
Mach-O Runtime Architecture



August 7, 2003



Apple Computer, Inc.
© 2003, 2004 Apple Computer, 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 Computer, 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 Computer, 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 or Apple-licensed computers.

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

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

Apple, the Apple logo, Carbon, Cocoa, Mac, Mac OS, Macintosh, and MPW are trademarks of Apple Computer, Inc., registered in the United States and other countries.

Finder and Velocity Engine are trademarks of Apple Computer, Inc.

Objective-C is a trademark of NeXT Software, Inc.

AIX is a trademark of IBM Corp., registered in the U.S. and other countries, and is being used under license.

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

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

Simultaneously published in the United States and Canada.

Even though Apple has reviewed this manual, APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THIS MANUAL, ITS QUALITY, ACCURACY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS MANUAL IS SOLD "AS IS," AND YOU, THE PURCHASER, 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 MANUAL, 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 Runtime Architecture	7
	Who Should Read This Book	8
	Where to Find Things in This Book	8
	Where to Find More Information	8
Chapter 1	Mach-O Runtime Architecture	11
	Building Mach-O Files	11
	The Tools—Building and Running Mach-O Files	12
	The Products—Types of Mach-O Files You Can Build	13
	Modules—The Smallest Unit of Code	14
	Static Archive Libraries	15
	Executing Mach-O Files	15
	Launching an Application	16
	Forking and Executing the Process	16
	Finding Imported Symbols	17
	Loading Code At Runtime	21
	Using Shared Libraries and Frameworks	21
	Loading Plug-In Code With Bundles	25
Chapter 2	Mach-O Runtime Conventions for PowerPC	27
	PowerPC Data Types	27
	PowerPC Data Alignment	29
	PowerPC Stack Structure	31
	Prologs and Epilogs	33
	The Red Zone	34
	PowerPC Calling Conventions	35
	Parameter Passing	35
	Function Return	39
	Register Preservation	39
	PowerPC Dynamic Code Generation	41
	Position-Independent Code	41
	Indirect Addressing	44

Chapter 3 **Mach-O File Format Reference** 49

- Mach-O Types and Data Structures 53
 - Mach-O Header Data Structure 53
 - Load Command Data Structures 56
 - Symbol Table and Related Data Structures 72
 - Relocation Data Structures 80
 - Static Archive Libraries 84
 - Multi-CPU Architecture Files 86

Chapter 4 **Mach-O Dynamic Linking Functions Reference** 89

- Dynamic Linker Functions 89
 - Object File Image Functions 89
 - Section and Segment Accessors 100
 - Low-Level Dynamic Linking Functions 103
 - Glue Functions for Indirect Addressing 111
- Dynamic Linker Data Types 112
 - Boolean Return Value 112

Document Revision History 113

Glossary 115

Index 119

Tables, Figures, and Listings

Chapter 1 Mach-O Runtime Architecture 11

- Listing 1-1 Building a framework 23
- Listing 1-2 Building a private framework 24
- Listing 1-3 Building a simple umbrella framework 25

Chapter 2 Mach-O Runtime Conventions for PowerPC 27

- Figure 2-1 The PowerPC stack 32
- Figure 2-2 The red zone 34
- Figure 2-3 The organization of the parameter area of the stack 37
- Figure 2-4 Parameter layout in registers and the parameter area 38
- Figure 2-5 Passing a variable number of parameters 38
- Listing 2-1 Sample PowerPC assembler prolog code 33
- Listing 2-2 Sample PowerPC routine epilog 34
- Listing 2-3 A variable-argument routine 38
- Listing 2-4 C source code example for position-independent code 42
- Listing 2-5 Position-independent code generated from the C example (with addresses in the left column) 42
- Listing 2-6 Sample C code for indirect function calls 45
- Listing 2-7 Example of an indirect function call 45
- Listing 2-8 Example of a position-independent indirect function call 46
- Table 2-1 Scalar data types in the Mach-O PowerPC Runtime Environment 28
- Table 2-2 Vector data types in the Mach-O PowerPC runtime environment 29
- Table 2-3 Embedded alignment modes (in bytes) 31
- Table 2-4 Volatile and nonvolatile registers 41

Chapter 3 Mach-O File Format Reference 49

- Figure 3-1 Mach-O file format basic structure 50
- Table 3-1 Typical sections in a Mach-O file 53
- Table 3-2 Mach-O load commands 58

Introduction to Runtime Architecture

A **runtime architecture** is a set of rules that define the software environment. Authors of compilers and other development tools must follow the definition of the runtime architecture to guarantee that programs released by different developers will work with each other. A runtime architecture typically specifies

- how to address code and data
- how to load and keep track of portions of program code in memory
- how compilers should generate code
- how to invoke certain system services, such as loading of application plug-ins

Mac OS X supports a number of different application environments, each with its own runtime rules, conventions, and file formats. The only executable format the Mac OS X kernel reads directly is the Mach-O file format, which gives the Mach-O runtime architecture its name. In Mac OS X, kernel extensions, command-line tools, applications, frameworks, and libraries (shared and static) are all implemented using Mach-O files.

The following list describes other runtime environments supported by Mac OS X:

- **Classic** is a Mac OS X application that runs Mac OS 9 within its address space and provides bridging services that allow Mac OS X to interact with the Mac OS 9 applications. Both classic 68K applications and PowerPC Code Fragment Manager (CFM) applications can run under Mac OS 9 in Classic. (Mac OS 9 does not support the 68K variant of the Code Fragment Manager, so you cannot run CFM-68K applications in Mac OS X.)
- **LaunchCFMApp** is a command-line tool that runs programs created for the PowerPC Code Fragment Manager. The file format used by such programs is called the Preferred Executable Format (PEF). Carbon provides bridging for Code Fragment Manager applications that allows them to link to Mach-O-based code, but, for ease of debugging if for no other reason, it's generally a good idea to use Mach-O for Carbon applications.
- The HotSpot **Java virtual machine** is a Mac OS X application that executes Java bytecode applications and applets.
- The **Mac OS X kernel** supports kernel extensions (kexts), static Mach-O executable files that are loaded directly into the address space of the kernel. Because errant code can write directly to memory used by the kernel, kernel extensions have the potential to crash the operating system, and you should generally avoid implementing functionality as kernel extensions if possible.

The Code Fragment Manager is documented in the book *Mac OS Runtime Architectures*, available from the Apple Developer Connection website.

This book describes the dynamic Mach-O runtime environment, including the data formats and calling conventions to which applications must adhere to successfully interoperate with other Mach-O applications and code libraries.

Who Should Read This Book

If you write development tools for Mac OS X, you will need to understand the information presented in this book.

This book is also useful for developers of shared libraries and frameworks, and for developers of applications that need to load code at runtime.

Where to Find Things in This Book

This document describes the Mach-O runtime linking architecture and calling conventions as of Mac OS X 10.2. The individual chapters discuss the following topics:

- [“Mach-O Runtime Architecture”](#) (page 11) describes the basics of the Mac OS X runtime architecture, including detailed conceptual information about how Mach-O executable files are built, linked, and executed. This chapter also explains how dynamic linking works with shared libraries, frameworks, bundles and plug-ins. All programmers who create shared libraries or who dynamically load code at runtime should read this chapter.
- [“Mach-O Runtime Conventions for PowerPC”](#) (page 27) describes the Mac OS X application binary interface for PowerPC microprocessors, which specifies low-level routine calling conventions and data formats. Writers of assembly language code and authors of developer tools should read this chapter.
- [“Mach-O File Format Reference”](#) (page 49) describes the layout of the Mach-O file format, used for executables, shared libraries, bundles, and all other native executable machine code in the Mach-O runtime architecture. Authors of developer tools should read this chapter.
- [“Mach-O Dynamic Linking Functions Reference”](#) (page 89) describes the Mach-O low-level programming interface. If you are loading code at runtime but cannot or do not wish to use CFBundle or NSBundle, you should refer to this chapter.

Where to Find More Information

You can access full reference documentation for the standard command line development tools using the `man` tool on the command line, or by selecting Open Man Page... from Project Builder’s Help menu in Mac OS X 10.2 and later.

This book documents effectively what a developer tool vendor might need to create a Mach-O-based development environment for a procedural language such as C. It does not document the following:

- The GCC C++ application binary interface—the specification of C++ class member layout, function/method name mangling, and related C++ issues. This information is documented for GCC 3.0 and later at <http://www.codesourcery.com/cxx-abi/abi.html>.
- The GCC Objective-C data structures and dynamic runtime functions. For this information, see the book *Inside Mac OS X: The Objective-C Programming Language*.
- The runtime environment of the Mac OS X kernel, Darwin. Darwin documentation is available at <http://developer.apple.com/documentation>.

For additional documentation on the standard Mac OS X developer tools, see the Tools documentation at <http://developer.apple.com/documentation>.

Source code from the Darwin project can be downloaded from <http://developer.apple.com/darwin/>. The source code for all of the Mac OS X linking and compiling tools is available in the following Darwin subprojects.

- `gcc3`—The Mac OS X compiler for the C, C++, and Objective-C languages, based on the GNU Compiler Collection version 3.1 (as of this writing). This is the standard compiler for Mac OS X 10.2 and later versions.
- `gcc`—The Mac OS X compiler for the C, C++, and Objective-C languages, based on the GNU Compiler Collection version 2.95.2. This is the standard compiler for Mac OS X 10.1 and earlier versions.
- `cctools`—The Mac OS X static linker, dynamic linker, and related tools for examining and manipulating Mach-O files and static archive libraries.

You might also find the following books useful in conjunction with this one.

- *Linkers and Loaders*, John R. Levine, Morgan Kaufmann, 2000, ISBN 1-55860-496-0. Describes the workings and operation of standard linkers from the earliest program loaders to the present dynamic link editors. Among the contents of this book are discussions of the classic BSD `a.out` format, the ELF format preferred by many current operating systems, the IBM System/360 linker output format, and the Microsoft Portable Executable (PE) format.
- *Mac OS Runtime Architectures*, Apple Computer, Inc. Available at <http://developer.apple.com/documentation>. Documents the classic 68K segment loader architecture, as well as the Code Fragment Manager Preferred Executable executable format used with classic PowerPC applications and with many Carbon applications.

I N T R O D U C T I O N
Introduction to Runtime Architecture

Mach-O Runtime Architecture

This chapter discusses how you use the Mach-O runtime architecture. It describes the types of programs you can build; how programs are loaded and executed; the ways in which you can change the way programs are loaded executed; how to load code at runtime; how to load and link code at runtime. If you create or load bundles, shared libraries or frameworks, you'll probably want to read and understand everything in this chapter.

The Mach-O file format provides both intermediate (during the build process) and final (after linking the final product) storage of machine code and data. It was designed as a flexible replacement for the BSD `a.out` format, to be used by the compiler and the static linker and to contain statically-linked executable code at runtime. Features for dynamic linking were added as the goals of the system evolved, resulting in a single file format for both statically and dynamically linked code.

A Mach-O file contains three primary regions of data: a header, a set of load commands, and raw segment data. The header and load commands describe the features, layout, and linking characteristics of the file. The segment data contains raw data for the segments that are defined in the load commands. The complete format is described in [“Mach-O File Format Reference”](#) (page 49).

The following sections describe the concepts behind the Mach-O runtime architecture:

- [“Building Mach-O Files”](#) (page 11) discusses how Mach-O programs are built and the types of build products that are supported.
- [“Executing Mach-O Files”](#) (page 15) discusses how Mach-O programs are loaded and linked.
- [“Loading Code At Runtime”](#) (page 21) discusses how you can load Mach-O code, or set up code to be loaded, at runtime.

Building Mach-O Files

The following sections loosely describe how Mac OS X programs are built, and discusses, in depth, the types of programs that you can build.

- [“The Tools—Building and Running Mach-O Files”](#) (page 12) describes the tools involved in the Mach-O build process.

- [“The Products—Types of Mach-O Files You Can Build”](#) (page 13) describes the types of Mach-O files that you can build.
- [“Modules—The Smallest Unit of Code”](#) (page 14) explains the role of the smallest indivisible unit of code within a Mach-O shared library.

The Tools—Building and Running Mach-O Files

To perform the work of actually loading and binding a program at runtime, the kernel uses the **dynamic linker** (a specially-marked dynamic shared library located at `/usr/lib/dyld`). The kernel loads the dynamic linker into the new process and then executes it. The dynamic linker loads the program and all of the frameworks and shared libraries that the program uses.

Throughout this book, the following tools are discussed abstractly:

- A **compiler** is a tool that translates from source code written in a high-level language into intermediate object files that contain machine binary code and data. Unless otherwise specified, this book considers a machine-language assembler to be a compiler.
- A **static linker** is a tool that combines intermediate object files into final products (in the next section, [“The Products—Types of Mach-O Files You Can Build”](#) (page 13)).

The Mac OS X Developer Tools CD contains several command-line tools (which this book refers collectively to as the **standard tools**) for building and analyzing your application, include compilers and `ld`, the standard static linker. Whether you use Project Builder, the standard command-line tools, or a third-party tool set to develop your application, understanding the role of each of the following tools can enhance your understanding of the Mach-O runtime and facilitate communication about these topics with other Mac OS X developers. The standard tools include the following:

- The compiler driver, `/usr/bin/cc` (or, on Mac OS X 10.2, `/usr/bin/gcc`), contains support for compiling, assembling, and linking modules of source code from the C, C++, and Objective-C languages. The compiler driver calls several other tools that implement the actual compiling, assembling, and static linking functionality. The actual compiler tools for each language dialect are normally hidden from view by the compiler driver; their role is to transform input source code into assembly language for input to the assembler.
- The C++ compiler driver, `/usr/bin/c++`, is like `/usr/bin/cc`, but automatically links C++ runtime functions (to support exceptions, runtime type information and other advanced language features) into the output file.
- The assembler, `/usr/bin/as`, creates intermediate object files from assembly-language code. It is primarily used by the compiler driver, which feeds it the assembly language source generated by the actual compiler.
- The static linker, `/usr/bin/ld`, is used by the compiler driver (and as a standalone tool to combine Mach-O executable files). You can use the static linker to bind programs either statically or dynamically. Statically bound programs are complete systems in and of themselves; they cannot make calls, other than system calls, to frameworks or shared libraries. In Mac OS X, kernel extensions are statically bound, while all other program types are dynamically bound, even traditional UNIX and BSD command line tools. All calls to the Mac OS X kernel by programs outside of the kernel are made through shared libraries, and only dynamically bound programs can access shared libraries.

- The library creation tool, `/usr/bin/libtool`, creates either static archive libraries or dynamic shared libraries, depending on the parameters given. `libtool` supersedes an older tool called `ranlib`, which was used in conjunction with the `ar` tool to create static libraries. When building shared libraries, `libtool` calls the static linker (`ld`).

Note: There is also a GNU tool named `libtool`, which allows portable source code to build libraries on various different UNIX systems. Don't confuse it with Mac OS X `libtool`; while they serve similar purposes, they are not related and they do not accept the same parameters.

Tools for analyzing Mach-O files include the following:

- The Mach-O file analyzer, `/usr/bin/otool`, lists the contents of specific sections and segments within a Mach-O file. It includes symbolic disassemblers for each supported CPU architecture and it knows how to format the contents of many common section types.
- The symbol table display tool, `/usr/bin/nm`, allows you to view the contents of a Mach-O file's symbol table.

The following sections describe how to create typical (and some atypical) Mach-O files. The section [“Searching for Symbols”](#) (page 18) discusses how to analyze existing Mach-O files. For more documentation on the tools, see [“Where to Find More Information”](#) (page 8).

The Products—Types of Mach-O Files You Can Build

In Mac OS X, a typical application executes code that originates from many different files. All of these different types of files contain code that conforms to the Mach-O file format and runtime calling conventions.

The main executable file usually contains the core logic of the program, including the entry point `main` function. The primary functionality of a program is usually implemented in the main executable file's code. See [“Executing Mach-O Files”](#) (page 15) for more information. Other Mach-O files that contain executable code include these:

- **Intermediate object files** are not final products; they are the basic building blocks of larger Mach-O files. Usually, a compiler creates one intermediate object file on output for the code and data generated from each input source code file. You can then use the static linker to combine the object files into dynamic linkers. Integrated development environments such as Project Builder usually hide this level of detail, and some development tools may not use the Mach-O format to store intermediate code and data.
- **Dynamic shared libraries** are files that contain modules of reusable executable code that your application references dynamically, and that are loaded by the dynamic linker when the application is launched. Shared libraries are typically used to store large amounts of code that is usable by many applications. See [“Using Shared Libraries and Frameworks”](#) (page 21) for more information.
- **Frameworks** are shared libraries that are packaged with associated resources, such as graphics files, developer documentation, and programming interfaces. See [“Using Shared Libraries and Frameworks”](#) (page 21) for more information.
- **Umbrella frameworks** are special types of frameworks that themselves contain more than one subframework. For example, the Cocoa umbrella framework consists of the Application Kit (user interface classes) and Foundation (non-user-interface classes) frameworks. See [“Using Shared Libraries and Frameworks”](#) (page 21) for more information.

- **Static archive libraries** contain modules of reusable code that the static linker can add to your application at build time. Static archive libraries generally contain very small amounts of code that is usable to only a few applications, or code that is difficult to maintain in a shared library for some reason. See [“Static Archive Libraries”](#) (page 15) for more information.
- **Bundles** are executable files that your program can load at runtime using dynamic linking functions. Bundles implement plug-in functionality, such as file format importers for a word processor. The term “bundle” has two related meanings in Mac OS X:
 - the actual Mach-O file containing the executable code
 - a file package—a folder containing the Mach-O bundle file and associated resources. A file package bundle need not contain a Mach-O bundle file.

The latter usage is the more common. However, unless otherwise specified, this book refers to the former.

See [“Loading Plug-In Code With Bundles”](#) (page 25) for more information.

- **Kernel extensions** are statically-bound Mach-O object files that are packaged similarly to bundles. Kernel extensions are loaded into the kernel address space and must therefore be built differently than other Mach-O file types; see the kernel documentation for more information. The kernel’s runtime environment is very different from the user space Mac OS X runtime, so it is not covered in this book.

To function properly in Mac OS X, all Mach-O files except kernel extensions must be **dynamically bound**—that is, built with code that allows dynamic references to shared libraries.

By default, the static linker searches for frameworks and umbrella frameworks in the directory `/System/Library/Frameworks` and for shared libraries and static archive libraries in the directory `/usr/lib`. Bundles are usually located in the `Resources` directory of an application package. However, you can specify the pathname for a different location at link time (and, for development purposes, at runtime as well).

The next section describes the role and usage of modules.

Modules—The Smallest Unit of Code

At the highest level, you can view a Mach-O shared library as a collection of modules. A **module** is the smallest unit of machine code and data that can be linked independently of other units of code. Usually, a module is an object file generated by compiling a single C source file. For example, given the source files `main.c`, `thing.c`, and `foo.c`, the compiler might generate the Mach-O object files `main.o`, `thing.o`, and `foo.o`. Each of these output object files is one module. When the static linker is used to combine all three files into a dynamic shared library, each of the object files is retained as an individual unit of code and data. When linking applications and bundles, the static linker always combines all of the object files into one module.

The static linker can also reduce several input modules into a single module. When building most dynamic shared libraries, it’s usually a good idea to do this before creating the final shared library, because function calls between modules are subject to a small amount of additional overhead. With `ld`, you can perform this optimization by using the command line as follows:

```
ld -r -o things.o thing1.o thing2.o thing3.o
```

Project Builder performs this optimization by default.

Static Archive Libraries

To group a set of modules, you can use a **static archive library**, which is an archive file with a table of contents entry. The format is that used by the `ar` command. You can use the `libtool` command to build a static archive library, and you can use the `ar` command to manipulate individual modules in the library.

Note: Again, please note that Mac OS X `libtool` is not GNU `libtool`.

In addition to Mach-O files, the static linker and other development tools accept static archive libraries as input. You might use a static archive library to distribute a set of modules that you do not want to include in a shared library, but that you want to make available to multiple programs.

Although an `ar` archive can contain any type of file, the typical purpose is to group several object files together with a table of contents file, forming a static archive library. The static linker can link the object files stored in a static archive library into a Mach-O executable or dynamic library. Note that you must use the `libtool` command to create the static library table of contents before an archive can be used as a static archive library.

Note: For historical reasons, the `tar` file format is different from the `ar` file format. The two formats are not interchangeable.

The `ar` archive file format is described in [“Static Archive Libraries”](#) (page 84).

With the standard tools, you can pass the `-static` option to `libtool` to create a static archive library. The following command creates a static archive library named `libthing.a` from a set of intermediate object files, `thing1.o` and `thing2.o`:

```
libtool -static thing1.o thing2 -o libthings.a
```

Note that if you pass neither `-static` nor `-dynamic`, `libtool` will assume `-static`. It is, however, considered good style to explicitly pass `-static` when creating static archive libraries.

Executing Mach-O Files

The next sections provide an overview of the Mac OS X dynamic loading process. The process of loading and linking a program in Mac OS X mainly involves two entities: the Mac OS X kernel and the dynamic linker. When you execute a program, the kernel creates a new process for the program, then loads and executes the dynamic linker shared library, `/usr/lib/dyld`, in the program’s address space. The dynamic linker then loads the program and the libraries it references. This process is described in detail in the next sections, as follows:

- [“Launching an Application”](#) (page 16)
- [“Forking and Executing the Process”](#) (page 16)
- [“Finding Imported Symbols”](#) (page 17)

Launching an Application

When you launch an application from the Finder or the Dock, or when you run a program in a shell using Terminal, the system ultimately calls two functions on your behalf, `fork` and `execve`. `fork` creates a new process; `execve` loads and executes the program. There are several variant `exec` functions, such as `execl`, `execv`, and `exec`, each providing a slightly different way of passing arguments and environment variables to the program. In Mac OS X, each of these other `exec` routines eventually calls the kernel routine `execve`.

When writing a Mac OS X application, you should use the Launch Services framework to launch other applications. Launch Services understands application packages, and you can use it to open both applications and documents. The Finder and Dock use Launch Services to maintain the database of mappings from document types to the applications that can open them. Cocoa applications can use the class `NSWorkspace` to launch applications and documents; `NSWorkspace` itself uses Launch Services. Launch Services ultimately calls `fork` and `execve` to do the actual work of creating and executing the new process.

Forking and Executing the Process

To create a new process using BSD system calls, your process must call the `fork` system call. `fork` creates a new logical copy of your process, then returns the new process ID to your process. Both the original process and the new process continue executing from the call to `fork`; the only difference is that `fork` returns the ID of the new process to the original process, and zero to the new process. (`fork` returns `-1` to the original process and sets `errno` to a specific error value if the new process could not be created.)

To run a different executable, your process must call the `execve` system call with a pathname specifying the location of the new executable. The `execve` call replaces the program currently in memory with a new executable file.

A Mach-O executable file contains a header consisting of a set of load commands. For programs that use shared libraries or frameworks, one of these commands specifies the location of the linker to be used to load the program. If you are using the standard Mac OS X development tools, this is always `/usr/bin/dyld`, the standard Mac OS X dynamic linker.

When you call the `execve` routine, the kernel first loads the specified program's file and examines the `mach_header` structure at the start of the file. The kernel verifies that the file appears to be a valid Mach-O file, and then interprets the load commands stored in the header. The kernel then loads the dynamic linker specified by the load commands into memory and executes the dynamic linker on the program file.

The dynamic linker loads all of the shared libraries that the main program links against (the **dependent libraries**) and binds enough of the symbols to start the program. It then calls the entry point function. At build time, the static linker adds the standard entry point function into the main executable file from the object file `/usr/lib/crt1.o`. This function sets up the runtime environment state for the kernel and calls static initializers for C++ objects, initializes the Objective-C runtime, and then calls the program's `main` function.

Finding Imported Symbols

When the dynamic linker loads a Mach-O file (which, for the purposes of this section, is called the **client program**), it connects the file's imported symbols to their definitions in a shared library or framework. The settings used to build the client program affect this process, as explained in the following sections:

- [“Binding Symbols”](#) (page 17) describes the process of binding the imported symbols in one Mach-O file to their definitions in other Mach-O files. The static linker allows you to specify a number of different ways to perform the binding process.
- [“Searching for Symbols”](#) (page 18) describes the process of finding a symbol. For each imported symbol the dynamic linker finds, it uses this process to determine the location of the symbol. Programs can also explicitly find symbols using the one of the APIs described in [“Loading Plug-In Code With Bundles”](#) (page 25).

