
Porting UNIX/Linux Applications to Mac OS X

Cross Platform: UNIX



2008-04-08



Apple Inc.
© 2002, 2008 Apple Inc.
All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, mechanical, electronic, photocopying, recording, or otherwise, without prior written permission of Apple Inc., with the following exceptions: Any person is hereby authorized to store documentation on a single computer for personal use only and to print copies of documentation for personal use provided that the documentation contains Apple's copyright notice.

The Apple logo is a trademark of Apple Inc.

Use of the "keyboard" Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

No licenses, express or implied, are granted with respect to any of the technology described in this document. Apple retains all intellectual property rights associated with the technology described in this document. This document is intended to assist application developers to develop applications only for Apple-labeled computers.

Every effort has been made to ensure that the information in this document is accurate. Apple is not responsible for typographical errors.

Apple Inc.
1 Infinite Loop
Cupertino, CA 95014
408-996-1010

.Mac is a registered service mark of Apple Inc.

Apple, the Apple logo, AppleScript, AppleScript Studio, Aqua, Bonjour, Carbon, Cocoa, ColorSync, eMac, FireWire, Mac, Mac OS, Macintosh, Objective-C, Pages, Panther, Quartz, QuickTime, Rosetta, and Xcode are trademarks of Apple Inc., registered in the United States and other countries.

Finder is a trademark of Apple Inc.

NeXT and OPENSTEP are trademarks of NeXT Software, Inc., registered in the United States and other countries.

Intel and Intel Core are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

OpenGL is a registered trademark of Silicon Graphics, Inc.

PowerPC and the PowerPC logo are trademarks of International Business Machines Corporation, used under license therefrom.

UNIX is a registered trademark of The Open Group

X Window System is a trademark of the Massachusetts Institute of Technology.

Simultaneously published in the United States and Canada.

Even though Apple has reviewed this document, APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THIS DOCUMENT, ITS QUALITY, ACCURACY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS DOCUMENT IS PROVIDED "AS IS," AND YOU, THE READER, ARE ASSUMING THE ENTIRE RISK AS TO ITS QUALITY AND ACCURACY.

IN NO EVENT WILL APPLE BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM ANY DEFECT OR INACCURACY IN THIS DOCUMENT, even if advised of the possibility of such damages.

THE WARRANTY AND REMEDIES SET FORTH ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESS OR IMPLIED. No Apple dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty.

Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.

Contents

Introduction **Introduction to Porting UNIX/Linux Applications to Mac OS X** 9

Who Should Read This Document? 9
Organization of This Document 10
See Also 10

Chapter 1 **Overview of Mac OS X** 13

The Family Tree 13
• BSD 13
• Mach 13
• NEXTSTEP 13
• Earlier Version of the Mac OS 14
Mac OS X and Darwin 14
What Macintosh Users Expect 15
 Benefits of Porting to Mac OS X 15
 Responsibilities of Porting to Mac OS X 16

Chapter 2 **Preparing to Port** 17

The Mac OS X Command-Line Environment 17
Installing the Mac OS X Developer Tools 18
Building Makefile Projects With Xcode 19
Windowing Environment Considerations 20
Working with 64-bit Software 20

Chapter 3 **Compiling Your Code in Mac OS X** 21

Using GNU Autoconf, Automake, and Autoheader 21
Compiling for Multiple CPU Architectures 22
 Cross-Compiling a Self-Bootstrapping Tool 25
Conditional Compilation on Mac OS X 26
Choosing a Default Compiler 27
 Using the gcc_select Command 28
Setting Compiler Flags 28
Understanding Two-Level Namespaces 29
Executable Format 29
Dynamic Libraries and Plug-ins 29
 Using Dynamic Libraries at Link Time 30
 Using Dynamic Libraries Programmatically 30
 Compiling Dynamic Libraries and Plugins 30
Bundles 31

- Application Bundles 31
- Frameworks 31
- For More Information 32

Chapter 4 **Using Mac OS X Native Application Environments 33**

- Choosing a Native Application Environment 33
 - Cocoa 34
 - Java 37
 - Carbon 38
- Lower-Level Graphics Technologies 39
 - Quartz 39
 - OpenGL 40
 - QuickTime 40

Chapter 5 **Choosing a Graphical Environment for Your Application 41**

- What Kind of Application Are You Porting? 41
- How Well Does It Need to Integrate With Mac OS X? 42
- Does Your Application Require Cross-Platform Functionality? 42

Chapter 6 **Using Traditional UNIX Graphical Environments 45**

- Choosing a UNIX Graphical Environment 45
 - X11R6 46
 - Benefits of X11R6 Development 46
 - Drawbacks of X11R6 Development 46
 - Tcl/Tk 47
 - Benefits of Tk Development 47
 - Drawbacks of Tk Development 47
 - Qt 47
 - Benefits of Qt Development 47
 - Drawbacks in Qt Development 47
 - GTK+/GTKmm 48
 - Benefits of GTK+/GTKmm Development 48
 - Drawbacks in GTK+/GTKmm Development 48
 - wxWidgets 48
- Wrapping a Command Line Application with a GUI Interface 48
 - Choosing a Wrapping Method 49
 - File Access Considerations 49
 - For More Information 50

Chapter 7 **Porting File, Device, and Network I/O 51**

- How Mac OS X File I/O Works 51
 - Using Carbon APIs for File I/O 52

- Using Cocoa APIs for File I/O 52
- Presenting File Open and Save Dialog Boxes 53
- How Mac OS X Device I/O Works 53
- File System Organization 54
- How Mac OS Networking Works 54

Chapter 8 **Distributing Your Application 57**

- Bundles vs. Installers 57
- Packaging Basics 58
 - Disk Utility (or Disk Copy) 58
- Creating Disk Images Programmatically Using hdiutil 60
- Tell the World About It 60

Chapter 9 **Additional Features 61**

- AppleScript and AppleScript Studio 61
- Audio Architecture 61
- Boot Sequence 62
- Configuration Files 62
- Device Drivers 62
- The File System 62
 - File-System Structure and Visibility 63
 - Supported File-System Types 63
- The Kernel 64
- Open Directory and the dscl Tool 65
 - Example: Adding a User From the Command Line 65
- Printing (CUPS) 66
- Bonjour 66
- Scripting Languages 67
- Security and Security Services 67
 - Role-Based Authentication 68

Chapter 10 **(Re)designing for Portability 69**

- What is Portability? 69
- Using Abstraction Layers 69
 - Avoid Conditionalizing Code 70
 - GUI Abstraction Issues 70
 - Other General Rules 71
- Using Plug-Ins and Libraries Effectively 72
 - Designing a Plug-In Architecture 72
- Architectural Portability 75

[Glossary](#) 77

[Document Revision History](#) 79

Figures, Tables, and Listings

Chapter 1 **Overview of Mac OS X 13**

Figure 1-1 Darwin's relation to Mac OS X 14

Chapter 4 **Using Mac OS X Native Application Environments 33**

Figure 4-1 High-level graphics technologies 33

Figure 4-2 Low-level graphics technologies 39

Listing 4-1 `main.m` 35

Listing 4-2 `HelloController.m` 35

Listing 4-3 `HelloController.h` 35

Listing 4-4 `SayHello.mm` 36

Listing 4-5 `SayHello.h` 36

Listing 4-6 `FooClass.cpp` 37

Listing 4-7 `FooClass.h` 37

Chapter 5 **Choosing a Graphical Environment for Your Application 41**

Figure 5-1 Graphical environments 42

Table 5-1 Platforms of cross-platform technologies 43

Chapter 7 **Porting File, Device, and Network I/O 51**

Table 7-1 POSIX versus Carbon APIs 52

Chapter 8 **Distributing Your Application 57**

Figure 8-1 Disk Utility options 59

Listing 8-1 Automatic Disk Image Creation using `hdiutil` 60

Chapter 9 **Additional Features 61**

Figure 9-1 XNU personalities 64

Introduction to Porting UNIX/Linux Applications to Mac OS X

The UNIX Porting Guide is a first stop for UNIX developers coming to Mac OS X. This document helps guide developers in bringing applications written for UNIX-based operating systems to Mac OS X. It provides the background needed to understand the operating system. It touches on some of the design decisions, and it provides a listing and discussion of some of the main areas that you should be concerned with in bringing UNIX applications to Mac OS X. It also points out some of the advanced features of Mac OS X not available in traditional UNIX applications that you can add to your ported applications.

This document also provides an entry point for other documentation on various subjects that may be of interest if you are porting an application from a UNIX environment to Mac OS X.

This document is an overview, not a tutorial. In many regards it is a companion to the more extensive *Mac OS X Technology Overview*, but with a bias toward the UNIX developer.

This document also does not cover porting shell scripts to Mac OS X. For more information about shell scripts and Mac OS X, you should read *Shell Scripting Primer*.

Bringing UNIX Apps to Mac OS X

The introduction of UNIX-like operating systems such as FreeBSD and Linux for personal computers was a great step in bringing the power and stability of UNIX to the mass market. Generally though, these projects were driven by power users and developers for their own use, without making design decisions that would make UNIX palatable to consumers. Mac OS X, on the other hand, was designed from the beginning with end users in mind.

With this operating system, Apple builds its well-known strengths in simplicity and elegance of design on a UNIX-based foundation. Rather than reinventing what has already been done well, Apple is combining their strengths with the strengths brought about by many years of advancement by the UNIX community.

Who Should Read This Document?

Any UNIX developers can benefit from reading this book.

- In-house corporate application developers
- Commercial UNIX developers
- Open source developers
- Open source porters
- Higher education faculty, staff, and students
- Science and research developers

INTRODUCTION

Introduction to Porting UNIX/Linux Applications to Mac OS X

If you're a commercial UNIX developer, you are already familiar with other UNIX-based systems and may want to understand the differences between other systems and Mac OS X. You might be interested in porting the GUI from an X11 environment into a native graphics environment using Carbon or Cocoa. You may also have special needs such as direct hardware access, exclusive file access guarantees, and so on.

If you're a corporate in-house developer (developing applications for internal use), you probably want to port applications with minimal code divergence.

If you're an open source developer, you might want information about how to incorporate new technologies into your software, and may be interested in GUI porting, depending on your level of interest. Alternately, you might be interested only in quickly porting code to a new platform with minimal changes so that you can easily get your changes back into the official code base. If so, you may be more likely to use compatibility shims than to use new APIs.

No matter what "flavor" of developer you are, this book will provide information that is helpful to you and provide pointers to additional documents that may be of interest.

Important: This document is not designed for pure Java developers. Mac OS X has a full and robust Java 2 Platform, Standard Edition (J2SE) implementation. If you have a pure Java application already, it should run in Mac OS X. More information on Java development can be found at <http://developer.apple.com/documentation/Java/Java.html>.

Organization of This Document

This document is a first step for UNIX developers coming to Mac OS X. It contains many links to more extensive documentation. Specific details of implementation are covered here only in cases where it is not adequately covered in other places in the documentation set.

To use this document most effectively, first read "[Overview of Mac OS X](#)" (page 13) to find out the basics about the Mac OS X platform and to get some of the high-level information you need to begin your port. If you already have an application that builds on other UNIX-based platforms, "[Compiling Your Code in Mac OS X](#)" (page 21) will help you find out how to compile your code on Mac OS X.

Most of your effort, however, should be spent towards making decisions concerning which, if any, graphical user interface to implement with your application. "[Choosing a Graphical Environment for Your Application](#)" (page 41) helps you with this.

If you want to refactor your application to take advantage of the rich feature set of Mac OS X, see "[Additional Features](#)" (page 61) for examples of features available in Mac OS X.

Once you have a complete application, read "[Distributing Your Application](#)" (page 57) for information on getting your application to Mac OS X users.

See Also

Developer documentation can be found at Apple's developer documentation website at <http://developer.apple.com/documentation/>. This site contains reference, conceptual, and tutorial material for the many facets of development on Mac OS X. The Mac OS X Developer Tools CD includes a snapshot of the developer

INTRODUCTION

Introduction to Porting UNIX/Linux Applications to Mac OS X

documentation, which is installed in `/Developer/ADC Reference Library/documentation` and can be searched for and viewed in Xcode's doc viewer. They are also accessible from Apple's Technical Publications website. The `man` pages are also included with the Mac OS X Developer Tools.

Apple Developer Connection (ADC) hosts a website full of information useful to UNIX developers targeting Mac OS X, <http://developer.apple.com/unix>. ADC also offers a variety of membership levels to help you in your development. These range from free memberships that give you access to developer software, to paid memberships that provide support incidents as well as the possibility of software seeds. More information on memberships is available at <http://developer.apple.com/membership/>.

Once a year in early Summer, Apple hosts the Worldwide Developers Conference (WWDC) in the San Francisco, California Bay area. This is an extremely valuable resource for developers trying to get an overall picture of Mac OS X as well as specific implementation details of individual technologies. Information on WWDC is available on the ADC website.

Apple hosts an extensive array of public mailing lists. These are available for public subscription and searching at <http://lists.apple.com>. The *unix-porting* list is highly recommended. The *darwin-development* and *darwinos-users* lists also offer much help but less specific to the task of porting.

In addition to Apple's own resources, many external resources exist. Two recommended websites are O'Reilly's Mac DevCenter, <http://www.oreillynet.com/mac/>, and Stepwise, at <http://www.stepwise.com>.

INTRODUCTION

Introduction to Porting UNIX/Linux Applications to Mac OS X

Overview of Mac OS X

Mac OS X is a modern operating system that combines the power and stability of UNIX-based operating systems with the simplicity and elegance of the Macintosh. For years, power users and developers have recognized the strengths of UNIX and its offshoots. While UNIX-based operating systems are indispensable to developers and power users, consumers have rarely been able to enjoy their benefits because of the perceived complexity. Instead consumers have lived with a generation of desktop computers that could only hope to achieve the strengths that UNIX-based operating systems have had from the beginning.

This chapter is for anyone interested in an overview of Mac OS X—its lineage and its open source core, called Darwin. Here you will find background information about Mac OS X and how your application fits in.

The Family Tree

Although this document covers the basic concepts in bringing UNIX applications to Mac OS X, it is by no means comprehensive. This section is provided to give you a hint on where to look for additional documentation by outlining how Mac OS X came to be. Knowing a little about the lineage of Mac OS X will help you to find more resources as the need arises.

- **BSD**

Part of the history of Mac OS X goes back to Berkeley Software Distributions (BSD) UNIX of the late seventies and early eighties. Specifically, it is based in part on BSD 4.4 Lite. On a system level, many of the design decisions are made to align with BSD-style UNIX systems. Most libraries and utilities are from FreeBSD (<http://www.freebsd.org/>), but some are derived from NetBSD (<http://www.netbsd.org/>). For future development, Mac OS X has adopted FreeBSD as a reference code base for BSD technology. Work is ongoing to more closely synchronize all BSD tools and libraries with the FreeBSD-stable branch..

- **Mach**

Although Mac OS X must credit BSD for most of the underlying levels of the operating system, Mac OS X also owes a major debt to Mach. The kernel is heavily influenced in its design philosophy by Carnegie Mellon's Mach project. The kernel is not a pure microkernel implementation, since the address space is shared with the BSD portion of the kernel and the I/O Kit.

- **NEXTSTEP**

In figuring out what makes Mac OS X tick, it is important to recognize the influences of NEXTSTEP and OPENSTEP in its design. Apple's acquisition of NeXT in 1997 was a major key in bringing Mac OS X from the drawing board into reality. Many parts of Mac OS X of interest to UNIX developers are enhancements and

offshoots of the technology present in NEXTSTEP. From the file system layer to the executable format and from the high-level Cocoa API to the kernel itself, the lineage of Mac OS X as a descendant of NEXTSTEP is evident.

• Earlier Version of the Mac OS

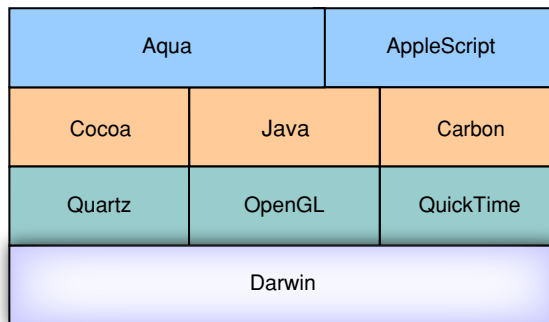
Although it shares its name with earlier versions of the Mac OS, Mac OS X is a fundamentally new operating system. This does not mean that all that went before has been left out. Mac OS X still includes many of the features that Mac OS 9 and earlier versions included. Although your initial port to Mac OS X may not use any of the features inherited from Mac OS 9, as you enhance the application, you might take advantage of some of the features provided by technologies like ColorSync or the Carbon APIs. Mac OS 9 is also the source of much of the terminology used in Mac OS X.

