DragonFlyBSD.org>.

- <ftp://ftp.theshell.com/pub/DragonFly/> (ftp / [http \(http://theshell.com/pub/DragonFly/\)](http://theshell.com/pub/DragonFly/) / rsync)

A.3 Using CVSup

A.3.1 Introduction

CVSup is a software package for distributing and updating source trees from a master CVS repository on a remote server host. The DragonFly sources are maintained in a CVS repository on a central development machine in California. With **CVSup**, DragonFly users can easily keep their own source trees up to date.

CVSup uses the so-called *pull* model of updating. Under the pull model, each client asks the server for updates, if and when they are wanted. The server waits passively for update requests from its clients. Thus all updates are instigated by the client. The server never sends unsolicited updates. Users must either run the **CVSup** client manually to get an update, or they must set up a `cron` job to run it automatically on a regular basis.

The term **CVSup**, capitalized just so, refers to the entire software package. Its main components are the client `cvsup` which runs on each user's machine, and the server `cvsupd` which runs at each of the DragonFly mirror sites that use **CVSup**.

A.3.2 Installation

CVSup is installed by default on all DragonFly systems.

A.3.3 CVSup Configuration

CVSup's operation is controlled by a configuration file called the `supfile`. There are some sample `supfiles` in the directory `/usr/share/examples/cvsup/`.

The information in a `supfile` answers the following questions for **CVSup**:

- Which files do you want to receive?
- Which versions of them do you want?
- Where do you want to get them from?
- Where do you want to put them on your own machine?
- Where do you want to put your status files?

In the following sections, we will construct a typical `supfile` by answering each of these questions in turn. First, we describe the overall structure of a `supfile`.

A `supfile` is a text file. Comments begin with `#` and extend to the end of the line. Lines that are blank and lines that contain only comments are ignored.

Each remaining line describes a set of files that the user wishes to receive. The line begins with the name of a "collection", a logical grouping of files defined by the server. The name of the collection tells the server which files you want. After the collection name come zero or more fields, separated by white space. These fields answer the questions listed above. There are two types of fields: flag fields and value fields. A flag field consists of a keyword standing alone, e.g., `delete` or `compress`. A value field also begins with a keyword, but the keyword is followed without intervening white space by `=` and a second word. For example, `release=cvs` is a value field.

A `supfile` typically specifies more than one collection to receive. One way to structure a `supfile` is to specify all of the relevant fields explicitly for each collection. However, that tends to make the `supfile` lines quite long, and it

is inconvenient because most fields are the same for all of the collections in a `supfile`. **CVSup** provides a defaulting mechanism to avoid these problems. Lines beginning with the special pseudo-collection name `*default` can be used to set flags and values which will be used as defaults for the subsequent collections in the `supfile`. A default value can be overridden for an individual collection, by specifying a different value with the collection itself. Defaults can also be changed or augmented in mid-supfile by additional `*default` lines.

With this background, we will now proceed to construct a `supfile` for receiving and updating the main source tree of DragonFly.

- Which files do you want to receive?

The files available via **CVSup** are organized into named groups called “collections”. The collections that are available are described in the following section. In this example, we wish to receive the entire main source tree for the DragonFly system. There is a single large collection `cvs-src` which will give us all of that. As a first step toward constructing our `supfile`, we simply list the collections, one per line (in this case, only one line):

```
cvs-src
```

- Which version(s) of them do you want?

With **CVSup**, you can receive virtually any version of the sources that ever existed. That is possible because the **cvsupd** server works directly from the CVS repository, which contains all of the versions. You specify which one of them you want using the `tag=` and `date=` value fields.

Warning: Be very careful to specify any `tag=` fields correctly. Some tags are valid only for certain collections of files. If you specify an incorrect or misspelled tag, **CVSup** will delete files which you probably do not want deleted. In particular, use *only* `tag=.` for the `ports-*` collections.

The `tag=` field names a symbolic tag in the repository. There are two kinds of tags, revision tags and branch tags. A revision tag refers to a specific revision. Its meaning stays the same from day to day. A branch tag, on the other hand, refers to the latest revision on a given line of development, at any given time. Because a branch tag does not refer to a specific revision, it may mean something different tomorrow than it means today.

Section A.4 contains branch tags that users might be interested in. When specifying a tag in **CVSup**'s configuration file, it must be preceded with `tag=` (`RELENG_4` will become `tag=RELENG_4`). Keep in mind that only the `tag=.` is relevant for the `ports` collection.

Warning: Be very careful to type the tag name exactly as shown. **CVSup** cannot distinguish between valid and invalid tags. If you misspell the tag, **CVSup** will behave as though you had specified a valid tag which happens to refer to no files at all. It will delete your existing sources in that case.

When you specify a branch tag, you normally receive the latest versions of the files on that line of development. If you wish to receive some past version, you can do so by specifying a date with the `date=` value field. The `cvsup(1)` manual page explains how to do that.