Binding Symbols

Binding is the process of resolving a module's references to functions and data in other modules (the **undefined external symbols**, sometimes called **imported symbols**). The modules may be in the same Mach-O file or in different Mach-O files; the semantics are identical in either case. When the application is first loaded, the dynamic linker loads the imported shared libraries into the address space of the program. When binding is performed, the linker replaces each of the program's imported references with the address of the actual definition from one of the shared libraries.

The dynamic linker can bind a program at several different points during loading and execution, depending on the options you specify at build time.

- With **just-in-time binding** (also called lazy binding), the dynamic linker binds a reference (and all of the other references in the same module) when the program first uses the reference. The dynamic linker loads any particular shared library the first time it binds a reference from that shared library.
- With **load-time binding**, the dynamic linker binds all of the imported references immediately upon loading the program, or, for bundles, upon loading the bundle. To use load-time binding with the standard tools, specify the `-bind_at_load` option to `ld` to specify that the dynamic linker should immediately bind all external references when the file is loaded. Without this option, `ld` will setup the output file for just-in-time binding.
- With **prebinding**, a form of load-time binding, the shared libraries referenced by the program are each prebound at a specified address. The static linker sets the address of each undefined references in the program to default to these addresses. At runtime, the dynamic linker needs only to verify that none of the addresses have changed since the program was built (or since the prebinding was recomputed). If the addresses have changed, the dynamic linker must undo the prebinding by clearing the prebound addresses for all of the undefined references and then proceed as if the program had been just-in-time bound. Otherwise, it does not need to perform any action to bind the program. Prebinding can greatly speed up application launch times.
Prebinding requires that each framework specify its desired base virtual memory address and that none of the prebound addresses of the loaded frameworks overlap. To prebind a file with the standard tools, specify the `-prebind` option to `ld`.

- **Weak references**, a feature introduced in Mac OS X 10.2, is useful for selectively implementing features that may be available on some systems, but not on others. This mode of binding allows a program to optionally bind to specified shared libraries. If the dynamic linker cannot find definitions for weak references, it sets them to null and continues to load the program. The program can check at runtime to see whether or not a reference is null, and if so, avoid using the reference. You can specify both libraries and individual symbols to be weakly-referenced.

Note: The Mach-O weak linking design is derived from the classic Mac OS Code Fragment Manager implementation of weak linking. If you are familiar with the ELF executable format, you may be used to a different meaning for the terms “weak symbol” or “weak linking,” where a “weak” symbol may be overridden by a non-weak symbol. The equivalent Mach-O feature is the “weak definition”—see [“Scope and Treatment of Symbol Definitions”](#) (page 19) for more information

If no other type of binding is specified for a given library, the static linker sets up the program’s undefined references to that library to use just-in-time binding.

Searching for Symbols

A **symbol** is a generic representation of the location of a function, data variable, or constant in an executable file. References to functions and data in a program are references to symbols. To refer to a symbol when using the dynamic linking routines, you usually pass the name of the symbol, although some functions also accept a number representing the ordering of the symbol in the executable file. The name of a symbol representing a function that conforms to standard C calling conventions is the name of the function with an underscore prefix. Thus, the name of the symbol representing the function `main` would be `_main`.

Programs created by the original Mac OS X 10.0 development tools add all symbols from all loaded shared libraries into a single global list. Any symbol that your program references can be located in any shared library, as long as that shared library is one of the program’s dependent libraries (or one of the dependent libraries of the dependent libraries).

Mac OS X 10.1 introduces a two-level symbol namespace feature. The first level of the two-level namespace is the name of the library that contains the symbol, and the second is the name of the symbol. With the two-level namespace enabled, when the static linker records references to imported symbols, it records a reference to the name of the library that contains the symbol and the name of the symbol. Linking your programs with the two level namespace feature offers two benefits over the flat namespace:

- **Enhanced performance when searching for symbols.** With the two-level namespace, the dynamic linker knows exactly where to start looking for the implementation of a symbol. With a flat namespace, the dynamic linker must search all of the loaded libraries for the one that contains the symbol.
- **Enhanced forward compatibility.** In the flat namespace, two or more libraries cannot contain symbols with different implementations that share the same name, because the dynamic linker cannot know which library contains the preferred implementation. This is not initially a problem, because the static linker catches any such problems when you first build the application. However, if the vendor of one of your dependent shared libraries later releases a new version of the library that contains a symbol with the same name as one in your program or in another dependent shared library, your program will fail to run.

Your application must link directly to the shared library that contains the symbol (or, if the library is part of an umbrella framework, to the umbrella framework that contains it).

When obtaining symbols in a program built with the two-level namespace feature enabled, you must specify a reference to the shared library that contains the symbols.

By default, the Mac OS X 10.1 static linker defaults to a two-level namespace for all Mach-O files.

Note: The Mach-O two-level namespace feature is loosely based on the design of the Code Fragment Manager's namespace. A two-level namespace is approximately equivalent to the namespace used to look up symbols in code fragments. Because the Code Fragment Manager always requires an explicit reference to a the library in which a symbol should be found, there is no Code Fragment Manager equivalent to a flat namespace search.

For programs that do not have a two-level namespace, you can tell the linker to define references to undefined symbols even if the linker cannot find the library that contains them. When you build an executable with such undefined symbols, you are making the assumption that one of the other files loaded as part of the executable file at runtime contains those symbols. Bundles and shared libraries sometimes use this option to reference symbols defined in the main executable. However, this causes you to lose the performance and compatibility benefits of two-level namespaces. It's usually better to explicitly link against an executable that defines the references. However, if you must link with undefined references, you can do it by enabling the flat namespace and suppressing undefined reference warnings, using the options `-flat_namespace` and `-undefined suppress` as in the following command line:

```
ld -o my_tool -flat_namespace -undefined suppress peace.o love.o
```

To build executables with a two-level namespace, the static linker must be able to find the source library for each of the symbols. This can present difficulties for authors of bundles and dynamic shared libraries that assume a flat, global symbol namespace. To build successfully with the two-level namespace, keep the following points in mind:

- Bundles that need to reference symbols defined in the program's main executable must use the `-bundle_loader` static linker option. The static linker can then search the main executable for the undefined symbols.
- Shared libraries that need to reference symbols defined in the program main executable must load the symbol dynamically using a function that does not require a library reference, such as `NSLookupAndBindSymbol` (page 96).

A two-level symbol namespace can still be searched using functions for doing flat symbol searches.

Scope and Treatment of Symbol Definitions

Symbols in a Mach-O file may exist at several different levels of scope. This section describes each of the possible scopes that a symbol may be defined at, and provides examples of C code used to create each symbol type. These examples work with the standard developer tools; a third party tool set may have different conventions.

A **defined external symbol** is any symbol defined in the current Mach-O file, including functions and data. The following C code defines an external symbol:

```
int x = 0;
```

An **undefined external symbol** is any symbol defined in a file outside of the current file. The following C code defines two external symbols, a variable and a function:

```
extern int x;
extern void SomeFunction(void);
```

A **common symbol** is a symbol that may appear in multiple intermediate object files. The static linker permits multiple common symbol definitions with the same name in input files, and copies the one with the largest size to the final product. If there is another symbol with the same name as a common symbol, the static linker will ignore the common symbol instead.

The standard C compiler generates a common symbol when it sees a **tentative definition**—a global variable that has no initializer and is not marked `extern`. The following line is an example of a tentative definition:

```
int x;
```

A shared library cannot have common symbols. To eliminate common symbols in an existing shared library, you must either explicitly define the symbol (with an initialized value, for example) in one of the modules of the shared library, or pass the `-fno-common` flag to the compiler.

A **private defined symbol** is a symbol that is not visible to other modules. The following C code defines a private symbol:

```
static int x;
```

A **private external symbol** is an external defined symbol that is visible only to other modules within the same Mach-O file as the module that contains it. The standard static linker changes private external symbols into private defined symbols unless the application developer specifies otherwise (using the `-keep_private_externs` parameter).

You can mark a symbol as private external by using the `__private_extern__` attribute, as in the following C example:

```
__attribute__((__private_extern__)) int x;
```

A **weak reference** is an undefined external symbol that need not be found in order for the client program to successfully link. If the symbol does not exist, the dynamic linker sets the address of the symbol to zero. Files with weak references can only be used on Mac OS X 10.2 and later. The following C code demonstrates conditionalizing an API call using a weak reference:

```
/* Only call this API if it exists */
if( SomeNewFunction != NULL )
    SomeNewFunction();
```

To specify that a function should be treated as a weak reference, an application developer should use the `weak_import` attribute on a function prototype, as demonstrated by the following code:

```
void SomeNewFunction(void) __attribute(weak_import);
```

A **coalesced symbol** is a symbol that may be defined in multiple object files, but that the static linker generates only one copy of in the output file. This can save a lot of memory with certain C++ language features that the compiler must generate for each individual object file, such as virtual function tables, runtime type information (RTTI) and C++ template instantiations. The compiler determines which constructs should be coalesced; no work on the part of the application developer is required.

Note: Programmers who use other operating systems may be familiar with the concept of symbols that are marked with a COMDAT flag; a coalesced symbol is the Mach-O equivalent feature.

A **weak definition** is a symbol that will be ignored by the linker if an otherwise identical but non-weak definition exists. This is used by the standard C++ compiler to support C++ template instantiations. The compiler marks implicit—and not explicit—template instantiations as weak definitions. The static linker will then prefer any explicit template instantiation to an implicit one for the same symbol, which provides correct C++ linking semantics. As with coalesced symbols, the compiler determines the constructs that require the weak definitions feature; no work on the part of the application developer is required.

Note: Files with weak definitions can only be used on Mac OS X 10.2 and later. The static linker changes any remaining weak definitions into non-weak definitions, so this is only a concern for intermediate object files and static libraries that you wish to deploy on older system releases.

A **debugging symbol** is a symbol generated by the compiler that allows the debugger to map from addresses in machine code to locations in source code. The standard compilers currently generate debugging symbols in the **stabs** debugging format, which is documented in the GDB debugger internals documentation (see [“Where to Find More Information”](#) (page 8)). Debugging symbols, like other symbols, are stored in the symbol table (see [“Symbol Table and Related Data Structures”](#) (page 72)).

Loading Code At Runtime

This section describes how you can load code at runtime.

- [“Using Shared Libraries and Frameworks”](#) (page 21)
- [“Static Archive Libraries”](#) (page 15)
- [“Loading Plug-In Code With Bundles”](#) (page 25)

Using Shared Libraries and Frameworks

Programmers often refer to dynamic shared libraries using different names, such as dynamically linked shared libraries, dynamic libraries, DLLs, dylibs, or just shared libraries; in Mac OS X, all of these names refer to the same thing: a library of code dynamically loaded into a process at runtime.

Shared libraries allow the operating system as a whole to use memory more efficiently. Each process in Mac OS X has its own virtual address space. The Mac OS X kernel allows regions of logical memory to be mapped into multiple processes at different addresses. The dynamic linker takes advantage of this feature by mapping the same read-only copy of the shared library code into the address space of each process. The result is that only one physical copy of a shared library is in memory at any time, even though many different processes may use it at the same time. Data, such as variables and constants, contained by a shared library is mapped into each client process using the kernel’s copy-on-write optimization capability. With copy-on-write, the data is shared among processes until one of the processes attempts to change the data. At that point,

the kernel creates a writable copy of the data private to that process. The other processes continue to use the read-only shared copy. Thus, additional memory for data is allocated only when absolutely necessary.

Shared libraries also provide a way for programs to seamlessly benefit from system upgrades. When the system is upgraded, the shared libraries are updated, but the programs need not be; since they are dynamically bound to the shared libraries, the programs can continue to call the same functions and the updated implementation of the shared libraries will be executed.

Client Program Compatibility

This section describes various parameters that affect compatibility with client programs. You can set these parameters at build time.

Shared libraries have two version numbers, which allow you to create new versions of a shared library that are **binary compatible** (that is, they do not require client programs to be recompiled) with the functions exported by the old versions of a library. These version numbers are as follows:

- The **current version** of the library specifies the current revision number of the library's implementation. A client program can examine this version number to find out the exact version of the library, which can be useful for checking for bug fixes and feature additions. The shared library can also examine the version number that the client program originally linked against, which can be useful for maintaining backwards compatibility.
- The **compatibility version** of the library specifies the version of the library's API that the shared library claims to be backward-compatible with. If the compatibility version of the shared library is newer than the version recorded with the client program, the program fails to link and an error occurs.

The **install name** is the pathname used by the dynamic linker to find a shared library at runtime. The install name is defined by the shared library and recorded into the client program by the static linker.

You can locate private frameworks and shared libraries in an application package using a relative-path install name beginning with `@executable_path`, such as `@executable_path/../Frameworks/MyFramework.framework`. This is useful for sharing functionality with plug-ins (bundles).

You can pass the `-dynamic` option to `libtool` to create a dynamic shared library. The following command creates a dynamic shared library named `libthing.dylib` from a set of intermediate object files, `thing1.o` and `thing2.o`:

```
libtool -dynamic thing1.o thing2.o -o libthing.dylib
```

Packaging a Shared Library as a Framework

A **framework** is a shared library packaged with associated resources, such as headers, localized strings, and documentation. installed in a standard folder hierarchy, usually in a standard location in the file system. The folders usually contain related header files, documentation in any format, and resource files. A framework may contain multiple versions of itself, and each version may have its own set of resources, documentation, and header files.

From a tools perspective, a framework is a shared library whose install name ends in the form `frameworkName.framework/Versions/versionName/frameworkName` or the form `frameworkName.framework/frameworkName`.

You compile a framework by building a normal dynamic shared library into a folder with the same name and a framework extension. For example, to create a framework named `Chaos`, place a dynamic shared library named `Chaos` into a folder called `Chaos.framework`. You can create other folders inside this folder to store related resources, such as header files, documentation, and graphics (the standard folder names for these are called `Headers`, `Documentation`, and `Resources`, respectively).

Apple follows a standard framework versioning convention, different from the shared library version numbering system. By versioning your framework, you can ship older versions of your framework alongside newer versions, to allow older clients to continue functioning, while still allowing you to advance the design of the framework in ways not compatible with older clients.

To version your framework, create a parent folder inside the framework called `Versions`, create a subfolder in `Versions` using a naming scheme of your choice, and build the framework shared library and other folders in this subfolder. Then create symbolic links in the framework's root folder to point to the shared library and folders. When you build a new, incompatible version of your framework, build it into a new directory in the versions directory and update the symlinks to point to the new version. When a client links to a versioned framework, the install name recorded in the client executable includes the full path to the shared library executable, and the dynamic linker will thus only load that version.

For example, a client links to a framework called `Peace.framework`, and the symlinks in `Peace.framework` point to the latest version, which is named "B." The install name of the framework ends with `Peace.framework/Versions/B/Peace`. The static linker records this install name in the client. When the client is loaded, the dynamic linker will attempt to load the shared library with this install name. Note that, while frameworks that ship with the system usually name successive versions with consecutive letters of the English alphabet (A through Z), you can use any name you wish.

A framework developer can build a simple, versioned framework in four steps.

1. Create the framework version directory.
2. Compile the framework executable into the framework version directory.
3. Create a symbolic link named `Current` that points to the framework version directory.
4. Create a symbolic link to the framework executable in the parent framework directory.

The shell commands in [Listing 1-1](#) (page 23) build a framework named `Peace` from the C source files `peace.c` and `love.c`. The resulting framework has the install name `Peace.framework/Versions/A/Peace`

Listing 1-1 Building a framework

```
mkdir -p Peace.framework/Versions/A
cc -dynamiclib -o Peace.framework/Versions/A/Peace peace.c love.c
ln -s Peace.framework/Versions/A Peace.framework/Versions/Current
ln -s Peace.framework/Versions/A/Peace Peace.framework/Peace
```


[Listing 1-2](#) (page 24) demonstrates how to create a private framework—that is, a framework located in an application’s package. Specify the install name explicitly during the linking phase and prefix it with `@executable_path`. The install name of the resulting framework is `@executable_path/Frameworks/Peace.framework/Versions/A/Peace`.

Listing 1-2 Building a private framework

```
mkdir -p Peace.framework/Versions/A
cc -c peace.c love.c
libtool -dynamic -install_name
@executable_path/Frameworks/Peace.framework/Versions/A/Peace -o
Peace.framework/Versions/A/Peace peace.o love.o -framework System
```

Packaging Frameworks and Libraries Under an Umbrella

An **umbrella framework** is a framework that serves as the “parent” of a group of frameworks and shared libraries that implement related functionality. Umbrella frameworks are useful to help manage extremely large development projects with complex interdependencies, such as subsystems of Mac OS X itself; for all other projects, a single framework should suffice (and is better for load-time performance).

To create an umbrella framework, you can take a normal framework and designate a subset of its imported frameworks as subframeworks. The subframeworks themselves need not be aware that they are part of the umbrella. With `ld`, you can use the `-sub_umbrella` option to designate a subframework.

When your program links against an umbrella framework, it also implicitly links against all of the subframeworks. Symbols located in subframeworks of umbrella frameworks are recorded in the client program as if they were implemented directly in the umbrella framework. This feature allows the contents of the umbrella framework to change over time while preserving compatibility with older client programs.

To ensure that developers link to the “parent” umbrella framework and not one of the subframeworks, the subframework can be built with a special load command to prevent unauthorized linking. When a client tries to link directly to such a subframework, the static linker produces an error. However, the subframework can authorize specific clients to link against it, and all subframeworks of an umbrella framework are implicitly authorized to link against each other. (Load commands are explained in [“Mach-O File Format Reference”](#) (page 49); the particular load commands referenced here are documented as `sub_framework_command` (page 70) and `sub_client_command` (page 71), and `ld` will generate them if given the `-sub_framework parent_umbrella_name` and `-sub_client client_name` options) Note that these conventions are enforced at build time by the static linker, but ignored by the dynamic linker at runtime.

You can also include libraries in umbrella frameworks. For example the Foundation framework includes both the Objective-C runtime library (`libobjc`) as a sublibrary and the CoreFoundation framework as a subframework. You might build Foundation using a variation on the commands listed in [Listing 1-3](#) (page 25).

Listing 1-3 Building a simple umbrella framework

```
mkdir Foundation.framework
ld -dylib -o Foundation.framework/Foundation -sub_umbrella CoreFoundation
  -sub_library libobjc -framework CoreFoundation -lobjc Foundation.o
```

By convention, subframeworks of an umbrella framework live within a `Frameworks` directory in the root directory of the umbrella framework, although this is obviously not a technical requirement. For example, the Cocoa framework is an umbrella framework that includes the `UIKit` framework; the `UIKit` framework is itself an umbrella framework that includes the `Foundation` and `ApplicationServices` frameworks as subframeworks.

Because an umbrella framework is a framework, you can use the same directory-based versioning strategy described in [“Packaging a Shared Library as a Framework”](#) (page 22).

Loading Plug-In Code With Bundles

Bundles provide the Mach-O mechanism for loading extension (or plug-in) code into an application at runtime. Typically, a bundle links against the application binary to gain access to the application’s exported API. Bundles can be—but are not required to be—packaged with resources, using the same folder hierarchy as that of an application package. In some cases (depending on the code in the bundle), bundles can also be unloaded.

Mac OS X supports several different schemes that allow other developers to extend the capabilities of your application by writing plug-in code that your program can load at runtime. Although you can use any one of these different plug-in schemes in any type of application, some are more suited to particular situations than others. For example:

- To load Objective-C classes at runtime, use the Foundation framework class **NSBundle**. `NSBundle` provides general services for referring to a packaged program, whether the program is an application or a plug-in.
- To load C functions at runtime, use the Core Foundation framework object **CFBundle**, which, like `NSBundle`, provides general services for referring to a packaged program, whether the program is an application or a plug-in.
- The Core Foundation framework object **CFPlugin**, implements a small subset of the Microsoft Component Object Model (COM) standard. COM allows you to instantiate C functions and data in an object-oriented manner at runtime.
- Carbon developers can also use the **Code Fragment Manager** to load code fragments updated for Carbon from PEF files. For more information, see the Code Fragment Manager documentation.
- For simpler needs, use the `dyld` library **object file image** functions and the `dyld` **low-level functions** to load and link bundle files. When porting UNIX tools that support plug-ins to Mac OS X, you usually want to use these two sets of functions. See [“Object File Image Functions”](#) (page 89) and [“Low-Level Dynamic Linking Functions”](#) (page 103) for more information.

Note: The dynamic linker in Mac OS X 10.0 will cause your program to crash if you ask it to load programs that are built with a two-level namespace hint table, so, by default, the static linker creates bundles that are compatible with Mac OS X 10.0 by not including the two-level namespace

hint table. You can use the `-twolevel_namespace_hints` option to ask the static linker to include the two-level hint table; the resulting bundle can be used only with Mac OS X 10.1 and later.

NSBundle and CFPlugin can both be used from Carbon applications running in both Mac OS 9 and Mac OS X. Both NSBundle and CFPlugin allow you to package plug-in code with the resources associated with the plug-in (such as graphics files and documentation), similar to the packaging for an application. To load COM objects in Mac OS 9, CFPlugin uses the Code Fragment Manager, and on Mac OS X, CFPlugin uses the object file image `dylld` library functions.

For more information on NSBundle, see the Cocoa documentation for NSBundle—Loadable Bundles and Dynamic Linking. For more information on the Code Fragment Manager, see the book *Mac OS Runtime Architectures*. For more information on CFPlugin and COM, see the Core Foundation framework documentation at <http://developer.apple.com/documentation>.

Mach-O Runtime Conventions for PowerPC

This chapter covers specific low-level details of the Mac OS X PowerPC runtime architecture, including the following:

- [“PowerPC Data Types”](#) (page 27)
- [“PowerPC Data Alignment”](#) (page 29)
- [“PowerPC Stack Structure”](#) (page 31)
- [“PowerPC Calling Conventions”](#) (page 35)
- [“PowerPC Dynamic Code Generation”](#) (page 41)

Together, these details specify the **Mach-O PowerPC Application Binary Interface** (ABI). If you are writing PowerPC assembly code or creating Mac OS X development tools, you should understand and conform to these conventions to ensure compatibility with the other programs running in the Mach-O runtime architecture.

PowerPC Data Types

[Table 2-1](#) (page 27) lists the scalar binary data types and their sizes in the Mach-O PowerPC runtime environment.

Table 2-1 Scalar data types in the Mach-O PowerPC Runtime Environment

C or C++ type	Size (in bytes)	Value range
unsigned char	1	0 to 255
char signed char	1	-128 to 127
unsigned short	2	0 to 65,535
signed short	2	-32,768 to 32,767

C or C++ type	Size (in bytes)	Value range
<code>_Bool bool</code>	4	0 or 1 (false or true)
<code>unsigned int</code> <code>unsigned long</code>	4	0 to 4,294,967,296
<code>int</code> <code>signed int</code> <code>signed long</code>	4	-2,147,483,648 to 2,147,483,647
<code>unsigned long long</code>	8	0 to 18,446,744,073,709, 551,616
<code>signed long long</code>	8	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
<code>float</code>	4	See <i>PPC Numerics</i>
<code>double</code>	8	See <i>PPC Numerics</i>
<code>long double</code>	see notes below	
<code>pointer</code>	4	0 to 0xFFFFFFFF

Table 2-2 (page 28) lists the binary vector types available in the Mach-O PowerPC runtime environment with the Velocity Engine (AltiVec).

Table 2-2 Vector data types in the Mach-O PowerPC runtime environment

AltiVec C or C++ type	Size (bytes)	Value range for each unit
<code>vector unsigned char</code>	16 (1 byte each)	0 to 255
<code>vector char</code> <code>vector signed char</code>	16 (1 byte each)	-128 to 127
<code>vector unsigned short</code>	16 (2 bytes each)	0 to 65,535
<code>vector signed short</code>	16 (2 bytes each)	-32,768 to 32,767
<code>vector unsigned int</code>	16 (4 bytes each)	0 to 4,294,967,296
<code>vector int</code> <code>vector signed int</code>	16 (4 bytes each)	-2,147,483,648 to 2,147,483,647
<code>vector bool char</code>	16 (1 bytes each)	0 (false), 1 (true)
<code>vector bool short</code>	16 (2 bytes each)	0 (false), 1 (true)
<code>vector bool int</code>	16 (4 bytes each)	0 (false), 1 (true)

Altivec C or C++ type	Size (bytes)	Value range for each unit
<code>vector float</code>	16 (4 bytes each)	See <i>PPC Numerics</i>
<code>vector pixel</code>	16 (2 bytes each)	1/5/5/5 pixel format

Here are some things to note about PowerPC data types:

- As in most CPU architectures, a byte is 8 bits long, and a `NULL` pointer has a value of zero.
- All floating point types conform to the IEEE-754 standard representation. For the value range and precise format of floating point data types, see the book *Inside Macintosh: PPC Numerics*.
- PowerPC processors use big-endian format to store numeric and pointer data types—most significant bytes first, then least significant bytes.
- PowerPC processors use two’s-complement binary representation for signed integer types.
- On 32-bit PowerPC processors, arithmetic for the 64-bit integer data types (`long long`) must be implemented by the compiler (using math library routines), since the CPU itself does not implement 64-bit integer math operations.
- The `long double` extended-precision type is 16 bytes on classic Mac OS, but GCC for PowerPC currently treats it as 8 bytes, equivalent to a `double`. A future revision of the compiler may extend `long double` to 128 bytes. For this reason, it is not currently recommended that you use `long double` on Mac OS X.
- Vector types are available only on CPUs that implement Altivec execution units.