Mac OS X and Darwin

The word *Darwin* is often used to refer to the underpinnings of Mac OS X. In fact, in some circles Mac OS X itself is rarely mentioned at all. It is important to understand the distinction between the two—how they are related and how they differ.

Darwin is the core of the Mac OS X operating system. Although Darwin can stand alone as an independent operating system, it includes only a subset of the features available in Mac OS X. [Figure 1-1](#) (page 14) shows how Darwin is related to Mac OS X as a whole.

Figure 1-1 Darwin's relation to Mac OS X



Darwin is an open source project. With it, you as a developer gain access to the foundation of Mac OS X. Its openness also allows you to submit changes that you think should be reflected in Mac OS X as a whole. Darwin has been released as a separate project that runs on PowerPC-based Macintosh computers as well as x86-compatible computers. Although it could be considered a standalone operating system in its own right, many of the fundamental design decisions of Darwin are governed by its being embedded within Mac OS X. In bringing your applications to the platform, you should target Mac OS X version 10.1.4 (Darwin 5.4) or later.

Mac OS X itself is not an Open Source project. As you can see from [Figure 1-1](#) (page 14), there are many parts of Mac OS X that are not included in the Open Source Darwin components. Part of your job while porting is deciding where your application will fit in Mac OS X.

If you are a developer whose tool is a command-line tool (or has a useful subset that is a command-line tool), you can, of course, simply port your application as a command-line tool or service, which is usually not that complicated. By doing this you gain a small benefit, in that it is now available to Mac OS X users who are familiar with the UNIX command-line environment. You will not be able to market it to Mac OS X users as a whole though, since many users do not even know how to access the command line on their computers.

The basic steps in porting a UNIX application to Mac OS X typically include:

1. Port to the command line.
2. Provide a graphical user interface (GUI).

What Macintosh Users Expect

In bringing your UNIX application to Mac OS X, you are entering a world where great emphasis is placed on user interactions. This brings you many opportunities and benefits as a developer, but also some responsibilities.

Benefits of Porting to Mac OS X

Porting to Mac OS X has three primary benefits:

- stable long-term customer base
- good inroad into education
- powerful developer tools

Bringing UNIX applications to Mac OS X can be very profitable if done correctly. Well-designed Macintosh applications of years past are the standards of today. PhotoShop, Illustrator, and Excel are all applications that first made their name on the Macintosh. Now is the time to win the hearts of Macintosh users with the next great application. In a word, possibly millions of paying customers!

Macintosh users are willing to spend their money on great applications because they know that Apple strives to give them a high-quality user environment. Apple developers are known for providing great applications for that environment.

For years, Apple has been known for its commitment to education. Mac OS X targets the education market for developers and is an ideal platform for learning for students. With its standards-based technologies as well as home-grown technologies, you have an ideal platform for use in educational application deployment and development.

Mac OS X also provides benefits in a development environment. Apple strives for standards first, then it adds that little bit that makes it better on a Mac. As a developer, you have access to many of the development tools and environments that you have on other platforms, like Java, OpenGL, POSIX libraries, and the BSD TCP/IP stack, but you also have built-in benefits like the Apache Web server on every computer, the Cocoa object-oriented development environment, a PDF-based display system (Quartz), Kerberos, QuickTime, a dynamic core audio implementation, and a suite of world-class developer tools. By adding a native Mac OS X front end to your application, you can achieve a cost-effective new deployment platform with minimal additional development effort.

Mac OS X adds tremendous value both to you and your customers on a standards-based operating system.

Responsibilities of Porting to Mac OS X

Along with benefits come responsibilities. If you have decided to make a full-featured Mac OS X application, here are some guidelines to keep in mind.

- A Mac OS X user should never have to resort to the command line to perform any task in an application with a graphical user interface. This is especially important to remember since the BSD user environment may not even be installed on a user's system. The libraries and kernel environment are of course there by default, but the tools may not be.
- If you are making graphical design decisions, you need to become familiar with the *Apple Human Interface Guidelines*, available from the Apple developer website. These are the standards that Macintosh users expect their applications to live up to. Well-behaved applications from Apple and third-party developers give the Macintosh its reputation as the most usable interface on the planet.

The responsibilities boil down to striving for an excellent product from a user's perspective. Mac OS X gives you the tools to make your applications shine.

Preparing to Port

A seasoned UNIX developer recognizes that no matter how similar two UNIX-based operating systems may be, there are always details that set one apart from another.

This chapter highlights areas to be aware of when compiling your code base for Mac OS X. It notes details about compiler flags that are important to you, and gives insight into how to link different parts of your code in Mac OS X. Many of these topics are covered more extensively in other resources as noted.

With few exceptions, this chapter applies to all varieties of UNIX developers. However, there are a few things that are specific to a given audience. These are pointed out as they appear.

If you are porting an open source application, you should first check to see if any of the open source port collections have already ported the application to avoid duplicating effort. You should also consider including your port in one of these collections to make it easier for others to take advantage of it. Finally, you should always make your changes available upstream to the original open source project for inclusion in future releases.

Some of the Darwin-based open source port collections include:

- MacPorts (<http://www.macports.org/>)
- Fink (<http://fink.sourceforge.net/>)
- GNU-Darwin (<http://www.gnu-darwin.org/>)

For more information on these port collections, visit their websites.

Important: This information is provided solely for your convenience, and that mention of these projects in no way represents any recommendation or endorsement by Apple.

Before you bring the basic port of your code to Mac OS X, make sure that you have the requisite tools for the task. It helps to know what is and isn't available to you by default.

The Mac OS X Command-Line Environment

The most important thing to do before porting an application is to familiarize yourself with the Mac OS X environment. In particular, you should familiarize yourself with the System Preferences application (and customize your computer to suit) and the Terminal application (to obtain a command-line interface).

The System Preferences application is located in the `Applications` folder at the root level of your hard drive. The Terminal application is located in the `Utilities` folder (which is within the `Applications` folder at the root level of your hard drive).

Once you have a Terminal window open, you can take advantage of a basic selection of common tools, assuming that they are installed. Before using Mac OS X as a development platform, you should make sure that the BSD subsystem is installed. On Mac OS X v10.4, this option is installed by default, but on previous versions, it was not.

You can check for this by looking for the BSD package receipt, `BSD.pkg`, in `/Library/Receipts`. If this receipt is not present on your system, you must reinstall Mac OS X. When you do, you must customize the installation by checking the BSD Subsystem checkbox.

Note: If you have already installed software updates, you cannot simply go back and reinstall over the existing system because of the potential for version conflicts with the software updates. If the BSD subsystem package is not installed, you will have to do an “archive and install” installation, then reinstall the software updates. For more information, see [Mac OS X: About the Archive and Install feature](#).

With the BSD subsystem installed, a look through `/bin` and `/usr/bin` should reveal a familiar environment. Welcome home.

Installing the Mac OS X Developer Tools

Because Mac OS X has a BSD core, you have access to the numerous open source tools that you are already familiar with (like the GNU tools, for example).

The Mac OS X Xcode Tools package provides additional tools that you will need to install to round out your development environment. These are not part of the default installation, but are essential to you. They contain some of the most important tools, such as the compiler (`gcc`) and debugger (`gdb`).

Most (if not all) of the tools described in this document are part of the Xcode Tools installation, so if you don't find `autoconf` or `make`, for example, that probably means you don't have Xcode Tools installed.

Important: If you are using makefiles for compilation, you should install Xcode in the default location (`/usr`). If you do not, you may have to do extra work to get your scripts to run the compiler, linker, and so on in a nonstandard location.

If you are compiling UNIX applications, be sure to select the optional BSD SDK package when installing Xcode tools. This tells the installer to install a number of BSD-level headers, without which many UNIX applications will not compile.

You can find the Xcode Tools Installer on your Mac OS X DVD (or as one of the CDs in the CD version). In addition, you can always obtain the very latest version online from the Apple Developer Connection (ADC) website, <http://connect.apple.com>. You need an ADC account to download the Developer Tools. Free accounts are available for those who just need access to the tools.

After installing the Mac OS X Developer Tools, you will have a selection of new tools to take advantage of:

- Inside `/Developer/Tools` are many Mac OS X-specific command-line tools. Documentation is provided in the form of man pages.
- `/usr/bin` now contains more command-line tools than were supplied by the BSD package alone, most notably `gcc` and `gdb`. `/usr/bin/cc` defaults to `gcc 3` in Mac OS X version 10.2 and later, but `gcc 2.95.2` is also available. See “[Choosing a Default Compiler](#)” (page 27) for more information.

- `/Developer/Applications` contains a wide assortment of graphical tools and utilities. Key among these are the following:
 - Xcode is a graphical integrated development environment for applications in multiple programming languages.
 - Interface Builder provides a simple way to build complex user interfaces.
 - FileMerge lets you graphically compare and merge files and directories.
 - IORegistryExplorer helps you determine which devices are registered with the I/O KIT. See “[Device Drivers](#)” (page 62) for a discussion of the I/O Kit.
 - MallocDebug analyzes all allocated memory in an application. It can measure the memory allocated since a given point in time and detect memory leaks
 - PackageMaker builds easily distributable Mac OS X packages.

Documentation for these tools is available from the Apple Technical Publications website.

Building Makefile Projects With Xcode

Mac OS X-native applications can be built in a number of environments. Most UNIX utilities are built using makefiles, though some are built using build scripts, imakefiles, or various other mechanisms.

If you are porting a UNIX command-line tool for personal use, you probably want to continue using the existing build environment, because keeping multiple build environments synchronized can be tricky at best. However if you are porting a more extensive tool and you plan to add a Mac OS X GUI to the tool, you may find it more convenient to work with the project in a development environment such as Xcode.

Note: These instructions apply to the latest version of Xcode.

Although Xcode keeps track of build settings in its own preferences files for information beyond what could normally be maintained in a makefile, it can also work closely with your project’s makefiles. If you want to use Xcode for development in Mac OS X, you can include a makefile in a Xcode project as follows:

1. Launch Xcode.
2. Choose New Project from the File menu.
3. Select whatever project type you are targeting. If you ultimately want an application, select something like Cocoa Application. If you are just trying to build a command-line utility, select one of the tools—for example, Standard Tool.
4. Follow the prompts to name and save your project. A new default project of that type is opened.
5. Open the Targets disclosure triangle and delete any default targets that may exist.
6. From the Project menu, Choose New Target.
7. Select “External Target” from the list. If this is not shown in the “Special Targets” list, you are not running the latest version of Xcode. Upgrade first.
8. Follow the prompts to name that target. When you have done this, a target icon with the name you just gave it appears in the Targets pane of the open Xcode window.

9. Double-click that new target. You should now see a new window with the build information for this target. This is *not* the same thing as clicking info. You *must* double-click the target itself.
10. In the “Custom Build Command” section of the target inspector, change the field called “Directory” to point to the directory containing your makefile, and change any other settings as needed. For example, in the Custom Build Settings pane, you could change Build Tool from `/usr/bin/gnumake` to `/usr/bin/bsdmake`. More information on the fields is available in Xcode Help.
11. Change the active target to your new target by choosing “Set Active Target” from the Project menu.
12. Add the source files to the project. To do this, first open the disclosure triangle beside the “Source” folder in the left side of the project window. Next, drag the folder containing the sources from the Finder into that “Source” folder in Xcode. Tell Xcode not to copy files. Xcode will recursively find all of the files in that folder. Delete anything you don’t want listed.
13. When you are ready to build the project, click the Build and Run button in the toolbar, select Build from the Build menu, or just press Command-B.
14. Once the project is built, tell Xcode where to find the executable by choosing “New Custom Executable” from the Project menu. Choose the path where the executable is located, then add the name of the executable.
15. Run the resulting program by pressing Command-R.

This should get you started in bringing your application into the native build environment of Mac OS X.

Windowing Environment Considerations

Before you compile an application in Mac OS X, be aware that the Mac OS X native windowing and display subsystem, Quartz, is based on the Portable Document Format (PDF). Quartz consists of a lightweight window server as well as a graphics rendering library for two-dimensional shapes. The window server features device-independent color and pixel depth, layered compositing, and buffered windows. The rendering model is PDF based.

Quartz is not an X Window System implementation. If you need an X11R6 implementation, you can easily install one. For more information, see “[X11R6](#)” (page 46).

Information on the graphical environments available to you can be found in “[Choosing a Graphical Environment for Your Application](#)” (page 41) along with some things to consider when choosing a graphical environment.

Working with 64-bit Software

Beginning in version 10.4, Mac OS X supports execution of 64-bit command-line tools. It is also possible to use a 32-bit Mac OS X front-end application to provide a user interface through careful use of client-server communication.

This document primarily covers general porting issues. Issues specific to 64-bit computing are beyond the scope of this document. For additional information on porting 64-bit software to Mac OS X, see the document *64-Bit Transition Guide*.

Compiling Your Code in Mac OS X

Now that you have the basic pieces in place, it is time to build your application. This section covers some of the more common issues that you may encounter in bringing your UNIX application to Mac OS X. These issues apply largely without regard to what type of development you are doing.

Using GNU Autoconf, Automake, and Autoheader

If you are bringing a preexisting command-line utility to Mac OS X that uses GNU `autoconf`, `automake`, or `autoheader`, you will probably find that it configures itself without modification (though the resulting configuration may be insufficient). Just run `configure` and `make` as you would on any other UNIX-based system.

If running the `configure` script fails because it doesn't understand the architecture, try replacing the project's `config.sub` and `config.guess` files with those available in `/usr/share/automake-1.6`. If you are distributing applications that use `autoconf`, you should include an up-to-date version of `config.sub` and `config.guess` so that Mac OS X users don't have to do anything extra to build your project.

If that still fails, you may need to run `/usr/bin/autoconf` on your project to rebuild the `configure` script before it works. Mac OS X includes `autoconf` in the BSD tools package. Beyond these basics, if the project does not build, you may need to modify your makefile using some of the tips provided in the following sections. After you do that, more extensive refactoring may be required.

Some programs may use `autoconf` macros that are not supported by the version of `autoconf` that shipped with Mac OS X. Because `autoconf` changes periodically, you may actually need to get a new version of `autoconf` if you need to build the very latest sources for some projects. In general, most projects include a prebuilt `configure` script with releases, so this is usually not necessary unless you are building an open source project using sources obtained from CVS or from a daily source snapshot.

However, if you find it necessary to upgrade `autoconf`, you can get a current version from <http://www.gnu.org/software/autoconf/>. Note that `autoconf`, by default, installs in `/usr/local/`, so you may need to modify your `PATH` environment variable to use the newly updated version. Do *not* attempt to replace the version installed in `/usr/`.

For additional information about using the GNU autotoolset, see <http://autotoolset.sourceforge.net/tutorial.html> and the manual pages `autoconf`, `automake`, and `autoheader`.

Compiling for Multiple CPU Architectures

Because the Macintosh platform includes more than one processor family, it is often important to compile software for multiple processor architectures. For example, libraries should generally be compiled as universal binaries even if you are exclusively targeting an Intel-based Macintosh computer, as your library may be used by a PowerPC binary running under Rosetta. For executables, if you plan to distribute compiled versions, you should generally create universal binaries for convenience.

When compiling programs for architectures other than your default host architecture, such as compiling for a ppc64 or Intel-based Macintosh target on a PowerPC-based build host, there are a few common problems that you may run into. Most of these problems result from one of the following mistakes:

- Assuming that the build host is architecturally similar to the target architecture and will thus be capable of executing intermediate build products
- Trying to determine target-processor-specific information at configuration time (by compiling and executing small code snippets) rather than at compile time (using macro tests) or execution time (for example, by using conditional byte swap functions)

Whenever cross-compiling occurs, extra care must be taken to ensure that the target architecture is detected correctly. This is particularly an issue when generating a binary containing object code for more than one architecture.

In many cases, binaries containing object code for more than one architecture can be generated simply by running the normal configuration script, then overriding the architecture flags at compile time.

For example, you might run

```
./configure
```

followed by

```
make CFLAGS="-isysroot /Developer/SDKs/MacOSX10.4u.sdk -arch ppc \  
-arch i386" LDFLAGS="-syslibroot /Developer/SDKs/MacOSX10.4u.sdk \  
-arch ppc -arch i386 -arch ppc -arch i386"
```

to generate a universal binary (for Intel-based and PowerPC-based Macintosh computers). To generate a 4-way universal binary that includes 64-bit versions, you would add `-arch ppc64` and `-arch x86_64` to the `CFLAGS` and `LDFLAGS`.