For our example, we wish to receive the current release of DragonFly. We add this line at the beginning of our `supfile`:

```
*default tag=.
```

There is an important special case that comes into play if you specify neither a `tag=` field nor a `date=` field. In that case, you receive the actual RCS files directly from the server's CVS repository, rather than receiving a particular version. Developers generally prefer this mode of operation. By maintaining a copy of the repository itself on their systems, they gain the ability to browse the revision histories and examine past versions of files. This gain is achieved at a large cost in terms of disk space, however.

- Where do you want to get them from?

We use the `host=` field to tell `cvsup` where to obtain its updates. Any of the CVSup mirror sites will do, though you should try to select one that is close to you in cyberspace. In this example we will use a fictional DragonFly distribution site, `cvsup666.dragonflybsd.org`:

```
*default host=cvsup666.dragonflybsd.org
```

You will need to change the host to one that actually exists before running **CVSup**. On any particular run of `cvsup`, you can override the host setting on the command line, with `-h hostname`.

- Where do you want to put them on your own machine?

The `prefix=` field tells `cvsup` where to put the files it receives. In this example, we will put the source files directly into our main source tree, `/usr/src`. The `src` directory is already implicit in the collections we have chosen to receive, so this is the correct specification:

```
*default prefix=/usr
```

- Where should `cvsup` maintain its status files?

The **CVSup** client maintains certain status files in what is called the “base” directory. These files help **CVSup** to work more efficiently, by keeping track of which updates you have already received. We will use the standard base directory, `/usr/local/etc/cvsup`:

```
*default base=/usr/local/etc/cvsup
```

This setting is used by default if it is not specified in the `supfile`, so we actually do not need the above line.

If your base directory does not already exist, now would be a good time to create it. The `cvsup` client will refuse to run if the base directory does not exist.

- Miscellaneous `supfile` settings:

There is one more line of boiler plate that normally needs to be present in the `supfile`:

```
*default release=cvs delete use-rel-suffix compress
```

`release=cvs` indicates that the server should get its information out of the main DragonFly CVS repository. This is virtually always the case, but there are other possibilities which are beyond the scope of this discussion.

`delete` gives **CVSup** permission to delete files. You should always specify this, so that **CVSup** can keep your source tree fully up-to-date. **CVSup** is careful to delete only those files for which it is responsible. Any extra files you happen to have will be left strictly alone.

`use-rel-suffix` is ... arcane. If you really want to know about it, see the `cvsup(1)` manual page. Otherwise, just specify it and do not worry about it.

`compress` enables the use of `gzip`-style compression on the communication channel. If your network link is T1 speed or faster, you probably should not use compression. Otherwise, it helps substantially.

- Putting it all together:

Here is the entire `supfile` for our example:

```
*default tag=.
```

```
*default host=cvsup666.dragonflybsd.org
*default prefix=/usr
*default base=/usr/local/etc/cvsup
*default release=cvs delete use-rel-suffix compress

src-all
```

A.3.3.1 The refuse File

As mentioned above, **CVSup** uses a *pull method*. Basically, this means that you connect to the **CVSup** server, and it says, “Here is what you can download from me...”, and your client responds “OK, I will take this, this, this, and this.” In the default configuration, the **CVSup** client will take every file associated with the collection and tag you chose in the configuration file. However, this is not always what you want, especially if you are syncing the `doc`, `ports`, or `www` trees — most people cannot read four or five languages, and therefore they do not need to download the language-specific files. If you are **CVSup**ing the `ports` collection, you can get around this by specifying each collection individually (e.g., *ports-astrology*, *ports-biology*, etc instead of simply saying *ports-all*). However, since the `doc` and `www` trees do not have language-specific collections, you must use one of **CVSup**’s many nifty features: the `refuse` file.

The `refuse` file essentially tells **CVSup** that it should not take every single file from a collection; in other words, it tells the client to *refuse* certain files from the server. The `refuse` file can be found (or, if you do not yet have one, should be placed) in `base/sup/`. `base` is defined in your `supfile`; by default, `base` is `/usr/local/etc/cvsup`, which means that by default the `refuse` file is `/usr/local/etc/cvsup/sup/refuse`.

The `refuse` file has a very simple format; it simply contains the names of files or directories that you do not wish to download. For example, if you cannot speak any languages other than English and some German, and you do not feel the need to use the German applications (or applications for any other languages, except for English), you can put the following in your `refuse` file:

```
ports/chinese
ports/french
ports/german
ports/hebrew
ports/hungarian
ports/japanese
ports/korean
ports/polish
ports/portuguese
ports/russian
ports/ukrainian
ports/vietnamese
doc/da_*
doc/de_*
doc/el_*
doc/es_*
doc/fr_*
doc/it_*
doc/ja_*
doc/nl_*
doc/no_*
doc/pl_*
doc/pt_*
```

```
doc/ru_*
doc/sr_*
doc/zh_*
```

and so forth for the other languages (you can find the full list by browsing the FreeBSD CVS repository (<http://www.FreeBSD.org/cgi/cvsweb.cgi/>)).