PowerPC Data Alignment

The PowerPC runtime environment supports multiple data alignment modes. Alignment of data types falls into two categories:

- the **natural alignment**, which is the alignment of a data type when allocated in memory or assigned a memory address
- the **embedding alignment**, which is the alignment of a data type within a composite data structure

For example, the alignment of an unsigned short variable on the stack may differ from that of an unsigned short data item embedded in a data structure.

Note: Data items passed as parameters in a function call have their own special alignment rules. See [“PowerPC Calling Conventions”](#) (page 35), for more information.

The natural alignment of a data type is the size of the type; [Table 2-1](#) (page 27) shows the size of each data type supported by the PowerPC runtime architecture.

In data structures, you can specify an embedding alignment that varies depending on the alignment mode selected. You can typically select the alignment mode using compiler options or pragmas. Your particular choice of mode is determined by compatibility and performance concerns, as detailed for each mode in the following list:

- **Power alignment mode** is derived from the alignment rules used by the IBM `xlc` compiler for the AIX operating system. It is the default alignment mode for Apple's PPC and MrC compilers for classic Mac OS, as well as the default for the PowerPC version of GCC used on AIX and Mac OS X. Because this mode is most likely to be compatible between PowerPC compilers from different vendors, you typically use it with data structures that will be shared between different programs. The rules for power alignment are as follows:
 - the embedding alignment of the first element in a data structure is equal to the element's natural alignment
 - for subsequent elements with a natural alignment less than 4, the embedding alignment of each element is equal to its natural alignment
 - for subsequent elements that have a natural alignment greater than 4 bytes, the embedding alignment is 4, unless the element is a vector data type
 - the embedding alignment for vector data types is always 16 bytes
 - the embedding alignment of a composite type (array or data structure) is determined by the largest embedding alignment of its members
 - the total size of a composite type is rounded up to a multiple of its embedding alignment, and is padded with null bytes

Be careful when defining data structures with `double` and `long long` data types in power alignment mode. Because these types have natural alignments greater than 4 bytes, they may not be appropriately aligned, which will impair performance when such data members are accessed. If you use these data types for any element after the first element, be sure to place padding in the data structure to align these elements to their natural alignment, or use natural alignment mode instead.

- **Mac68k alignment mode** is derived from the alignment rules used by the MPW compilers for classic Mac OS. This alignment mode is usually used with legacy data structures inherited from classic Mac OS. New code should not need to use this alignment mode except to preserve compatibility with older data structures. The rules for mac68k alignment are as follows:
 - the embedding alignment of a `char` type is 1 byte
 - the embedding alignment of all other types other than vector types is 2 bytes
 - the embedding alignment for vector data types is 16 bytes
 - the total size of a composite data type is rounded up to a multiple of 2 bytes
- **Natural alignment mode** uses the natural alignment of each data type as its embedding alignment. Use this mode for highest performance when working with `double`, `long long`, and `long double` data types.
- **Packed alignment mode** contains no alignment padding between elements (the embedding alignment for all elements is 1 byte). Use this mode when you need a data structure to be compressed as small as possible in memory.

[Table 2-3](#) (page 31) compares the embedding alignment for each data in each of the alignment modes.

Table 2-3 Embedded alignment modes (in bytes)

C data type	Power	68K	Packed	Natural
char	1	1	1	1
short	2	2	1	2
long	4	2	1	4
_Bool	4	2	1	4
float	4	2	1	4
double	4 or 8	2	1	8
long long	8	2	1	8
all vector types	16	16	1	16
composite (data structure or array)	4, 8, or 16	2	1	1, 2, 4, 8, or 16

With the standard C compiler, you can control data structure alignment by adding pragmas to your source code, or by using parameters on the command line. The power alignment mode will be used if you do not specify otherwise.

To change the default alignment at the command line, you can pass the options `-malign-power`, `-malign-mac68k`, and `-malign-natural` to the compiler. To enable a particular alignment mode for a data structure, place an alignment pragma of the following form before the data structure:

```
#pragma option align=mode
```

where *mode* is `power`, `mac68k`, `natural`, or `packed`.

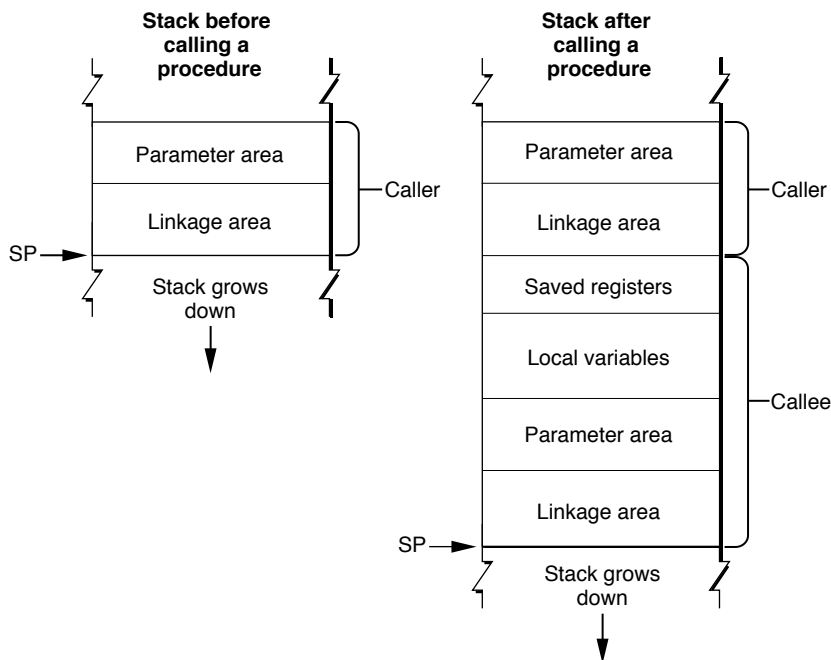
Alignment modes are nested. To restore the previous alignment mode, use `reset`, as follows:

```
#pragma option align=reset
```

PowerPC Stack Structure

The PowerPC runtime environment uses a grow-down stack that contains linkage information, local variables, and a routine's parameter information, as shown in [Figure 2-1](#) (page 32).

Figure 2-1 The PowerPC stack



The PowerPC stack conventions use only a stack pointer (held in register GPR1) and no frame pointer. This configuration assumes a fixed stack frame size, which is known at compile time. Parameters are not passed by pushing them onto the stack.

The calling routine's stack frame includes a parameter area and some linkage information. The **parameter area** has space for the parameters of any routines the caller calls (not the parameters of the caller itself). Since the calling routine might call several different routines, the parameter area must be large enough to accommodate the largest parameter list of all the routines the caller calls. It is the calling routine's responsibility for setting up the parameter area before each call to some other routine, and the called routine's responsibility for accessing the parameters placed within it.

The calling routine's **linkage area** holds a number of values, some of which are saved by the calling routine and some by the called routine. The elements within the linkage area are as follows:

- The Link Register (LR) value is saved at $8(SP)$ by the called routine if it chooses to do so.
- The Condition Register (CR) value may be saved at $4(SP)$ by the called routine. As with the Link Register value, the called routine is not required to save this value.
- The stack pointer is always saved by the calling routine as part of its stack frame.

Note that the linkage area is at the top of the stack, adjacent to the stack pointer. This positioning is necessary so the calling routine can find and restore the values stored there and also to enable the called routine to find the caller's parameter area. This placement means that a routine cannot push and pop parameters from the stack once the stack frame is set up.

The stack frame also includes space for the called routine's local variables. In general, the general-purpose registers GPR13 through GPR31, the floating-point registers FPR14 through FPR31, and vector registers v0, v1, and v14 through v31 are reserved for the routine's local

variables. However, if the routine contains more local variables than would fit in the registers, it uses additional space on the stack. The size of the local variable area is determined at compile time; once a stack frame is allocated, the size of the local variable area cannot change.

Prologs and Epilogs

The called routine is responsible for allocating its own stack frame, making sure to preserve 16-byte alignment on the stack. This action is accomplished by a section of code called the **prolog**, which the compiler places before the body of the routine. After the body of the routine, the compiler generates an **epilog** to restore the processor to the state it was prior to the prolog.

The compiler-generated prolog code does the following:

- Decrements the stack pointer to account for the new stack frame.
- Writes the previous value of the stack pointer to its own linkage area. This procedure ensures that the stack can be restored to its original state after returning from the call.
- Saves all nonvolatile general-purpose and floating-point registers into the saved-registers area. Note that if the called routine does not change a particular nonvolatile register, it does not save it.
- Saves the Link Register and Condition Register values in the caller's linkage area, if needed.

These actions need not be executed in any particular order. [Listing 2-1](#) (page 33) shows a sample routine prolog. Note that the order of these actions differs from the order previously described.

Listing 2-1 Sample PowerPC assembler prolog code

```
linkageArea: set 24                # size in PowerPC environment
params: set 32                    # callee parameter area
localVars: set 0                 # callee local variables
numGPRs: set 0                   # volatile GPRs used by callee
numFPRs: set 0                   # volatile FPRs used by callee

spaceToSave: set linkageArea + params + localVars
spaceToSave: set spaceToSave + 4*numGPRs + 8*numFPRs

.functionName:                   # PROLOG
    mflr      r0,                 # extract return address
    stw      r0,8(SP)            # save the return address
    stwu     SP, -spaceToSave(SP) # skip over caller save
```

At the end of the function, the compiler-generated epilog does the following:

- Restores the nonvolatile general-purpose and floating-point registers that were saved in the stack frame.
- Restores the Condition Register and Link Register values that were stored in the linkage area.
- Restores the stack pointer to its previous value.
- Returns to the calling routine using the address stored in the Link Register.

Again, these actions need not be executed in any particular order. [Listing 2-2](#) (page 34) shows a sample PowerPC routine epilog.

Listing 2-2 Sample PowerPC routine epilog

```


; EPILOG
lhz      r0,spaceToSave(SP)+8    # get the return address
mtlr     R0                      # into the Link Register
addic   SP,SP,spaceToSave       # restore stack pointer
blr                               # and branch to the return address

```

The Red Zone

The space beneath the stack pointer, where a new stack frame would normally be allocated, is called the **red zone**. This area, as shown in [Figure 2-2](#) (page 34), may be used for any purpose as long as a new stack frame does not need to be added to the stack.

Figure 2-2 The red zone



For example, the red zone may be used by a **leaf procedure**. A leaf procedure is a routine that does not call any other routines. Since it does not call any other routines, it does not need to allocate a parameter area on the stack. Furthermore, if it does not need to use the stack to store local variables, it need save and restore only the nonvolatile registers that it uses for local variables. Since by definition no more than one leaf procedure is active at any time, there is no possibility of multiple leaf procedures competing for the same red zone space.

A leaf procedure does not allocate a stack frame nor does it decrement the stack pointer. Instead it stores the Link Register and Condition Register values in the linkage area of the routine that calls it (if necessary) and stores the values of any nonvolatile registers it uses in the red zone. This streamlining means that a leaf procedure's prolog and epilog do only minimal work; they do not have to set up and take down a stack frame.

Note: The value of 224 bytes is the space occupied by nineteen 32-bit general-purpose registers plus eighteen 64-bit floating-point registers, rounded up to the nearest 16-byte boundary. If a leaf procedure's red zone usage would exceed 224 bytes, then it must set up a stack frame just like routines that call other routines.

PowerPC Calling Conventions

This section details the process of passing parameters to a routine in the PowerPC runtime environment. See the section [“PowerPC Dynamic Code Generation”](#) (page 41) for information about generating indirect calls and position-independent code.

Note: These parameter passing conventions are part of Apple’s standard for procedural interfaces. Object-oriented languages may use different rules for their own method calls. For example, the conventions for C++ virtual function calls may be different from those for C functions.

Parameter Passing

A routine can have a fixed or variable number of arguments. In an ANSI-style C syntax definition, a routine with a variable number of arguments typically appears with ellipsis points (...) at the end of its input parameter list.

A variable-argument routine may have several required (that is, fixed) parameters preceding the variable parameter portion. For example, the routine definition

```
(void) fooColor(int number, ...)
```

gives no restriction on the number of arguments after *number*, but you precede them with a *number* argument. Therefore, *number* is a fixed parameter.

Typically the calling routine passes parameters in registers. However, the compiler generates a parameter area in the caller’s stack frame that is large enough to hold all parameters passed to the called routine, regardless of how many of the parameters are actually passed in registers. There are several reasons for this scheme:

- It provides the callee with space to store a register-based parameter if it wants to use one of the parameter registers for some other purpose (for instance, to pass parameters to a subroutine).
- Routines with variable-length parameter lists must often access their parameters from RAM, not from registers. Such routines must reserve eight registers (32 bytes) in the parameter area to hold the parameter values.
- To simplify debugging, some compilers may write parameters from the parameter registers into the parameter area in the stack frame; this allows you to see all the parameters by looking only at that parameter area.

You can think of the parameter area as a data structure that has space to hold all the parameters in a given call. The parameters are placed in the structure from left to right according to the following rules:

- All non-vector parameters are aligned on 4-byte (word) boundaries.
- Noncomposite parameters (that is, parameters that are not arrays or data structures) smaller than 4 bytes occupy the high-order bytes of their word.
- Composite parameters (arrays or data structures) larger than 4 bytes are followed by padding to make a multiple of 4 bytes, with the padding bytes being undefined. (It is recommended that you use zero, however).

- Composite parameters smaller than 4 bytes are *preceded* by padding to 4 bytes.

Note: The latter rule is inconsistent with other PowerPC ABIs. In AIX and classic Mac OS, padding bytes always follow the structure data even in the case of composite parameters smaller than 4 bytes.

For a routine with fixed parameters, the first 8 words (32 bytes) of the parameters, no matter the size of the individual parameters, are passed in registers according to the following rules:

- The first 8 words are placed in GPR3 through GPR10 unless a floating-point parameter is encountered.
- Floating-point parameters are placed in the floating-point registers FPR1 through FPR13.
- If a floating-point parameter appears before all the general-purpose registers are filled, the corresponding GPRs that match the size of the floating-point parameter are skipped. For example, a float item causes one (4-byte) GPR to be skipped, while an item of type double causes two GPRs to be skipped.
- If the number of parameters exceeds the number of usable registers, the calling routine writes the excess parameters into the parameter area of its stack frame.
- The caller places vector parameters in vector registers v2 through v13. For routines with a fixed number of parameters, the presence of vectors does not affect the allocation of GPR and FPR registers. The caller should not allocate space for vector register values in the parameter area of the stack unless the number of vector parameters exceeds the number of available vector registers.

For example, consider a routine `fooFunc` with this declaration:

```
void fooFunc (SInt32 i1, float f1, double d1, SInt16 s1, double d2,
             UInt8 c1, UInt16 s2, float f2, SInt32 i2);
```

To see how the parameters of `fooFunc` are arranged in the parameter area on the stack, first convert the parameter list into a structure, as follows:


```
struct params {
    SInt32          p_i1;
    float          p_f1;
    double         p_d1;
    SInt16         p_s1;
    double         p_d2;
    UInt8          p_c1;
    UInt16         p_s2;
    float          p_f2;
    SInt32         p_i2;
};
```

This structure serves as a template for constructing the parameter area on the stack. (Remember that, in actual practice, many of these variables are passed in registers; nonetheless, the compiler still allocates space for all of them on the stack, for the reasons just mentioned.)

The “top” position on the stack is for the field `p_i1` (the structure field corresponding to parameter `i1`). The floating-point field `p_f1` is assigned to the next word in the parameter area. The 64-bit double field `p_d1` is assigned to the next two words in the parameter area. Next, the short integer field `p_s1` is placed into the following 32-bit word; the original value of `p_s1` is in the lower half of the word, and the padding is in the upper half. The remaining fields of the `params` structure

are assigned space on the stack in exactly the same way, with unsigned values being extended to fill each field to make it a 32-bit word. The final arrangement of the stack is illustrated in [Figure 2-3](#) (page 37). (Because the stack grows down, it looks as though the fields of the params structure are upside down.)


Figure 2-3 The organization of the parameter area of the stack



To see which parameters are passed in registers and which are passed on the stack, you need to map the stack, as illustrated in [Figure 2-3](#) (page 37), to the available general-purpose and floating-point registers. Therefore, the parameter *i 1* is passed in GPR3, the first available general-purpose register. The floating-point parameter *f 1* is passed in FPR1, the first available floating-point register. This action causes GPR4 to be skipped.


The parameter *d 1* is placed into FPR2 and the corresponding general-purpose registers GPR5 and GPR6 are unused. The parameter *s 1* is placed into the next available general-purpose register, GPR7. Parameter *d 2* is placed into FPR3, with GPR8 and GPR9 masked out. Parameter *c 1* is placed into GPR10, which fills out the first 8 words of the data structure. Parameter *s 2* is then passed in the parameter area of the stack. Parameter *f 2* is passed in FPR4, since there are still floating-point registers available. Finally, parameter *i 2* is passed on the stack. [Figure 2-4](#) (page 38) shows the final layout of the parameters in the registers and the parameter area.

Figure 2-4 Parameter layout in registers and the parameter area



If you have a C routine with a variable number of parameters (that is, one that does not have a fixed prototype), the compiler cannot know whether to pass a parameter in the variable portion of the routine in the general-purpose (that is, fixed-point) registers or in the floating-point registers. Therefore, the compiler passes the parameter in both the floating-point and the general-purpose registers, as shown in [Figure 2-5](#) (page 38).

Figure 2-5 Passing a variable number of parameters



The called routine can access parameters in the fixed portion of the routine definition as usual. However, in the variable-argument portion of the routine, the called routine must copy the GPRs to the parameter area and access the values from there. [Listing 2-3](#) (page 38) shows a routine that accesses values by walking through the stack.

Listing 2-3 A variable-argument routine

```
double dsum (int count, ...)
{
    double sum = 0.0;
    double * arg = (double *) (&count + 1 /* pointer arithmetic */);
    while (count > 0 ) {
        sum += *arg;
        arg += 1; /* pointer arithmetic */
        count -= 1;
    }
    return sum;
}
```

Vector parameters in the fixed portion of the routine definition are passed in v2 through v13. For functions with variable arguments only, they are also shadowed on the stack, where they must be aligned to a 16-byte boundary. Vector parameters that appear in the variable-argument portion of the routine must also be shadowed in the GPRs.

Function Return

In the PowerPC runtime environment, functions return floating-point values in register FPR1. Other values are returned as follows:

- Functions returning simple values smaller than 4 bytes (such as the GCC C compiler types `char` or `short`) place the return value in the least significant byte or bytes of GPR3. The most significant bytes in GPR3 are undefined.
- Functions returning 4-byte values (such as pointers, including array pointers, or GCC C compiler types `long` and `int`) return them normally in GPR3.
- Functions returning `long long` values place the return value in GPR3 (the 4 high-order bytes) and GPR4 (the 4 low-order bytes).
- If a function returns a composite value (for example, a struct or union data type) or a value larger than 4 bytes, a pointer must be passed as an implicit left-most parameter before passing all the user-visible arguments (that is, the address is passed in GPR3, and the actual parameters begin with GPR4). The address of the pointer must be a memory location large enough to hold the function return value. Since GPR3 is treated as a parameter in this case, its value is not guaranteed on return. Note that Mac OS X differs from PPC Linux in how they handle 64-bit composite values. In PPC Linux, those values are stored in GPR3 and GPR4.

Register Preservation

Table 2-4 (page 39) lists registers used in the PowerPC runtime environment and their volatility in routine calls. Registers that retain their value after a routine call are called **nonvolatile**. All registers are 4 bytes long.

Table 2-4 Volatile and nonvolatile registers

Type	Register	Preserved by a routine call (nonvolatile)?	Notes
General-purpose register	GPR0	No	
	GPR1	Yes	Used as the stack pointer to store parameters and other temporary data items.

Type	Register	Preserved by a routine call (nonvolatile)?	Notes
	GPR2	No	On other PowerPC platforms, including Mac OS 9, GPR2 is usually a pointer to the current TOC. Mac OS X uses a different indirect addressing scheme, and GPR2 is thus considered a volatile register available for general use.
	GPR3	No	The caller passes parameter values to the called function in GPR3 through GPR10. The called function should place the return value, if any, in GPR3.
	GPR4–GPR10	No	Used to pass parameter values in routine calls (see previous row).
	GPR11	No	
	GPR12	No	Set to the address of the branch target before an indirect call for dynamic code generation. This register is not set for a routine that has been called directly, so routines that may be called directly should not depend on this register being set up correctly. See “Indirect Addressing” (page 44) for more information.
	GPR13–GPR31	Yes	
Floating-point register	FPR0	No	
	FPR1–FPR13	No	Used to pass floating-point parameters in routine calls.

Type	Register	Preserved by a routine call (nonvolatile)?	Notes
	FPR14–FPR31	Yes	
Vector register	v0–v19	No	The caller passes vector parameters in v2 to v13 during a routine call.
	v20–31	Yes	
Vector special purpose	VRSAVE	Yes	32-bit special purpose register; set bits for each vector register that must be saved during a thread or process context switch.
Link Register	LR	No	Stores the return address of the calling routine during a routine call.
Count Register	CTR	No	
Fixed-point exception register	XER	No	
Condition Registers	CR0–CR1	No	
	CR2–CR4	Yes	
	CR5–CR7	No	

PowerPC Dynamic Code Generation

To support dynamically-bound shared libraries, applications, and bundles, the compiler tools and the dynamic linker support two features: **position-independent code** (abbreviated **PIC**) and **indirect addressing**.

Position-Independent Code

Position-independent code, or **PIC**, is the name of the code generation technique that allows the dynamic linker to load a region of code at a different virtual memory addresses. Without some form of position-independent code generation, the operating system would need to place all code you wanted to be shared at fixed addresses in virtual memory, which would make

maintenance of the operating system remarkably difficult. For example, it would be nearly impossible to support shared libraries and frameworks, because each one would need to be preassigned an address, which could never change.

Mach-O position-independent code design is based on the observation that the `__DATA` segment is always located at a constant offset from the `__TEXT` segment. That is, the dynamic linker, when loading any Mach-O file, will never move a file's `__TEXT` segment relative to its `__DATA` segment. Therefore, a function can use its own current address plus a fixed offset to determine the location of the data it wishes to access. All segments of a Mach-O file are at fixed offsets relative to the other segments.

Note: If you are familiar with the Executable Linking Format (ELF), you may note that Mach-O position-independent code is similar to the GOT (global offset table) scheme. The primary difference is that Mach-O code references data using a direct offset, while ELF indirections all data access through the global offset table.

Position-independent code is typically required for shared libraries and bundles, to allow the dynamic linker to relocate them to different addresses at load time. However, it is not typically required for applications, which typically reside at the same address in virtual memory. Apple's version of GCC 3.1 introduces a new option, called `-dynamic-no-pic`, to reduce the code size of application executables by eliminating position-independent code references, while preserving indirect calls to shared libraries. If you are using Project Builder to create your application, this option is enabled by default. For an example of dynamic code generated without PIC, see [Listing 2-7](#) (page 45).

Note: Dynamic code generation without PIC is new for GCC 3.1, the standard compiler shipped with Mac OS X 10.2. However, executables generated with this option will run on older versions of Mac OS X, as long as they do not rely on incompatible new features of the Mac OS X 10.2 tools (such as weak references).

[Listing 2-5](#) (page 42) shows an example of the position-independent code generated for the C code in [Listing 2-4](#) (page 42).