Note: If you are using an older version of `gcc` and your makefile passes `LDFLAGS` to `gcc` instead of passing them directly to `ld`, you may need to specify the linker flags as `-Wl, -syslibroot, /Developer/SDKs/MacOSX10.4u.sdk`. This tells the compiler to pass the unknown flags to the linker without interpreting them. Do not pass the `LDFLAGS` in this form to `ld`, however; `ld` does not currently support the `-Wl` syntax.

If you need to support an older version of `gcc` and your makefile passes `LDFLAGS` to both `gcc` and `ld`, you may need to modify it to pass this argument in different forms, depending on which tool is being used. Fortunately, these cases are rare; most makefiles either pass `LDFLAGS` to `gcc` or `ld`, but not both. Newer versions of `gcc` support `-syslibroot` directly.

If your makefile does not explicitly pass the contents of `LDFLAGS` to `gcc` or `ld`, they may still be passed to one or the other by a make rule. If you are using the standard built-in make rules, the contents of `LDFLAGS` are passed directly to `ld`. If in doubt, assume that it is passed to `ld`. If you get an invalid flag error, you guessed incorrectly.

If your makefile uses `gcc` to run the linker instead of invoking it directly, you *must* specify a list of target architectures to link when working with universal binary object (`.o`) files even if you are not using all of the architectures of the object file. If you don't, you will not create a universal binary, and you may also get a linker error. For more information about 64-bit executables, see *64-Bit Transition Guide*.

However, applications that make configuration-time decisions about the size of data structures will generally fail to build correctly in such an environment (since those sizes may need to be different depending on whether the compiler is executing a `ppc` pass, a `ppc64` pass, or an `i386` pass). When this happens, the tool must be configured and compiled for each architecture as separate executables, then glued together manually using `lipo`.

In rare cases, software not written with cross-compilation in mind will make configure-time decisions by executing code on the build host. In these cases, you will have to manually alter either the configuration scripts or the resulting headers to be appropriate for the actual target architecture (rather than the build architecture). In some cases, this can be solved by telling the configure script that you are cross-compiling using the `--host`, `--build`, and `--target` flags. However, this may simply result in defaults for the target platform being inserted, which doesn't really solve the problem.

The best fix is to replace configure-time detection of endianness, data type sizes, and so on with compile-time or run-time detection. For example, instead of testing the architecture for endianness to obtain consistent byte order in a file, you should do one of the following:

- Use C preprocessor macros like `__BIG_ENDIAN__` and `__LITTLE_ENDIAN__` to test endianness at compile time.
- Use functions like `htonl`, `htons`, `ntohl`, and `ntohs` to guarantee a big-endian representation on any architecture.
- Extract individual bytes by bitwise masking and shifting (for example, `lowbyte=word & 0xff`; `nextbyte = (word >> 8) & 0xff`; and so on).

Similarly, instead of performing elaborate tests to determine whether to use `int` or `long` for a 4-byte piece of data, you should simply use a standard sized type such as `uint32_t`.

Note: Not all script execution is incompatible with cross-compiling. A number of open source tools (GTK, for example) use script execution to determine the presence or absence of libraries, determine their versions and locations, and so on.

In those cases, you must be certain that the info script associated with the universal binary installation (or the target platform installation if you are strictly cross-compiling) is the one that executes during the configuration process, rather than the info script associated with an installation specific to your host architecture.

There are a few other caveats when working with universal binaries:

- The library archive utility, `ar`, cannot work with libraries containing code for more than one architecture (or single-architecture libraries generated with `lipo`) after `ranlib` has added a table of contents to them. Thus, if you need to add additional object files to a library, you must keep a separate copy without a TOC.
- The `-M` switch to `gcc` (to output dependency information) is not supported when multiple architectures are specified on the command line. Depending on your makefile, this may require substantial changes to your makefile rules. For autoconf-based configure scripts, the flag `--disable-dependency-tracking` should solve this problem.

For projects using `automake`, it may be necessary to run `automake` with the `-i` flag to disable dependency checks or put `no-dependencies` in the `AUTOMAKE_OPTIONS` variable in each `Makefile.am` file.

- If you run into problems building a universal binary for an open source tool, the first thing you should do is to get the latest version of the source code. This does two things:
 - Ensures that the version of autoconf and automake used to generate the configuration scripts is reasonably current, reducing the likelihood of build failures, execution failures, backwards or forwards compatibility problems, and other idiosyncratic or downright broken behavior.
 - Reduces the likelihood of building a version of an open source tool that contains known security holes or other serious bugs.
- Older versions of autoconf do not handle the case where `--target`, `--host`, and `--build` are not handled gracefully. Different versions also behave differently when you specify only one or two of these flags. Thus, you should always specify all three of these options if you are running an autoconf-generated configure script with intent to cross-compile.
- Some earlier versions of autoconf handle cross-compiling poorly. If your tool contains a configure script generated by an early autoconf, you may be able to significantly improve things by replacing some of the `config.*` files (and `config.guess` in particular) with updated copies from the version of autoconf that comes with Mac OS X.

This will not always work, however, in which case it may be necessary to actually regenerate the configure script by running autoconf. To do this, simply change into the root directory of the project and run `/usr/bin/autoconf`. It will automatically detect and use the `configure.in` file and use it to generate a new configure script. If you get warnings, you should first try a web search for the error message, as someone else may have already run into the problem (possibly on a different tool) and found a solution.

If you get errors about missing `AC_` macros, you may need to download a copy of libraries on which your tool depends and copy their `.m4` autoconf configuration files into `/usr/share/autoconf`. Alternately, you can add the macros to the file `acinclude.m4` in your project's main directory and autoconf should automatically pick up those macros.

You may, in some cases, need to rerun `automake` and/or `autoheader` if your tool uses them. Be prepared to run into missing `AM_` and `AH_` macros if you do, however. Because of the added risk of missing macros, this should generally only be done if running `autoconf` by itself does not correct a build problem.

Important: Be sure to make a backup copy of the original scripts, headers, and other generated files (or, ideally, the entire project directory) before running `autoheader` or `automake`.

- Different makefiles and configure scripts handle command-line overrides in different ways. The most consistent way to force these overrides is to specify them prior to the command. For example:

```
make CFLAGS="-isysroot /Developer/SDKs/MacOSX10.4u.sdk -arch i386 -arch ppc"
```

should generally result in the above being added to `CFLAGS` during compilation. However, this behavior is not completely consistent across makefiles from different projects.

For additional information about `autoconf`, `automake`, and `autoheader`, you can view the `autoconf` documentation at <http://www.gnu.org/software/autoconf/manual/index.html>.

For additional information on compiler flags for Intel-based Macintosh computers, modifying code to support little-endian CPUs, and other porting concerns, you should read *Universal Binary Programming Guidelines, Second Edition*, available from the ADC Reference Library.

Cross-Compiling a Self-Bootstrapping Tool

Probably the most difficult situation you may experience is that of a self-bootstrapping tool—a tool that uses a (possibly stripped-down) copy of itself to either compile the final version of itself or to construct support files or libraries. Some examples include TeX, Perl, and gcc.

Ideally, you should be able to build the executable as a universal binary in a single build pass. If that is possible, everything “just works,” since the universal binary can execute on the host. However, this is not always possible. If you have to cross-compile and glue the pieces together with `lipo`, this obviously will not work.

If the build system is written well, the tool will bootstrap itself by building a version compiled for the host, then use that to build the pieces for the target, and finally compile a version of the binary for the target. In that case, you should not have to do anything special for the build to succeed.

In some cases, however, it is not possible to simultaneously compile for multiple architectures and the build system wasn’t designed for cross-compiling. In those cases, the recommended solution is to pre-install a version of the tool for the host architecture, then modify the build scripts to rename the target’s intermediate copy of the tool and copy the host’s copy in place of that intermediate build product (for example, `mv miniperl miniperl-target; cp /usr/bin/perl miniperl`).

By doing this, later parts of the build script will execute the version of the tool built for the host architecture. Assuming there are no architecture dependencies in the dependent tools or support files, they should build correctly using the host’s copy of the tool. Once the dependent build is complete, you should swap back in the original target copy in the final build phase. The trick is in figuring out when to have each copy in place.

Conditional Compilation on Mac OS X

You will sometimes find it necessary to use conditional compilation to make your code behave differently depending on whether certain functionality is available.

Older code sometimes used conditional statements like `#ifdef __MACH__` or `#ifdef __APPLE__` to try to determine whether it was being compiled on Mac OS X or not. While this seems appealing as a quick way of getting ported, it ultimately causes more work in the long run. For example, if you make the assumption that a particular function does not exist in Mac OS X and conditionally replace it with your own version that implements the same functionality as a wrapper around a different API, your application may no longer compile or may be less efficient if Apple adds that function in a later version.

Apart from displaying or using the name of the OS for some reason (which you can more portably obtain from the `uname(3)` API), code should never behave differently on Mac OS X merely because it is running on Mac OS X. Code should behave differently because Mac OS X behaves differently in some way—offering an additional feature, not offering functionality specific to another operating system, and so on. Thus, for maximum portability and maintainability, you should focus on that difference and make the conditional compilation dependent upon detecting the *difference* rather than dependent upon the OS itself. This not only makes it easier to maintain your code as Mac OS X evolves, but also makes it easier to port your code to other platforms that may support different but overlapping feature sets.

The most common reasons you might want to use such conditional statements are attempts to detect differences in:

- processor architecture
- byte order
- file system case sensitivity
- other file system properties
- compiler, linker, or toolchain differences
- availability of application frameworks
- availability of header files
- support for a function or feature

Instead it is better to figure out *why* your code needs to behave differently in Mac OS X, then use conditional compilation techniques that are appropriate for the *actual* root cause.

The misuse of these conditionals often causes problems. For example, if you assume that certain frameworks are present if those macros are defined, you might get compile failures when building a 64-bit executable. If you instead test for the availability of the framework, you might be able to fall back on an alternative mechanism such as X11, or you might skip building the graphical portions of the application entirely.

For example, Mac OS X provides preprocessor macros to determine the CPU architecture and byte order. These include:

- `__i386__`—Intel (32-bit)
- `__x86_64__`—Intel (64-bit)
- `__ppc__`—PowerPC (32-bit)

- `__ppc64__`—PowerPC (64-bit)
- `__BIG_ENDIAN__`—Big endian CPU
- `__LITTLE_ENDIAN__`—Little endian CPU
- `__LP64__`—The LP64 (64-bit) data model

In addition, using tools like `autoconf`, you can create arbitrary conditional compilation on nearly any practical feature of the installation, from testing to see if a file exists to seeing if you can successfully compile a piece of code.

For example, if a portion of your project requires a particular application framework, you can compile a small test program whose `main` function calls a function in that framework. If the test program compiles and links successfully, the application framework is present for the specified CPU architecture.

You can even use this technique to determine whether to include workarounds for known bugs in Apple or third-party libraries and frameworks, either by testing the versions of those frameworks or by providing a test case that reproduces the bug and checking the results.

For example, in Mac OS X, `poll(2)` does not support device files such as `/dev/tty`. If you just avoid `poll` if your code is running on Mac OS X, you are making two assumptions that you should not make:

- You are assuming that what you are doing will always be unsupported. Mac OS X is an evolving operating system that adds new features on a regular basis, so this is not necessarily a valid assumption.
- You are assuming that Mac OS X is the only platform that does not support using `poll` on device files. While this is probably true for most device files, not all device files support `poll` in all operating systems, so this is also not necessarily a valid assumption.

A better solution is to use a configuration-time test that tries to use `poll` on a device file, and if the call fails, disables the use of `poll`. If using `poll` provides some significant advantage, it may be better to perform a runtime test early in your application execution, then use `poll` only if that test succeeds. By testing for support at runtime, your application can use the `poll` API if it is supported by a particular version of any operating system, falling back gracefully if it is not supported.

A good rule is to always test for the most specific thing that is guaranteed to meet your requirements. If you need a framework, test for the framework. If you need a library, test for the library. If you need a particular compiler version, test the compiler version. And so on. By doing this, you increase your chances that your application will continue to work correctly without modification in the future.

Choosing a Default Compiler

Mac OS X ships two compilers and their corresponding toolchains. The default compiler is based on gcc 4. In addition, a compiler based on gcc 3.3 is provided. Compiling for 64-bit PowerPC and Intel-based Macintosh computers is only supported in version 4.0 and later.

Always try to compile your software using gcc 4 because future toolchains will be based on gcc version 4 or later. However, because gcc 4 is a relatively new toolchain, you may find bugs that prevent compiling certain programs.

Use of the legacy gcc 2.95.2-based toolchain is strongly discouraged unless you have to maintain compatibility with Mac OS X version 10.1.

Using the `gcc_select` Command

If you run into a problem that looks like a compiler bug, the command `sudo gcc_select 3.5` changes the default compiler to gcc 3.5. You can change it back at any time by typing `sudo gcc_select 4.0`. Do this only as a last resort, though. If possible, you should make whatever code and compiler flag changes are needed to get gcc 4.0 to compile your application.

Setting Compiler Flags

When building your projects in Mac OS X, simply supplying or modifying the compiler flags of a few key options is all you need to do to port most programs. These are usually specified by either the `CFLAGS` or `LDFLAGS` variable in your makefile, depending on which part of the compiler chain interprets the flags. Unless otherwise specified, you should add these flags to `CFLAGS` if needed.

Note: The 64-bit toolchain in Mac OS X v10.4 and later has additional compiler flags (and a few deprecated flags). These are described in more detail in *64-Bit Transition Guide*.

Some common flags include:

`-flat_namespace` (in `LDFLAGS`)

Changes from a two-level to a single-level (flat) namespace. By default, Mac OS X builds libraries and applications with a two-level namespace where references to dynamic libraries are resolved to a definition in a specific dynamic library when the image is built. Use of this flag is generally discouraged, but in some cases, is unavoidable. For more information, see [“Understanding Two-Level Namespaces”](#) (page 29).

`-bundle` (in `LDFLAGS`)

Produces a Mach-O bundle format file, which is used for creating loadable plug-ins. See the `ld` man page for more discussion of this flag.

`-bundle_loader executable` (in `LDFLAGS`)

Specifies which executable will load a plug-in. Undefined symbols in that bundle are checked against the specified executable as if it were another dynamic library, thus ensuring that the bundle will actually be loadable without missing symbols.

`-framework framework` (in `LDFLAGS`)

Links the executable being built against the listed framework. For example, you might add `-framework vecLib` to include support for vector math.

`-mmacosx-version-min version`

Specifies the version of Mac OS X you are targeting. You must target your compile for the *oldest* version of Mac OS X on which you want to run the executable. In addition, you should install and use the cross-development SDK for that version of Mac OS X. For more information, see *SDK Compatibility Guide*.

More extensive discussion for the compiler in general can be found at <http://developer.apple.com/releasesnotes/DeveloperTools/>.

Understanding Two-Level Namespaces

By default, Mac OS X builds libraries and applications with a two-level namespace. In a two-level namespace environment, when you compile a new dynamic library, any references that the library might make to other dynamic libraries are resolved to a definition in those *specific* dynamic libraries.

The two-level namespace design has many advantages for Carbon applications. However, it can cause problems for many traditional UNIX applications if they were designed to work in a flat namespace environment.

For example, suppose one library, call it `libfoo`, uses another library, `libbar`, for its implementation of the function `barIt`. Now suppose an application wants to override the use of `libbar` with a compressed version, called `libzbar`. Since `libfoo` was linked against `libbar` at compile time, this is not possible without recompiling `libfoo`.

To allow the application to override references made by `libfoo` to `libbar`, you would use the flag `-flat_namespace`. The `ld` man page has a more detailed discussion of this flag.

If you are writing libraries from scratch, it may be worth considering the two-level namespace issue in your design. If you expect that someone may want to override your library's use of another library, you might have an initializer routine that takes pointers to the second library as its arguments, and then use those pointers for the calls instead of calling the second library directly.

Alternately, you might use a plug-in architecture in which the calls to the outside library are made from a plug-in that could be easily replaced with a different plug-in for a different outside library. See [“Dynamic Libraries and Plug-ins”](#) (page 29) for more information.

For the most part, however, unless you are designing a library from scratch, it is not practical to avoid using `-flat_namespace` if you need to override a library's references to another library.