With this very useful feature, those users who are on slow links or pay by the minute for their Internet connection will be able to save valuable time as they will no longer need to download files that they will never use. For more information on `refuse` files and other neat features of **CVSup**, please view its manual page.

A.3.4 Running CVSup

You are now ready to try an update. The command line for doing this is quite simple:

```
# cvsup supfile
```

where `supfile` is of course the name of the `supfile` you have just created. Assuming you are running under X11, `cvsup` will display a GUI window with some buttons to do the usual things. Press the `go` button, and watch it run.

Since you are updating your actual `/usr/src` tree in this example, you will need to run the program as `root` so that `cvsup` has the permissions it needs to update your files. Having just created your configuration file, and having never used this program before, that might understandably make you nervous. There is an easy way to do a trial run without touching your precious files. Just create an empty directory somewhere convenient, and name it as an extra argument on the command line:

```
# mkdir /var/tmp/dest
# cvsup supfile /var/tmp/dest
```

The directory you specify will be used as the destination directory for all file updates. **CVSup** will examine your usual files in `/usr/src`, but it will not modify or delete any of them. Any file updates will instead land in `/var/tmp/dest/usr/src`. **CVSup** will also leave its base directory status files untouched when run this way. The new versions of those files will be written into the specified directory. As long as you have read access to `/usr/src`, you do not even need to be `root` to perform this kind of trial run.

If you are not running X11 or if you just do not like GUIs, you should add a couple of options to the command line when you run `cvsup`:

```
# cvsup -g -L 2 supfile
```

The `-g` tells **CVSup** not to use its GUI. This is automatic if you are not running X11, but otherwise you have to specify it.

The `-L 2` tells **CVSup** to print out the details of all the file updates it is doing. There are three levels of verbosity, from `-L 0` to `-L 2`. The default is 0, which means total silence except for error messages.

There are plenty of other options available. For a brief list of them, type `cvsup -h`. For more detailed descriptions, see the manual page.

Once you are satisfied with the way updates are working, you can arrange for regular runs of **CVSup** using `cron(8)`. Obviously, you should not let **CVSup** use its GUI when running it from `cron(8)`.

A.3.5 CVSup File Collections

The most commonly used collections are `cvs-src`, and `cvs-dfports`.

`cvs-src`

The DragonFly source code.

`cvs-doc`

Documentation. This does not include the DragonFly website.

`cvs-dfports`

Overrides for the FreeBSD Ports Collection.

`cvs-site`

The DragonFly website code.

`cvs-root`

Basic CVS data. This is only needed if you are pulling the RCS data.

A.3.6 For More Information

For the **CVSup** FAQ and other information about **CVSup**, see The CVSup Home Page (<http://www.polstra.com/projects/freeware/CVSup/>).

Questions and bug reports should be addressed to the author of the program at [<cvsup-bugs@polstra.com>](mailto:cvsup-bugs@polstra.com).

A.3.7 CVSup Sites

CVSup servers for DragonFly are running at the following sites:

Primary Mirror Sites, Australia, USA.

(as of 2005/06/27 23:37:47 UTC)

Primary Mirror Sites

- chlamydia.fs.ei.tum.de
- cvsup.allbsd.org
- grappa.unix-ag.uni-kl.de
- mirror.isp.net.au
- alxl.info
- dragonflybsd.delphij.net
- fred.acm.cs.rpi.edu

Australia

- mirror.isp.net.au

USA

- fred.acm.cs.rpi.edu

A.4 CVS Tags

When obtaining or updating sources from **cv**s and **CVS**up a revision tag (reference to a date in time) must be specified.

A revision tag refers to either a particular line of DragonFly development, or a specific point in time. The first type are called “branch tags”, the second type are called “release tags”.

A.4.1 Branch Tags

The DragonFly tree has no branch tags at the current time.

A.4.2 Release Tags

These tags correspond to the DragonFly `src/` tree at a specific point in time, when a particular version was released.

HEAD

The latest bleeding-edge DragonFly code. Be warned that this is considered unstable and possibly may not build or compile at any time.

DragonFly_Preview

A "preview" of the latest bleeding-edge DragonFly code. The main purpose of the Preview tag is to support those users who want a fairly recent snapshot at a "reasonably stable" point in development. Under normal conditions, one should consider syncing Preview at least once a month.

DragonFly_RELEASE_1_2

DragonFly 1.2

DragonFly_1_0_REL

DragonFly 1.0

Appendix B. Bibliography

While the manual pages provide the definitive reference for individual pieces of the DragonFly operating system, they are notorious for not illustrating how to put the pieces together to make the whole operating system run smoothly. For this, there is no substitute for a good book on UNIX system administration and a good users' manual.