Listing 2-4 C source code example for position-independent code

```
struct s { int member1; int member2; };

struct s bar = {1,2};

int foo(void)
{
    return bar.member2;
}
```

Listing 2-5 Position-independent code generated from the C example (with addresses in the left column)

```

      .text
      ; The function foo
      .align 2
      .globl _foo
```

```

0x0    _foo:      mflr r0                ; save the link register
(LR)
0x4    bcl 20,31,L1$pb        ; Use the branch always
instruction                                ; that does not affect
the link                                    ; register stack to get
the address                                ; of L1$pb into the LR.
0x8    L1$pb:     mflr r10        ; then move LR to r10
0xc    mtlr r0        ; restore the previous LR
                                ; bar is located at L1$pc
+ distance
0x10   addis r9,r10,ha16(_bar-L1$pb); L1$pb plus high 16 bits
of distance
0x14   la r9,lo16(_bar-L1$pb)(r9) ; plus low 16 of distance
                                ; => r9 now contains
address of bar
0x18   lwz r3,4(r9)          ; return bar.member2
0x1c   blr
.data
                                ; The initialized structure bar
                                .align 2
                                .globl _bar
0x20   _bar:      .long 1          ; member1's initialized
value
0x24   .long 2          ; member2's initialized
value

```

To calculate the address of `_bar`, the generated code adds the address of the `L1$pb` symbol (0x8) to the distance to `bar`. The distance to `bar` from the address of `L1$pb` is the value of the expression `_bar - L1$pb`, which is 0x18 (0x20 - 0x8).

Relocating Position-Independent Code

To support relocation of code in intermediate object files, Mach-O supports a section difference relocation entry format. Relocation entries are described in [“Relocation Data Structures”](#) (page 80).

Each of the add-immediate instructions is represented by two relocation entries. For the `addis` instruction (at address 0x10 in the example), the following tables list the two relocation entries. The fields of the first relocation entry (of type `scattered_relocation_info` (page 81)) are as follows:

<code>r_scattered</code>	1—true
<code>r_pcrel</code>	0—false
<code>r_length</code>	2—indicating 4 bytes
<code>r_type</code>	PPC_RELOC_HA16_SECTDIFF
<code>r_address</code>	0x10—the address of the <code>addis</code> instruction
<code>r_value</code>	0x20—the address of the symbol <code>_bar</code>

The values of the second relocation entry are as follows:

r_scattered	1—true
r_pcrel	0—false
r_length	2—indicating 4 bytes
r_type	PPC_RELOC_PAIR
r_address	0x18—the low sixteen bits of the expression (<code>_bar - L1\$pb</code>)
r_value	0x8—the address of the symbol <code>L1\$pb</code>

The first relocation entry for the `la` instruction (at address 0x14 in the example) is as follows:

r_scattered	1—true
r_pcrel	0—false
r_length	2—indicating 4 bytes
r_type	PPC_RELOC_L016_SECTDIFF
r_address	0x14—the address of the <code>addi</code> instruction
r_value	0x20—the address of the symbol <code>_bar</code>

The values of the second relocation entry are as follows:

r_scattered	1—true
r_pcrel	0—false
r_length	2—indicating 4 bytes
r_type	PPC_RELOC_PAIR
r_address	0x0—the high sixteen bits of the expression (<code>_bar - L1\$pb</code>)
r_value	0x8—the address of the symbol <code>L1\$pb</code>

Indirect Addressing

Indirect addressing is the name of the code generation technique, separate from position-independent code, that allows symbols defined in one file to be referenced from another file, without requiring the first file to have explicit knowledge of the layout of the second. This allows the second file to be modified independently of the first.

When generating calls to functions that are defined in other files, the compiler creates a **symbol stub** and a **lazy symbol pointer**. The symbol stub is a small amount of code that directly dereferences and jumps to the lazy symbol pointer. The lazy symbol pointer is an address that is initially set to glue code that calls the linker glue function `dyld_stub_binding_helper`. `dyld_stub_binding_helper` calls the dynamic linker function that performs the actual work of binding the stub. On return from `dyld_stub_binding_helper`, the lazy pointer points to the actual address of the external function.

The simple example code in [Listing 2-6](#) (page 45) might produce two different types of symbol stubs, depending on whether or not it is compiled with position-independent code generation. [Listing 2-7](#) (page 45) shows indirect addressing without position-independent code, while [Listing 2-8](#) (page 46) shows both indirect addressing and position-independent code.

Listing 2-6 Sample C code for indirect function calls

```
extern void bar(void);
void foo(void)
{
    bar();
}
```

Listing 2-7 Example of an indirect function call

```
.text
    ; The function foo
    .align 2
    .globl _foo
_foo:
    mflr r0          ; move the link register into r0
    stw r0,8(r1)     ; save the link register value on the stack
    stwu r1,-64(r1) ; set up the frame on the stack
    bl L_bar$stub    ; branch and link to the symbol stub for _bar
    lwz r0,72(r1)    ; load the link register value from the stack
    addi r1,r1,64    ; removed the frame from the stack
    mtlr r0          ; restore the link register
    blr              ; branch to the link register to return

.symbol_stub        ; the standard symbol stub section
L_bar$stub:
    .indirect_symbol _bar          ; identify this symbol stub
for the                          ; symbol _bar
                                ; load r11 with the high 16
    lis r11,ha16(L_bar$lazy_ptr)   ; address of bar's lazy
bits of the                       ; pointer
                                ; load the value of bar's lazy
pointer                             ; into r12
    lwz r12,lo16(L_bar$lazy_ptr)(r11) ; move r2 to the count register
                                ; load r11 with the address
    mtctr r12                    ; pointer
    addi r11,r11,lo16(L_bar$lazy_ptr) ; of bars lazy
```

```

        bctr                                ; jump to the value in bar's
lazy pointer

.lazy_symbol_pointer    ; the lazy pointer section
L_bar$lazy_ptr:
        .indirect_symbol _bar              ; identify this lazy pointer
for symbol
                                                ; bar
        .long dyld_stub_binding_helper    ; initialize the lazy pointerp
to the stub
                                                ; binding helper address

```

Listing 2-8 Example of a position-independent indirect function call

```

.text
        ; The function foo
        .align 2
        .globl _foo
_foo:
        mflr r0                ; move the link register into r0
        stw r0,8(r1)           ; save the link register value on the stack
        stwu r1,-80(r1)       ; set up the frame on the stack
        bl L_bar$stub         ; branch and link to the symbol stub for _bar
        lwz r0,88(r1)         ; load the link register value from the stack
        addi r1,r1,80         ; removed the frame from the stack
        mtlr r0               ; restore the link register
        blr                   ; branch to the link register to return

.picsymbol_stub              ; the standard pic symbol stub section
L_bar$stub:
        .indirect_symbol _bar    ; identify this symbol stub for the
symbol _bar
        mflr r0                 ; save the link register (LR)
        bcl 20,31,L0$_bar       ; Use the branch-always instruction
that does not
                                ; affect the link register stack to
get the
                                ; address of L0$_bar into the LR.
L0$_bar:
        mflr r11                ; then move LR to r11
                                ; bar's lazy pointer is
located at
                                ; L1$_bar + distance
        addis r11,r11,ha16(L_bar$lazy_ptr-L0$_bar); L0$_bar plus high 16
bits of
                                ; distance
        mtlr r0                 ; restore the previous LR
        lwz r12,lo16(L_bar$lazy_ptr-L0$_bar)(r11); ...plus low 16 of
distance
        mtctr r12               ; move r12 to the count
register
        addi r11,r11,lo16(L_bar$lazy_ptr-L0$_bar); load r11 with the
address of bar's
                                ; lazy pointer
        bctr                    ; jump to the value in
bar's lazy

```

```

                                ; pointer

.lazy_symbol_pointer    ; the lazy pointer section
L_bar$lazy_ptr:
    .indirect_symbol _bar    ; identify this lazy pointer for
    symbol bar
    .long dyld_stub_binding_helper ; initialize the lazy pointer to
    the stub
                                ; binding helper address.

```

As you can see, the `__picsymbol_stub` code in [Listing 2-8](#) (page 46) resembles the position-independent code generated for [Listing 2-5](#) (page 42). For any position-independent Mach-O file, symbol stubs must obviously be position-independent, too.

The static linker performs two optimizations when writing output files:

- it removes symbol stubs for references to symbols that are defined in the same module, modifying branch instructions that were calling through stubs to branch directly to the call
- it removes duplicates of the same symbol stub, updating branch instructions as necessary

Note that a routine that branches indirectly to another routine must store the target of the call in the GPR12 register. Standardizing the register used by the compiler to store the target address makes it possible to optimize dynamic code generation. Because the target address needs to be stored in a register in any event, this convention simply standardizes what register to use. Routines that may have been called directly should not depend on the value of GR12, because in the case of a direct call, its value is not defined.

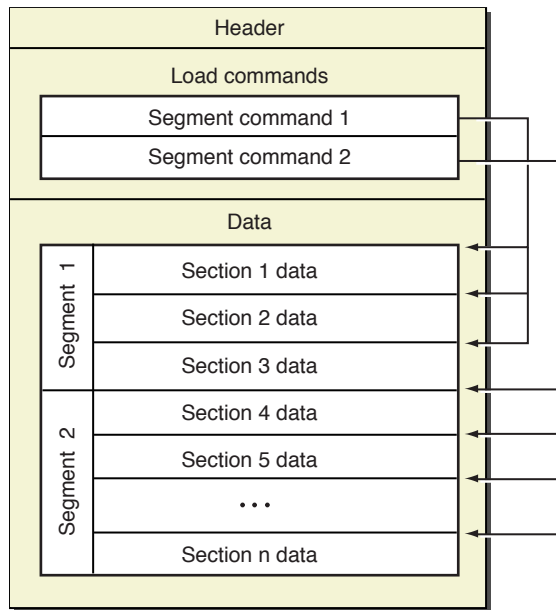
Mach-O File Format Reference

This chapter describes the structure of the Mach-O executable file format, which is the standard used to store programs on disk in the Mach-O runtime architecture. To understand how the development tools work with Mach-O files, and to perform low-level debugging tasks, you need to understand this information.

A Mach-O file contains three major regions (as shown in [Figure 3-1](#) (page 50)):

- At the beginning of every Mach-O file is a **header structure** that identifies the file as a Mach-O executable file. The header also contains other basic file type information, indicates the target CPU architecture, and contains flags specifying options that affect the interpretation of the rest of the file.
- Directly following the header are a series of variable-size **load commands** that specify the layout and linkage characteristics of the Mach-O file. Among other information, the load commands can specify
 - the initial layout of the file in virtual memory
 - the location of the symbol table (used for both dynamic linking and debugging information)
 - the initial execution state of the main thread of the program
 - the names of shared libraries that contain definitions for the main executable's imported symbols
- Following the load commands, all Mach-O files contain the data of one or more **segments**. Each segment contains zero or more **sections**. Each section of a segment contains code or data of some particular type. Each segment defines a region of virtual memory that the dynamic linker will map into the address space of the process. The exact number and layout of segments and sections is specified by the load commands and the file type.

Figure 3-1 Mach-O file format basic structure



Various tables within a Mach-O file refer to sections by number. Section numbering begins at 1 (not zero) and continues across segment boundaries. Thus, the first segment in a file may contain sections 1 and 2 and the second segment may contain sections 3 and 4.

A Mach-O file contains code and data for one CPU architecture. The header structure of a Mach-O file specifies the target CPU architecture, which allows the kernel to ensure that, for example, binary machine code intended for PowerPC processors is not executed on an x86 processor. You can store Mach-O files for multiple CPU architectures in one file using the format described in [“Multi-CPU Architecture Files”](#) (page 86).

Segments and sections are normally accessed by name. Segments, by convention, are named using all uppercase letters preceded by two underscores (for example, `__TEXT`); sections should be named using all lowercase letters preceded by two underscores (for example, `__text`). This naming convention is standard, though not required for the tools to operate correctly.

A segment defines a range of bytes in a Mach-O file and the addresses and memory protection attributes at which those bytes are mapped into virtual memory when the dynamic linker loads the application. As such, segments are always virtual memory page-aligned.

Segments that require more memory at runtime than they do at build time can specify a larger in-memory size than they actually have on disk. For example, the `__PAGEZERO` segment generated by the linker for PowerPC executable files has a virtual memory size of one page, but an on-disk size of zero. Because `__PAGEZERO` contains no data, there is no need for it to occupy any space in the executable file.

A segment contains zero or more sections. For performance reasons, sections that are to be filled with zeros should always be placed at the end of the segment.

For compactness, an intermediate object file contains only one segment. This segment has no name; it contains all of the sections destined ultimately for different segments in the final Mach-O file. The data structure that defines a section ([section](#) (page 60)) contains the name of the segment the section is intended for, and the static linker will place each section in the final Mach-O file accordingly.

For best performance, segments should be aligned on virtual memory page boundaries—4096 bytes for PowerPC processors and 8192 bytes for x86 processors. To calculate the size of a segment, add up the size of each section, then round up the sum to the next virtual memory page boundary (4096 bytes, or 4 kilobytes). Using this algorithm, the minimum size of a segment is 4 kilobytes, and thereafter it is sized at 4 kilobyte increments.

The header and load commands are considered part of the first segment of the file for paging purposes. In an executable file, this generally means that the headers and load commands live at the start of the `__TEXT` segment, because that is the first segment that contains data. The `__PAGEZERO` segment contains no data on disk, and so is ignored for this purpose.

The standard Mac OS X development tools add five segment types to a typical Mac OS X executable:

- The static linker creates a `__PAGEZERO` segment as the first segment of an executable file. This segment is located at virtual memory location zero and has no protection rights assigned, the combination of which causes accesses to NULL, a common C programming error, to immediately crash. The `__PAGEZERO` segment is the size of one full VM page for the current CPU architecture (for x86 and PowerPC, this is 4096 bytes or 0x1000 in hexadecimal). Because there is no data in the `__PAGEZERO` segment, it occupies no space in the file (the file size in the segment command is zero).
- The `__TEXT` segment contains executable code and other read-only data. To allow the kernel to map it directly from the executable into sharable memory, the static linker sets this segment's virtual memory permissions to disallow writing. When the segment is mapped into memory, it can be shared among all processes interested in its contents. (This is primarily used with frameworks, bundles, and shared libraries, but it is possible to run multiple copies of the same executable in Mac OS X, and this applies in that case as well.) The read-only attribute also means that the pages that make up the `__TEXT` segment never need to be written back to disk. When the kernel needs to free up physical memory, it can simply discard one or more `__TEXT` pages and re-read them from disk when they are next needed.
- The `__DATA` segment contains writable data. The static linker sets the virtual memory permissions of this segment to allow both reading and writing. Because it is writable, the `__DATA` segment of a framework or other shared library is logically copied for each process linking with the library. When memory pages such as those making up the `__DATA` segment are readable and writable, the kernel marks them copy-on-write; therefore when a process writes to one of these pages, that process receives its own private copy of the page.
- The `__OBJC` segment contains data used by the Objective-C language runtime support library.
- The `__LINKEDIT` segment contains raw data used by the dynamic linker, such as symbol, string, and relocation table entries.

The `__TEXT` and `__DATA` segments may contain a number of standard sections, listed in [Table 3-1](#) (page 52). The `__OBJC` segment contains a number of sections which are private to the Objective-C compiler. Note that the static linker and file analysis tools typically use the section type and attributes (instead of the section name) to determine how they should treat the section. The section name, type and attributes are explained further in the description of the [section](#) (page 60) data type.

Table 3-1 Typical sections in a Mach-O file

Segment and Section Name	Contents
__TEXT,__text	Executable machine code. The compiler places only executable code in this section; no tables or data of any sort are stored here.
__TEXT,__cstring	Constant C strings. A C string is a sequence of non-null bytes that ends with a null byte (‘\0’). The static linker coalesces constant C string values, removing duplicates, when building the final product.
__TEXT,__picsymbol_stub	Position -independent indirect symbol stubs. See “Indirect Addressing” (page 44) for more information.
__TEXT,__symbol_stub	Indirect symbol stubs. See “Indirect Addressing” (page 44) for more information.
__TEXT,__const	Initialized constant variables. The compiler places all data declared <code>const</code> in this section.
__TEXT,__literal4	4-byte literal values. The compiler places single-precision floating point constants in this section. The static linker coalesces these values, removing duplicates, when building the final product. With some CPU architectures, it is more efficient for the compiler to use immediate load instructions rather than adding to this section.
__TEXT,__literal8	8-byte literal values. The compiler places double-precision floating point constants in this section. The static linker coalesces these values, removing duplicates, when building the final product. With some CPU architectures, it is more efficient for the compiler to use immediate load instructions rather than adding to this section.
__DATA,__data	Initialized mutable variables, such as writable C strings and data arrays.
__DATA,__la_symbol_ptr	Lazy symbol pointers, which are indirect references to functions imported from a different file. See “Indirect Addressing” (page 44) for more information.
__DATA,__nl_symbol_ptr	Non-lazy symbol pointers, which are indirect references to data items imported from a different file. See “Indirect Addressing” (page 44) for more information.
__DATA,__dyld	Information used by the static linker.

Segment and Section Name	Contents
<code>__DATA,__const</code>	Uninitialized constant variables.
<code>__DATA,__mod_init_func</code>	Module initialization functions. The C++ compiler places static constructors here.
<code>__DATA,__mod_term_func</code>	Module termination functions.
<code>__DATA,__bss</code>	Data for uninitialized static variables (for example, <code>static int i;</code>).
<code>__DATA,__common</code>	Uninitialized imported symbol definitions (for example, <code>int i;</code>) located in the global scope (outside of a function declaration).

Each section in a Mach-O file has both a type and a set of attribute flags. In intermediate object files, the type and attributes determine how the static linker copy the sections from intermediate object files into the final product. Object file analysis tools (such as `otool`) use the type and attributes to determine how to read and display the sections. The section type and attributes are not used by the dynamic linker. Descriptions for important variants of the symbol type and attributes as they apply to static linking follow:

- **Regular sections.** In a regular section, only one definition of an external symbol may exist in intermediate object files. The static linker returns an error if it finds any duplicate external symbol definitions.
- **Coalesced sections.** In the final product, the static linker retains only one instance of each symbol defined in coalesced sections. Some complex language features (such as C++ vtables and RTTI) require a definition of a particular symbol to be duplicated in every intermediate object file. This wastes a lot of space in the final product; to reduce the memory occupied by a program, the compiler can place symbol definitions in coalesced sections.

Mach-O Types and Data Structures

This section describes the data types that compose a Mach-O file. Values for integer types in all Mach-O data structures are written using the host CPU's byte ordering scheme, except for `fat_header` (page 86) and "`fat_arch`" (page 87), which are written in big-endian byte order. All of these data types can be found in `/usr/include/mach-o/loader.h`, unless otherwise specified in the description.

Mach-O Header Data Structure

`mach_header`

Specifies general attributes of the file.

```
struct mach_header
```

```

{
    unsigned long magic;
    cpu_type_t cputype;
    cpu_subtype_t cpusubtype;
    unsigned long filetype;
    unsigned long ncmds;
    unsigned long sizeofcmds;
    unsigned long flags;
};
/* Constant for the magic field of the mach_header */
#define MH_MAGIC 0xfeedface /* the mach magic number */
#define MH_CIGAM NXSwapInt(MH_MAGIC)

```

Field Descriptions

magic

An integer containing a value identifying this file as a Mach-O executable file. Use the constant `MH_MAGIC` if the file is intended for use on a CPU with the same endianness as the machine on which the compiler is running. The constant `MH_CIGAM` can be used when the byte ordering scheme of the target machine is the reverse of the host CPU.

cputype

An integer indicating the CPU architecture you intend to use the file on. Appropriate values include

- `CPU_TYPE_POWERPC` for PowerPC-architecture CPUs
- `CPU_TYPE_I386` for x86-architecture CPUs

cpusubtype

An integer specifying the exact model of the CPU. To run on all PowerPC or x86 processors supported by the Mac OS X kernel, this should be set to `CPU_SUBTYPE_POWERPC_ALL` or `CPU_SUBTYPE_I386_ALL`.

`filetype`

An integer indicating the usage and alignment of the file. Valid values for this field include the following:

- The `MH_OBJECT` file type is the format used for intermediate object files. It is a very compact format containing all of its sections in only one segment. The compiler and assembler usually create one `MH_OBJECT` file for each source code file. By convention, the file name extension for this format is `.o`.
- The `MH_EXECUTE` file type is the format used by standard executable programs.
- The `MH_BUNDLE` file type is the type typically used by code that you load at runtime (typically called bundles or plug-ins). By convention, the file name extension for this format is `.bundle`.
- The `MH_DYLIB` file type is for dynamic shared libraries. It contains some additional tables to support multiple modules. By convention, the file name extension for this format is `.dylib`, except for the main shared library of a framework, which does not usually have a file name extension.
- The `MH_PRELOAD` file type is an executable format used for special-purpose programs that are not loaded by the Mac OS X kernel, such as programs burned into programmable ROM chips. Do not confuse this file type with the `MH_PREBOUND` flag, which is a flag that the static linker sets in the header structure to mark a prebound image.
- The `MH_CORE` file type is used to store core files, which are traditionally created when a program crashes. Core files store the entire address space of a process at the time it crashed; you can later run `gdb` on the core file to figure out why the crash occurred.
- The `MH_DYLINKER` file type is the type of a dynamic linker shared library. This is the type that `dylld` is constructed from.

`ncmds`

An integer indicating the number of load commands following the header structure.

`sizeofcmds`

An integer indicating the number of bytes occupied by the load commands following the header structure.

flags

An integer containing a set of bit flags that indicate the state of certain optional features of the Mach-O file format. Masks that you can use to manipulate this field are as follows:

- `MH_NOUNDEFS`—the object file contained no undefined references when it was built.
- `MH_INCRLINK`—the object file is the output of an incremental link against a base file and can't be linked again.
- `MH_DYLDLINK`—the file is input for the dynamic linker and can't be statically linked again.
- `MH_BINDATLOAD`—the dynamic linker should bind the undefined references when the file is loaded.
- `MH_PREBOUND`—the file's undefined references are prebound.
- `MH_SPLIT_SEGS`—the file has its read only and read-write segments split.
- `MH_TWOLEVEL`—the image is using two-level namespace bindings.
- `MH_FORCE_FLAT`—the executable is forcing all images to use flat namespace bindings.

Special Considerations

For all except the `MH_OBJECT` file type, segments must be aligned on page boundaries for the given CPU architecture; 4096 bytes for PowerPC processors and 8192 bytes for x86 processors. This allows the kernel to page virtual memory directly from the segment into the address space of the process. The header and load commands must be aligned as part of the data of the first segment stored on disk (which would be the `__TEXT` segment, in all current file types).

Load Command Data Structures

The load command structures are located directly after the header of the Mach-O file, and they specify both the logical structure of the file and the layout of the file in virtual memory. Each load command begins with fields that specify the command type and the size of the command data.

`load_command`

Contains fields that are common to all load commands.

```
struct load_command
{
    unsigned long cmd;
    unsigned long cmdsize;
};
```

Field Descriptions

`cmd`

An integer indicating the type of load command. [Table 3-2](#) (page 57) lists the valid load command types.

`cmdsize`

An integer specifying the total size in bytes of the load command data structure. Each load command structure contains a different set of data depending on the load command type, so each might have a different size. The size must always be a multiple of 4. This means the `cmdsize` field must always divide evenly into 4. If the load command data does not divide evenly into four, add bytes containing zeros to the end until it does.

Discussion

[Table 3-2](#) (page 57) lists the valid load command types, with links to the full data structures for each type.

Table 3-2 Mach-O load commands

Command	Data structure	Purpose
LC_SEGMENT	segment_command (page 58)	Defines a segment of this file to be mapped into the address space of the process that loads this file. It also includes all of the sections contained by the segment. See “Mach-O File Format Reference” (page 49).
LC_SYMTAB	symtab_command (page 72)	Specifies the symbol table for this file. This information is used by both static and dynamic linkers when linking the file, and also by debuggers to map symbols to the original source code files from which the symbols were generated.
LC_DYSYMTAB	dysymtab_command (page 76)	Specifies additional symbol table information used by the dynamic linker.
LC_THREAD LC_UNIXTHREAD	thread_command (page 68)	For an executable file, the <code>LC_UNIXTHREAD</code> command defines the initial thread state of the main thread of the process. <code>LC_THREAD</code> is like <code>LC_UNIXTHREAD</code> , but does not cause the kernel to allocate a stack.
LC_LOAD_DYLIB	dylib_command (page 66)	Defines the name of a dynamic shared library that this file links against.
LC_ID_DYLIB	dylib_command (page 66)	Specifies the install name of a dynamic shared library.
LC_PREBOUND_DYLIB	prebound_dylib_command (page 67)	For a shared library that this executable is linked prebound against, specifies the modules in the shared library that are used.