Executable Format

The only executable format that the Mac OS X kernel understands is the Mach-O format. Some bridging tools are provided for classic Macintosh executable formats, but Mach-O is the native format. It is very different from the commonly used Executable and Linking Format (ELF). For more information on Mach-O, see *Mac OS X ABI Mach-O File Format Reference*.

Dynamic Libraries and Plug-ins

Dynamic libraries and plug-ins behave differently in Mac OS X than in other operating systems. This section explains some of those differences.

Using Dynamic Libraries at Link Time

When linking an executable, Mac OS X treats dynamic libraries just like libraries in any other UNIX-based or UNIX-like operating system. If you have a library called `libmytool.a`, `libmytool.dylib`, or `libmytool.so`, for example, all you have to do is this:

```
ld a.o b.o c.o ... -L/path/to/lib -lmytool
```

As a general rule, you should avoid creating static libraries (`.a`) except as a temporary side product of building an application. You *must* run `ranlib` on any archive file before you attempt to link against it.

Using Dynamic Libraries Programmatically

Mac OS X makes heavy use of dynamically linked code. Unlike other binary formats such as ELF and xcoff, Mach-O treats plug-ins differently than it treats shared libraries.

The preferred mechanism for dynamic loading of shared code, beginning in Mac OS X v10.4 and later, is the `dlopen` family of functions. These are described in the man page for `dlopen`. The `ld` and `dyld` man pages give more specific details of the dynamic linker's implementation.

Note: By default, the names of dynamic libraries in Mac OS X end in `.dylib` instead of `.so`. You should be aware of this when writing code to load shared code in Mac OS X.

Libraries that you are familiar with from other UNIX-based systems might not be in the same location in Mac OS X. This is because Mac OS X has a single dynamically loadable framework, `libSystem`, that contains much of the core system functionality. This single module provides the standard C runtime environment, input/output routines, math libraries, and most of the normal functionality required by command-line applications and network services.

The `libSystem` library also includes functions that you would normally expect to find in `libc` and `libm`, RPC services, and a name resolver. Because `libSystem` is automatically linked into your application, you do not need to explicitly add it to the compiler's link line. For your convenience, many of these libraries exist as symbolic links to `libSystem`, so while explicitly linking against `-lm` (for example) is not needed, it will not cause an error.

To learn more about how to use dynamic libraries, see *Dynamic Library Programming Topics*.

Compiling Dynamic Libraries and Plugins

For the most part, you can treat dynamic libraries and plugins the same way as on any other platform if you use GNU `libtool`. This tool is installed in Mac OS X as `glibtool` to avoid a name conflict with NeXT `libtool`. For more information, see the manual page for `glibtool`.

You can also create dynamic libraries and plugins manually if desired. As mentioned in “[Using Dynamic Libraries Programmatically](#)” (page 30), dynamic libraries and plugins are not the same thing in Mac OS X. Thus, you must pass different flags when you create them.

To create dynamic libraries in Mac OS X, pass the `-dynamiclib` flag.

To create plugins, pass the `-bundle` flag.

Because plugins can be tailored to a particular application, the Mac OS X compiler provides the ability to check these plugins for loadability at compile time. To take advantage of this feature, use the `-bundle_loader` flag. For example:

```
gcc -bundle a.o b.o c.o -o mybundle.bundle \  
    -bundle_loader /Applications/MyApp.app/Contents/MacOS/MyApp
```

If the compiler finds symbol requests in the plugin that cannot be resolved in the application, you will get a link error. This means that you must use the `-l` flag to link against any libraries that the plugin requires as though you were building a complete executable.

Important: Mac OS X does not support the concept of weak linking as it is found in systems like Linux. If you override one symbol, you must override all of the symbols in that object file.

To learn more about how to create and use dynamic libraries, see *Dynamic Library Programming Topics*.

Bundles

In the Mac OS X file system, some directories store executable code and the software resources related to that code in one discrete package. These packages, known as bundles, come in two varieties: application bundles and frameworks.

There are two basic types of bundles that you should be familiar with during the basic porting process: application bundles and frameworks. In particular, you should be aware of how to use frameworks, since you may need to link against the contents of a framework when porting your application.

Application Bundles

Application bundles are special directories that appear in the Finder as a single entity. Having only one item allows a user to double-click it to get the application with all of its supporting resources. If you are building Mac OS X applications, you should make application bundles. Xcode builds them by default if you select one of the application project types. More information on application bundles is available in [“Bundles vs. Installers”](#) (page 57) and in *Mac OS X Technology Overview*.

Frameworks

A framework is a type of bundle that packages a shared library with the resources that the library requires. Depending on the library, this bundle could include header files, images, and reference documentation. If you are trying to maintain cross-platform compatibility, you may not want to create your own frameworks, but you should be aware of them because you might need to link against them. For example, you might want to link against the Core Foundation framework. Since a framework is just one form of a bundle, you can do this by linking against `/System/Library/Frameworks/CoreFoundation.framework` with the `-framework` flag. A more thorough discussion of frameworks is in *Mac OS X Technology Overview*.

For More Information

You can find additional information about bundles in *Mac OS X Technology Overview* and in *Creating Bundles*, available from Apple's Technical Publications website.

Using Mac OS X Native Application Environments

Mac OS X includes three high-level native development environments that you can use for your application's graphical user interface: Carbon, Cocoa, and Java. These environments are full-featured development environments in their own right, and you can write complete applications in any one of these environments.

In addition to these technologies, you can also use OpenGL, X11, Qt, Tcl/Tk, wxWidgets, and a number of other traditional UNIX graphics technologies when writing applications for Mac OS X. In each case, there are tradeoffs. This chapter explains these tradeoffs and shows how they affect you as a programmer as well as your users.

Choosing a Native Application Environment

In the context of this document, Carbon, Cocoa, and Java are presented as a way to write a GUI wrapper for a UNIX-based tool.

Writing to these environments lets you build an application that is indistinguishable from any other native Mac OS X application. The Java or Cocoa frameworks are probably the environments that you will use in bringing UNIX applications to Mac OS X, although the Carbon frameworks are used by some developers. The following sections outline all three.

A good general rule is that if you decide to port your application from X11 to a native GUI environment, Carbon is a good choice for C and other procedural X11 applications, while Cocoa is a good choice if your application uses high-level toolkits such as Tcl/Tk or Qt. Cocoa is also a good choice for wrapping command-line applications with a GUI.

Figure 4-1 High-level graphics technologies

Aqua		AppleScript	
Cocoa	Java	Carbon	
Quartz	OpenGL	QuickTime	
Darwin			

If you decide to port your cross-platform application to a native Mac OS X GUI environment, you should carefully consider how best to do this port in a way that is maintainable. For some helpful tips, you should also read [“\(Re\)designing for Portability”](#) (page 69).

Cocoa

Cocoa is an object-oriented framework that incorporates many of Mac OS X's greatest strengths. Based on the highly-respected OpenStep frameworks, it allows for rapid development and deployment with both its object-oriented design and integration with the Mac OS X development tools. Cocoa is divided into two major parts: Foundation and the Application Kit. Foundation provides the fundamental classes that define data types and collections; it also provides classes to access basic system information like dates and communication ports. The Application Kit builds on that by giving you the classes you need to implement graphical event-driven user interfaces.

Cocoa also provides file system abstraction that makes things such as file browsers look like Mac users expect. If you are developing a commercial application, you should use these. If you are developing an in-house application that will largely be used by people familiar with UNIX, the layout may be confusing, so these APIs may not be appropriate. See “[Presenting File Open and Save Dialog Boxes](#)” (page 53) for more information.

The native language for Cocoa is Objective-C, which provides object-oriented extensions to standard C and Objective-C++. *The Objective-C Programming Language* describes the grammar of Objective-C and presents the concepts behind it. Objective-C is supported in gcc 2.95 and 3. Most of the Cocoa API is also accessible through Java. Cocoa is also the basis for AppleScript Studio, an application environment for script-based GUI development.

Additional Cocoa information, including sample code, can be found at <http://developer.apple.com/cocoa>. Cocoa documentation including tutorials is available at <http://developer.apple.com/documentation/Cocoa/Cocoa.html>.

Benefits of Cocoa Development

- Rapid development environment
- Object-oriented framework design
- Excellent integration with Mac OS X developer tools
- Very robust feature set
- Can take advantage of existing C, C++, Objective-C, and Java code
- Similar level of abstraction to high level toolkits such as Qt, Tcl/Tk, and so on

Drawbacks of Cocoa Development

- Cross-platform deployment requires having a separate, non-Cocoa code base for the GUI portion
- Requires integrating Java, Objective C, or Objective C++ into your code base for the GUI layer
- Requires knowledge of Java, Objective C, or Objective C++
- Very different level of abstraction from raw X11 programming
- Potential performance penalties if used incorrectly

Example: Calling C or C++ Code With Cocoa

When designing an application from the ground up in Cocoa, you are in an object-oriented environment. You can also take advantage of Cocoa's object-oriented nature when converting preexisting code, as you can use the Cocoa frameworks to wrap the functionality of C or C++ code.

The Objective-C language has been extended to understand C++. Often this is called Objective-C++, but the functionality remains basically the same. Because Cocoa understands Objective-C++, you can call native C and C++ code from Cocoa. This is one example of how you can take advantage of your code base while adding a Macintosh front end. An example is provided below of how Objective-C wraps together C, C++, and Objective-C++ functionality. [Listing 4-1](#) (page 35) shows the Objective-C main class.

Listing 4-1 main.m

```
#import <AppKit/AppKit.h>

int main(int argc, const char *argv[]) {
    return NSApplicationMain(argc, argv);
}
```

This gets everything started when the user double-clicks the application icon. A call is then sent to invoke a `HelloController` object by the NIB, a file that holds interface information. The listings for `HelloController.m` and `HelloController.h` follow.

Listing 4-2 HelloController.m

```
#import "HelloController.h"

@implementation HelloController

- (void)doAbout:(id)sender
{
    NSRunAlertPanel(@"About",@"Welcome to Hello World!",@"OK",NULL,NULL);
}

- (IBAction)switchMessage:(id)sender
{
    int which=[sender selectedRow]+1;
    [helloButton setAction:NSSelectorFromString([NSString stringWithFormat:@"%@@%d:",@"message",which])];
}

- (void)awakeFromNib
{
    [[helloButton window] makeKeyAndOrderFront:self];
}

@end
```

Listing 4-3 HelloController.h

```
#import <AppKit/AppKit.h>

@interface HelloController : NSObject
{
```

```

        id helloButton;
        id messageRadio;
    }
    - (void)doAbout:(id)sender;
    - (void)switchMessage:(id)sender;
@end

```

The communication between the C, C++, and the Objective-C code is handled as shown in [Listing 4-4](#) (page 36). The header file `SayHello.h` is shown in [Listing 4-5](#) (page 36).

Listing 4-4 SayHello.mm

```

#import "SayHello.h"
#include "FooClass.h"
#include <Carbon/Carbon.h>

@implementation SayHello

- (void)message1:(id)sender
{
    NSRunAlertPanel(@"Regular Obj-C from Obj-C",@"Hello, World! This is a regular
old NSRunAlertPanel call in Cocoa!",@"OK",NULL,NULL);
}

- (void)message2:(id)sender
{
    int howMany;
    NSString *theAnswer;
    Foo* myCPlusPlusObj;
    myCPlusPlusObj=new Foo();
    howMany=myCPlusPlusObj->getVariable();
    delete myCPlusPlusObj;
    theAnswer=[NSString stringWithFormat:@"Hello, World! When our C++ object is
queried, it tells us that the number is %i!",howMany];
    NSRunAlertPanel(@"C++ from Obj-C",theAnswer,@"OK",NULL,NULL);
}

- (void)message3:(id)sender
{
    Alert(128,NULL); //This calls into Carbon
}

@end

```

Listing 4-5 SayHello.h

```

#import <AppKit/AppKit.h>

@interface SayHello : NSObject
{
}

- (void)message1:(id)sender;
- (void)message2:(id)sender;
- (void)message3:(id)sender;
@end

```

The C++ class wrapped by these Cocoa calls is shown in [Listing 4-6](#) (page 37). The header file, `FooClass.h`, is shown in [Listing 4-7](#) (page 37).

Listing 4-6 `FooClass.cpp`

```
#include "FooClass.h"
Foo::Foo()
{
    variable=3;
}

int Foo::getVariable()
{
    return variable;
}
```

Listing 4-7 `FooClass.h`

```
class Foo {
public:
    Foo();
    int getVariable();
    void * objcObject;
private:
    int variable;
};
```

You should be careful when writing code using Cocoa, because the same constructs that make it easy for the developer tend to degrade performance when overused. In particular, heavy use of message passing can result in a sluggish application.

The ideal use of Cocoa is as a thin layer on top of an existing application. Such a design gives you not only good performance and ease of GUI design, but also minimizes the amount of divergence between your UNIX code base and your Mac OS X code base.

Java

Mac OS X includes a complete implementation of Java 2 Platform Standard Edition with both the 1.4.x JDK and 1.4.x SDK. Pure Java development in Mac OS works just as it would on other platforms. More information on Java development in Mac OS X can be found online at <http://developer.apple.com/java>.

Benefits of Java Development

- Installed by default with Mac OS X
- Swing elements are displayed with the native Aqua look and feel of Mac OS X
- Cross-platform deployment

Drawbacks of Java Development

- Usually not as fast as a C or C++ implementation

- Complexities of communicating between two different languages if your code is C-based
- Limited access to hardware and OS-specific capabilities

Carbon

Carbon is an environment designed to bring existing Mac OS applications to Mac OS X. It can be used to bring an X11 application to a native Mac OS X environment, since the basic drawing primitives are at a similar level of abstraction.

Carbon also offers relatively straightforward file I/O with a few enhancements that are unavailable through POSIX APIs, such as aliases (which behave like a cross between a symbolic link and a hard link). Carbon presents the file system differently to the user than UNIX applications do. This is the layout that Mac users are used to seeing, and thus you should use it if possible when developing a commercial application for broad-scale deployment.

However, if your users are predominantly UNIX users (for example, if you are an in-house developer in a corporate environment), the use of the Carbon file API may not be appropriate, since volumes appear as they do in the Finder, rather than being organized according to the file system hierarchy.

Benefits of Carbon Development

- Well-documented feature set
- Integration with Mac OS X developer tools
- Very robust feature set
- Simple C and C++ integration
- Similar abstraction level to X11
- Procedural design rather than object-oriented

Drawbacks of Carbon Development

- No cross-platform deployment without a separate non-Carbon code base
- Slightly more effort required to take advantage of native Mac OS X technologies
- Procedural design rather than object-oriented

In case you're wondering, that isn't a typo. Procedural design in a GUI can be both a benefit and a drawback.

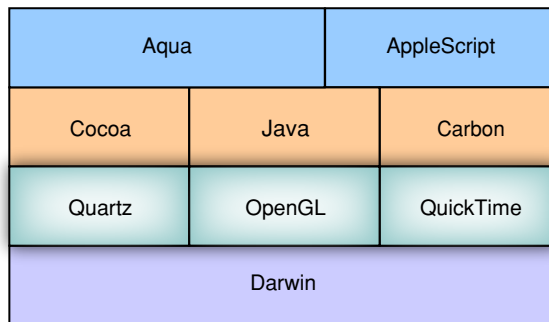
Procedural design can be a benefit in terms of being able to easily integrate into any environment, regardless of language. It is also a benefit because it fits many styles of graphical programming well.

Procedural design can be a drawback, however, when dealing with more complex GUI designs where it would be more convenient to have the methods more closely tied to the data on which they operate.

Lower-Level Graphics Technologies

With some code it is simple to abstract the display code from the underlying computational engine, but often this isn't the case. Especially for graphics-intensive applications, you may want to take advantage of directly calling the core graphic functionality of your targeted operating system. You still need to wrap this functionality in a higher-level API to display it to the screen and allow user interaction, but for times where you need to push pixels in very specific ways, Mac OS X offers you access to three first-class graphics technologies: Quartz, OpenGL, and QuickTime.

Figure 4-2 Low-level graphics technologies



Quartz

Quartz is the graphical system that forms the foundation of the imaging model for Mac OS X. Quartz gives you a standards-based drawing model and output format. Quartz provides both a two-dimensional drawing engine and the Mac OS X windowing environment. Its drawing engine leverages the Portable Document Format (PDF) drawing model to provide professional-strength drawing. The windowing services provide low-level functionality such as window buffering and event handling as well as translucency. Quartz is covered in more detail in *Mac OS X Technology Overview* and in Apple's Quartz website (<http://developer.apple.com/quartz/>).