B.1 Books & Magazines Specific to BSD

International books & Magazines:

- Using FreeBSD (<http://jdli.tw.FreeBSD.org/publication/book/freebsd2/index.htm>) (in Chinese).
- FreeBSD for PC 98'ers (in Japanese), published by SHUWA System Co, LTD. ISBN 4-87966-468-5 C3055 P2900E.
- FreeBSD (in Japanese), published by CUTT. ISBN 4-906391-22-2 C3055 P2400E.
- Complete Introduction to FreeBSD (<http://www.shoeisha.com/book/Detail.asp?bid=650>) (in Japanese), published by Shoeisha Co., Ltd (<http://www.shoeisha.co.jp/>). ISBN 4-88135-473-6 P3600E.
- Personal UNIX Starter Kit FreeBSD (<http://www.ascii.co.jp/pb/book1/shinkan/detail/1322785.html>) (in Japanese), published by ASCII (<http://www.ascii.co.jp/>). ISBN 4-7561-1733-3 P3000E.
- FreeBSD Handbook (Japanese translation), published by ASCII (<http://www.ascii.co.jp/>). ISBN 4-7561-1580-2 P3800E.
- FreeBSD mit Methode (in German), published by Computer und Literatur Verlag (<http://www.cul.de/>)Vertrieb Hanser, 1998. ISBN 3-932311-31-0.
- FreeBSD 4 - Installieren, Konfigurieren, Administrieren (<http://www.cul.de/freebsd.html>) (in German), published by Computer und Literatur Verlag (<http://www.cul.de/>), 2001. ISBN 3-932311-88-4.
- FreeBSD 5 - Installieren, Konfigurieren, Administrieren (<http://www.cul.de/freebsd.html>) (in German), published by Computer und Literatur Verlag (<http://www.cul.de/>), 2003. ISBN 3-936546-06-1.
- FreeBSD de Luxe (<http://www.mitp.de/vmi/mitp/detail/pWert/1343/>) (in German), published by Verlag Modere Industrie (<http://www.mitp.de/>), 2003. ISBN 3-8266-1343-0.
- FreeBSD Install and Utilization Manual (<http://www.pc.mycom.co.jp/FreeBSD/install-manual.html>) (in Japanese), published by Mainichi Communications Inc. (<http://www.pc.mycom.co.jp/>).
- Onno W Purbo, Dodi Maryanto, Syahrial Hubbany, Widjil Widodo *Building Internet Server with FreeBSD* (<http://maxwell.itb.ac.id/>) (in Indonesia Language), published by Elex Media Komputindo (<http://www.elexmedia.co.id/>).

English language books & Magazines:

- Absolute BSD: The Ultimate Guide to FreeBSD (<http://www.AbsoluteBSD.com/>), published by No Starch Press (<http://www.nostarch.com/>), 2002. ISBN: 1886411743
- The Complete FreeBSD (<http://www.freebsdmail.com/cgi-bin/fm/bsdcomp>), published by O'Reilly (<http://www.oreilly.com/>), 2003. ISBN: 0596005164
- The FreeBSD Corporate Networker's Guide (<http://www.freebsd-corp-net-guide.com/>), published by Addison-Wesley (<http://www.awl.com/awl/>), 2000. ISBN: 0201704811

- FreeBSD: An Open-Source Operating System for Your Personal Computer (<http://andrsn.stanford.edu/FreeBSD/introbook/>), published by The Bit Tree Press, 2001. ISBN: 0971204500
- Teach Yourself FreeBSD in 24 Hours, published by Sams (<http://www.sampublishing.com/>), 2002. ISBN: 0672324245
- FreeBSD unleashed, published by Sams (<http://www.sampublishing.com/>), 2002. ISBN: 0672324563
- FreeBSD: The Complete Reference, published by McGrawHill (<http://books.mcgraw-hill.com/>), 2003. ISBN: 0072224096

B.2 Users' Guides

- Computer Systems Research Group, UC Berkeley. *4.4BSD User's Reference Manual*. O'Reilly & Associates, Inc., 1994. ISBN 1-56592-075-9
- Computer Systems Research Group, UC Berkeley. *4.4BSD User's Supplementary Documents*. O'Reilly & Associates, Inc., 1994. ISBN 1-56592-076-7
- *UNIX in a Nutshell*. O'Reilly & Associates, Inc., 1990. ISBN 093717520X
- Mui, Linda. *What You Need To Know When You Can't Find Your UNIX System Administrator*. O'Reilly & Associates, Inc., 1995. ISBN 1-56592-104-6
- Ohio State University (<http://www-wks.acs.ohio-state.edu/>) has written a UNIX Introductory Course (http://www-wks.acs.ohio-state.edu/unix_course/unix.html) which is available online in HTML and PostScript format.