Command	Data structure	Purpose
LC_LOAD_DYLINKER	dylinker_command (page 67)	Specifies the dynamic linker that the kernel executes to load this file.
LC_ID_DYLINKER	dylinker_command (page 67)	Identifies this file as a dynamic linker.
LC_ROUTINES	routines_command (page 69)	Contains the offset of the shared library initialization routine (specified by the linker's <code>-init</code> option).
LC_TWOLEVEL_HINTS	twolevel_hints_command (page 63)	Contains the two-level namespace lookup hint table.
LC_SUB_FRAMEWORK	sub_framework_command (page 70)	Identifies this file as the implementation of a subframework of an umbrella framework. The name of the umbrella framework is stored in the string parameter.
LC_SUB_UMBRELLA	sub_umbrella_command (page 70)	Specifies a file that is a subumbrella of this umbrella framework.
LC_SUB_LIBRARY	sub_library_command (page 71)	Identifies this file as the implementation of a sublibrary of an umbrella framework. The name of the umbrella framework is stored in the string parameter. Note that Apple has not defined a supported location for sublibraries at this time.
LC_SUB_CLIENT	sub_client_command (page 71)	A subframework can explicitly allow another framework or bundle to link against it by including a <code>LC_SUB_CLIENT</code> load command containing the name of the framework or a client name for a bundle.

segment_command

Segments are defined by the `LC_SEGMENT` load command, which specifies a range of bytes in the file that are to be mapped by the loader into the address space of a program.

```
struct segment_command
{
    unsigned long cmd; /* LC_SEGMENT */
    unsigned long cmdsize; /* includes sizeof section structs */
    char segname[16]; /* segment name */
    unsigned long vmaddr; /* memory address of this segment */
    unsigned long vmsize; /* memory size of this segment */
    unsigned long fileoff; /* file offset of this segment */
    unsigned long filesize; /* amount to map from the file */
};
```

```

    vm_prot_t maxprot; /* maximum VM protection */
    vm_prot_t initprot; /* initial VM protection */
    unsigned long nsects; /* number of sections in segment */
    unsigned long flags; /* flags */
};

```

Field Descriptions

cmd

Common to all load command structures. Set to LC_SEGMENT for this structure.

cmdsize

Common to all load command structures. For this structure, set this field to `sizeof(segment_command)` plus the size of all of the section data structures that follow (`sizeof(segment_command) + (sizeof(section) * segment->nsect)`).

segname

A C string specifying the name of the segment. The value of this field can be any sequence of ASCII characters, although segment names defined by Apple begin with two underscores and consist of capital letters (as in `__TEXT` and `__DATA`). This field is fixed at 16 bytes in length.

vmaddr

Indicates the starting virtual memory address of this segment.

vmsize

Indicates the number of bytes of virtual memory occupied by this segment. See also the description of `filesize`, below.

fileoff

Indicates the offset in this file of the data to be mapped at `vmaddr`.

filesize

Indicates the number of bytes occupied by this segment on disk. For segments that require more memory at runtime than they do at build time, `vmsize` can be larger than `filesize`. For example, the `__PAGEZERO` segment generated by the linker for MH_EXECUTABLE files has a `vmsize` of 0x1000 but a `filesize` of zero. Because `__PAGEZERO` contains no data, there is no need for it to occupy any space until runtime. Also, the static linker often allocates uninitialized data at the end of the `__DATA` segment; in this case, the `vmsize` will be larger than the `filesize`. The loader guarantees that any memory of this sort is initialized with zeros.

maxprot

Specifies the maximum permitted virtual memory protections of this segment.

initprot

Specifies the initial virtual memory protections of this segment.

nsects

Indicates the number of section data structures following this load command.

flags

Defines a set of flags which affect the loading of this segment.

- `SG_HIGHVM`—The file contents for this segment are for the high part of the virtual memory space, the low part is zero filled (for stacks in core files).
- `SG_NORELOC`—This segment has nothing that was relocated in it and nothing relocated to it. It may be safely replaced without relocation.

section

Directly following a `segment_command` data structure is an array of `section` data structures, with the exact count determined by the `nsect` field of the `segment_command` structure.

```
struct section
{
    char sectname[16]; /* name of this section */
    char segname[16]; /* segment this section goes in */
    unsigned long addr; /* memory address of this section */
    unsigned long size; /* size in bytes of this section */
    unsigned long offset; /* file offset of this section */
    unsigned long align; /* section alignment (power of 2) */
    unsigned long reloff; /* file offset of relocation entries */
    unsigned long nreloc; /* number of relocation entries */
    unsigned long flags; /* flags (section type and attributes)*/
    unsigned long reserved1; /* reserved */
    unsigned long reserved2; /* reserved */
};
```

Field Descriptions

`sectname`

A string specifying the name of this section. The value of this field can be any sequence of ASCII characters, although section names defined by Apple begin with two underscores and consist of lowercase letters (as in `__text` and `__data`). This field is fixed at 16 bytes in length.

`segname`

A string specifying the name of the segment that should eventually contain this section. For compactness, intermediate object files—files of type `MH_OBJECT`—contain only one segment, in which all sections are placed. The static linker places each section in the named segment when building the final product (any file that is not of type `MH_OBJECT`).

`addr`

An integer specifying the virtual memory address of this section.

`size`

An integer specifying the size in bytes of the virtual memory occupied by this section.

`offset`

An integer specifying the offset in this file to the section.

`align`

An integer specifying the section's byte alignment. Specify this as a power of two; for example, a section with 8-byte alignment would have an `align` value of 3 (2 to the 3rd power equals 8).

`reloff`

An integer specifying the file offset of the first relocation entry for this section.

`nreloc`

An integer specifying the number of relocation entries located at `reloff` for this section.

flags

An integer divided into two parts. The least significant 8 bits contain the section type, while the most significant 24 bits contain a set of flags that specify other attributes of the section. These types and flags are primarily used by the static linker and file analysis tools such as `otool` to determine how to modify or display the section. These are the possible types:

- `S_REGULAR`—This section has no particular type. The standard tools create a `__TEXT, __text` section of this type.
- `S_ZEROFILL`—Zero fill-on-demand section—when this section is first read from or written to, each page within is automatically filled with bytes containing zero.
- `S_CSTRING_LITERALS`—This section contains only constant C strings. The standard tools create a `__TEXT, __cstring` section of this type.
- `S_4BYTE_LITERALS`—This section contains only constant values that are 4 bytes long. The standard tools create a `__TEXT, __literal4` section of this type.
- `S_8BYTE_LITERALS`—This section contains only constant values that are 8 bytes long. The standard tools create a `__TEXT, __literal8` section of this type.
- `S_LITERAL_POINTERS`—This section contains only pointers to constant values.
- `S_NON_LAZY_SYMBOL_POINTERS`—This section contains only nonlazy pointers to symbols. The standard tools create a section of the `__DATA, __nl_symbol_ptrs` section of this type.
- `S_LAZY_SYMBOL_POINTERS`—This section contains only lazy pointers to symbols. The standard tools create a `__DATA, __la_symbol_ptrs` section of this type.
- `S_SYMBOL_STUBS`—This section contains symbol stubs. The standard tools create `__TEXT, __symbol_stub` and `__TEXT, __picsymbol_stub` sections of this type. See [“Indirect Addressing”](#) (page 44) for more information.
- `S_MOD_INIT_FUNC_POINTERS`—This section contains pointers to module initialization functions. The standard tools create `__DATA, __mod_init_func` sections of this type.
- `S_MOD_TERM_FUNC_POINTERS`—This section contains pointers to module termination functions. The standard tools create `__DATA, __mod_term_func` sections of this type.
- `S_COALESCED`—This section contains symbols that are coalesced by the static linker and possibly the dynamic linker. More than one file may contain coalesced definitions of the same symbol without causing multiple-defined-symbol errors.

The following are the possible attributes of a section:

- `S_ATTR_PURE_INSTRUCTIONS`—This section contains only executable machine instructions. The standard tools set this flag for the sections `__TEXT, __text, __TEXT, __symbol_stub`, and `__TEXT, __picsymbol_stub`.
- `S_ATTR_NO_TOC`—This section contains coalesced symbols that must not be placed in the table of contents (SYMDEF member) of a static archive library.
- `S_ATTR_SOME_INSTRUCTIONS`—This section contains executable machine instructions and other data.

- `S_ATTR_EXT_RELOC`—This section contains references that must be relocated. These references refer to data that exists in other files (undefined symbols). To support external relocation, the maximum virtual memory protections of the segment that contains this section must allow both reading and writing.
- `S_ATTR_LOC_RELOC`—This section contains references that must be relocated. These references refer to data within this file.

`reserved1`

An integer reserved for use with certain section types. For symbol pointer sections and symbol stubs sections which refer to indirect symbol table entries, this is the index into the indirect table for this section's entries. The number of entries is based on the section size divided by the size of the symbol pointer or stub. Otherwise this field is set to zero.

`reserved2`

For sections of type `S_SYMBOL_STUBS`, an integer specifying the size (in bytes) of the symbol stub entries contained in the section. Otherwise, this field is reserved for future use and should be set to zero.

Discussion

Each section in a Mach-O file has both a type and a set of attribute flags. In intermediate object files, the type and attributes determine how the static linker copy the sections from intermediate object files into the final product. Object file analysis tools (such as `otool`) use the type and attributes to determine how to read and display the sections. The section type and attributes are not used by the dynamic linker. Descriptions for important variants of the symbol type and attributes as they apply to static linking follow:

- **Regular sections.** In a regular section, only one definition of an external symbol may exist in intermediate object files. The static linker returns an error if it finds any duplicate external symbol definitions.
- **Coalesced sections.** In the final product, the static linker retains only one instance of each symbol defined in coalesced sections. Some complex language features (such as C++ vtables and RTTI) require a definition of a particular symbol to be duplicated in every intermediate object file. This wastes a lot of space in the final product; to reduce the memory occupied by a program, the compiler can place symbol definitions in coalesced sections.
- **Coalesced weak definition sections** (available in Mac OS X 10.2 and later). When the static linker finds duplicate definitions for a symbol, it will discard any that are in a section that has the weak definitions attribute. If there are no non-weak definitions, the first weak definition is used instead. This feature is designed to support C++ templates; it allows explicit template instantiations to override implicit ones. The C++ compiler places the explicit definitions in a section that is not marked weak, and places the implicit ones in a coalesced section marked weak. Intermediate object files (and thus static archive libraries) built with weak definitions cannot be used with static linker versions prior to Mac OS X 10.2. Final products (applications and shared libraries) should not contain weak definitions, so they can usually be used on prior versions of Mac OS X.

`twolevel_hints_command`

The data structure of a `LC_TWOLEVEL_HINTS` load command.

```
struct twolevel_hints_command
{
    unsigned long cmd; /* LC_TWOLEVEL_HINTS */
};
```

```
    unsigned long cmdsize; /* sizeof(struct twolevel_hints_command) */  
    unsigned long offset; /* offset to the hint table */  
    unsigned long nhints; /* number of hints in the hint table */  
};
```

Field Descriptions

cmd

Common to all load command structures. Set to LC_TWOLEVEL_HINTS for this structure.

cmdsize

Common to all load command structures. For this structure, set to sizeof(twolevel_hints_command).

offset

An integer specifying the byte offset from the start of this file to an array of twolevel_hint data structures, known as the two-level namespace hint table.

nhints

The number of twolevel_hint data structures located at offset.

Discussion

The static linker adds the LC_TWOLEVEL_HINTS load command and the two-level namespace hint table to the output file when building a two-level namespace image.

Special Considerations

By default, `ld` does not include the LC_TWOLEVEL_HINTS command or the two-level namespace hint table in an MH_BUNDLE file, because the presence of this load command causes the version of the dynamic linker shipped with Mac OS X 10.0 to crash. If you know the code will run only on Mac OS X 10.1 and later, you should explicitly enable the two-level namespace hint table. See the `ld` man page, specifically about the `-twolevel_namespace_hints` option, for more information.

twolevel_hint

Specifies an entry in the two-level namespace hint table.

```
struct twolevel_hint  
{  
    unsigned long isub_image:8,  
                itoc:24;  
};
```

Field Descriptions

isub_image

The subimage in which the symbol is defined. It is an index into the subimage list of the symbol's image (which is specified by the high eight bits of the `n_desc` field of the symbol data structure—see [nlist](#) (page 73)). If this field is zero, the symbol is assumed to be in the umbrella image itself. If the symbol is not from an umbrella framework or library, `isub_image` must be zero.

`itoc`

The symbol index into the table of contents of the image specified by the `isub_image` field

Discussion

The two-level namespace hint table provides the dynamic linker with suggested positions to start searching for symbols in the libraries the current image is linked against.

Every undefined symbol (that is, every symbol of type `N_UNDF` or `N_PBUD`) in a two-level namespace image must have a corresponding entry in the two-level hint table, at the same index.

The static linker adds the `LC_TWOLEVEL_HINTS` load command and the two-level namespace hint table to the output file when building a two-level namespace image.

By default, the linker does not include the `LC_TWOLEVEL_HINTS` command or the two-level namespace hint table in an `MH_BUNDLE` file, because the presence of this load command causes the version of the dynamic linker shipped with Mac OS X 10.0 to crash. If you know the code will run only on Mac OS X 10.1 and later, you should explicitly enable the two-level namespace hints. See the linker documentation for more information.

`lc_str`

Defines a variable-length string.

```
union lc_str
{
    unsigned long offset;
    char* ptr;
};
```

Field Descriptions

`offset`

A long integer. A byte offset from the start of the load command that contains this string to the start of the string data.

`ptr`

A pointer to an array of bytes. At runtime, this pointer contains the virtual memory address of the string data.

Discussion

Load commands store variable-length data such as library names using the `lc_str` data structure. Unless otherwise specified, the data consists of a C string.

The data pointed to is stored just after the load command, and the size is added to the size of the load command. You can determine the size of the string by subtracting the size of the load command data structure from the `cmdsize` field of the load command data structure.

`dylib`

Defines the data used by the dynamic linker to match a shared library against the files that have linked to it. Used exclusively in the `dylib_command` data structure.

```
struct dylib
{
```

```
union lc_str name;  
unsigned long timestamp;  
unsigned long current_version;  
unsigned long compatibility_version;  
};
```

Field Descriptions

name

A data structure of type `lc_str` (page 65). Specifies the name of the shared library.

timestamp

The date and time when the shared library was built.

current_version

The current version of the shared library.

compatibility_version

The compatibility version of the shared library.

Discussion

In order for the dynamic linker to successfully link an file to a dynamic library at runtime, two criteria must be met:

- The name for the shared library must match exactly the install name previously recorded by the static linker in the file, or as modified by the various dynamic linker environment variables (see the `dyld` man page for more information.)
- The compatibility version for the shared library must be less than or equal to the compatibility version recorded by the static linker in the image.

The dynamic linker uses the timestamp to determine whether it can use the prebinding information. The current version is returned by the function `NSVersionOfRunTimeLibrary` to allow you to determine the version of the library your program is using.

dylib_command

The data structure for the `LC_LOAD_DYLIB` and `LC_ID_DYLIB` load commands.

```
struct dylib_command  
{  
    unsigned long cmd; /* LC_ID_DYLIB or LC_LOAD_DYLIB */  
    unsigned long cmdsize; /* includes pathname string */  
    struct dylib dylib; /* the library identification */  
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to either `LC_LOAD_DYLIB` or `LC_ID_DYLIB`.

`cmdsize`

Common to all load command structures. For this structure, set to `sizeof(dylib_command)`, plus the size of the data pointed to by the `name` field of the `dylib` field.

`dylib`

A data structure of type `dylib` (page 65). Specifies the attributes of the shared library.

Discussion

The static linker adds a `LC_ID_DYLIB` load command to shared libraries to identify the linking attributes of the library.

The static linker adds one `LC_LOAD_DYLIB` load command for each shared library that a file links against. All the `LC_LOAD_DYLIB` commands together form a list that is ordered according to location in the file, earliest `LC_LOAD_DYLIB` command first. For two-level namespace files, undefined symbol entries in the symbol table refer to their parent shared libraries by index into this list. The index is called a *library ordinal*, and it is stored in the `n_desc` field of the `nlist` (page 73) data structure.

dylinker_command

The data structure for the `LC_LOAD_DYLINKER` and `LC_ID_DYLINKER` load commands.

```
struct dylinker_command
{
    unsigned long cmd; /* LC_ID_DYLINKER or LC_LOAD_DYLINKER */
    unsigned long cmdsize; /* includes pathname string */
    union lc_str name; /* dynamic linker's path name */
};
```

Field Descriptions

`cmd`

Common to all load command structures. For this structure, set to either `LC_ID_DYLINKER` or `LC_LOAD_DYLINKER`.

`cmdsize`

Common to all load command structures. For this structure, set to `sizeof(dylinker_command)`, plus the size of the data pointed to by the `name` field.

`name`

A data structure of type `lc_str` (page 65). Specifies the name of the dynamic linker.

Discussion

Every executable file that is dynamically linked contains a `LC_LOAD_DYLINKER` command that specifies the name of the dynamic linker that the kernel must load in order to execute the file. The dynamic linker itself specifies its name using the `LC_ID_DYLINKER` load command.

prebound_dylib_command

The data structure for the `LC_PREBOUND_DYLIB` load command. For every library that a prebound executable file links to, the static linker adds one `LC_PREBOUND_DYLIB` command.

```
struct prebound_dylib_command
{
    unsigned long cmd; /* LC_PREBOUND_DYLIB */
    unsigned long cmdsize; /* includes strings */
    union lc_str name; /* library's path name */
    unsigned long nmodules; /* number of modules in library */
    union lc_str linked_modules; /* bit vector of linked modules */
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to LC_PREBOUND_DYLIB

cmdsize

Common to all load command structures. For this structure, set to `sizeof(prebound_dylib_command)`, plus the size of the data pointed to by the `name` and `linked_modules` fields.

name

A data structure of type `lc_str` (page 65). Specifies the name of the prebound shared library.

nmodules

An integer. Specifies the number of modules the prebound shared library contains. The size of the `linked_modules` string is $(nmodules / 8) + (nmodules \% 8)$.

linked_modules

A data structure of type `lc_str` (page 65). Usually, this data structure defines the offset of a C string; in this usage, it is a variable-length bitset, containing one bit for each module. Each bit represents whether the corresponding module is linked to a module in the current file, 1 for yes, zero for no. The bit for the first module is the low bit of the first byte.

thread_command

The data structure for the LC_THREAD and LC_UNIXTHREAD load commands. The data of this command is specific to each CPU architecture, and appears in the header `thread_status.h` located in the CPU's directory in `/usr/include/mach/`.

```
struct thread_command
{
    unsigned long cmd;
    unsigned long cmdsize;
    /* unsigned long flavor; */
    /* unsigned long count; */
    /* struct cpu_thread_state state; */
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to LC_THREAD or LC_UNIXTHREAD.

`cmdsize`

Common to all load command structures. For this structure, set to `sizeof(thread_command)` plus the size of the `flavor` and `count` fields, plus the size of the CPU-specific thread state data structure.

`flavor`

Integer specifying the particular flavor of the thread state data structure. See the `thread_status.h` header file for your CPU architecture.

`count`

Size of the thread state data, in number of 32-bit integers. The thread state data structure must be fully padded to 32-bit alignment.

`routines_command`

The data structure for the `LC_ROUTINES` load command. Describes the location of the shared library initialization function, which is a function that the dynamic linker calls before allowing any of the routines in the library to be called.

```
struct routines_command
{
    unsigned long cmd;
    unsigned long cmdsize;
    unsigned long init_address;
    unsigned long init_module;
    unsigned long reserved1;
    unsigned long reserved2;
    unsigned long reserved3;
    unsigned long reserved4;
    unsigned long reserved5;
    unsigned long reserved6;
};
```

Field Descriptions

`cmd`

Common to all load command structures. For this structure, set to `LC_ROUTINES`

`cmdsize`

Common to all load command structures. For this structure, set to `sizeof(routines_command)`.

`init_address`

An integer specifying the virtual memory address of the initialization function.

`init_module`

An integer specifying the index into the module table of the module containing the initialization function.

`reserved1`

Reserved for future use. Set this field to zero.

reserved2
Reserved for future use. Set this field to zero.

reserved3
Reserved for future use. Set this field to zero.

reserved4
Reserved for future use. Set this field to zero.

reserved5
Reserved for future use. Set this field to zero.

reserved6
Reserved for future use. Set this field to zero.

Discussion

The static linker adds an `LC_ROUTINES` command when you specify a shared library initialization function using the `-init` parameter.

sub_framework_command

The data structure for the `LC_SUB_FRAMEWORK` load command. Identifies the umbrella framework of which this file is a subframework.

```
struct sub_framework_command
{
    unsigned long cmd;
    unsigned long cmdsize;
    union lc_str umbrella;
};
```

Field Descriptions

cmd
Common to all load command structures. For this structure, set to `LC_SUB_FRAMEWORK`.

cmdsize
Common to all load command structures. For this structure, set to `sizeof(sub_framework_command)` plus the size of the data pointed to by the `umbrella` field.

umbrella
A data structure of type `lc_str` (page 65). Specifies the name of the umbrella framework of which this file is a member.

sub_umbrella_command

The data structure for the `LC_SUB_UMBRELLA` load command. Identifies the named framework as a subumbrella of this framework. Unlike a subframework, any client may link to a subumbrella.

```
struct sub_umbrella_command
{
    unsigned long cmd;
```

```
    unsigned long cmdsize;  
    union lc_str sub_umbrella;  
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to LC_SUB_UMBRELLA

cmdsize

Common to all load command structures. For this structure, set to `sizeof(sub_umbrella_command)` plus the size of the data pointed to by the `sub_umbrella` field.

sub_umbrella

A data structure of type `lc_str` (page 65). Specifies the name of the umbrella framework of which this file is a member.

sub_library_command

The data structure for the LC_SUB_LIBRARY load command. Identifies a sublibrary of this framework, and marks this framework as an umbrella framework. Unlike a subframework, any client may link to a sublibrary.

```
struct sub_library_command  
{  
    unsigned long cmd;  
    unsigned long cmdsize;  
    union lc_str sub_library;  
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to LC_SUB_LIBRARY.

cmdsize

Common to all load command structures. For this structure, set to `sizeof(sub_library_command)` plus the size of the data pointed to by the `sub_library` field.

sub_library

A data structure of type `lc_str` (page 65). Specifies the name of the framework of which this file is a member.

sub_client_command

The data structure for the LC_SUB_CLIENT load command. Specifies the name of a file that is allowed to link to this subframework; this file would otherwise be required to link to the umbrella framework of which this file is a component.

```
struct sub_client_command  
{  
    unsigned long cmd;
```

```
    unsigned long cmdsize;  
    union lc_str sub_client;  
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to LC_SUB_CLIENT

cmdsize

Common to all load command structures. For this structure, set to `sizeof(sub_client_command)` plus the size of the data pointed to by the `sub_client` field.

sub_client

A data structure of type `lc_str` (page 65). Specifies the name of a client authorized to link to this library.

Special Considerations

`ld` generates a `sub_client_command` load command in the built product if you pass the option `-allowable_client_name name`, where `name` is the install name of a framework or the client name of a bundle. See the `ld` man page, specifically about the options `-allowable_client_name` and `-client_name`, for more information.

Symbol Table and Related Data Structures

Two load commands, LC_SYMTAB and LC_DYSYMTAB, describe the size and location of the symbol tables, along with additional metadata. The other data structures listed in this section represent the symbol tables themselves.

symtab_command

The data structure for the LC_SYMTAB load command. Describes the size and location of the symbol table data structures.

```
struct symtab_command  
{  
    unsigned long cmd; /* LC_SYMTAB */  
    unsigned long cmdsize; /* sizeof(struct symtab_command) */  
    unsigned long symoff; /* symbol table offset */  
    unsigned long nsyms; /* number of symbol table entries */  
    unsigned long stroff; /* string table offset */  
    unsigned long strsize; /* string table size in bytes */  
};
```

Field Descriptions

cmd

Common to all load command structures. For this structure, set to LC_SYMTAB

cmdsize

Common to all load command structures. For this structure, set to `sizeof(symtab_command)`.

symoff

An integer containing the byte offset from the start of the file to the location of the symbol table entries. The symbol table is an array of [nlist](#) (page 73) data structures.

nsyms

An integer indicating the number of entries in the symbol table.

stroff

An integer containing the byte offset from the start of the image to the location of the string table.

strsize

An integer indicating the size (in bytes) of the string table.