Benefits of using Quartz

- PostScript-like drawing features
- PDF-based
- Included color management tools
- Unified print and imaging model

Drawbacks to using Quartz

- No cross-platform deployment

OpenGL

OpenGL is an industry-standard 2D and 3D graphics technology. It provides functionality for rendering, texture mapping, special effects, and other visualization functions. It is a fundamental part of Mac OS X and is implemented on many other platforms.

Given its integration into many modern graphics chipsets and video cards, it is of special interest for programs that require intricate graphic manipulation. OpenGL's homepage gives more information on the technology in general at <http://www.opengl.org>.

Apple's OpenGL page, at <http://developer.apple.com/opengl/>, gives more information on how it is integrated into Mac OS X and explains the enhancements that the Mac OS X OpenGL implementation provides.

Benefits of using OpenGL

- Cross-platform technology
- Native Mac OS X integration
- Included with every installation of Mac OS X
- Very robust feature set for handling graphics

Drawbacks to using OpenGL

- Some level of integration at an application level is required

QuickTime

QuickTime is a powerful multimedia technology for manipulating, enhancing, storing, and delivering graphics, video, and sound. It is a cross-platform technology that provides delivery on Mac OS as well as Windows. More information on the technology can be found at <http://developer.apple.com/quicktime/>.

Benefits of using QuickTime

- Robust feature set for manipulating sound and video
- Cross-platform development

Drawbacks to using QuickTime

- Not supported on other UNIX-based systems.

Note that while QuickTime itself is not supported on UNIX-based systems, the QuickTime file format is supported by various third-party utilities.

Choosing a Graphical Environment for Your Application

Mac OS X offers many options for transforming your applications with a graphical user interface from a UNIX code base to a native Mac OS X code base, or even for wrapping preexisting command-line tools or utilities with a graphical front end, making them available to users who never want to go to the command line.

This chapter describes some of the issues you will face when porting a GUI application to Mac OS X or adding a GUI wrapper around a command line application. It also describes the various GUI environments supported by Mac OS X and gives advantages and disadvantages of each.

In choosing a graphical environment to use in bringing a UNIX-based application to Mac OS X, you will need to answer the questions posed in the following sections:

- [“What Kind of Application Are You Porting?”](#) (page 41)
- [“How Well Does It Need to Integrate With Mac OS X?”](#) (page 42)
- [“Does Your Application Require Cross-Platform Functionality?”](#) (page 42)

These questions should all be evaluated as you weigh the costs and benefits of each environment. You may already be using a cross-platform toolkit API. If you aren't doing so, you may want to port your application to a native API such as Carbon or Cocoa.

If you are a commercial developer adding a new graphical user interface to a command-line application and want to take advantage of the greatest strengths of Mac OS X, you will probably want to use the Cocoa API. In some cases, you may want to use a different API for reasons such as cross-platform compatibility.

If you decide to use a nonnative API, like X11R6, to provide a user interface for your Mac OS X application, it is important to remember that users and developers with a UNIX background might be perfectly content to just have the application running in Mac OS X. Traditional Macintosh users, however, will pass up an application with a traditional UNIX interface for a more integrated and modern interface. Whether you target a straight port to the Darwin layer or a more robust transformation of your application to take advantage of other Mac OS X technologies (like the Cocoa frameworks) is your decision.

What Kind of Application Are You Porting?

Are you bringing a preexisting code base to Mac OS X, or are you adding new functionality—for example a graphical interface—to a command-line application? If you already have a code base written to a particular API, and that API is supported in Mac OS X, you probably want to continue using that API for any large, complex application unless you desire features of another API.

For simple applications, or for applications where you are wrapping a command-line utility with a graphical user interface, you need to evaluate what API to use. Reading the next two sections will help you recognize the benefits and drawbacks of each technology.

How Well Does It Need to Integrate With Mac OS X?

Who are you marketing your application to? If they are traditional UNIX users that just want to run, for example, a gene-sequencing application alongside Microsoft Office, then it may be sufficient to just install an X Window System on their Mac OS X computer. You would simply port your X11R6-based application to Mac OS X, leaving your code as it stands (aside from the little changes you may need to make to make it compile in Mac OS X). For more information on X11 in Mac OS X, see “X11R6” (page 46).

If you are an in-house developer of UNIX applications, this may be as far as you want to go, particularly if you want to maintain the same user experience across multiple platforms. However, you may still want to use Carbon or Cocoa APIs to improve the overall look of the UI to make it easier to use.

If you sell an application, some customers might at first be happy just to have it on their platform. However, if a competing product is released using Mac OS X native functionality, customers are likely to gravitate to that product.


A hot topic in the science and technology industries is not only bringing a code base to Mac OS X, but also giving that application a Mac OS X native user interface. This is not a decision to be made trivially for an application with a complex GUI and a large code base, but it is one that can make or break a product’s success in the market. The individual discussions of the APIs that follow should help you to make a well-informed decision.

If you’re an open source software developer, you will probably gravitate toward a basic port of the X11 application. However, you should consider creating a native GUI if your application is likely to be used by consumers such as a word processor or a web browser.

Does Your Application Require Cross-Platform Functionality?

If you have an application that must work on multiple platforms, Mac OS X has you set up for success. You have many options; some are built in and shipped with every version of the operating system; others require the installation of additional components. [Figure 5-1](#) (page 42) depicts the distinction between the cross-platform APIs that are native and those that aren’t.

Figure 5-1 Graphical environments



You can see that Mac OS X includes some standard cross-platform APIs: Java, OpenGL, and QuickTime. There are also commercial and free implementations of some of the traditional UNIX technologies. If you are building a cross-platform application, you should evaluate which platforms you are targeting with your application and determine which API allows you to bring your UNIX-based application to Mac OS X. [Table 5-1](#) (page 43) lists the platforms on which Mac OS X cross-platform technologies run.

Table 5-1 Platforms of cross-platform technologies

API	Platforms
"OpenGL" (page 40)	Mac OS X, UNIX-based systems, Windows
"Java" (page 37)	Mac OS X, UNIX-based systems, Windows
"QuickTime" (page 40)	Mac OS X, Windows
"Qt" (page 47)	Mac OS X (Native & X11), UNIX-based systems, Windows
"Tcl/Tk" (page 47)	Mac OS X (Native & X11), UNIX-based systems, Windows
"wxWidgets" (page 48)	Mac OS X (Native & X11), UNIX-based systems, Windows
"X11R6" (page 46)	Mac OS X, UNIX-based systems

Note: Although some of the technologies in [Table 5-1](#) (page 43) are supported in Mac OS 9, support for these technologies in Mac OS 9 is not considered here because the rest of your UNIX code base will not run in Mac OS 9.

In the next two chapters, you'll find brief descriptions of each of the technologies available to you for your application's graphical user interface.

Using Traditional UNIX Graphical Environments

UNIX-based operating systems have grown to include many environments for providing a graphical interface to users. The X Window System, also known as X11, is probably the most well known. As more specific needs have arisen, other architectures have been developed that assume an X11 implementation.

Mac OS X does not use the X Window System by default, but implementations of X11 are available from Apple and from third-party developers. This means that X Window System–based applications can be run, as can many of the alternative UNIX-style graphical environments. This section gives more details on some of these environments.

If your application uses the X Window System you need to either port the user interface to a native Mac OS X environment, provide an X Window System implementation with your application, or require that your users install the Apple X11 package (an installation option in the Mac OS X installer beginning in Mac OS X version 10.3).

If you are a commercial software developer or if your software will be used by traditional end users (including word processors, web browsers, and so on), you may want to consider porting to a native graphical environment instead of requiring X11. However, for internal tools or for tools with a specialized user base, it may be more reasonable to continue using X11. The sections in this chapter will help you decide how to choose an appropriate X11-based environment for your application.

Choosing a UNIX Graphical Environment

Many UNIX applications are designed using high-level widget toolkits such as Tcl/Tk, Motif (or OpenMotif or Lesstif), and Qt. These environments work well across many platforms. However, that cross-platform nature comes with a price in flexibility. In essence, to be cross-platform, they can only support capabilities that are generic to all of the potential operating environments, and as such, they tend to provide only the lowest common denominator in terms of functionality.

Such toolkits are often weak in areas such as performance and flexibility, and many do not allow precise control over how GUI elements are laid out.

If they provide everything you need, then they are good choices. Otherwise, a more flexible alternative such as raw Xlib programming, or a more Mac-tuned alternative such as Carbon or Cocoa may be more appropriate.

Other UNIX solutions exist for more extreme cross-platform environments. One example is MicroWindows, which is designed to allow ease of code sharing between X11 and Windows by implementing a subset of the Win32 graphics APIs on X11. Because this is really just an X11 application, it should be possible to use it on Mac OS X. However, due to the amount of abstraction involved, such a solution will not perform as well as a native application.

If you are trying to design an application for such an environment, it is generally better to rearchitect your code to have multiple front-ends—one for Windows, one for Mac OS X, and one for UNIX (X11). See [“\(Re\)designing for Portability”](#) (page 69) for more information on creating a graphics abstraction layer.

X11R6

Mac OS X does not include an X11R6 implementation by default. If you are intent on only porting your application without adding any Mac OS X functionality, you can still run X11 on Mac OS X. Along with your application, you should instruct your users on how to install one.

The easiest way to get an X11 implementation is to install the Apple X11 implementation. Beginning in Mac OS X v10.3, the X11 package is an optional installation that comes on the install DVD or CD.

For *older versions* of Mac OS X (before v10.4), you can download X11 from http://www.apple.com/downloads/macosx/apple/macosx_updates/x11formacosx.html.

Important: Do *not* install the downloadable version of X11 on Mac OS X v10.4. It will not work. The Apple X11 version compatible with v10.4 is available *only* on the Mac OS X install DVD.

In addition, several other commercial and free alternatives exist.

XTools by Tenon Intersystems (<http://www.tenon.com>) and eXodus from Powerlan USA both provide X11 servers that coexist with Aqua.

If you do not need a commercial implementation of X11, XFree86 offers a very robust free X11R6 implementation. The Mac OS X version of XFree86 is an active project with integral support from the XDarwin project. More information on the XDarwin project including tips for installing X11 is available at <http://www.xdarwin.org>. XFree86 itself can be downloaded from <http://xfree86.org/>.

With an X Window System implementation, you are now free to use toolkits that you are already familiar with—for example, GTK+ or KDE.

Benefits of X11R6 Development

- The de facto standard in display technology for UNIX-based operating systems
- Open source implementation is available
- Somewhat portable to certain embedded systems

Drawbacks of X11R6 Development

- Complicated development environment
- Requires substantial disk space for a full installation
- Is not native to Mac OS X. This means your application is treated as a second-class citizen to some extent, particularly in the areas of drag-and-drop and system services.

Tcl/Tk

There is a native Aqua version of Tk available at <http://tcl.sourceforge.net>. With this you can easily add graphical elements to your existing Perl, Python, or Tcl scripts.

Benefits of Tk Development

- Cross-platform development environment
- Easily integrates with Tcl, Perl, and Python scripts
- Open source implementation is available

Drawbacks of Tk Development

- Not supported by default in Mac OS X
- As with most high-level toolkits, flexibility is limited

Qt

Qt is a C++ toolkit for building graphical user interfaces. It is very popular in the UNIX world, as it is used by the very popular K Desktop Environment, KDE. Trolltech has a native version of Qt, Qt/Mac, available for Mac OS X. If you already have a C++ application or are considering building one, Qt lets you build applications that run natively in Mac OS X as well as Linux and Windows. Information about Qt/Mac is available at <http://www.trolltech.com>. Qt/Mac does not run on top of X11 in Mac OS X, but the source code is compatible with Qt's X11 implementation.

Benefits of Qt Development

- Cross-platform development environment
- Integrates easily with C++ code
- Robust feature set

Drawbacks in Qt Development

- Requires licensing for commercial development
- As with most high-level toolkits, flexibility is limited

GTK+/GTKmm

GTK+ and GTKmm are the C and C++ interfaces (respectively) to another toolkit for building graphical user interfaces. It is popular in the Gimp and Gnome development communities. There is a beta native version of the GTK+ 2.x toolkits available for Mac OS X. If you already have a C or C++ application or are considering building one, GTK+ lets you build applications that run natively in Mac OS X as well as Linux and Windows.

Information about Gtk-MacOSX (an implementation of the GTK+ 2.x toolkit) is available at <http://developer.imendio.com/projects/gtk-macosx>.

These toolkits do not run on top of X11 in Mac OS X, but are source-code-compatible with the equivalent X11 implementations.

Benefits of GTK+/GTKmm Development

- Cross-platform development environment
- Integrates easily with C/C++ code
- Robust feature set

Drawbacks in GTK+/GTKmm Development

- LGPL license reverse engineering clause may raise concerns for some commercial developers
- As with most high-level toolkits, flexibility is limited

wxWidgets

The wxWidgets toolkit is another toolkit for building graphical user interfaces. It is a C++ framework for building native look-and-feel applications in a cross-platform way. It currently supports UNIX/Linux (X11), Windows, and Mac OS X. Support for other platforms is in development.

More information about wxWidgets can be found at <http://www.wxwidgets.org>.

Wrapping a Command Line Application with a GUI Interface

Command line applications are a good first step in writing software for Mac OS X. However, most Mac OS X users tend to shy away from the command line (former UNIX and Linux users notwithstanding). For this reason, if you want your command line application to have mass appeal, you should consider wrapping it in a GUI interface.

Mac OS X provides a number of methods for wrapping a command line application with a GUI interface. These include:

- Do Shell Script (AppleScript Studio)
- NSTask (Cocoa)

Each provides different capabilities for different purposes. This section points you to existing documentation on the topic and describes common issues that arise when wrapping a command line application with a GUI interface.

This section is primarily intended for open source developers, shareware/freeware developers, and in-house UNIX application developers, as these techniques do not offer the same level of functionality that you can get through tighter integration between the GUI and the underlying application.

Choosing a Wrapping Method

If you just need to run a shell script by double-clicking it, and have no need for a GUI, you can simply rename the shell script to end with `.command` and it will open in Terminal automatically.

If you need more control over arguments but the command-line tool is non-interactive—that is, the GUI wrapper will merely start it and display the results—then either `NSTask` or "Do Shell Script" is acceptable, and you should use whichever one is most comfortable.

If the application needs to interact with the tool in any way, you should use the function `NSTask`. With it, you can chain arbitrary numbers of tasks using the `NSPipe` function.

File Access Considerations

Carbon, Cocoa, AppleScript, and other Mac OS X development environments deal with files by using file references rather than referring to them by their path. As a result, you can move files anywhere on the disk and the file reference will still be valid even though the path has changed.

Using file references presents two problems for the developer. First, you have to explicitly convert them to POSIX paths to use them as the argument to a shell script. Second, the POSIX paths can contain spaces, so you must construct the command line carefully to prevent the script from failing if a directory name (or the name of your hard drive) contains a space.

Because of the potential side effects, you should always use the quoted form of the POSIX path when passing the path to a shell script. This is described in detail in *Technote #2065: Issuing Commands*.

The issue of spaces in pathnames is also a slight issue for Cocoa developers using the `NSTask` API, usually because of bugs in the shell script itself. (You do not need to escape strings passed as arguments using `NSTask`.)

You should always test any GUI-wrapped command-line application with filenames that contain spaces and other strange characters such as double quotes (") and trademark (™) to make sure that there are no unexpected failures.

For More Information

For more information on AppleScript's `do shell script` command, see *Technote #2065: Issuing Commands*, at <http://developer.apple.com/technotes/tn2002/tn2065.html> or the “simple sheet” example in AppleScript Studio.

For more information on Cocoa's `NSTask` API, see *Interacting with the Operating System*.

Porting File, Device, and Network I/O

Mac OS X offers most of the standard UNIX mechanisms for file and device I/O. There are, however, differences to be aware of when porting your application to Mac OS X from other UNIX-based and UNIX-like platforms.

This chapter describes file I/O and device I/O in Mac OS X, including APIs that will enhance the user experience such as the file manager APIs for file access.

Note: This document does not cover device driver porting. For information on device driver porting, read *Porting Drivers to Mac OS X*.

If you are a commercial software developer or if your application will be used by end users, you should read this chapter.

If you are writing a port of an open source application or an in-house UNIX application, you should read this chapter only if your application already uses or plans to use alternate file APIs for other platforms or if you need to do device I/O.