An Italian translation ([../..../it_IT.ISO8859-15/books/unix-introduction/index.html](http://www-wks.acs.ohio-state.edu/~it_IT.ISO8859-15/books/unix-introduction/index.html)) of this document is available as part of the FreeBSD Italian Documentation Project.

- Jpman Project, Japan FreeBSD Users Group (<http://www.jp.FreeBSD.org/>). FreeBSD User's Reference Manual (<http://www.pc.mycom.co.jp/FreeBSD/urm.html>) (Japanese translation). Mainichi Communications Inc. (<http://www.pc.mycom.co.jp/>), 1998. ISBN4-8399-0088-4 P3800E.
- Edinburgh University (<http://www.ed.ac.uk/>) has written an Online Guide (<http://unixhelp.ed.ac.uk/>) for newcomers to the UNIX environment.

B.3 Administrators' Guides

- Albitz, Paul and Liu, Cricket. *DNS and BIND*, 4th Ed. O'Reilly & Associates, Inc., 2001. ISBN 1-59600-158-4
- Computer Systems Research Group, UC Berkeley. *4.4BSD System Manager's Manual*. O'Reilly & Associates, Inc., 1994. ISBN 1-56592-080-5
- Costales, Brian, et al. *Sendmail*, 2nd Ed. O'Reilly & Associates, Inc., 1997. ISBN 1-56592-222-0
- Frisch, Eelen. *Essential System Administration*, 2nd Ed. O'Reilly & Associates, Inc., 1995. ISBN 1-56592-127-5
- Hunt, Craig. *TCP/IP Network Administration*, 2nd Ed. O'Reilly & Associates, Inc., 1997. ISBN 1-56592-322-7
- Nemeth, Evi. *UNIX System Administration Handbook*. 3rd Ed. Prentice Hall, 2000. ISBN 0-13-020601-6

- Stern, Hal *Managing NFS and NIS* O'Reilly & Associates, Inc., 1991. ISBN 0-937175-75-7
- Jpman Project, Japan FreeBSD Users Group (<http://www.jp.FreeBSD.org/>). FreeBSD System Administrator's Manual (<http://www.pc.mycom.co.jp/FreeBSD/sam.html>) (Japanese translation). Mainichi Communications Inc. (<http://www.pc.mycom.co.jp/>), 1998. ISBN4-8399-0109-0 P3300E.
- Dreyfus, Emmanuel. Cahiers de l'Admin: BSD (<http://www.editions-eyrolles.com/php.informatique/Ouvrages/9782212112443.php3>) (in French), Eyrolles, 2003. ISBN 2-212-11244-0

B.4 Programmers' Guides

- Asente, Paul, Converse, Diana, and Swick, Ralph. *X Window System Toolkit*. Digital Press, 1998. ISBN 1-55558-178-1
- Computer Systems Research Group, UC Berkeley. *4.BSD Programmer's Reference Manual*. O'Reilly & Associates, Inc., 1994. ISBN 1-56592-078-3
- Computer Systems Research Group, UC Berkeley. *4.BSD Programmer's Supplementary Documents*. O'Reilly & Associates, Inc., 1994. ISBN 1-56592-079-1
- Harbison, Samuel P. and Steele, Guy L. Jr. *C: A Reference Manual*. 4rd ed. Prentice Hall, 1995. ISBN 0-13-326224-3
- Kernighan, Brian and Dennis M. Ritchie. *The C Programming Language*. PTR Prentice Hall, 1988. ISBN 0-13-110362-9
- Lehey, Greg. *Porting UNIX Software*. O'Reilly & Associates, Inc., 1995. ISBN 1-56592-126-7
- Plauger, P. J. *The Standard C Library*. Prentice Hall, 1992. ISBN 0-13-131509-9
- Spinellis, Diomidis. *Code Reading: The Open Source Perspective* (<http://www.spinellis.gr/codereading/>). Addison-Wesley, 2003. ISBN 0-201-79940-5
- Stevens, W. Richard. *Advanced Programming in the UNIX Environment*. Reading, Mass. : Addison-Wesley, 1992. ISBN 0-201-56317-7
- Stevens, W. Richard. *UNIX Network Programming*. 2nd Ed, PTR Prentice Hall, 1998. ISBN 0-13-490012-X
- Wells, Bill. "Writing Serial Drivers for UNIX". *Dr. Dobb's Journal*. 19(15), December 1994. pp68-71, 97-99.