Discussion

LC_SYMTAB should exist in both statically linked and dynamically linked file types.

nlist

Describes an entry in the symbol table. Declared in the header `/usr/include/mach-o/nlist.h`.

```
struct nlist
{
    union {
        char *n_name;
        long n_strx;
    } n_un;
    unsigned char n_type;
    unsigned char n_sect;
    short n_desc;
    unsigned long n_value;
};
```

Field Descriptions

n_un

A union that holds an index into the string table, `n_strx`. To specify an empty string (""), set this value to zero. The `n_name` field is not currently used in Mach-O.

n_type

A byte value consisting of data accessed using four bit masks. The masks are used as follows:

- **N_STAB (0xe0)** If any of these 3 bits are set, the symbol is a symbolic debugging table (stab) entry. In that case, the entire n_type field is interpreted as a stab value. See the header `/usr/include/mach-o/stab.h` for valid stab values.
- **N_PEXT (0x10)**. If this bit is on, this symbol is marked as having limited global scope. When the file is fed to the static linker, it clears the N_EXT bit for each symbol with the N_PEXT bit set. (The ld option `-keep_private_externs` disables this behavior.) With Mac OS X GCC, you can use the `__private_extern__` function attribute to set this bit.
- **N_TYPE (0x0e)** These bits define the type of the symbol.
- **N_EXT (0x01)** If this bit is on, this symbol is an external symbol, a symbol that is either defined outside of this file or that is defined in this file, but can be referenced by other files.

Values for the N_TYPE field include the following:

- **N_UNDF (0x0)**—the symbol is undefined. Undefined symbols are symbols referenced in this module but are defined in a different module. Set the n_sect field to NO_SECT.
- **N_ABS (0x2)** —the symbol is absolute. The linker does not update the value of an absolute symbol. Set the n_sect field to NO_SECT.
- **N_SECT (0xe)** —the symbol is defined in the section number given in n_sect.
- **N_PBUD (0xc)** —the symbol is undefined and the image is using a prebound value for the symbol. Set the n_sect field to NO_SECT.
- **N_INDR (0xa)** —the symbol is defined to be the same as another symbol. The n_value field is an index into the string table specifying the name of the other symbol. When that symbol is linked, both this and the other symbol point to the same defined type and value.

n_sect

An integer specifying the number of the section that this symbol can be found in, or NO_SECT if the symbol is not to be found in any section of this image. The sections are contiguously numbered across segments, starting from 1, according to the order they appear in the LC_SEGMENT load commands.

`n_desc`

A 16-bit value providing additional information about the nature of this symbol. The reference flag can be accessed using the `REFERENCE_TYPE` mask (0xF) and are defined as follows:

- `REFERENCE_FLAG_UNDEFINED_NON_LAZY` (0x0)—This symbol is a reference to an external non-lazy (data) symbol.
- `REFERENCE_FLAG_UNDEFINED_LAZY` (0x1)—This symbol is a reference to an external lazy symbol—that is, to a function call.
- `REFERENCE_FLAG_DEFINED` (0x2)—This symbol is defined in this module.
- `REFERENCE_FLAG_PRIVATE_DEFINED` (0x3)—This symbol is defined in this module and is visible only to modules within this shared library.
- `REFERENCE_FLAG_PRIVATE_UNDEFINED_NON_LAZY` (0x4)—This symbol is defined in another module in this file, is a non-lazy (data) symbol, and is visible only to modules within this shared library.
- `REFERENCE_FLAG_PRIVATE_UNDEFINED_LAZY` (0x5)—This symbol is defined in another module in this file, is a lazy (function) symbol, and is visible only to modules within this shared library.

Additionally, the following bits might also be set:

- `REFERENCED_DYNAMICALY` (0x10) must be set for any symbol that might be referenced by another image. The `strip` tool uses this bit to avoid removing symbols that must exist: If the symbol has this bit set, `strip` does not strip it.
- `N_DESC_DISCARDED` (0x20) is used by the dynamic linker at runtime; do not set this bit.
- `N_WEAK_REF` (0x40) indicates that this symbol is a weak reference; if the dynamic linker cannot find a definition for this symbol, it set the address of this symbol to zero. The static linker sets this symbol given the appropriate weak-linking flags.
- `N_WEAK_DEF` (0x80) indicates that this symbol is a weak definition; if the static linker or the dynamic linker finds another (non-weak) definition for this symbol, the weak definition is ignored.

If this file is a two-level namespace image (that is, if the `MH_TWOLEVEL` flag of the `mach_header` structure is set), the high 8 bits of `n_desc` specify the number of the library in which this symbol is defined. Use the macro `GET_LIBRARY_ORDINAL` to obtain this value, and the macro `SET_LIBRARY_ORDINAL` to set it. A zero value indicates the current image. 1 through 254 specify the library number according to the order of `LC_LOAD_DYLIB` commands in the file. For plug-ins that load symbols from the executable program they are linked against, 255 specifies the executable image. For flat namespace images, the high 8 bits must be zero.

`n_value`

An integer that contains the value of the symbol. This format of this value is different for each type of symbol table entry (as specified by the `n_type` field). For the `N_SECT` symbol type, `n_value` is the address of the symbol. See the description of the `n_type` field for information on other possible values.

Discussion

Common symbols must be of type `N_UNDF` and must have the `N_EXT` bit set. The `n_value` for a common symbol is the size (in bytes) of the data of the symbol. In C, a common symbol is a variable that is declared but not initialized in this file. Common symbols can only appear in `MH_OBJECT`-type Mach-O files.

dysymtab_command

The data structure for the `LC_DYSYMTAB` load command. Describes the sizes and locations of the parts of the symbol table used for dynamic linking.

```
struct dysymtab_command
{
    unsigned long cmd; /* LC_DYSYMTAB */
    unsigned long cmdsize; /* sizeof(struct dysymtab_command) */
    unsigned long ilocalsym; /* index to local symbols */
    unsigned long nlocalsym; /* number of local symbols */
    unsigned long iextdefsym; /* index to externally defined symbols */
    unsigned long nextdefsym; /* number of externally defined symbols */
    unsigned long iundefsym; /* index to undefined symbols */
    unsigned long nundefsym; /* number of undefined symbols */
    unsigned long tocoff; /* file offset to table of contents */
    unsigned long ntoc; /* # of entries in table of contents */
    unsigned long modtaboff; /* file offset to module table */
    unsigned long nmodtab; /* # of module table entries */
    unsigned long extrefoff; /* offset to referenced symbol table */
    unsigned long nextrefsyms; /* # of referenced symbol table entries */
    unsigned long indirectsymoff; /* f-offset to the indirect symbol table */
    unsigned long nindirectsyms; /* #of indirect symbol table entries */
    unsigned long extreloff; /* offset to external relocation entries */
    unsigned long nextrel; /* number of external relocation entries */
    unsigned long locreloff; /* offset to local relocation entries */
    unsigned long nlocalrel; /* number of local relocation entries */
};
```

Field Descriptions

`cmd`

Common to all load command structures. For this structure, set to `LC_DYSYMTAB`

`cmdsize`

Common to all load command structures. For this structure, set to `sizeof(dysymtab_command)`.

`ilocalsym`

An integer indicating the index of the first symbol in the group of local symbols.

`nlocalsym`

An integer indicating the total number of symbols in the group of local symbols.

`iextdefsym`

An integer indicating the index of the first symbol in the group of defined external symbols.

nextdefsym

An integer indicating the total number of symbols in the group of defined external symbols.

iundefsym

An integer indicating the index of the first symbol in the group of undefined external symbols.

nundefsym

An integer indicating the total number of symbols in the group of undefined external symbols.

tocoff

An integer indicating the byte offset from the start of the file to the table of contents data.

ntoc

An integer indicating the number of entries in the table of contents.

modtaboff

An integer indicating the byte offset from the start of the file to the module table data.

nmodtab

An integer indicating the number of entries in the module table.

extrefsymoff

An integer indicating the byte offset from the start of the file to the external reference table data.

nextrefsyms

An integer indicating the number of entries in the external reference table.

indirectsymoff

An integer indicating the byte offset from the start of the file to the indirect symbol table data

nindirectsyms

An integer indicating the number of entries in the indirect symbol table

extreloff

An integer indicating the byte offset from the start of the file to the external relocation table data

nextrel

An integer indicating the number of entries in the external relocation table

locreloff

An integer indicating the byte offset from the start of the file to the local relocation table data.

nlocrel

An integer indicating the number of entries in the local relocation table.

Discussion

The `LC_DYSYMTAB` load command contains a set of indexes into the symbol table and a set of file offsets that define the location of several other tables. Fields for tables not used in the file should be set to zero. These tables are described in the section [“Indirect Addressing”](#) (page 44).

dylib_table_of_contents

Describes an entry in the table of contents of a dynamic shared library.

```
struct dylib_table_of_contents
{
    unsigned long symbol_index;
    unsigned long module_index;
};
```

Field Descriptions

`symbol_index`

An index into the symbol table indicating the defined external symbol to which this entry refers.

`module_index`

An index into the module table indicating the module in which this defined external symbol is defined.

dylib_module

Describes a module table entry for a dynamic shared library.

```
struct dylib_module
{
    unsigned long module_name;
    unsigned long iextdefsym;
    unsigned long nextdefsym;
    unsigned long irefsym;
    unsigned long nrefsym;
    unsigned long ilocalsym;
    unsigned long nlocalsym;
    unsigned long iextrel;
    unsigned long nextrel;
    unsigned long iinit_iterm;
    unsigned long ninit_nterm;
    unsigned long objc_module_info_addr;
    unsigned long objc_module_info_size;
};
```

Field Descriptions

`module_name`

An index to an entry in the string table indicating the name of the module.

`iextdefsym`

The index into the symbol table of the first defined external symbol provided by this module.

nextdefsym

The number of defined external symbols provided by this module.

irefsym

The index into the external reference table of the first entry provided by this module.

nrefsym

The number of external reference entries provided by this module.

ilocalsym

The index into the symbol table of the first local symbol provided by this module.

nlocalsym

The number of local symbols provided by this module.

ixtrel

The index into the external relocation table of the first entry provided by this module.

nextrel

The number of entries in the external relocation table that are provided by this module.

init_iterm

Contains both the index into the module initialization section (the low 16 bits) and the index into the module termination section (the high 16 bits) to the pointers for this module.

ninit_termin

Contains both the number of pointers in the module initialization (the low 16 bits) and the number of pointers in the module termination section (the high 16 bits) for this module.

objc_module_info_addr

The statically linked address of the start of the data for this module in the `__module_info` section of the `__OBJC` segment.

objc_module_info_size

The number of bytes of data for this module that are used in the `__module_info` section of the `__OBJC` segment.

dylib_reference

The structure of an external reference table entry for the external reference entries provided by a module in a shared library.

```
struct dylib_reference
{
    unsigned long isym:24, /* index into the symbol table */
                 flags:8; /* flags to indicate the type of reference */
};
```

Field Descriptions

isym

An index into the symbol table for the symbol being referenced.

flags

A constant for the type of reference being made. Use the same `REFERENCE_FLAG` constants as described in the `nlist` (page 73) structure description.

Relocation Data Structures

Relocation is the process of moving symbols to a different address. When the static linker moves a symbol (a function or an item of data) to a different address, it needs to change all of the references to that symbol to use the new address. The **relocation entries** in a Mach-O file contain offsets in the file to addresses that need to be relocated when the contents of the file are relocated. The addresses are usually relative offsets stored in CPU instructions; the exact format of the address is specified in each relocation entry. When creating the intermediate object file, the compiler generates one relocation entry for every instruction that contains a relative address. The static linker typically removes the relocation entries when building the final product, as relocation local to a single Mach-O file does not usually occur at runtime.

`relocation_info`

Describes an item in the file that uses an address that needs to be updated if the address is changed. This data structure is declared in the header `/usr/include/mach-o/reloc.h`.

```
struct relocation_info
{
    long r_address;
    unsigned int r_symbolnum:24,
               r_pcrel:1,
               r_length:2,
               r_extern:1,
               r_type:4;
};
#define R_ABS 0 /* absolute relocation type for Mach-O files */
```

Field Descriptions

`r_address`

In `MH_OBJECT` files, this is an offset from the start of the section to the item containing the address requiring relocation. If the high bit of this field is set (which you can check using the `R_SCATTERED` bit mask), the `relocation_info` structure is actually a `scattered_relocation_info` (page 81) structure.

In images used by the dynamic linker, this is an offset from the virtual memory address of the data of the first `segment_command` (page 58) that appears in the file (not necessarily the one with the lowest address). For images with the `MH_SPLIT_SEGS` flag set, this is an offset from the virtual memory address of data of the first read/write `segment_command` (page 58).

`r_symbolnum`

Indicates either an index into the symbol table (when the `r_extern` field is set to 1) or a section number (when the `r_extern` field is set to zero). As previously mentioned, sections are ordered from 1 to 255 in the order in which they appear in the `LC_SEGMENT` load commands.

r_pcrel

Indicates whether the item containing the address to be relocated is part of a CPU instruction that uses PC-relative addressing.

For addresses contained in PC-relative instructions, the CPU adds the address of the instruction to the address contained in the instruction.

r_length

Indicates the length of item containing the address to be relocated. A value of zero indicates a single byte; a value of 1 indicates a 2-byte address, and a value of 2 indicates a 4-byte address.

r_extern

Indicates whether the `r_symbolnum` field is an index into the symbol table (1) or a section number (zero).

r_type

Indicates the type of relocation to be performed. Possible values for this field are shared between this structure and the `scattered_relocation_info` (page 81) data structure; see the description of the `r_type` field in the `scattered_relocation_info` (page 81) data structure for more details.

`scattered_relocation_info`

Describes an item in the file that uses a non-zero constant in its relocatable expression or two addresses in its relocatable expression that needs to be updated if the addresses being used are changed. This information is needed to reconstruct the addresses that make up the relocatable expression's value in order to change the addresses independently of each other. This data structure is declared in the header `/usr/include/mach-o/reloc.h`.

```
struct scattered_relocation_info
{
#ifdef __BIG_ENDIAN__
    unsigned int r_scattered:1,
                r_pcrel:1,
                r_length:2,
                r_type:4,
                r_address:24;
    long r_value;
#endif /* __BIG_ENDIAN__ */
#ifdef __LITTLE_ENDIAN__
    unsigned int r_address:24,
                r_type:4,
                r_length:2,
                r_pcrel:1,
                r_scattered:1;
    long r_value;
#endif /* __LITTLE_ENDIAN__ */
};
```

Field Descriptions

`r_address`

In `MH_OBJECT` files, this is an offset from the start of the section to the item containing the address requiring relocation. If the high bit of this field is clear (which you can check using the `R_SCATTERED` bit mask), this structure is actually a `relocation_info` (page 80) structure.

In images used by the dynamic linker, this is an offset from the virtual memory address of the data of the first `segment_command` (page 58) that appears in the file (not necessarily the one with the lowest address). For images with the `MH_SPLIT_SEGS` flag set, this is an offset from the virtual memory address of data of the first read/write `segment_command` (page 58).

Since the `r_address` field is only 24 bits long, the offset in this field can never be larger than `0x00FFFFFF`, thus limiting the size of the relocatable contents of this image to 16 megabytes.

`r_value`

The address of the relocatable expression for the item in the file that needs to be updated if the address is changed. For relocatable expressions with the difference of two section addresses, the address from which to subtract (in mathematical terms, the minuend) is contained in the first relocation entry and the address to subtract (the subtrahend) is contained in the second relocation entry.

For x86 processors, the `r_type` field may contain any of the values listed below.

- `GENERIC_RELOC_VANILLA`. A generic relocation entry for both addresses contained in data and addresses contained in CPU instructions.
- `GENERIC_RELOC_PAIR`. The second relocation entry of a pair.
- `GENERIC_RELOC_SECTDIFF`. A relocation entry for an item that contains the difference of two section addresses. This is generally used for position-independent code generation. `GENERIC_RELOC_SECTDIFF` contains the address from which to subtract; it must be followed by a `GENERIC_RELOC_PAIR` containing the address to subtract.
- `GENERIC_RELOC_PB_LA_PTR`. A relocation entry for a prebound lazy pointer. This is always a scattered relocation entry. The `r_value` field contains the non-prebound value of the lazy pointer.

For PowerPC processors, the `r_type` field is usually `PPC_RELOC_VANILLA` for addresses contained in data. Relocation entries for addresses contained in CPU instructions are described by other `r_type` values, as follows.

- `PPC_RELOC_PAIR`. The second relocation entry of a pair. A `PPC_RELOC_PAIR` entry must follow each of the other relocation entry types, except for `PPC_RELOC_VANILLA`, `PPC_RELOC_BR14`, `PPC_RELOC_BR24`, and `PPC_RELOC_PB_LA_PTR`.
- `PPC_RELOC_BR14`—The instruction contains a 14-bit branch displacement.
- `PPC_RELOC_BR24`—The instruction contains a 24-bit branch displacement.
- `PPC_RELOC_HI16`—The instruction contains the high 16 bits of a relocatable expression. The next relocation entry must be a `PPC_RELOC_PAIR` specifying the low 16 bits of the expression in the low 16 bits of the `r_value` field.
- `PPC_RELOC_L016`—The instruction contains the low 16 bits of an address. The next relocation entry must be a `PPC_RELOC_PAIR` specifying the high 16 bits of the expression in the low (not the high) 16 bits of the `r_value` field.
- `PPC_RELOC_HA16`—Same as the `PPC_RELOC_HI16` except the low 16 bits and the high 16 bits are added together with the low 16 bits sign-extended first. This means if bit 15 of the low 16 bits is set, the high 16 bits stored in the instruction will be adjusted.
- `PPC_RELOC_L014`—Same as `PPC_RELOC_L016` except that the low 2 bits are not stored in the CPU instruction and are always zero. `PPC_RELOC_L014` is used in 64-bit load/store instructions.
- `PPC_RELOC_SECTDIFF`—A relocation entry for an item that contains the difference of two section addresses. This is generally used for position-independent code generation. `PPC_RELOC_SECTDIFF` contains the address from which to subtract; it must be followed by a `PPC_RELOC_PAIR` containing the section address to subtract.

- `PPC_RELOC_PB_LA_PTR`—A relocation entry for a prebound lazy pointer. This is always a scattered relocation entry. The `r_value` field contains the non-prebound value of the lazy pointer.
- `PPC_RELOC_HI16_SECTDIFF`—Section difference form of `PPC_RELOC_HI16`.
- `PPC_RELOC_L016_SECTDIFF`—Section difference form of `PPC_RELOC_L016`.
- `PPC_RELOC_HA16_SECTDIFF`—Section difference form of `PPC_RELOC_HA16`.
- `PPC_RELOC_JBSR`—A relocation entry for the assembler synthetic opcode `jbsr`, which is a 24-bit branch-and-link instruction using a branch island. The branch displacement is assembled to the branch island address and the relocation entry indicates the actual target symbol. If the linker can make the branch reach the actual target symbol that is done. Otherwise, the branch is relocated to the branch island.

Discussion

Mach-O relocation data structures support two different types of relocatable expressions in machine code and data:

- **symbol address + constant.** The most typical form of relocation, adding a simple constant value to the existing address.
- **address of section *y* - address of section *x* + constant.** The section difference form of relocation. This form of relocation supports position-independent code.

Static Archive Libraries

This section describes the file format used for static archive libraries, which are described in the section [“Static Archive Libraries”](#) (page 15). Mac OS X uses a format derived from the original BSD static archive library format, with a few minor additions. See the discussion for the `ranlib` data structure for more information.

`ranlib`

The data structure of a static archive library symbol table entry. Declared in the header file `/usr/include/mach-o/ranlib.h`.

```
struct ranlib
{
    union
    {
        unsigned long ran_strx; /* string table index of */
        char * ran_name; /* symbol defined by */
    } ran_un;
    unsigned long ran_off; /* library member at this offset */
};
```

Field Descriptions

`ran_strx`

The index number (zero-based) of the string in the string table that follows the array of `ranlib` data structures.

`ran_name`

The byte offset, from the start of the file, at which the symbol name can be found.

`ran_off`

The byte offset, from the start of the file, at which the header line for the member containing this symbol can be found.

Discussion

A static archive library begins with the file identifier string `!<arch>`, followed by a newline character (ASCII value 0x0A). The file identifier string is followed by a series of member files. Each member consists of a fixed-length header line followed by the file data. The header line is 60 bytes long and is divided into five fixed-length fields, as shown in this sample header line:

```
grapple.c      999514211   501   20   100644  167   `
```

The last 2 bytes of the header line are a backquote (```) character (ASCII value 0x60) and a newline character. All header fields are defined in ASCII and padded with spaces to the full length of the field. All fields are defined in decimal notation, except for the file mode field, which is defined in octal. The fields are described in depth as follows:

- The name field (16 bytes) contains the name of the file. If the name is either longer than 16 bytes, or contains a space character, the actual name should be written directly after the header line and the name field should contain the string `#1/` followed by the length. To keep the archive entries aligned to 4 byte boundaries, length of the name that follows the `#1/` is rounded to 4 bytes and the name that follows the header is padded with null bytes.
- The modified date field (12 bytes) is taken from the `st_time` field returned by the `stat` system call.
- The user ID field (6 bytes) is taken from the `st_uid` field returned by the `stat` system call.
- The group ID field (6 bytes) is taken from the `st_gid` field returned by the `stat` system call.
- The file mode field (8 bytes) is taken from the `st_mode` field returned by the `stat` system call. This field is written in octal notation.
- The file size field (8 bytes) is taken from the `st_size` field returned by the `stat` system call.

The first member in a static archive library is always the symbol table describing the contents of the rest of the member files. This member is always called either `__SYMDEF` or `__SYMDEF SORTED`. (note the two leading underscores and period). The name used depends on the sort order of the symbol table. The older variant—`__SYMDEF`—contains entries in the same order that they appear in the object files. The newer variant—`__SYMDEF SORTED`—contains entries in alphabetical order, which allows the static linker to load the symbols faster.

The `__SYMDEF` and `__SORTED SYMDEF` archive members contain an array of `ranlib` data structures, preceded by the length (in bytes) (a long integer, 4 bytes) of the number of items in the array. The array is followed by a string table of null-terminated strings, which are preceded by the length in bytes of the entire string table (again, a 4-byte long integer).

The string table is an array of C strings, each terminated by a null byte.

The `ranlib` declarations can be found in the header `/usr/include/mach-o/ranlib.h`.

Special Considerations

Prior to the advent of `libtool`, a tool called `ranlib` was used to generate the symbol table. `ranlib` has since been integrated into `libtool`. See the man page for `libtool` for more information.

Multi-CPU Architecture Files

The standard development tools normally accept multiple-CPU (or “fat”) files as parameters wherever a normal Mach-O file or static archive library is accepted. The dynamic linker will load the correct data for the currently-running CPU from fat shared libraries, frameworks, and bundles.

A fat file is not really a Mach-O file at all. It is a simple archive format that contains the data of either multiple Mach-O files (one for each CPU architecture you wish to support) or multiple static archive libraries (again, one for each CPU architecture you wish to support). You can have a fat file containing data for only one CPU architecture, although it’s not very useful to do so.

A multiple-architecture, or “fat” file is one that contains compiled code and data for more than one CPU architecture. A fat file contains a set of single-CPU files, one for each CPU architecture, with a special header at the beginning of the file to allow the various runtime tools to quickly find a particular CPU architecture. Each single-CPU file is stored as a continuous set of bytes at an offset in the fat file. The single-CPU files may currently be either Mach-O files or static archive libraries. For example, a fat static archive library might contain the data of one static archive library containing PowerPC modules and also the data for one static archive library containing x86 modules.

A fat file always begins with a `fat_header` (page 86) data structure, followed by a set of `fat_arch` (page 87) data structures, and then the actual data for all of the CPU architectures. All data in these data structures is stored in big-endian byte order.

`fat_header`

Describes the layout of a “fat” file, which is a file that can contain code and data for more than one CPU architecture. This data structure is declared in the header `/usr/include/mach-o/fat.h`.

```
struct fat_header
{
    unsigned long magic; // FAT_MAGIC
    unsigned long nfat_arch; // count structs that follow
};
```

Field Descriptions

`magic`

An integer containing the value `0xCAFEBABE` in big-endian byte order format. On a big-endian host CPU, this can be validated using the constant `FAT_MAGIC`; on a little-endian host CPU, it can be validated using the constant `FAT_CIGAM`.

`nfat_arch`

An integer specifying the number of `fat_arch` (page 87) data structures that follow. This is the number of CPU architectures contained in this file.

Discussion

The `fat_header` data structure is placed at the start of a file that contains Mach-O files for multiple CPU architectures. Directly following the `fat_header` data structure is a set of `fat_arch` (page 87) data structures, one for each CPU architecture included in the file.