How Mac OS X File I/O Works

Mac OS X contains all of the standard UNIX and POSIX file I/O functionality. For a basic port of an application to Mac OS X, you generally do not need to make any changes in this area unless your application assumes that the file system stores files in a case-sensitive manner. (HFS+, the default Mac file system, is case-preserving, but not case-sensitive.)

However, traditional Mac applications have some enhanced behavior that you may want to incorporate in your application. The most powerful of these is the use of aliases to find files when their location has changed. It is not possible to use this capability through standard UNIX APIs. To obtain this functionality, you must use either Carbon or Cocoa APIs.

For example, some Mac applications also take advantage of the HFS+ file system's ability to handle multi-forked files. This could be used, for example, to store disposable information about a file such as last window position. It is not recommended that crucial information be stored in resource forks when writing new applications. However, for compatibility reasons, you may need to access data stored in resource forks if your application needs to read files created by older Mac applications, or if you are writing a backup tool that needs to retain the entire contents of a file.

Also, Mac applications present the file system in a different way than POSIX applications in open and save dialogs. Mac applications have a directory structure multiply rooted from the individual volumes instead of singly rooted from the root volume. Using the Carbon or Cocoa APIs achieves this view of the file system automatically. These APIs also present standard file open and save dialogs that match the standard user experience that Mac users have come to expect.

Before moving to a Carbon or Cocoa file browser interface, you should consider what kinds of users are likely to use your application on the Mac platform. Are they mostly experienced in using another UNIX-based platform, or are they largely Mac users?

If they are mostly UNIX users, the Mac file dialog may be confusing to them. If they are mostly Mac users, a traditional UNIX file dialog may be equally confusing. You should generally choose which file API to use based on long-term expectations rather than the current user base to avoid problems down the line.

Note: Before considering the use of a non-POSIX file API such as those in Carbon or Cocoa, you should read the chapter “(Re)designing for Portability” (page 69) and consider ways to rearchitect your file accesses to minimize the number of platform-conditional pieces of code, both for readability and for debuggability.

Using Carbon APIs for File I/O

Carbon APIs are useful when writing file I/O entirely in C. They provide very basic access much like the traditional POSIX APIs.

Table 7-1 POSIX versus Carbon APIs

POSIX function	Carbon Function	Description
open, fopen	FSOpenFork	Takes a filesystem reference (FSRef) as an argument and returns a fork reference number.
close, fclose	FSClose	Takes a fork reference number as an argument.
create	FSCreateFileUnicode	Takes the FSRef of the parent directory and the name of the file to be created.
mkdir	FSCreateDirectoryUnicode	Takes the FSRef of the parent directory and the name of the directory to be created.
rmdir	FSDeleteObject	Deletes a file or folder (by FSRef).
unlink	FSDeleteObject or FSDeleteFork	Deletes a file or folder (by FSRef) or delete a single fork of a file.

Note that these functions use `FSSpec` and `FSRef` structures rather than file names. The functions `FSRefMakePath` and `FSPathMakeRef` convert an `FSRef` to a path and vice versa. Similarly, `FSPathMakeRef` converts an `FSSpec` to an `FSRef`.

Using Cocoa APIs for File I/O

Cocoa file APIs are very different from traditional POSIX APIs. An explanation of the Cocoa APIs is beyond the scope of this book. For more information on Cocoa file APIs, see <http://developer.apple.com/reference/library/Cocoa/idxFileManagement-date.html> in Cocoa Documentation.

Presenting File Open and Save Dialog Boxes

File open and save dialog boxes can be presented in Mac OS X using Carbon or Cocoa file dialog APIs. If you choose to use the standard dialog boxes, use the same API you used for the rest of your GUI.

The Carbon file dialog API, Navigation Services, is described in *Navigation Services Reference* in Carbon Documentation.

The Cocoa file dialog API consists of the functions `NSOpenPanel` and `NSSavePanel`, for open and save dialogs. These are described in *Application File Management* in Cocoa Documentation.

How Mac OS X Device I/O Works

Mac OS X provides many of the traditional UNIX mechanisms for device I/O. However, individual device driver designs determine whether or not to use these mechanisms.

In Mac OS X, disk devices, serial port devices, the random-number generator pseudo-device, and pseudo-tty devices (pttys) are traditionally accessed through standard UNIX-style block and character devices. Other pseudo-devices (devices without hardware backing) can also be implemented as BSD devices. You can use these device files in the same way as you would use them on any other UNIX-based or UNIX-like system.

Most actual hardware devices, however, such as USB and FireWire devices are not handled with block or character devices in Mac OS X. Instead, the primary mechanism for accessing hardware devices is through I/O Kit user clients, which are basically remote procedure call or system call interfaces that allow you to call functions within a kernel driver from an application.

To access a hardware device, you must use APIs to search the device tree for a device matching various parameters. You then call functions within the driver. The mechanism used to call these functions depends on the design of the device interface, as do the functions themselves. You can also get information about the device using a standardized mechanism provided by the I/O Kit to examine its properties in the device tree.

A device cannot be controlled from an application unless there is a driver in the kernel. In many cases, the driver may be a simple pass-through that presents a device interface that allows the real device driver to be part of the application itself. In other cases, it may be a full driver for the device that presents a device interface for configuration. The device interface may be provided by Mac OS X itself (such as the user space device interface for USB HID devices) or it may be provided by the individual hardware vendor (such as an audio card driver).

The implementation of a user client is beyond the scope of this document. It is described in detail in the book *Accessing Hardware From Applications*. Since the design of a user client is also heavily dependent on the design of the device interface to which it is connecting, you may also need to read documentation specific to the technology area in question. In the case of device interfaces created by hardware vendors, you may also need to contact the hardware vendor for additional programming information.

Note: For information on porting actual device drivers, read *Porting Drivers to Mac OS X*.

File System Organization

The Mac OS X file system organization is slightly different from that of most UNIX environments, and is described in detail in the man page `hier`. This section presents only the most important differences between a typical UNIX environment and Mac OS X.

First, Mac OS X has a number of folders intended primarily for use with GUI applications or for parts of Mac OS X itself. These paths generally start with a capital letter, and include:

- `/Applications`—contains GUI applications
- `/System`—contains system frameworks and parts of the OS itself
- `/Library`—contains primarily user-installed frameworks
- `/Developer`—contains the developer tools (if installed)
- `/Users`—contains the home directories of local users

In addition, various ports collections are available, and install files in various locations, such as `/sw` (fink), `/usr/local` (GNU-Darwin), and `/opt/local/` (Darwin Ports).

A few directories, including `/etc` and `/tmp` are actually symbolic links into `/private`. You should be careful not to stomp on these symbolic links.

Finally, you should be aware that by default, a number of directories are hidden from users when viewing the file system using the Mac OS X GUI. These directories include most of the standard system directories, including `/bin`, `/dev`, `/etc`, `/sbin`, `/tmp`, and `/usr`. These directories still appear and behave in the expected manner when accessed from the command line. For more information on making these visible from GUI applications, see “[File-System Structure and Visibility](#)” (page 63).

Mac OS X has different privilege sets for file system access. Users by default have write access to their home directory and certain other directories created in previous versions of Mac OS. Admin users have write access to additional parts of the system, including various application directories and configuration files, without the need to become the root user. Finally, the directories containing the OS itself are read-only except as root.

How Mac OS Networking Works

Mac OS X provides the usual POSIX and BSD networking functionality. For more information on the actual APIs, see *Mac OS X Man Pages*.

For the most part, this networking behaves just like it does on any other UNIX system. Mac OS X differs in one crucial way, however. It does not use `/etc/hosts`, `/etc/resolv.conf`, and other similar configuration files for network configuration. (More precisely, the file `/etc/resolv.conf` is provided for read-only use, but should not be modified directly. The file `/etc/hosts` is provided, but is not used by default.)

CHAPTER 7

Porting File, Device, and Network I/O

If you need to configure name server settings (or other network settings), you should use the System Configuration framework (described in *System Configuration Programming Guidelines* and *System Configuration Framework Reference*).

If you need to add static host entries, you should use Directory Services (described in *Directory Service Framework Reference*, the `dsc1(1)` manual page, and [“Open Directory and the dscl Tool”](#) (page 65) elsewhere in this document).

Distributing Your Application

Developing an application is only part of the story. You must now get it out there for people to use. Given that this is a UNIX-based operating system, you could just `tar` and `gzip` your file. That's fine if your end users are mostly UNIX users, but this doesn't meet the needs of more general users. To complete the transition, you should package your application like other Mac OS X applications. This chapter walks you through some of those details since they are probably new to you as a UNIX developer.

This chapter is important for all non-command-line developers, whether your application is an end-user commercial suite or an open source tool.

Bundles vs. Installers

Most applications in Mac OS X do not need to use an installer. To make installation and removal easy, Mac OS X provides bundles.

A bundle is basically a directory that contains an application. Unlike normal folders, however, it appears to the user like an ordinary file. The user can double-click a bundle, and the application will launch. Since the bundle is really a directory, it can contain all the support files that are needed for the application.

The reason to use bundles is somewhat obvious if you have installed many applications in Mac OS X: applications in the form of a bundle can be installed simply by dragging the application to the destination folder in the Finder.

There are, however, some situations where bundles are problematic. An application that installs kernel extensions, startup items, system-wide preference panes, or any other system-wide resources cannot be installed by dragging the application, since those resources need to be in a separate place on the drive.

If your application requires installing a startup item, the only practical way to install your application is the use of an installer. Mac OS X makes this easy using PackageMaker. Other commercial solutions are also available from various third parties such as Stuffit InstallerMaker from Aladdin Systems and Installer VISE from MindVision.

In most cases, however, it is preferable to install system-wide components the first time your application is launched. You can do this using Authorization Services, as described in the book *Authorization Services Programming Guide*, available from the Apple Technical Publications website.

For more information about how to create a bundle, see *Bundle Programming Guide*.

Packaging Basics

There are two main applications for compressing your application: Disk Utility (or Disk Copy in older versions of Mac OS X) and PackageMaker. Disk Utility allows you to create compressed disk images, while PackageMaker creates packages that can be installed using the Mac OS X installer.

The recommended form of application distribution is a compressed disk image. A compressed disk image preserves resource forks that may be present, allows drag-and-drop installation, allows license display, and even allows encryption of data, if required.

If your application is a single application bundle, simply put it and any relevant documentation on a disk image with Disk Utility, then compress it and distribute it.

If you have an application that requires administrator privileges to install into privileged directories or requires more than a simple drag-and-drop installation, use PackageMaker (/Developer/Applications/PackageMaker) to build installer packages for Apple's Installer application.

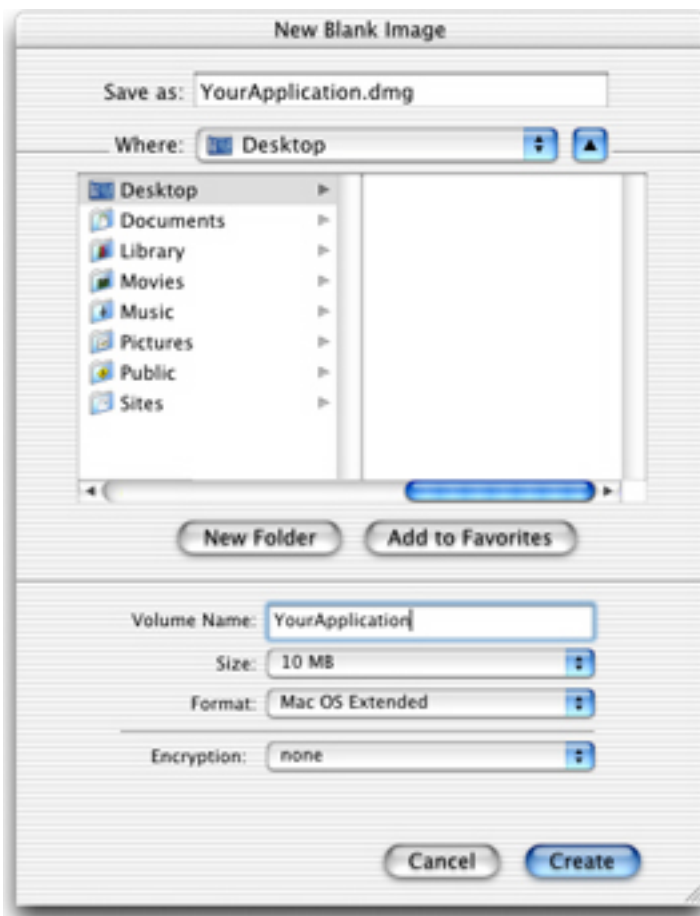
The basics of using Disk Utility to make a disk image are given in the next section. For help using PackageMaker, choose PackageMaker Help from the PackageMaker Help menu.

Disk Utility (or Disk Copy)

The following steps help you package your application as a disk image (.dmg file) for distribution on Mac OS X.

1. Open /Applications/Utilities/Disk Utility.app by double-clicking it.
2. From the Image menu, choose New Blank Image. Disk Utility opens a new window with customization options as in [Figure 8-1](#) (page 59).
3. In the "Save as" text box, enter the name of the compressed file that you will distribute. By default, a .dmg suffix is appended to the name you enter. Although it is not required, it is a good idea to retain this suffix for clarity and simplicity.
4. In the Volume Name text field, enter the name of the volume that you want to appear in the Finder of a user's computer. Usually this name is the same as the name of the compressed file without the .dmg suffix.
5. In the file browser, set the location to save the file on your computer. This has nothing to do with the installation location on the end user's computer, only where it saves it on your computer.
6. Set the Size pop-up menu to a size that is large enough to hold your application.
7. Leave the Format set to Mac OS Extended (the HFS+ file format).
8. Leave Encryption set to none. If you change it, the end user must enter a password before the image can be mounted, which is not the normal way to distribute an application.
9. Click Create.

Figure 8-1 Disk Utility options



Once you have a disk image, mount it by double-clicking it. You can now copy your files to that mounted image. When you have everything on the image that you want, you should make your image read-only. Again from Disk Utility, perform these steps:

1. Unmount the disk image by dragging the volume to the Trash, clicking the eject button next to the device in a Finder window, or selecting the mounted volume and choosing Eject from the Finder's File menu.
2. Choose Convert Image from the Image menu.
3. In the file browser, select the disk image you just modified and click Convert.
4. Choose a location to save the resulting file, change the image format to read-only, and click Convert.

You now have a disk image for your application that is easy to distribute.

Creating Disk Images Programmatically Using hdiutil

If you find yourself regularly creating disk images, you may find it helpful to automate this process. While a complete script is beyond the scope of this document, this section includes a code snippet to get you started.

For more information on hdiutil, see `hdiutil`.

Listing 8-1 Automatic Disk Image Creation using hdiutil

```
# Create an initial disk image (32 megs)
hdiutil create -size 32m -fs HFS+ -volname "My Volume" myimg.dmg

# Mount the disk image
hdiutil attach myimg.dmg

# Obtain device information
DEVS=$(hdiutil attach myimg.dmg | cut -f 1)
DEV=$(echo $DEVS | cut -f 1 -d ' ')

# Unmount the disk image
hdiutil detach $DEV

# Convert the disk image to read-only
hdiutil convert myimg.dmg -format UDZO -o myoutputimg.dmg
```

Tell the World About It

Once you have an application, how do you get the word out? First, let Apple know about it. To get your application listed on Apple's main download page for Mac OS X, go to <http://www.apple.com/downloads/macosx/submit/> and fill out the appropriate information about your application. You should also go to <http://guide.apple.com/> and at the bottom of the page, click Submit a Product to get your application listed in the Apple Guide. You might then also want to send notices to <http://www.versiontracker.com/> and Macintosh news sites like <http://www.maccentral.com/> and <http://www.macnn.com/>.

Additional Features

Although many parts of Mac OS X are the same as other UNIX-based operating systems, Mac OS X also includes many things that set it apart. This chapter lists some of the key details that distinguish Mac OS X from most other UNIX-based operating systems. These may not be important for a basic port of a simple application, but the more robust your application and the more tightly it integrates with the underlying operating system, the more important it is to understand the additional functionality provided by the operating system. Most of the information here is covered only as an overview, with references to more detailed documentation where appropriate.

This chapter is of interest to all varieties of developers, though different developers will have different interests, depending on the nature of your application.

AppleScript and AppleScript Studio

AppleScript is a technology that does for GUI applications what shell scripts do for command-line applications. The scripts control applications using a relatively straightforward messaging protocol (whose commands are application-dependent).

AppleScript Studio is a suite of tools—Xcode, Interface Builder, and the Cocoa application framework—that together allow you to create whole applications written entirely in AppleScript.