B.5 Operating System Internals

- Andleigh, Prabhat K. *UNIX System Architecture*. Prentice-Hall, Inc., 1990. ISBN 0-13-949843-5
- Jolitz, William. "Porting UNIX to the 386". *Dr. Dobb's Journal*. January 1991-July 1992.
- Leffler, Samuel J., Marshall Kirk McKusick, Michael J Karels and John Quarterman *The Design and Implementation of the 4.3BSD UNIX Operating System*. Reading, Mass. : Addison-Wesley, 1989. ISBN 0-201-06196-1
- Leffler, Samuel J., Marshall Kirk McKusick, *The Design and Implementation of the 4.3BSD UNIX Operating System: Answer Book*. Reading, Mass. : Addison-Wesley, 1991. ISBN 0-201-54629-9

- McKusick, Marshall Kirk, Keith Bostic, Michael J Karels, and John Quarterman. *The Design and Implementation of the 4.4BSD Operating System*. Reading, Mass. : Addison-Wesley, 1996. ISBN 0-201-54979-4
(Chapter 2 of this book is available online as part of the FreeBSD Documentation Project, and chapter 9 here (http://www.netapp.com/tech_library/nfsbook.print.)
- Stevens, W. Richard. *TCP/IP Illustrated, Volume 1: The Protocols*. Reading, Mass. : Addison-Wesley, 1996. ISBN 0-201-63346-9
- Schimmel, Curt. *Unix Systems for Modern Architectures*. Reading, Mass. : Addison-Wesley, 1994. ISBN 0-201-63338-8
- Stevens, W. Richard. *TCP/IP Illustrated, Volume 3: TCP for Transactions, HTTP, NNTP and the UNIX Domain Protocols*. Reading, Mass. : Addison-Wesley, 1996. ISBN 0-201-63495-3
- Vahalia, Uresh. *UNIX Internals -- The New Frontiers*. Prentice Hall, 1996. ISBN 0-13-101908-2
- Wright, Gary R. and W. Richard Stevens. *TCP/IP Illustrated, Volume 2: The Implementation*. Reading, Mass. : Addison-Wesley, 1995. ISBN 0-201-63354-X

B.6 Security Reference

- Cheswick, William R. and Steven M. Bellovin. *Firewalls and Internet Security: Repelling the Wily Hacker*. Reading, Mass. : Addison-Wesley, 1995. ISBN 0-201-63357-4
- Garfinkel, Simson and Gene Spafford. *Practical UNIX & Internet Security*. 2nd Ed. O'Reilly & Associates, Inc., 1996. ISBN 1-56592-148-8
- Garfinkel, Simson. *PGP Pretty Good Privacy* O'Reilly & Associates, Inc., 1995. ISBN 1-56592-098-8

B.7 Hardware Reference

- Anderson, Don and Tom Shanley. *Pentium Processor System Architecture*. 2nd Ed. Reading, Mass. : Addison-Wesley, 1995. ISBN 0-201-40992-5
- Ferraro, Richard F. *Programmer's Guide to the EGA, VGA, and Super VGA Cards*. 3rd ed. Reading, Mass. : Addison-Wesley, 1995. ISBN 0-201-62490-7
- Intel Corporation publishes documentation on their CPUs, chipsets and standards on their developer web site (<http://developer.intel.com/>), usually as PDF files.
- Shanley, Tom. *80486 System Architecture*. 3rd ed. Reading, Mass. : Addison-Wesley, 1995. ISBN 0-201-40994-1
- Shanley, Tom. *ISA System Architecture*. 3rd ed. Reading, Mass. : Addison-Wesley, 1995. ISBN 0-201-40996-8
- Shanley, Tom. *PCI System Architecture*. 4th ed. Reading, Mass. : Addison-Wesley, 1999. ISBN 0-201-30974-2
- Van Gilluwe, Frank. *The Undocumented PC*, 2nd Ed. Reading, Mass: Addison-Wesley Pub. Co., 1996. ISBN 0-201-47950-8
- Messmer, Hans-Peter. *The Indispensable PC Hardware Book*, 4th Ed. Reading, Mass: Addison-Wesley Pub. Co., 2002. ISBN 0-201-59616-4

B.8 UNIX History

- Lion, John *Lion's Commentary on UNIX, 6th Ed. With Source Code*. ITP Media Group, 1996. ISBN 1573980137
- Raymond, Eric S. *The New Hacker's Dictionary, 3rd edition*. MIT Press, 1996. ISBN 0-262-68092-0. Also known as the Jargon File (<http://www.catb.org/~esr/jargon/html/index.html>)
- Salus, Peter H. *A Quarter Century of UNIX*. Addison-Wesley Publishing Company, Inc., 1994. ISBN 0-201-54777-5
- Simon Garfinkel, Daniel Weise, Steven Strassmann. *The UNIX-HATERS Handbook*. IDG Books Worldwide, Inc., 1994. ISBN 1-56884-203-1
- Don Libes, Sandy Ressler *Life with UNIX* — special edition. Prentice-Hall, Inc., 1989. ISBN 0-13-536657-7
- *The BSD family tree*. <http://www.FreeBSD.org/cgi/cvsweb.cgi/src/share/misc/bsd-family-tree> or [/usr/share/misc/bsd-family-tree](http://usr/share/misc/bsd-family-tree) on a modern FreeBSD machine.
- *The BSD Release Announcements collection*. 1997. <http://www.de.FreeBSD.org/de/ftp/releases/>
- *Networked Computer Science Technical Reports Library*. <http://www.ncstrl.org/>
- *Old BSD releases from the Computer Systems Research group (CSRG)*. <http://www.mckusick.com/csrg/>: The 4CD set covers all BSD versions from 1BSD to 4.4BSD and 4.4BSD-Lite2 (but not 2.11BSD, unfortunately). As well, the last disk holds the final sources plus the SCCS files.