Regardless of the content it describes, all of the fields in this data structure are stored in big-endian byte order.

`fat_arch`

Describes the location within the file of data for a single CPU architecture. This data structure is declared in the header `/usr/include/mach-o/fat.h`.

```
struct fat_arch
{
    cpu_type_t cputype;
    cpu_subtype_t cpusubtype;
    unsigned long offset;
    unsigned long size;
    unsigned long align;
};
```

Field Descriptions

`cputype`

An enumeration value of type `cpu_type_t`. Specifies the CPU family.

`cpusubtype`

An enumeration value of type `cpu_subtype_t`. Specifies the specific member of the CPU family on which this entry may be used, or a constant specifying all members.

`offset`

Offset to the beginning of the data for this CPU.

`size`

Size of the data for this CPU.

`align`

The power of 2 alignment for the offset of the contents of this CPU architecture. This is required to ensure that, if this fat file is changed, the contents it retains are correctly aligned for virtual memory paging and other uses.

Discussion

An array of `fat_arch` data structures appears directly after the `fat_header` (page 86) data structure of a file that contains Mach-O files for multiple CPU architectures.

Regardless of the content it describes, all of the fields in this data structure are stored in big-endian byte order.

Mach-O Dynamic Linking Functions Reference

This chapter documents data structures that compose the Mach-O file format, as well as the low-level functions that you can use to manipulate Mach-O files at runtime.

Dynamic Linker Functions

The dynamic linker provides several different groups of functionality that allows your application to manipulate Mach-O files at runtime.

Object File Image Functions

NSAddImage

Adds the specified Mach-O image to the currently running process.

```
const struct mach_header* NSAddImage(
    char* image_name,
    unsigned long options);
```

Parameter Descriptions

image_name

A pointer to a C string. Pass the pathname to a shared library on disk. For best performance, specify the full pathname of the shared library—not a symlink.

options

A bit mask. Pass one or more of the following options or `NSADDIMAGE_OPTION_NONE` to specify no options.

`NSADDIMAGE_OPTION_RETURN_ON_ERROR`

If an error occurs and you have specified this option, `NSAddImage` returns `NULL`. You can then use the function `NSLinkEditError` to retrieve information about the error.

If an error occurs, and you have not specified this option, `NSAddImage` calls the `linkEdit` error handler you have installed using the `NSInstallLinkEditErrorHandlers` function. If you have not installed a link edit error handler, `NSAddImage` prints an error to `stderr` and causes a breakpoint trap to end the program.

`NSADDIMAGE_OPTION_WITH_SEARCHING`

With this option, the `image_name` passed for the library and all its dependents is affected by the various `dylld` environment variables as if this library were linked into the program.

`NSADDIMAGE_OPTION_RETURN_ONLY_IF_LOADED`

With this option, `NSAddImage` returns `NULL` if the shared library was not loaded prior to this call to `NSAddImage`.

function result A pointer to a `mach_header` (page 53) data structure. This is a pointer to the start of the loaded image.

Discussion

`NSAddImage` loads the shared library specified by `image_name` into the current process, returning a pointer to the `mach_header` (page 53) data structure of the loaded image. Any libraries that the specified library depends on are also loaded.

The `linkEdit` error handler is documented in the `NSModule(3)` man page.

Version Notes

This function was first introduced in Mac OS X 10.1.

NSCreateObjectFileImageFromFile

Creates an image reference for a given Mach-O file.

```
NSObjectFileImageReturnCode NSCreateObjectFileImageFromFile(
    const char* pathName,
    NSObjectFileImage* objectFileImage);
```

Parameter Descriptions*pathName*

A C string. Pass the pathname to a Mach-O executable file. You must have previously built this file with the `-bundle` linker option, or this function returns an error.

objectFileImage

On output, a pointer to an `NSObjectFileImage` opaque data structure.

function result A result code (see below).

Discussion

Given a pathname to a Mach-O executable, this function creates and returns an `NSObjectFileImage` reference. The current implementation works only with bundles, so you must build the Mach-O executable file using the `-bundle` linker option.

Result Codes

`NSObjectFileImageSuccess`

The operation was completed successfully.

`NSObjectFileImageFailure`

The operation was not successfully completed. When this result code is returned, an error message is printed to the standard error stream.

`NSObjectFileImageInappropriateFile`

The specified Mach-O file is not of a type this function can operate upon.

`NSObjectFileImageArch`

The specified Mach-O file is for a different CPU architecture.

`NSObjectFileImageFormat`

The specified file does not appear to be a Mach-O file.

`NSObjectFileImageAccess`

The access permissions for the specified file do not permit the creation of the image.

NSCreateObjectFileImageFromMemory

Creates an image reference for a Mach-O file currently in memory.

```
NSObjectFileImageReturnCode NSCreateObjectFileImageFromFile(
    void* address,
    unsigned long size,
    NSObjectFileImage* objectFileImage);
```

Parameter Descriptions

address

A pointer to the memory block containing the Mach-O file contents.

size

The size of the memory block, in bytes.

objectFileImage

On output, a pointer to an `NSObjectFileImage` opaque data structure.

function result A result code (see below).

Discussion

Given a pointer to a Mach-O file in memory, this function creates and returns an `NSObjectFileImage` reference. The current implementation works only with bundles, so you must build the Mach-O executable file using the `-bundle` linker option.

Availability

Available in Mac OS X version 10.3 and later.

Result Codes

`NSObjectFileImageSuccess`

The operation was completed successfully.

`NSObjectFileImageFailure`

The operation was not successfully completed. When this result code is returned, an error message is printed to the standard error stream.

`NSObjectFileImageInappropriateFile`

The Mach-O file in memory is not of a type this function can operate upon.

`NSObjectFileImageArch`

The specified Mach-O file is for a different CPU architecture.

`NSObjectFileImageFormat`

The memory block does not appear to point to a Mach-O file.

`NSObjectFileImageAccess`

The access permissions for the specified file do not permit the creation of the image.

NSDestroyObjectFileImage

Releases the given object file image.

```
enum DYLD_BOOL NSDestroyObjectFileImage(
    NSObjectFileImage objectFileImage);
```

Parameter Descriptions

objectFileImage

A reference to the object file image to destroy.

function result TRUE if the image was successfully destroyed, FALSE if not.

NSNameOfModule

Returns the name of the given module.

```
const char* NSNameOfModule(
    NSModule module);
```

Parameter Descriptions*module*

A module reference. Pass the module whose name you wish to retrieve.

function result A C string containing the name of the module. The string is owned by the dynamic linker and you should not free it.

Discussion

See [“Modules—The Smallest Unit of Code”](#) (page 14) for more information about modules.

NSLibraryNameForModule

Returns the name of the library that contains the given module.

```
const char* NSLibraryNameOfModule(
    NSModule module);
```

Parameter Descriptions*module*

A module reference. Pass the module whose library name you wish to retrieve.

function result A C string containing the name of the library that contains the module. The string is owned by the dynamic linker and you should not free it.

Discussion

See [“Modules—The Smallest Unit of Code”](#) (page 14) for more information about modules.

NSLinkModule

Links the given object file image as a module into the current program.

```
NSModule NSLinkModule(
    NSObjectFileImage objectFileImage,
    const char* moduleName,
    unsigned long options);
```

Parameter Descriptions*objectFileImage*

An object file image reference. Pass a reference created using the [NSCreateObjectFileImageFromFile](#) (page 90) function.

moduleName

A C string. Pass the absolute path to the object file image. GDB uses this path to retrieve debug symbol information from the library.

options

An unsigned long value. Pass one or more of the following bit masks or `NSLINKMODULE_OPTION_NONE` to specify no options.

`NSLINKMODULE_OPTION_BINDNOW`

The dynamic linker binds all of the undefined references immediately, rather than waiting until the references are actually used. All dependent libraries are also bound.

`NSLINKMODULE_OPTION_PRIVATE`

Do not add the global symbols from the module to the global symbol list. Instead, you must use the [NSLookupSymbolInModule](#) (page 99) function to obtain symbols from this module.

`NSLINKMODULE_OPTION_RETURN_ON_ERROR`

If an error occurs while binding the module, return `NULL`. You can then use the function `NSLinkEditError` to retrieve information about the error.

Without this option, `NSLinkModule` calls the `linkEdit` error handler you have installed using the `NSInstallLinkEditErrorHandlers` function. If you have not installed a link edit error handler, `NSLinkModule` prints a message to the standard error stream and causes a breakpoint trap to end the program.

function result A reference to the linked module.

Discussion

When you call `NSLinkModule`, all libraries referenced by the given module are added to the library search list. Unless you pass the `NSLINKMODULE_OPTION_PRIVATE`, `NSLinkModule` adds all global symbols in the module to the global symbol list.

See “[Modules—The Smallest Unit of Code](#)” (page 14) for more information about modules.

NSUnLinkModule

Unlinks the given module from the current program.

```
enum DYLD_BOOL NSUnLinkModule(
    NSModule module,
    unsigned long options);
```

Parameter Descriptions*module*

A module reference. Pass a reference to a module that you have previously linked using the [NSLinkModule](#) (page 93) function.

options

An unsigned long value. You can specify one or more of the following bit masks.

`NSUNLINKMODULE_OPTION_NONE`

Unlink the module and deallocate the memory it occupies.

`NSUNLINKMODULE_OPTION_KEEP_MEMORY_MAPPED`

Unlink the module, but do not deallocate the memory it occupies. Addresses that reside within the module will still be valid. You cannot unmap this memory later; it will be released when the process exits or is terminated.

`NSUNLINKMODULE_OPTION_RESET_LAZY_REFERENCES`

Unlink the module and reset lazy references from other modules that are bound to the module. You can then link a new module that implements the same symbols, and the function call references will then bind to the new module when accessed.

Discussion

See “[Modules—The Smallest Unit of Code](#)” (page 14) for more information about modules.

Special Considerations

As of this writing (Mac OS X 10.2), `NSUNLINKMODULE_OPTION_RESET_LAZY_REFERENCES` can be used only with PowerPC CPU executables.

`NSIsSymbolNameDefined`

Returns `true` if the given symbol is defined in the current program.

```
enum DYLD_BOOL NSIsSymbolNameDefined(
    const char* symbolName);
```

Parameter Descriptions

symbolName

A C string. Pass the name of the symbol whose definition status you wish to discover.

function result A boolean value indicating `true` if the symbol is defined by any image loaded in the current process, `false` if the symbol cannot be found.

Discussion

If you know the name of the library in which the symbol is likely to be located, you can use the [NSIsSymbolNameDefinedWithHint](#) (page 95) function, which may be faster than `NSIsSymbolNameDefined`. You should use the [NSIsSymbolNameDefinedInImage](#) (page 96) function to perform a two-level namespace lookup.

`NSIsSymbolNameDefinedWithHint`

Returns `true` if the given symbol is defined in the current program, with a hint specifying the name of the shared library likely to contain the symbol.

```
enum DYLD_BOOL NSIsSymbolNameDefinedWithHint(
    const char* symbolName,
```

```
const char* libraryNameHint);
```

Parameter Descriptions

symbolName

A C string. Pass the name of the symbol whose definition status you wish to discover.

libraryNameHint

A C string. Pass any part of the name of the shared library that is likely to contain the symbol. It searches only the first shared library that matches.

function result A boolean value indicating `true` if the symbol is defined by any image loaded in the current process, `false` if the symbol cannot be found.

Discussion

The library name you pass to `NSIsSymbolNameDefinedWithHint` allows it to determine a position in the list of loaded symbols from which to start the search. This can result in a considerably faster lookup search time than is possible using `NSIsSymbolNameDefined` (page 95).

Note that `NSIsSymbolNameDefinedWithHint` performs a flat lookup even if the symbol namespace of the current program has two levels. You should use the `NSIsSymbolNameDefinedInImage` (page 96) function to perform a two-level namespace lookup.

NSIsSymbolNameDefinedInImage

Returns `true` if the given image contains the named symbol.

```
enum DYLD_BOOL NSIsSymbolNameDefinedInImage(
    const struct mach_header* image,
    const char* symbolName);
```

Parameter Descriptions

image

A pointer to a `mach_header` (page 53) data structure.

symbolName

A C string. Pass the name of the symbol.

function result An enumeration value of type `DYLD_BOOL`. `true` if the image contains a symbol with the given name, `false` otherwise.

Version Notes

This function was first introduced in Mac OS X 10.1.

NSLookupAndBindSymbol

Given a symbol name, returns the corresponding symbol from the global symbol table.

```
NSSymbol NSLookupAndBindSymbol(
    const char* symbolName);
```


Parameter Descriptions*symbolName*

A pointer to a C string. Pass the name of the symbol you wish to find.

function result The symbol reference for the requested symbol.

Discussion

On error, if you have installed a link edit error handler, it will be called; otherwise, this function writes an error message to file descriptor 2 (usually the standard error stream, `stderr`) and causes a breakpoint trap to end the program.

If you know the name of the library in which the symbol is likely to be located, you can use the [NSLookupAndBindSymbolWithHint](#) (page 97) function, which may be faster than `NSLookupAndBindSymbol`. You should use the [NSLookupSymbolInImage](#) (page 97) function to perform a two-level namespace lookup.

NSLookupAndBindSymbolWithHint

Given a symbol name, returns the corresponding symbol from the global symbol table.

```
NSSymbol NSLookupAndBindSymbolWithHint(
    const char* symbolName,
    const char* libraryNameHint);
```

Parameter Descriptions*symbolName*

A pointer to a C string. Pass the name of the symbol you wish to find.

libraryNameHint

A pointer to a C string. Pass any part of the name of the library that the symbol is likely to be found in.

function result The symbol reference for the requested symbol.

Discussion

On error, if you have installed a link edit error handler, it will be called; otherwise, this function writes an error message to file descriptor 2 (usually the standard error stream, `stderr`), and causes a breakpoint trap to end the program.

Note that `NSLookupAndBindSymbolWithHint` performs a flat lookup even if the symbol namespace of the current program has two levels. You should use the [NSLookupSymbolInImage](#) (page 97) function to perform a two-level namespace lookup.

NSLookupSymbolInImage

Returns a reference to the specified symbol from the specified image.

```
NSSymbol NSLookupSymbolInImage(
    const struct mach_header* image,
    const char* symbolName)
```

```
unsigned long options);
```

Parameter Descriptions

image

A pointer to a `mach_header` data structure. Pass a pointer to the start of the image that contains the symbol. You can get this pointer from a shared library name using the `NSAddImage` (page 89) function.

If the process does not have a two-level namespace, `NSLookupSymbolInImage` ignores this argument and searches for the symbol in the global symbol table.

symbolName

A pointer to a C string. Pass the name of the symbol you wish to find.

options

A bit mask. Pass any of the following options.

`NSLOOKUPSYMBOLINIMAGE_OPTION_BIND`

Bind the nonlazy symbols of the module in the image that defines `symbolName` and let all lazy symbols in the module be bound on first call. You should pass this option when you expect the module to bind without errors (for example, a library supplied with the system). If, later, you call a lazy symbol, and the lazy symbol fails to bind, the runtime calls the link edit error handler you have installed using the `NSInstallLinkEditErrorHandlers` function.

If there is no link edit error handler installed, the runtime prints a message to the standard error stream and causes a breakpoint trap to end the program.

`NSLOOKUPSYMBOLINIMAGE_OPTION_BIND_NOW`

Bind all the non-lazy and lazy symbols of the module in the image that defines the symbol name, and bind symbols in the dependent libraries as needed.

Pass this option for a library that might not be expected to bind without errors but that links against only system-supplied libraries that are themselves expected to bind without any errors.

`NSLOOKUPSYMBOLINIMAGE_OPTION_BIND_FULLY`

Bind all the symbols of the module that defines `symbolName` and all the dependent symbols of all needed libraries.

Because it may take a long time to fully bind the image, you should pass this option only for libraries that cannot bind other symbols once executed, such as code that implements signal handlers.

`NSLOOKUPSYMBOLINIMAGE_OPTION_RETURN_ON_ERROR`

Return NULL if the symbol cannot be bound.

function result The symbol reference for the requested symbol, or NULL if the symbol cannot be found and you passed the option `NSLOOKUPSYMBOLINIMAGE_OPTION_RETURN_ON_ERROR`.

Discussion

On error, if you have installed a link edit error handler, it will be called; otherwise, this function writes an error message to file descriptor 2 (usually the standard error stream, `stderr`), and causes a breakpoint trap to end the program.

Version Notes

This function was first introduced in Mac OS X 10.1.

NSLookupSymbolInModule

Given a module reference, returns a reference to the symbol with the given name.

```
NSSymbol NSLookupSymbolInModule(
    NSModule module,
    const char* symbolName);
```

Parameter Descriptions

module

A module reference. Pass the module that contains the symbol.

symbolName

A pointer to a C string. Pass the name of the symbol to look up.

function result The symbol reference or NULL if the symbol cannot be found.

NSNameOfSymbol

Returns the name of the given symbol.

```
const char* NSNameOfSymbol(
    NSSymbol symbol);
```

Parameter Descriptions

symbol

A symbol reference. Pass the symbol whose name you wish to obtain.

function result A pointer to a C string containing the name of the reference. The dynamic linker owns this string and you should not free it.

NSAddressOfSymbol

Returns the address in the program's address space of the data represented by the given symbol. The data may be a variable, a constant, or the first instruction of a function.

```
void* NSAddressOfSymbol(
    NSSymbol symbol);
```

Parameter Descriptions*symbol*

A symbol reference. Pass the symbol whose address you wish to obtain.

function result A pointer to the data represented by the given symbol.

NSModuleForSymbol

Returns a reference to the module containing the given symbol.

```
NSModule NSModuleForSymbol(
    NSSymbol symbol);
```

Parameter Descriptions*symbol*

A symbol reference. Pass the symbol whose module you wish to obtain.

function result A reference to the module that contains the given symbol.

Section and Segment Accessors

getsectbyname

Returns a data structure representing a section of the Mach-O file that contains the main executable program of the current process.

```
const struct section* getsectbyname(
    const char* segname,
    const char* sectname);
```

Parameter Descriptions*segname*

A pointer to a C string. Pass the name of the segment in which the section resides.

sectname

A pointer to a C string. Pass the name of the section.

function result A pointer to a [section](#) (page 60) data structure.

getsegbyname

Returns a data structure representing a segment of the Mach-O file containing the main executable program of the current process.

```
const struct segment_command* getsegbyname(
```

```
const char* segname);
```

Parameter Descriptions

segname

A pointer to a C string. Pass the name of the segment.

function result A pointer to a [segment_command](#) (page 58) data structure.

getsectdata

Returns the data for a section from the Mach-O file of the main executable program of the current process.

```
extern char* getsectdata(
    const char* segname,
    const char* sectname,
    unsigned long* size);
```

Parameter Descriptions

segname

A pointer to a C string. Pass the name of the segment in which the section resides.

sectname

A pointer to a C string. Pass the name of the section.

size

A pointer to a long integer. On output, contains the length (in bytes) of the section.

function result A pointer to the data of the section.

getsectbynamefromheader

Returns the data structure representing a section of a specified Mach-O file.

```
const struct section* getsectbynamefromheader(
    const struct mach_header* mhp,
    const char* segname,
    const char* sectname);
```

Parameter Descriptions

mhp

A pointer to a [mach_header](#) (page 53) data structure. Pass the `mach_header` of the file containing the section data you wish to retrieve.

segname

A pointer to a C string. Pass the name of the segment in which the section resides.

sectname

A pointer to a C string. Pass the name of the section.

function result A pointer to the [section](#) (page 60) data structure.

getsectdatafromheader

Returns the data for a section of a specified Mach-O file.

```
extern char* getsectdatafromheader(
    const struct mach_header* mhp,
    const char* segname,
    const char* sectname,
    unsigned long* size);
```

Parameter Descriptions

mhp

A pointer to a [mach_header](#) (page 53) data structure. Pass the `mach_header` of the file containing the section data you wish to retrieve.

segname

A pointer to a C string. Pass the name of the segment in which the section resides.

sectname

A pointer to a C string. Pass the name of the section.

size

A pointer to a long integer. On output, contains the length (in bytes) of the section.

function result A pointer to the data of the section. If the Mach-O file is a dynamic shared library (MH_DYLIB), you will need to add the virtual memory slide amount to this address to get the true address of the data. See [_dyld_get_image_vmaddr_slide](#) (page 104) for more information.

getsectdatafromFramework

Returns the data for a section of the Mach-O file containing a specified framework.

```
extern char* getsectdatafromFramework(
    const char* FrameworkName,
    const char* segname,
    const char* sectname,
    unsigned long* size);
```

Parameter Descriptions

FrameworkName

A pointer to a C string. Pass the name of the framework in which the section resides.

segname

A pointer to a C string. Pass the name of the segment in which the section resides.

sectname

A pointer to a C string. Pass the name of the section.

size

A pointer to a long integer. On output, contains the length (in bytes) of the section.

function result A pointer to the data of the section. If the Mach-O file is a dynamic shared library (MH_DYLIB), you will need to add the virtual memory slide amount to this address to get the true address of the data. See [_dyld_get_image_vmaddr_slide](#) (page 104) for more information.

Low-Level Dynamic Linking Functions

_dyld_present

Indicates whether or not the dynamic linker is loaded into the current program

```
unsigned long _dyld_present (void);
```

function result A long integer indicating the presence of dyld. This value is zero if dyld is not loaded in the current process, and greater than zero if dyld is loaded in the current process.

_dyld_image_count

Returns the number of images that dyld has mapped into the address space of the current process.

```
unsigned long _dyld_image_count(void);
```

function result A long integer containing the number of images that dyld has mapped into the address space of the current process.

Discussion

This function provides you with a count of the number of the images in the image list. You can use this number to iterate the images loaded into the address space of the current process, using functions such as [_dyld_get_image_header](#) (page 103) and [_dyld_get_image_name](#) (page 104).

_dyld_get_image_header

Returns the data structure for the header of a specified image. The image is specified by index into the list of images maintained by dyld for the current process.

```
struct mach_header* _dyld_get_image_header(
```

```
unsigned long image_index);
```

Parameter Descriptions

image_index

A long integer. Pass a zero-based index indicating the position of the image in the list of images loaded into the address space of the current process.

function result A pointer to the [mach_header](#) (page 53) data structure of the specified image. If *image_index* is greater than the number of loaded images, this pointer is NULL.

_dyld_get_image_vmaddr_slide

Returns the virtual memory address slide amount of an image.

```
unsigned long _dyld_get_image_vmaddr_slide(
    unsigned long image_index);
```

Parameter Descriptions

image_index

A long integer. Pass a zero-based index indicating the position of the image in the list of images loaded into the address space of the current process.

function result If *image_index* is greater than or equal to the value returned by [_dyld_image_count](#) (page 103), zero. Otherwise, the *vmaddr_slide* value for the specified image.

Discussion

When the dynamic linker loads an image, the image must be mapped into the virtual address space of the process at an unoccupied address. The dynamic linker accomplishes this by adding a value—the virtual memory slide amount—to the base address of the image.

_dyld_get_image_name

Retrieves the name of an image.

```
char* _dyld_get_image_name(
    unsigned long image_index);
```

Parameter Descriptions

image_index

A long integer. Pass a zero-based index indicating the position of the image in the list of images loaded into the address space of the current process.

function result A pointer to a C string. If *image_index* is greater than the number of loaded images, the string pointer is NULL.

Discussion

Returns the name of the image located at the given index into the global image list.

_dyld_lookup_and_bind

Finds the given symbol name and binds it into the program.

```
void _dyld_lookup_and_bind(  
    char* symbol_name,  
    unsigned long* address  
    void ** module);
```

Parameter Descriptions

symbol_name

A pointer to a C string. Specify the name of the symbol to bind.

address

A pointer to a long integer. On output, the long integer contains the address of the symbol specified by *symbol_name*. This parameter is optional; pass NULL for this pointer on input if you do not want to retrieve this data.

module

A pointer to a module pointer. On output, the module pointer contains the module of the symbol specified by *symbol_name*. This parameter is optional; specify NULL for this pointer on input if you do not want to retrieve this data.

Discussion

You can use `_dyld_lookup_and_bind` to find a given symbol name in the global search list and bind it (and all other defined symbols in the same module) into the program.

If the program is prebound and you know the name of the library that contains the symbol, consider using `_dyld_lookup_and_bind_with_hint` (page 105) instead.