For more information on AppleScript and AppleScript Studio, see the AppleScript documentation at <http://developer.apple.com/documentation/ScriptingAutomation>.

Audio Architecture

Mac OS X provides much of the audio functionality normally associated with third-party MIDI and audio toolkits. This gives developers a simple platform for developing dynamic audio applications. For end users, it minimizes the configuration that is normally required to get high-end audio applications to work. As a UNIX developer, it means that if you have been looking for a platform to develop a robust audio application on, you now have one. Among the high-level features of the Mac OS X Core Audio subsystem are these:

- Native multi-channel audio with plug-in support
- Native MIDI
- Audio Hardware Abstraction Layer (HAL)
- A built-in USB class driver compliant with the USB audio specification
- A simplified driver model
- A direct relation with the I/O Kit through the `IOAudioDevice` class that enables rapid device-driver development

If you are developing applications that need access to the audio layer of Mac OS X, you can pursue the extensive resources available at <http://developer.apple.com/audio/>.

Boot Sequence

The boot sequence in Mac OS X offers a number of advanced features that are not yet common in other UNIX-based and UNIX-like operating systems. These differences are significant enough to warrant their own document: *System Startup Programming Topics*.

In that document, you will learn about legacy startup items (the recommended way to start daemons prior to v10.4), `launchd` (the recommended way in v10.4 and later), and why you should start your daemons with these mechanisms instead of modifying `rc` files.

Configuration Files

Often on a UNIX-based system, you find system configuration files in `/etc`. This is still true in Mac OS X, but configuration information is also found in other places. Networking, printing, and user system configuration details are, by default, regulated by the NetInfo database. Applications usually make use of XML property lists (plist) to hold configuration information. You can view many example property lists by looking in `~/Library/Preferences`.

If changing a configuration file in `/etc` does not have your desired effect, then that information is probably regulated by information in the NetInfo database or by an application's property list.

For example, users and groups are not handled with files in `/etc`. To create a user, you instead should use the Accounts preferences pane or NetInfo Manager. NetInfo Manager can also be used to create groups.

Device Drivers

Mac OS X implements an object-oriented programming model for developing device drivers. This technology is called the I/O Kit. It is a collection of system frameworks, libraries, tools, and other resources. This model is different from the model traditionally found on a BSD system. If your code needs to access any devices other than disks, you use the I/O Kit. I/O Kit programming is done using a restricted subset of C++ that omits features unsuitable for use within a multithreaded kernel. By modeling the hardware connected to a Mac OS X system and abstracting common functionality for devices in particular categories, the I/O Kit streamlines the process of device-driver development. I/O Kit information and documentation is available at <http://developer.apple.com/documentation/DeviceDrivers/Conceptual/IOKitFundamentals/index.html>.

The File System

The Mac OS X file system is similar to other UNIX-based operating systems, but there are some significant differences in the file system structure and in case sensitivity. These differences are described in the sections that follow.

File-System Structure and Visibility

The file-system structure of Mac OS X is similar to a BSD-style system. A quick glance at `hier` should comfort you. When in doubt as to where to put things, you can put them where you would in a BSD-style system. There are a few directories that you might not recognize. These differences are described in more detail in “[File System Organization](#)” (page 54).

The default behavior of the Mac OS X Finder is to hide the directories that users normally would not be interested in, as well as invisible files, such as those preceded by a dot (`.`). This appearance is maintained by the Finder to promote simplicity in the user interface. As a developer, you might want to see the dot files and your complete directory layout. The `/usr/bin/defaults` tool allows you to override the default behavior of hiding invisible files. To show all of the files that the Finder ordinarily hides, type the following command in the shell:

```
defaults write com.apple.Finder AppleShowAllFiles true
```

Then restart the Finder either by logging out and back in or by choosing Force Quit from the Apple Menu.

There are a couple of other simple ways to view the contents of hidden folders without modifying the default behavior of the Finder itself. You can use the `/usr/bin/open` command or the Finder Go to Folder command. With `open` you can open a directory in the Finder, hidden or not, from the shell. For example, `open /usr/include` opens the hidden folder in a new Finder window. If you are in the Finder and want to see the contents of an invisible folder hierarchy, choose Go to Folder from the Go menu, or just press `command-~`, and type in the pathname of your desired destination.

For information on how to lay out the directory structure of your completed Mac OS X applications, consult *Apple Human Interface Guidelines* and *Mac OS X Technology Overview*.

Supported File-System Types

Mac OS X supports Mac OS Extended (HFS+), the traditional Macintosh volume format, and the UNIX File System (UFS). HFS+ is recommended and is what most users have their system installed on. Some more server-centric installations have their system installed on UFS. If you develop on UFS, you should thoroughly test your code on HFS+ as well. Although the HFS+ file system preserves case, it is not case sensitive. This means that if you have two files whose names differ only by case, the HFS+ file system regards them as the same file. In designing your application, you should therefore not attempt to put two objects with names that differ only by case in the same directory—for example `Makefile` and `makefile`.

However, developing on HFS+ does not necessarily ensure that your application will work on UFS. It is far too easy to write code in which your program opens a file as `org.mklinux.formattool.prefs` one time and as `org.MkLinux.formattool.prefs` another time and gets completely different results.

Also, do not assume that a bug is unimportant simply because you have seen it only on UFS. Other file systems have similar properties, including potentially an NFS, SMB, or AFP share, particularly when those shares are being served by something other than a Mac. Thus, a bug that occurs on one file system will likely occur on others.

The Kernel

The core of any operating system is its kernel. Though Mac OS X shares much of its underlying architecture with BSD, the Mac OS X kernel, known as XNU, differs significantly. XNU is based on the Mach microkernel design, but it also incorporates BSD features. It is not technically a microkernel implementation, but still has many of the benefits of a microkernel, such as Mach interprocess communication mechanisms and a relatively clean API separation between various parts of the kernel.

Why is it designed like this? With pure Mach, the core parts of the operating system are user space processes. This gives you flexibility, but also slows things down because of message passing performance between Mach and the layers built on top of it. To regain that performance, BSD functionality has been incorporated in the kernel alongside Mach. The result is that the kernel combines the strengths of Mach with the strengths of BSD.

How does this relate to the actual tasks the kernel must accomplish? [Figure 9-1](#) (page 64) illustrates how the kernel's different personalities are manifested.

Figure 9-1 XNU personalities

BSD	Mach
<ul style="list-style-type: none"> - Users and permissions - Networking stack - Virtual File System - POSIX 	<ul style="list-style-type: none"> - Memory management - Messaging - I/O Kit

The Mach aspects of the kernel handle

- Memory management
- Mach messaging and Mach inter process communication (IPC)
- Device drivers

The BSD components

- Manage users and permissions
- Contain the networking stack
- Provide a virtual file system
- Maintain the POSIX compatibility layer

See Inside *Kernel Programming Guide* for more information on why you would (or wouldn't) want to program in the kernel space, including a discussion on the kernel extension (KEXT) mechanism.

Open Directory and the dscl Tool

Open Directory is the built-in Mac OS X directory system that system and application processes can use to store and find administrative information about resources and users. Open Directory includes components such as OpenLDAP and Kerberos for providing local and remote authentication.

By default, each Mac OS X computer runs client and server processes, but the server only serves to the local client. You can also bind client computers to servers other than the local server over a number of protocols including LDAP. Information is then accessed in a hierarchical scheme. In such a scheme, each client computer accesses the union of the information provided first by its local server and then by any higher-level servers it is bound to.

Directory Services is the default way that Mac OS X stores user and some network information. When a user is added, the system automatically adds their information to a local database using Open Directory. Traditional tools such as `adduser` either do not exist or do not work as you might expect. You can add users and groups in several ways:

- Through the Users pane of System Preferences
- Through `/Applications/Utilities/Directory` (for adding groups)
- From the command line (see “[Example: Adding a User From the Command Line](#)” (page 65))

You can find more information on NetInfo in the manual pages for `netinfo`, `netinfo`, `nidump`, `nicl`, `nifind`, `niload`, `niutil`, and `nireport` on a computer running Mac OS X v10.4 or earlier. NetInfo is no longer supported in Mac OS X v10.5 and later. You should use Directory Service functionality instead.

You can find more information on Directory Service in *Open Directory Programming Guide*, *Mac OS X Server Open Directory Administration*, and the manual pages `DirectoryService`, `dscl`, `dsconfigldap`, `dsexport`, `dsimport`, and `dsperfmonitor`.

Example: Adding a User From the Command Line

This section shows a simple example of using the Directory Service command-line tool, `dscl`, to add a user to the system. The example specifies some of the properties that you would normally associate with any user.

Note: These commands must be run as the root user. If you are executing them from the command line manually, you should do this with `sudo`. If you are using them in a script, you should use `sudo` when running the script.

1. Create a new entry in the local (`/`) domain under the category `/users`.

```
dscl / -create /Users/portingunix
```

2. Create and set the shell property to `bash`.

```
dscl / -create /Users/portingunix UserShell /bin/bash
```

3. Create and set the user's full name.

```
dsccl / -create /Users/portingunix RealName "Porting Unix Applications To Mac OS X"
```

4. Create and set the user's ID.

```
dsccl / -create /Users/portingunix UniqueID 503
```

5. Create and set the user's group ID property.

```
dsccl / -create /Users/portingunix PrimaryGroupID 1000
```

6. Create and set the user home directory. (Despite the name NFSHomeDirectory, this can also be used for a path to a home directory on a local volume.)

```
dsccl / -create /Users/portingunix NFSHomeDirectory
/Network/Servers/techno/Users/portingunix
```

or

```
dsccl / -create /Users/portingunix NFSHomeDirectory /Users/portingunix
```

7. Set the password.

```
dsccl / -passwd /Users/portingunix PASSWORD
```

or

```
passwd portingunix
```

8. To make that user useful, you might want to add them to the admin group.

```
dsccl / -append /Groups/admin GroupMembership portingunix
```

This is essentially what System Preferences does when it makes a new user, but the process is presented here so you can see more clearly what is going on behind the scenes with the database. A look through the hierarchies using `dsccl` interactively also helps you understand how the database is organized.

Printing (CUPS)

Mac OS X, beginning in version 10.2, uses the Common UNIX Printing System (CUPS) as the basis of its printing system. CUPS is a free software (open source) print system that is also popular in Linux and UNIX circles. For more information on CUPS, see the Mac OS X Printing website at <http://developer.apple.com/printing>.

Bonjour

Bonjour is an implementation of the IETF ZeroConf working draft. Bonjour has three basic parts:

- IP assignment
- Multicast DNS

- Service discovery

You can see the first part, IP assignment, simply by plugging two computers together using a crossover cable and setting them up to use DHCP to obtain an address. Upon failing to obtain an address, the computers choose an IP number randomly from within the range reserved for ZeroConf numbers, then ARP to see if anyone has that address already, repeating this process as needed until a free address is found. This means that a relatively large number of machines can choose their own IP numbers without any administration.

Multicast DNS is a service that allows you to look up the ZeroConf IP number assigned to a particular machine without an established DNS infrastructure. Your computer sends out a request using IP multicast, asking who responds for a particular host name. The host in question then responds. That host could be a computer, a printer, or any other device that might reasonably exist on an IP-based network.

Service discovery is a way to query local machines (using multicast DNS) or a DNS server to see what machines (computers, printers, and so on) support a given service. If you are writing a service, you can register the service using Bonjour. Then anyone else who makes a request sees that your machine supports that service.

The service discovery portion of Bonjour is independent of the IP assignment portion. Thus, Bonjour can also be used for service discovery on a more traditional, configured IP network.

For more information about Bonjour, see *DNS Service Discovery Programming Guide*.

Scripting Languages

With Mac OS X, you can run shell scripts in its native `sh` compatible shell, `bash`, or the included `csh` and `tcsh`. You can also run Ruby, Perl, Python, Tcl, PHP, and other scripts you have developed. In addition, Mac OS X provides an Apple-specific scripting language, AppleScript. Although AppleScript is immensely powerful and can be used to build applications itself, it is designed mainly to communicate with graphical components of the operating system. There are other uses you can find for it, but it is not a replacement for UNIX-style scripting languages. You can use it, though, to put a GUI front end onto your traditional scripts.

AppleScript conforms to the Open Scripting Architecture (OSA) language and can be used from the command line through the `osascript` command. Other languages can be made OSA compliant to enable interaction with the operating system.

Security and Security Services

Mac OS X implements the Common Data Security Architecture (CDSA). If you need to use Authorization Services, Secure Transport, or certificates within the scope of CDSA, online documentation is available at <http://developer.apple.com/documentation/Security/index.html>.

In addition, Mac OS X provides OpenSSL and PAM to ease porting of applications from other UNIX-based operating systems to Mac OS X.

Role-Based Authentication

By default, there is no root user password in Mac OS X. This is a deliberate design decision to maximize security and to simplify the user interface. Your applications should not assume that a user needs superuser access to the system. For power users and developers, `sudo` is provided to run privileged applications in the shell. Privileged applications can also be run by members of the admin group. By default, the admin group is included in the list of “sudoers” (which can be found in the file `/etc/sudoers`). You can assign users to the admin group in System Preferences:

1. Click the Users button.
2. Select a user from the list and click Edit User, or make a new user.
3. Click the Password tab.
4. Check Allow user to administer this computer.
5. That user can now use administrative applications as well as `sudo` in the shell.

Although it is generally considered unsafe practice to log in as the root user, it is mentioned here because the root user is often used to install applications during development. If during development you need to enable the root user for yourself.

In Mac OS X prior to version 10.5:

1. Launch `/Applications/Utilities/NetInfo Manager`.
2. Choose Domain > Security > Authenticate.
3. As prompted, enter your administrator name and password.
4. Now choose Domain > Security > Enable Root User. The first time you do this, you need to select a root password.

In Mac OS X version 10.5 and later:

1. Launch `/Applications/Utilities/Directory Utility`.
2. Click the lock icon to authenticate yourself.
3. Select Edit > Enable Root User. The first time you do this, you need to select a root password.

Alternatively, you can use `sudo passwd root` from the shell and set the appropriate root password.

Important: Do not assume that an end user of your software can enable the root user. This information is provided to help you in development work only; software you distribute should never require the user to enable the root user.

(Re)designing for Portability

When porting applications to Mac OS X, you might consider making architectural changes to make your port (and future ports) more maintainable.

If you are a new developer designing an application that you eventually hope to use on multiple platforms, many of the portability issues described in this chapter apply regardless of the platforms involved.

What is Portability?

There are many different definitions for portability in the context of code design. One definition of portability is limiting yourself to functions specified in a commonly accepted standard such as the Single UNIX Specification (SUS) or the Portable Operating System Interface (POSIX). This is useful for non-GUI software, particularly when moving among UNIX-based and UNIX-like systems, but it does not address graphical user interfaces and does not allow you to take advantage of operating-system-specific capabilities.

The alternative to such limitations is to design your code in a modular fashion so that additional features specific to a given OS can be “plugged into” your application. Whether you do this with a plug-in interface, a class hierarchy, or simple abstraction, this is the most effective way to allow your software to be easily ported to multiple platforms without sacrificing functionality.

There are many different ways to achieve this degree of portability, all of which have valid uses. The goal of this chapter is to describe some of these and to explain when it is appropriate (or inappropriate) to use each one. This should help you avoid some of the common pitfalls when modifying software to support multiple platforms.

Of course, the proper time to make such decisions is before you begin writing the first line of code. Since this is not always possible, this chapter will also describe ways to retrofit these concepts into an existing application in the most maintainable manner possible.

Using Abstraction Layers

Abstraction layers are the easiest, most general-purpose way of making code more portable. The basic concept is well discussed in computer science textbooks. However, most people still find themselves unclear on when it is appropriate to add an additional layer of abstraction.

The most straightforward rule of abstraction is that you should split functionality by use. If a block of code is likely to be reused, even if the code needs to be slightly modified or wrapped in different code, that block should be a separate function. If a block of code is only relevant to the function in which it is enclosed, it probably should not be, unless doing so represents a significant benefit to readability.

You'd be surprised how well this single rule fits most situations. Take file access, for example. Say you want to take advantage of the Carbon file API (to get alias support, for example). The code to open, read, and write a file is used in many places in most applications. Therefore, it should be abstracted into its own function and called from those places.

What you should generally avoid doing, though, is abstracting the open, read, and write calls individually, then using them inside large loops. This leads to needlessly unreadable code with lots of very small abstraction layer functions. You should instead present an interface that makes sense in the context of your application.