B.9 Magazines and Journals

- *The C/C++ Users Journal*. R&D Publications Inc. ISSN 1075-2838
- *Sys Admin — The Journal for UNIX System Administrators* Miller Freeman, Inc., ISSN 1061-2688
- *freeX — Das Magazin für Linux - BSD - UNIX* (in German) Computer- und Literaturverlag GmbH, ISSN 1436-7033

Appendix C. Resources on the Internet

The rapid pace of DragonFly progress makes print media impractical as a means of following the latest developments. Electronic resources are the best, if not often the only, way stay informed of the latest advances. Since DragonFly is a volunteer effort, the user community itself also generally serves as a “technical support department” of sorts, with electronic mail and USENET news being the most effective way of reaching that community.

The most important points of contact with the DragonFly user community are outlined below. If you are aware of other resources not mentioned here, please send them to the DragonFly Documentation project mailing list (<http://leaf.dragonflybsd.org/mailarchive/>) so that they may also be included.

C.1 Mailing Lists

Though many of the DragonFly development members read USENET, we cannot always guarantee that we will get to your questions in a timely fashion (or at all) if you post them only to one of the `comp.unix.bsd.*` groups. By addressing your questions to the appropriate mailing list you will reach both us and a concentrated DragonFly audience, invariably assuring a better (or at least faster) response.

The charters for the various lists are given at the bottom of this document. *Please read the charter before joining or sending mail to any list.* Most of our list subscribers now receive many hundreds of DragonFly related messages every day, and by setting down charters and rules for proper use we are striving to keep the signal-to-noise ratio of the lists high. To do less would see the mailing lists ultimately fail as an effective communications medium for the project.

Archives are kept for all of the mailing lists and can be searched using the mail archives (<http://leaf.dragonflybsd.org/mailarchive/>). The keyword searchable archive offers an excellent way of finding answers to frequently asked questions and should be consulted before posting a question.

C.1.1 List Summary

General lists: The following are general lists which anyone is free (and encouraged) to join:

List	Purpose
bugs (http://leaf.dragonflybsd.org/mailarchive/)	Bug reports
commits (http://leaf.dragonflybsd.org/mailarchive/)	Messages generated by code changes to DragonFly source, documentation, or the website.
docs (http://leaf.dragonflybsd.org/mailarchive/)	Discussion of DragonFly documentation
kernel (http://leaf.dragonflybsd.org/mailarchive/)	Ostensibly for discussion of kernel work, though this list also serves as a catch-all for any topic pertaining to DragonFly.
submit (http://leaf.dragonflybsd.org/mailarchive/)	Submission and discussion of new code or ideas for DragonFly.
Test (http://leaf.dragonflybsd.org/mailarchive/)	For testing your newsreader or mail software.
users (http://leaf.dragonflybsd.org/mailarchive/)	User related discussion about DragonFly.

C.1.2 How to Subscribe

To subscribe to a list, click on the list name above or send an email to <listname-request@dragonflybsd.org> and put 'subscribe' in the body of the message.

To actually post to a given list you simply send mail to <listname@dragonflybsd.org>. It will then be redistributed to mailing list members world-wide.

To unsubscribe yourself from a list, email <listname-request@dragonflybsd.org> and put 'unsubscribe' in the body of the message.

C.1.3 List Charters

All DragonFly mailing lists have certain basic rules which must be adhered to by anyone using them. Failure to comply with these guidelines may result in from all DragonFly mailing lists and filtered from further posting to them. We regret that such rules and measures are necessary at all, but today's Internet is a pretty harsh environment, it would seem, and many fail to appreciate just how fragile some of its mechanisms are.

Rules of the road:

- The topic of any posting should adhere to the basic charter of the list it is posted to, e.g. if the list is about technical issues then your posting should contain technical discussion. Ongoing irrelevant chatter or flaming only detracts from the value of the mailing list for everyone on it and will not be tolerated.
- No posting should be made to more than two mailing lists, and only to two when a clear and obvious need to post to both lists exists. For most lists, there is already a great deal of subscriber overlap. If a message is sent to you in such a way that multiple mailing lists appear on the Cc line then the Cc line should also be trimmed before sending it out again. *You are still responsible for your own cross-postings, no matter who the originator might have been.*
- Personal attacks and profanity (in the context of an argument) are not allowed, and that includes users and developers alike. Gross breaches of netiquette, like excerpting or reposting private mail when permission to do so was not and would not be forthcoming, are frowned upon but not specifically enforced. *However*, there are also very few cases where such content would fit within the charter of a list and it would therefore probably rate a warning (or ban) on that basis alone.
- Advertising of non-DragonFly related products or services is strictly prohibited and will result in an immediate ban if it is clear that the offender is advertising by spam.