_dyld_lookup_and_bind_with_hint

Finds the given symbol name and binds it into the program, with a hint to allow `dyld` to speed up the symbol search for a prebound program.

```
void _dyld_lookup_and_bind_with_hint(  
    char* symbol_name,  
    const char* library_name_hint,  
    unsigned long* address,  
    void** module);
```

Parameter Descriptions

symbol_name

A pointer to a C string. Specify the name of the symbol to bind.

library_name_hint

A pointer to a C string. Specify the name of the library in which the symbol is probably located. The dynamic linker compares this name with the actual library install names using the standard C library function `strstr`.

address

A pointer to a long integer. On output, the long integer contains the address of the symbol specified by `symbol_name`. This parameter is optional; pass `NULL` for this pointer on input if you do not want to retrieve this data.

module

A pointer to a module pointer. On output, the module pointer contains the module of the symbol specified by `symbol_name`. This parameter is optional; specify `NULL` for this pointer on input if you do not want to retrieve this data.

Discussion

You can use `_dyld_lookup_and_bind_with_hint` to quickly find a given symbol name in the global search list of a prebound program and bind the symbol (and all other defined symbols in the same module) into the program.

`_dyld_lookup_and_bind_fully`

Finds the module containing the specified symbol and fully bind all of the symbol references within it.

```
void _dyld_lookup_and_bind_fully(
    char* symbol_name,
    unsigned long* address,
    void** module);
```

Parameter Descriptions

symbol_name

A pointer to a C string. Specify the name of the symbol to bind.

address

A pointer to a long integer. On output, this is the address of the specified symbol.

module

A pointer to a pointer. On output, the pointer is set to the address of the module in which the specified symbol resides.

Discussion

You can use this function to bind modules containing signal handlers or other error handling code which cannot be initialized lazily.

Errors in binding are reported through the normal mechanisms.

`_dyld_bind_fully_image_containing_address`

Finds the image containing the specified address and fully binds all of the modules within it.

```
DYLD_BOOL _dyld_bind_fully_image_containing_address(
    unsigned long* address);
```

Parameter Descriptions*address*

A pointer to an address located somewhere within a loaded image.

function result A boolean value. If `true`, the address resides somewhere within a loaded image, and so `_dyld_bind_fully_image_containing_address` attempted to bind that image. If `false`, the address does not reside within a loaded image, and so `_dyld_bind_fully_image_containing_address` did nothing.

Discussion

You can use this function to bind error handling code like signal handlers when you have the address of a function, but not the symbol name. This may bind more symbols than are actually needed.

If the image containing the address is a flat namespace image, multiply-defined errors can occur even if the symbols are not really used. Errors in binding are reported through the normal error reporting mechanisms.

`_dyld_image_containing_address`

Returns whether or not a specified address is within any loaded image.

```
DYLD_BOOL _dyld_image_containing_address(
    unsigned long address);
```

Parameter Descriptions*address*

An unsigned long integer. Pass the address that you wish to obtain status about.

function result `TRUE` if the address is located within an image loaded by the dynamic linker and `FALSE` otherwise.

`_dyld_launched_prebound`

Returns whether or not the dynamic linker was able to launch the program with the prebinding optimization enabled.

```
DYLD_BOOL _dyld_launched_prebound(void);
```

function result A boolean value. `true` if the program was launched successfully using the prebound state and `false` if either the program was not prebound or the prebinding couldn't be used for some reason.

Discussion

If the program was not successfully launched with the prebinding optimization, either the linker did not prebind the program, or the addresses of some images overlapped and so the linker could not use the prebound addresses, or some other problem occurred. In any case, the program continues to launch, but it will be slower than with prebinding enabled.

_dyld_func_lookup

Obtains the address of the implementation of a dyld library function.

```
int _dyld_func_lookup(
    char* dyld_func_name,
    unsigned long* address);
```

Parameter Descriptions

dyld_func_name

A pointer to a C string. Pass the name of a dyld library function.

address

A pointer to a long integer. On output, the long integer value is set to the address of the function if the function is found, otherwise the value is undefined.

function result An integer value. Nonzero if the function was found. Zero if the function was not found.

Discussion

This function is used by the library code that implements the dyld functions.

_dyld_bind_objc_module

Binds the module that contains a given Objective-C address.

```
void _dyld_bind_objc_module(
    void* address);
```

Parameter Descriptions

address

A pointer. Pass any address residing within the `__OBJC,__module` section of a loaded Mach-O file.

Discussion

This function is used by the Objective-C runtime library.

_dyld_get_objc_module_sect_for_module

Obtains the size and starting location of an Objective-C module.

```
extern void _dyld_get_objc_module_sect_for_module(
    NSModule module,
    void** moduleStart,
    unsigned long* moduleSize);
```

Parameter Descriptions

module

A module reference from an image.

moduleStart

A pointer to a pointer. On output, contains a pointer to the start of the `__OBJC,__module` section for the specified module.

moduleSize

A pointer to a long integer. On output, the long integer contains the a value indicating the size of the output module.

Discussion

This function is used by the Objective-C runtime library.

`_dyld_lookup_and_bind_objc`

Obtains and binds the an Objective-C module that contains the specified symbol.

```
void _dyld_lookup_and_bind_objc(
    const char* symbol_name,
    unsigned long* address,
    void** module);
```

Parameter Descriptions

symbol_name

A pointer to a C string. Specify the name of the symbol to bind, such as `.objc_class_name_Foo`.

address

A pointer to a long integer. On output, the long integer contains the address of the symbol specified by `symbol_name`. This parameter is optional; pass `NULL` for this pointer on input if you do not want to retrieve this data.

module

A pointer to a module pointer. On output, the module pointer contains the module of the symbol specified by `symbol_name`. This parameter is optional; specify `NULL` for this pointer on input if you do not want to retrieve this data.

Discussion

This routine is used by the Objective-C runtime library. It performs the same function as [_dyld_lookup_and_bind](#) (page 105) but, for performance reasons, does not update the symbol pointers if the symbol is in a bound module. An Objective-C symbol such as `.objc_class_name_Object` is never used by a symbol pointer, and updating the symbol pointers is a relatively expensive operation, so this provides a way for the Objective-C runtime to avoid that overhead.

`_dyld_moninit`

This function is used by the profiling routine `moninit` to allow images other than the main executable to be profiled.

```
void _dyld_moninit( void (*monaddition)(
    char* lowpc,
```

```
char* highpc);
```

Parameter Descriptions

monaddition

A pointer to a callback function. The callback is called when an image is first mapped in.

Discussion

This function is usually called by the profiling runtime (specifically, from the `moninit` function). It is documented here for completeness. See the man page for `moninit` and `monaddtion` for further information.

_dyld_register_func_for_add_image

Registers a function to be called by the dynamic linker runtime when an image is added to the program.

```
void _dyld_register_func_for_add_image(
    void (*func)(struct mach_header* mh, unsigned long vmaddr_slide));
```

Parameter Descriptions

func

A pointer to a callback function that accepts a pointer to a `mach_header` (page 53) data structure and a virtual memory slide amount. The virtual memory slide amount specifies the difference between the address at which the image was linked and the address at which the image is loaded.

Discussion

When you call `_dyld_register_func_for_add_image`, the dynamic linker runtime calls the specified callback (`func`) once for each of the images that is currently loaded into the program. When a new image is added to the program, your callback is called again with the `mach_header` (page 53) for the new image, and the virtual memory slide amount of the new image.

You might use this, for example, in implementing a runtime system, such as the Objective-C runtime, to discover when new images are added to the program.

_dyld_register_func_for_remove_image

Registers a function to be called by the dynamic linker runtime when an image is removed from the program.

```
void _dyld_register_func_for_remove_image(
    void (*func) (struct mach_header* mh),
    unsigned long vmaddr_slide );
```

Parameter Descriptions*func*

A pointer to a callback function that accepts a pointer to a [mach_header](#) (page 53) data structure and a virtual memory slide amount. The virtual memory slide amount specifies the difference between the address at which the image was linked and the address at which the image is loaded.

_dyld_register_func_for_link_module

Registers a function to be called by the dynamic linker runtime when a module is linked into the program.

```
void _dyld_register_func_for_link_module(
    void (*func)(NSModule module));
```

Parameter Descriptions*func*

A pointer to a callback that accepts a module reference.

Discussion

When you call `_dyld_register_func_for_link_module`, the dynamic linker runtime calls the specified callback (*func*) once for each module that is currently linked into the program. When a new module is linked into the program, the *func* callback is called again for that module.

Glue Functions for Indirect Addressing

dyld_stub_binding_helper

Assembly-language glue code that performs binding for a lazy function symbol.

```
.private_extern dyld_stub_binding_helper
```

Parameter Descriptions*PowerPC: r11 x86: stack-based parameter*

A pointer to the lazy symbol pointer for the function to be bound.

Discussion

The `dyld stub binding helper` is a glue function that assists the dynamic linker in lazily binding an external function. When the compiler sees a call to an external function, it generates a symbol stub and a lazy pointer for the function. At the call site, the compiler generates a call to the symbol stub. The symbol stub is a sequence of code that loads the lazy pointer and jumps to it. Initially, the sequence of code and the contents of the lazy pointer call `dyld_stub_binding_helper`, which calls the dynamic linker to bind the symbol. After the symbol is bound, the lazy pointer is set to the address of the symbol, and the symbol is reached directly by jumping to the lazy pointer.

Thereafter, because the address has been changed to the actual address of the function, all calls to the external function call the external function.

On entry, `dyld_stub_binding_helper` accepts the address of the lazy symbol pointer. On exit, the value of the lazy symbol pointer is set to the address of the external function. The `dyld` stub binding helper is assembly-language based and does not use standard calling conventions, and as such, the location of the parameters are specific to each CPU architecture. On PowerPC, the address of the lazy symbol pointer is expected to be in GPR11. On x86, the address of the lazy symbol pointer should be the pushed on the stack.

`dyld_stub_binding_helper` is located in the runtime startup files that are statically linked into the image. For executables, the file is `/lib/crt1.o`. For bundles, it is `/lib/bundle1.o`, and for shared libraries, it is `/lib/dylib1.o`.

Dynamic Linker Data Types

Boolean Return Value

The low-level dynamic functions return status information using the `DYLD_BOOL` data type.

DYLD_BOOL

An enumeration specifying a Boolean variable, which contains a value that can be either true or false.

```
enum DYLD_BOOL {FALSE, TRUE};
```

Constant Descriptions

TRUE

Boolean true. Equivalent to the integer value 1.

FALSE

Boolean false. Equivalent to the integer value 0.

Document Revision History

The table below describes the revisions to *Mach-O Runtime Architecture*.

Date	Notes
August 7, 2003	Added description of new API for Mac OS X version 10.3.
Jan 1, 2003	Incorporated developer feedback. Updated code-generation examples.
	Fixed bugs 2462895, 2749339, 2909989, 2910422, 2921574.
July 1, 2002	More developer feedback. Document weak definitions and weak references (new for 10.2). Substantially update the glossary. Other tweaks and additional material. Clarify common vs. coalesced symbol definitions.
	ABI: Rewrote position-independent and indirect code section, incorporating correct examples and separating PIC and indirect code generation. Add <code>C99_Bool</code> data type.
	Fixed bugs 2909989, 2910422, and 2921574.
May 1, 2002	This was a preliminary draft distributed with the WWDC 2002 developer tools.
	Incorporated many corrections from developer review. More to come.
	By popular demand, added some common usage scenarios to map runtime features to the options in the standard Mac OS X tools that implement those features. To satisfy a related popular demand, this information is collected in a separate chapter, which allows users of third-party tool sets to ignore it. This chapter is currently unfinished, and the overview chapter is yet to be modified to cross-reference it.
	Updated umbrella framework description to better match reality.
	Added <code>long double</code> and <code>long long</code> return value information. Removed last vestiges of CFM. Rewrote data alignment section, incorporating the correct rules (inherited from IBM's <code>xlc</code> compiler) for power alignment mode, and adding new natural alignment mode.

REVISION HISTORY

Document Revision History

Date	Notes
April 1, 2002	This was a preliminary draft distributed with the April 2002 Developer Tools CD.

Glossary

bundle (1) A Mach-O file that must be explicitly loaded by the client application before use. Bundles typically contain code and data that extend the capabilities of an application. Also called a *plug-in*, *application extension* or *drop-in addition*. Compare [dynamic shared library](#).

bundle (2) A file package containing loadable code and resources, in the format understood by the NSBundle and CFBundle classes.

client application In the context of a shared library or a bundle, the application that imports the shared library or that loads the bundle. Usually refers specifically to the main program file of the application.

coalesced symbol A symbol that may be defined in multiple intermediate object files. The symbols have the same name and occupy the same amount of space. The static linker ignores all but one copy of the symbol. Compare [common symbol](#).

Code Fragment Manager (CFM) The part of Mac OS 9 that loads code from [Preferred Executable Format \(PEF\)](#) files into memory and prepares it for execution. Supported in Carbon for compatibility with Mac OS 9.

common symbol A symbol that may be defined in multiple intermediate object files. The symbols have the same name but may occupy different amounts of space. The static linker uses only one definition—the largest—in the output file. Compare [coalesced symbol](#). See also [tentative definition](#).

debugging symbol A symbol generated by the compiler to enable the debugger to map machine code to source code.

defined external symbol A data item or executable routine within a fragment that is made available for use by other Mach-O files. Also called *imported symbol*. Compare [undefined external symbol](#).

dependent library Another name for [import library](#).

dynamic shared library An image that exports functions and global variables to other images. A shared library is not included with the application code at link time but is linked in dynamically at runtime. Also known as a *shared library*, *dynamic library*, *dylib*, *dynamic link library*, and *DLL*. Compare [framework](#).

ELF An executable file format commonly used in UNIX operating systems.

embedding alignment The alignment of a data item within a composite data item (such as a data structure). Compare [natural alignment](#).

entry point A location (offset) within a [module](#).

epilog A sequence of code that cleans up the stack after a procedure call (restoring registers, restoring the stack pointer, and so on).

executable file As a generic term, refers to any file containing binary machine code, including bundles, shared libraries and programs. Often used to specifically refer to the main program file of an application. See also [program](#).

exported symbol See [defined external symbol](#).

external reference A reference to a routine or variable defined in a separate compilation unit or assembly.

fat application An application that contains code for two or more CPU architectures. For example, a fat application may contain both x86 and PowerPC code. Compare [fat file](#).

fat file A file that contains an archive of one or more Mach-O files, each containing code and data tailored to a specific CPU architecture.

final product Any file output from the static linker that the static linker cannot perform further binding on; dynamic shared libraries, main program files, and bundles are all final products, while intermediate object files and static archive libraries are not.

fragment In the Code Fragment Manager runtime architecture, an executable unit of code and its associated data. No defined meaning for Mach-O.

frame pointer (FP) A pointer to the beginning of a stack frame.

framework A shared library packaged with associated resources, such as header files, localized strings, and documentation files.

image Any Mach-O file built for use with the dynamic linker. In Mac OS X, this currently includes shared libraries, bundles, and executables (that is, Mach-O files of types MH_DYLIB, MH_BUNDLE, and MH_EXECUTABLE).

imported symbol See [undefined external symbol](#)

import library A dynamic shared library referenced by a Mach-O file. Also called [dependent library](#).

initialization function A function contained in a shared library or bundle that is executed immediately after the file is loaded. Compare [termination function](#).

install name the pathname to a dynamic shared library, as recorded at build time. Mach-O files refer to shared libraries using the install name.

leaf procedure A routine that calls no other routines.

linkage area The area in the PowerPC stack that holds the calling routine's RTOC value and the saved values of the Condition Register and the Link Register. Compare [parameter area](#).

main symbol For applications, the main routine or main entry point. Shared libraries and bundles do not require a main symbol.

main program file The Mach-O file (of type MH_EXECUTABLE) that contains the main entry point of a program. Often called an *executable*, but *executable* can also be used to describe any file containing machine code that can be executed.

module The smallest indivisible unit of machine code and data that can be linked independently in a Mach-O file.

natural alignment The alignment of a data type when allocated in memory or assigned a memory address. Compare [embedding alignment](#).

parameter area The area in the PowerPC stack that holds the parameters for any routines called by a given routine. Compare [linkage area](#).

PEF See [Preferred Executable Format \(PEF\)](#).

plug-in See [bundle \(1\)](#).

PowerPC microprocessor Any member of the family of PowerPC microprocessors.

CFM runtime architecture The runtime architecture for PowerPC computers running classic Mac OS or using the [Preferred Executable Format \(PEF\)](#) with Carbon on Mac OS X.

Preferred Executable Format (PEF) The format of executable files used for Carbon applications and shared libraries that use the Code Fragment Manager.

private framework A framework that is a part of one or more applications, but not part of the system. Private frameworks are often installed in a Frameworks directory inside an application package. Private frameworks are often used to provide functionality that is shared between the main program and any bundles loaded by the main program.

process A running program. See also [program](#).

program An executable unit that contains executable code.

prolog A sequence of code that prepares the stack for a procedure call (by saving registers, adjusting the stack, and so on).

red zone In the PowerPC CPU architecture, the area of memory immediately above the address pointed to by the stack pointer. The red zone is reserved for temporary use by a routine's prolog and as an area to store a leaf procedure's nonvolatile registers.

reference The location within one module that contains the address of another module or entry point.

relocation The process of replacing references to symbols with actual addresses during fragment preparation.

runtime architecture A set of basic rules that define how software operates. It dictates how code and data are addressed, the form of generated code, how applications are handled, and how to enable system calls. The runtime architecture defines the core of the runtime environment. Compare [runtime environment](#).

runtime environment The execution environment provided by the Mac OS X kernel and dynamic linker. The runtime environment dictates how executable code is loaded into memory, where data is stored, and how routines call other functions. Compare [runtime architecture](#).

section A named storage unit in a segment of a Mach-O file that contains either code or data.

segment A named collection of sections in a Mach-O file.

shared library See [dynamic shared library](#).

stack An area of memory in the application partition that is used for temporary storage of data during the operation of that application or other software.

stack frame The area of the stack used by a routine for its parameters, return address, local variables, and temporary storage.

stack pointer (SP) A pointer to the top of the stack.

static archive library An archive of modules whose code is included in the application at link time. Also called a *static library*.

symbol A reference to a function or data item.

termination function A function contained in a shared library or bundle that is executed just before the file is unloaded. Compare [initialization function](#).

tentative definition A symbol that has no initializer and is not marked with the ANSI C `extern` keyword. The standard compiler transforms tentative definitions into [common symbols](#). Compare [coalesced symbol](#).

transition vector In the CFM runtime architecture, an 8-byte data structure that describes the entry point and base register address of a routine.

undefined external symbol A data item or executable routine referenced by a fragment but not contained in it. An import is identified by name to the linker, but its actual address is bound at load time by the dynamic linker. Also called *imported symbol*. Compare [defined external symbol](#).

weak library A shared library that does not need to be present at runtime for the client application to run. Sometimes called a *soft library*.

weak definition A [defined external symbol](#) that the static linker may ignore in the presence of a non-weak defined external symbol. Used to support some features of C++. Compare [weak reference](#).

weak reference An [undefined external symbol](#) that does not need to be present at runtime for the client application to run. If the dynamic linker cannot find a definition of a weak reference, it sets the address of the symbol to zero. The client application can then test the weak symbol against NULL to see whether or not the symbol was found. Also known as a *weak import* or *soft import*. Compare [weak definition](#).

x86 microprocessor Any microprocessor capable of directly executing machine code for the IA-32 instruction set.

Index

Symbols

`_dyld_bind_fully_image_containing_address` function 106
`_dyld_bind_objc_module` function 108
`_dyld_func_lookup` function 108
`_dyld_get_image_header` function 103
`_dyld_get_image_name` function 104
`_dyld_get_image_vmaddr_slide` function 104
`_dyld_get_objc_module_sect_for_module` function 108
`_dyld_image_containing_address` function 107
`_dyld_image_count` function 103
`_dyld_launched_prebound` function 107
`_dyld_lookup_and_bind` function 105
`_dyld_lookup_and_bind_fully` function 106
`_dyld_lookup_and_bind_objc` function 109
`_dyld_lookup_and_bind_with_hint` function 105
`_dyld_moninit` function 109
`_dyld_present` function 103
`_dyld_register_func_for_add_image` function 110
`_dyld_register_func_for_link_module` function 111
`_dyld_register_func_for_remove_image` function 110

A

application binary interface (ABI) 8
application package 14

B

binding 17
bundles 14, 25

C

CFPlugin object 25
Classic runtime environment 7
coalesced symbol 20
Cocoa framework 16
Code Fragment Manager 8
Code Fragment Manager, using in Mac OS X 7
Code Fragment Manager, using with Carbon 25
COM objects 26
copy-on-write (COW) 51

D

dependent libraries 16
DLL. *See* dynamic shared libraries
dyld tool 15
DYLD_BOOL data type 112
`dyld_stub_binding_helper` function 111
dylib structure 65
`dylib_command` structure 66
`dylib_module` structure 78
`dylib_reference` structure 79
`dylib_table_of_contents` structure 78
`dlinker_command` structure 67
dynamic linker 15, 16
dynamic shared libraries 21
`dysymtab_command` structure 76

E

`errno` variable 16
`execve` function 16
external symbol 20

F

FALSE constant [112](#)
 fat_arch structure [87](#)
 fat_header structure [86](#)
 file package [14](#)
 fork function [16](#)
 frameworks [13](#)
 function value return
 PowerPC [39](#)

G

getsectbyname function [100](#)
 getsectbynamefromheader function [101](#)
 getsectdata function [101](#)
 getsectdatafromFramework function [102](#)
 getsectdatafromheader function [102](#)
 getsegbyname function [100](#)

H

HotSpot Java virtual machine [7](#)

I

install name [22](#)
 intermediate object files [13](#)

J

Java virtual machine [7](#)
 just-in-time binding [17](#)

L

LaunchCFMApp tool [7](#)
 lazy binding [17](#)
 lc_str union [65](#)
 load commands [16](#)
 load-time binding [17](#)
 load_command structure [56](#)

M

Mach-O [7](#)
 mach_header data structure [16](#)
 mach_header structure [53](#)
 main function [16](#)
 memory
 freeing [51](#)
 module [14](#)

N

nlist structure [73](#)
 NSAddImage function [89](#)
 NSAddressOfSymbol function [99](#)
 NSBundle class [25](#)
 NSCreateObjectFileImageFromFile function [90](#)
 NSCreateObjectFileImageFromMemory function
 [91](#)
 NSDestroyObjectFileImage function [92](#)
 NSIsSymbolNameDefined function [95](#)
 NSIsSymbolNameDefinedInImage function [96](#)
 NSIsSymbolNameDefinedWithHint function [95](#)
 NSLibraryNameForModule function [93](#)
 NSLinkModule function [93](#)
 NSLookupAndBindSymbol function [96](#)
 NSLookupAndBindSymbolWithHint function [97](#)
 NSLookupSymbolInImage function [97](#)
 NSLookupSymbolInModule function [99](#)
 NSModuleForSymbol function [100](#)
 NSNameOfModule function [92](#)
 NSNameOfSymbol function [99](#)
 NSUnLinkModule function [94](#)

O

object file image functions [25](#)

P

parameter area
 in PowerPC stack
 plug-in. *See* bundle
 PowerPC implementation of CFM-based
 architecture
 routine calling conventions [39](#)
 prebinding [17](#)
 prebound_dylib_command structure [67](#)

Preferred Executable Format (PEF) [7, 25](#)
 private defined symbol [20](#)

R

ranlib structure [84](#)
 registers, PowerPC environment
 and function value return [39](#)
 preservation [39](#)
 saving and restoring values in
 relocation entries [80](#)
 relocation_info structure [80](#)
 routine calling conventions
 function value return
 PowerPC [39](#)
 PowerPC [39](#)
 register preservation
 PowerPC [39](#)
 routines_command structure [69](#)
 runtime architecture, defined [7](#)

S

scattered_relocation_info structure [81](#)
 section structure [60](#)
 segment_command structure [58](#)
 shared libraries [13](#)
 shared libraries, and versioning [22](#)
 static archive libraries [14](#)
 static archive library [15](#)
 sub_client_command structure [71](#)
 sub_framework_command structure [70](#)
 sub_library_command structure [71](#)
 sub_umbrella_command structure [70](#)
 symbol [18](#)
 symtab_command structure [72](#)

T

tentative symbol [20](#)
 thread_command structure [68](#)
 TRUE constant [112](#)
 two-level namespace hint table [18, 19](#)
 two-level symbol namespace [18](#)
 twolevel_hint structure [64](#)
 twolevel_hints_command structure [63](#)
 -twelevel_hint_table linker option [26](#)

U

umbrella framework [24](#)
 umbrella frameworks [13](#)
 /usr/lib/crt1.o [16](#)

W

weak binding [18](#)