For example, if your application performs streaming, you might have the following overall structure:

- A function that opens a file
- A function that reads the next n bytes that returns to a buffer containing the data
- A function that rewinds the stream by n bytes
- A function that closes the file

Your open function might return an opaque handle that can be passed around within the core of the application but whose structure is unknown outside the file functions.

Similarly, you might have functions that read and write the preferences file using a large data structure, a function to read the header of a file, and so on.

Avoid Conditionalizing Code

Don't fall into the trap of conditionalizing hundreds of bits of code with `#ifdef` directives. (An occasional `#ifdef` is OK.) This quickly leads to unmanageable code, particularly when you support Mac OS X, multiple UNIX-based and UNIX-like systems, Classic Mac OS, and Windows.

If you study the spots in your code that you need to special-case, more often than not, you will find that their issues are closely related. For example, they might all be graphics routines. If so, you might pull all the functions that call graphics routines into their own file and conditionalize the inclusion of the entire file. For example, you might use using one file for X11 routines, one for Win32 routines, and one for Carbon or Cocoa routines. It is easier to conditionalize code in a few core functions than to conditionalize it in a hundred random places in the code. Similarly, it is easier to replace an entire file than to replace bits of a file.

GUI Abstraction Issues

Most *non-GUI* abstraction layers should be as thin as possible, since a thick abstraction layer tends to lead to significant platform divergence. However, with a GUI, platform divergence is desirable and is thus an exception to this rule.

If you want your application to look identical on all platforms, then you should make the abstraction layer thin. However, this tends to result in Mac OS X users throwing your application in a dumpster because it doesn't feel like a Mac application. Mac OS X has common GUI style rules that most X11 applications don't follow.

For example, X11 applications often have per-window menus, while Mac OS X has per-application menus (with the ability to add or remove menus when a different window is in the foreground). X11 applications tend to have very square features and small, space-efficient buttons, while Mac OS X applications tend to have rounded, large, easier-to-read buttons.

Similarly, other operating systems, such as Windows and Classic Mac OS, have different design rules. You must understand the GUI design rules for each computing platform or your application will not be accepted easily on those platforms.

For this reason, it is easiest to split your entire user interface into a separate file or directory for X11, then clone that and rewrite it for Carbon or Cocoa when porting to Mac OS X. That way, you can heavily redesign the Mac OS X interface to match what Mac users expect without annoying your UNIX users.

A good way to implement this is to write your GUI to operate in a separate thread (or multiple threads, if desired). It can then communicate with the core of the application using pipes or other platform-specific solutions.

Another way to implement this is to simply have the GUI call functions in the core of the code. This design is effective when the core of the application is entirely GUI driven, but tends to result in the GUI appearing to wedge if the core of the application can take a long time to complete an operation, and is particularly hard to implement if the core code needs to do something continuously in the background during normal operation. You should consider these issues in your redesign to make your application more palatable to end users.

Other General Rules

Don't fall into the trap of overly abstracting your code. If a block of code will not be used in multiple places, don't abstract it out unless it is platform-specific. Abstracting out code that appears in two or three places is, for readability reasons, usually not worth splitting into a separate function.

Do try to limit your functions to a reasonable length. Abstraction for readability alone is probably not a good idea, as splitting a function unnecessarily can end up making it harder to read. If the purpose of a function doesn't divide in an obvious way into multiple subtasks, it is probably better not to split it. In many cases, though, you may find that an inner loop within that function can be considered a distinct operation unto itself. If so, that loop may be split into a separate function if doing so improves readability.

Don't make abstraction layers more than five or six layers deep, as measured from the top level of a major functional unit. The deeper the nesting, the harder it is to follow the code when debugging.

Don't make functions shorter than about ten lines of code. The readability improvement of such a small change in the outer function is rarely significant. There are, of course, exceptions, such as inline assembly (which should always be abstracted). These exceptions are good candidates for an inline function (or even a macro, if you are so inclined).

Remember that these are guidelines, not rules. You should generally follow your instincts. If you feel like something might be reusable, go ahead and abstract it even if it is only being used in one place. If you feel that splitting an inner loop into its own function is pointless, don't split it. This applies to all guidelines, not just the ones in this chapter.

Using Plug-Ins and Libraries Effectively

Plug-ins can be an effective way to minimize platform-specific code in the core of an application. This section describes places where plug-ins are appropriate. The specifics of writing plug-ins for Mac OS X are described in depth in “[Dynamic Libraries and Plug-ins](#)” (page 29).

Effectively managing plug-ins can be tricky when dealing with multiple platforms. You may choose to have a plug-in for each platform, or for each service, or both. The choice of methodology depends largely on the amount of code that you need to put into an external module.

If you have only a few platform-specific bits, it is probably sufficient to have a single per-platform plug-in. A good place for such a design is in an application that needs, for example, to support a single feature in both Mac OS 9 and Mac OS X. By isolating this block of code into an external module, the appropriate piece can be loaded according to the platform in which the application is launched.

Another approach is to have a separate plug-in for each distinct service from the operating system. This is particularly effective if the number of services is significant. Separating services into separate modules makes debugging each service easier from a version control point of view.

Designing a Plug-In Architecture

There are several right ways to design a plug-in architecture and many, many more wrong ways. The most important features of a plug-in architecture are extensibility, simplicity, extensibility, robustness, and extensibility, with emphasis on extensibility.

Ensuring that an architecture is robust is the responsibility of the implementors, and is beyond the scope of any design document. You should consider extensibility and simplicity early and often during the design process.

Simplicity

In general, the simpler the plug-in API, the more likely people will use it rather than grafting hacks into other parts of the code. Of course, this becomes a problem if you fail to expose features that are needed for a given plug-in. To solve this problem, design your API based around message passing concepts rather than function calls. That way, you can easily add additional types of messages to the API without modifying the API itself.

Of course, if your plug-in API only needs to handle one particular type of communication and you are relatively certain that this will not change, a plug-in architecture based on functions is somewhat easier to use. As with all designs, there are tradeoffs.

Extensibility

Extensibility in API design is a tricky issue. When you are creating a plug-in architecture in general, you can rarely envision the sorts of plug-ins that might eventually be useful. Someone might want plug-ins to add additional functionality that no one had even thought of, much less invented, when you created the architecture

You can also use plug-ins to abstract the interface presented by platform-specific services into a more generic form to simplify your application. These designs are more straightforward than architectures intended for adding new functionality; creating plug-in architectures for adding features are beyond the scope of this document.

The first step in designing a plug-in API for service abstraction is to choose the level of abstraction at which the application ends and the plug-in begins. If you choose a level that is too close to the application, the plug-ins will contain redundant code. If you choose a level that is too far removed, you will eventually need to interface your application with an environment that does not meet your initial expectations. As a result, the interface code for that module will become excessively complex. In either case, a rewrite is probably in order, which wastes time and resources that could have been saved had you chosen the right level of abstraction to begin with.

Choosing the level of abstraction for a plug-in interface can range from straightforward to downright impossible, depending on the application. In general, you should choose to split out the smallest amount of code possible such that the following conditions are met:

- No assumptions are made about the data structures of the underlying service.
- Return values are standardized in a way that meets the needs of your application.
- The calling convention can be implemented on any system that provides the support necessary for your application to function—that is, no special features of the underlying service should be implied at the plug-in API level or in your application.
- The data types passed to and from the plug-in API are one of two kinds: either base types in the programming language of choice, or composites of those base types that are structurally relevant to the application as a whole rather than to the underlying service.

Designing a plug-in API in this fashion may, however, result in substantial duplication of code. In such environments, a nested plug-in approach may be more practical.

Nesting Modules for Shared Functionality

The concept of nesting plug-ins is straightforward. You should consider using nested plug-ins when you find numerous opportunities for dividing an application from its plug-ins—that is, when dealing with a general class of underlying services divided into a number of subclasses that contain multiple specific variants with common characteristics.

Nested modules are convenient for:

- Authentication—a generic authentication layer with a generic plug-in for UNIX-based systems and specific plug-ins below to handle platform-specific variations
- Databases—at the top level, you might have ODBC, JDBC, STET, and SQL, with submodules for other, more specific SQL implementations
- Printing—for example, you might have a plug-in that handles PostScript printers with narrower plug-ins that override generic assumptions for a particular printer model

In short, whenever you have a group of underlying services that are substantially similar in behavior, but where you want to be able to also support services with dramatically different behavior, a nested plug-in approach is an effective design (unless the differences in behavior at the lowest level are very minimal).

Example: Database Support

A good example of plug-in API design is database access. As an example, consider a project that uses a database to obtain song playlist information.

Such a program requires various specific pieces of information from the database. For a given song, it might need the title, the artist, the total length, the intro (talkover) time, the time at which the next song should begin playing (trigger time), and the time at which the song should be faded out, if applicable. In addition, a comments field could be included for composer or other genre-specific information. This is referred to as the internal data level.

To facilitate support for arbitrary databases, you add the first layer of abstraction at the internal data level. The core code calls functions with names like `getsongname` and `getsongtrigger`. Any arbitrary database, regardless of the query language used, can easily return a string or an integer (or at worst, be coerced into doing so). This is considered the minimum level of functionality needed to support this application, and thus forms the first plug-in split.

However, a surprising number of databases use a common syntax, SQL, for making requests. Although the syntax is the same, certain details (data types) are different, and the libraries used for accessing them are also different. For these reasons, supporting multiple SQL servers is a desirable goal, because you can use most common databases as storage.

Because the SQL instructions themselves are so similar between databases, the second split comes somewhat naturally. The code to actually execute a query can be placed in an implementation-specific module, and the code to generate the query—for a song's name, for example—can be placed in a generic SQL language module. The API for the implementation-specific module could, for example, include:

- `getdbc`—Returns a pointer to a database connection object. The SQL core code treats this as an opaque type, but this is necessary to allow many database implementations to be used in a multithreaded environment. This information should not leave the SQL core unless the larger design requires connections to multiple databases simultaneously.
- `opendb`—Opens a database connection.
- `closedb`—Closes a database connection.
- `dbquerystring`—Executes an SQL query and returns the first result as a string.
- `dblistsongs`—Returns an array of song IDs that match an SQL query.

It might strike you as odd that this design doesn't specify a generic SQL query call. This is largely a space-time tradeoff. Implementing a generic query would allow the GUI, for example, to request all of the information about a call in a single request. However, doing so would add a great deal of code for a relatively small time benefit. Since the GUI does not need to refresh frequently, and since the number of queries per second, therefore, tends to be relatively small, it makes sense to design the API to be as simple as possible. If low latency and high bandwidth are required for the specific application, a generic routine is desirable.

If you are familiar with SQL, though, you are probably aware that the similarities among SQL implementations end at inserting new tables into the database. Because this is such a small piece of the overall picture for this type of application (one or two lines of code total), it could easily be "special-cased" in the core code or (preferably) in the generic SQL code. This is one of those cases where you must decide whether the extra 1% of functional abstraction is worth the extra effort involved. In many cases, it is not.

If creating tables is more significant, you could create them in one of two ways: by adding a function in the implementation-specific plug-in, or by adding a function in the generic SQL plug-in, using a table for the data types (which could, if desired, reside in the implementation-specific plug-in).

To demonstrate the effectiveness of this design, MySQL database support was added to an application that used this exact design in about an hour. With the exception of table creation, no modifications to the generic SQL support routines were required.

In summary, proper plug-in design is much like proper abstraction. Code that can be reused should be reused, but the API to the plug-ins should not assume any knowledge of the underlying system.

Architectural Portability

As a good general rule, when writing abstraction layers, plug-in architectures, and so on, architectural portability should be considered. Most existing open source software is already fairly flexible in this regard. However, when dealing with OS-specific code, you may find that consistent support for things like endianness and alignment are not always considered.

When you encounter code specific to Mac OS X with architectural dependencies on a particular processor architecture, there are a number of ways to handle the situation. Here are some tips that should help:

- Altivec or SSE code should be special-cased with equivalent scalar versions that can be compiled in. This will ensure that it is easy for developers familiar with other chip architectures to understand what is happening and write equivalents for other architectures.
- Testing of alignment should generally occur at compile time, not configuration time, to avoid unnecessary problems when cross-compiling.
- Testing of endianness should generally occur at either compile time or execution time, at your option.
- Executing intermediate build products is a bad idea and should be avoided where possible.

For more detailed information, see [“Compiling for Multiple CPU Architectures”](#) (page 22) and the document *Universal Binary Programming Guidelines, Second Edition*.

Glossary

ADC See [Apple Developer Connection](#)

Apple Developer Connection The primary source for technical and business resources and information for anyone developing for Apple's software and hardware platforms anywhere in the world. It includes programs, products, and services and a website filled with up-to-date technical documentation for existing and emerging Apple technologies. The Apple Developer Connection is at <http://www.apple.com/developer/>.

Aqua The graphical user interface for Mac OS X.

bom (Bill Of Materials) A file in an installer package used by the Installer to determine which files to install, remove, or upgrade. It contains all the files within a directory, along with information about each file such as the file's permissions, its owner and group, size, its time of last modification, a checksum for each file, and information about hard links.

bundle A directory in the file system that stores executable code and the software resources related to that code. Applications, plug-ins, and frameworks are types of bundles. Except for frameworks, bundles are file packages, presented by the Finder as a single file.

Carbon An application environment for Mac OS X that features a set of programming interfaces derived from earlier versions of the Mac OS. The Carbon API has been modified to work properly with Mac OS X, especially with the foundation of the operating system, the kernel environment. Carbon applications can run in Mac OS X, Mac OS 9, and all versions of Mac OS 8 later than Mac OS 8.1.

Classic An application environment for Mac OS X that lets you run non-Carbon legacy Mac OS software. It supports programs built for both Power PC and 68K chip architectures and is fully integrated with the Finder and the other application environments.

Cocoa An advanced object-oriented development platform for Mac OS X. Cocoa is a set of frameworks with programming interfaces in both Java and Objective-C. It is based on the integration of OPENSTEP, Apple technologies, and Java.

Darwin Another name for the core of the Mac OS X operating system. The Darwin kernel is equivalent to the Mac OS X kernel plus the BSD libraries and commands essential to the BSD command-line environment. Darwin is open source technology.

.dmg file A Mac OS X disk image file.

Finder The system application that acts as the primary user interface for file-system interaction.

HFS (Hierarchical File System) The Mac OS Standard file-system format, used to represent a collection of files as a hierarchy of directories (folders), each of which may contain either files or folders themselves. HFS is a two-fork volume format.

HFS+ The Mac OS Extended file-system format. This file-system format was introduced as part of Mac OS 8.1, adding support for filenames longer than 31 characters, Unicode representation of file and directory names, and efficient operation on very large disks. HFS+ is a multiple-fork volume format.

Mach-O The executable format of Mach object files. This is the default executable format in Mac OS X.

NetInfo The network administrative information database and information retrieval system for Mac OS X. Many Mac OS X services consult the NetInfo database for their configuration information.

nib file An XML archive that describes the user interface of applications built with Interface Builder.

.pkg file A Mac OS X Installer file. May be grouped together into a metapackage (.mpkg).

plist See [property list](#).

property list A structured, textual representation of data that uses the Extensible Markup Language (XML) as the structuring medium. Elements of a property list represent data of certain types, such as arrays, dictionaries, and strings.

Xcode Apple's graphical integrated development environment. It is available free with the Mac OS X Developer Tools package.

XNU The Mac OS X kernel. The acronym stands for X is Not Unix. XNU combines the functionality of Mach and BSD with the I/O Kit, the driver model for Mac OS X.

Document Revision History

This table describes the changes to *Porting UNIX/Linux Applications to Mac OS X*.

Date	Notes
2008-04-08	Fixed minor typographical errors and omissions.
2008-02-08	Corrected instructions for building with an external Makefile target.
2006-11-07	Improved shared library and bundle explanation. Updated example from <code>niel</code> to <code>dscl</code> .
2006-10-03	Added an explanation of why one should not conditionalize code specifically for Mac OS X.
2006-03-08	Fixed typographical errors.
2006-01-10	Corrected external target build instructions.
2005-10-04	Added <code>hdiutil</code> example and made minor wording improvements.
2005-08-11	Added an explanation of universal binaries. Added a reference to "Dynamic Library Programming Topics."
2005-07-07	Added missing step in packaging instructions. Added link to GNU <code>autoconf</code> . Globally changed "Disk Copy" to "Disk Utility."
2005-04-29	Updated for Mac OS X v10.4.
2004-08-31	Minor wording changes.
2004-04-22	Minor wording changes.
2003-10-17	Updated for Panther

REVISION HISTORY

Document Revision History