C.1.4 Filtering on the Mailing Lists

The DragonFly mailing lists are filtered in multiple ways to avoid the distribution of spam, viruses, and other unwanted emails. The filtering actions described in this section do not include all those used to protect the mailing lists.

Only certain types of attachments are allowed on the mailing lists. All attachments with a MIME content type not found in the list below will be stripped before an email is distributed on the mailing lists.

- application/octet-stream
- application/pdf
- application/pgp-signature

- application/x-pkcs7-signature
- message/rfc822
- multipart/alternative
- multipart/related
- multipart/signed
- text/html
- text/plain
- text/x-diff
- text/x-patch

Note: Some of the mailing lists might allow attachments of other MIME content types, but the above list should be applicable for most of the mailing lists.

If an email contains both an HTML and a plain text version, the HTML version will be removed. If an email contains only an HTML version, it will be converted to plain text.

C.2 Usenet Newsgroups

All the DragonFly mailing lists are duplicated as newsgroups, served at nntp.dragonflybsd.org.

In addition to these newsgroups, there are many others in which DragonFly is discussed or are otherwise relevant to DragonFly users. Keyword searchable archives (http://minnie.tuhs.org/BSD-info/bsdnews_search.html) are available for some of these newsgroups from courtesy of Warren Toomey <wkt@cs.adfa.edu.au>.

C.2.1 BSD Specific Newsgroups

- comp.unix.bsd.freebsd.announce (news:comp.unix.bsd.freebsd.announce)
- de.comp.os.unix.bsd (news:de.comp.os.unix.bsd) (German)
- fr.comp.os.bsd (news:fr.comp.os.bsd) (French)
- it.comp.os.freebsd (news:it.comp.os.freebsd) (Italian)

C.2.2 Other UNIX Newsgroups of Interest

- comp.unix (news:comp.unix)
- comp.unix.questions (news:comp.unix.questions)
- comp.unix.admin (news:comp.unix.admin)
- comp.unix.programmer (news:comp.unix.programmer)

- comp.unix.shell (news:comp.unix.shell)
- comp.unix.user-friendly (news:comp.unix.user-friendly)
- comp.security.unix (news:comp.security.unix)
- comp.sources.unix (news:comp.sources.unix)
- comp.unix.advocacy (news:comp.unix.advocacy)
- comp.unix.misc (news:comp.unix.misc)
- comp.bugs.4bsd (news:comp.bugs.4bsd)
- comp.bugs.4bsd.ucb-fixes (news:comp.bugs.4bsd.ucb-fixes)
- comp.unix.bsd (news:comp.unix.bsd)

C.2.3 X Window System

- comp.windows.x.i386unix (news:comp.windows.x.i386unix)
- comp.windows.x (news:comp.windows.x)
- comp.windows.x.apps (news:comp.windows.x.apps)
- comp.windows.x.announce (news:comp.windows.x.announce)
- comp.windows.x.intrinsics (news:comp.windows.x.intrinsics)
- comp.windows.x.motif (news:comp.windows.x.motif)
- comp.windows.x.pex (news:comp.windows.x.pex)
- comp.emulators.ms-windows.wine (news:comp.emulators.ms-windows.wine)

C.3 World Wide Web Servers

Central Servers,

(as of 2005/06/27 23:37:47 UTC)

- Central Servers
 - <http://www.DragonFlyBSD.org/>

Appendix D. PGP Keys

In case you need to verify a signature or send encrypted email to one of the developers, a number of keys are provided here for your convenience.

D.1 Developers

D.1.1 Hiten Pandya <hmp at dragonflybsd.org>

```
pub 1024D/938CACA8 2004-02-13 Hiten Pandya (FreeBSD) <hmp@FreeBSD.org>
   Key fingerprint = 84EB C75E C75A 50ED 304E E446 D974 7842 938C ACA8
uid                               Hiten Pandya <hmp@backplane.com>
sub 2048g/783874B5 2004-02-13
```

Colophon

This book is the combined work of hundreds of contributors to “The FreeBSD Documentation Project” and the “The DragonFly BSD Documentation Project”. The text is authored in SGML according to the DocBook DTD and is formatted from SGML into many different presentation formats using **Jade**, an open source DSSSL engine. Norm Walsh’s DSSSL stylesheets were used with an additional customization layer to provide the presentation instructions for **Jade**. The printed version of this document would not be possible without Donald Knuth’s **TeX** typesetting language, Leslie Lamport’s **LaTeX**, or Sebastian Rahtz’s **JadeTeX** macro package.