



Linker and Libraries Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 817-3677-10
April 2004

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2004 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



031111@6671



Contents

Preface	15
1 Introduction to the Solaris Linkers	19
Link-Editing	20
Runtime Linking	21
Related Topics	22
Dynamic Linking	22
Application Binary Interfaces	22
32-Bit and 64-Bit Environments	22
Environment Variables	23
Support Tools	23
2 Link-Editor	25
Invoking the Link-Editor	26
Direct Invocation	26
Using a Compiler Driver	27
Specifying the Link-Editor Options	27
Input File Processing	28
Archive Processing	29
Shared Object Processing	30
Linking With Additional Libraries	31
Initialization and Termination Sections	36
Symbol Processing	38
Symbol Resolution	38
Undefined Symbols	42

Tentative Symbol Order Within the Output File	46
Defining Additional Symbols	46
Reducing Symbol Scope	52
External Bindings	56
String Table Compression	56
Generating the Output File	57
Relocation Processing	58
Displacement Relocations	58
Debugging Aids	60
3 Runtime Linker	63
Shared Object Dependencies	64
Locating Shared Object Dependencies	64
Directories Searched by the Runtime Linker	64
Configuring the Default Search Paths	67
Dynamic String Tokens	67
Relocation Processing	67
Symbol Lookup	68
When Relocations Are Performed	71
Relocation Errors	72
Loading Additional Objects	73
Lazy Loading of Dynamic Dependencies	74
Initialization and Termination Routines	76
Initialization and Termination Order	77
Security	80
Runtime Linking Programming Interface	81
Loading Additional Objects	82
Relocation Processing	83
Obtaining New Symbols	89
Feature Checking	92
Debugging Aids	93
Debugging Library	93
Debugger Module	96
4 Shared Objects	101
Naming Conventions	102
Recording a Shared Object Name	103

Shared Objects With Dependencies	105
Dependency Ordering	106
Shared Objects as Filters	107
Generating Standard Filters	108
Generating Auxiliary Filters	110
Filter Processing	111
Performance Considerations	112
Analyzing Files	112
Underlying System	114
Lazy Loading of Dynamic Dependencies	114
Position-Independent Code	115
Remove Unused Material	117
Maximizing Shareability	118
Minimizing Paging Activity	120
Relocations	120
Using <code>-B symbolic</code>	125
Profiling Shared Objects	125
5 Application Binary Interfaces and Versioning	129
Interface Compatibility	130
Internal Versioning	131
Creating a Version Definition	131
Binding to a Version Definition	136
Specifying a Version Binding	140
Version Stability	144
Relocatable Objects	145
External Versioning	145
Coordination of Versioned Filenames	146
6 Support Interfaces	149
Link-Editor Support Interface	149
Invoking the Support Interface	150
Support Interface Functions	151
Support Interface Example	153
Runtime Linker Auditing Interface	155
Establishing a Namespace	156
Creating an Audit Library	156

Invoking the Auditing Interface	157
Recording Local Auditors	157
Audit Interface Functions	158
Audit Interface Example	162
Audit Interface Demonstrations	163
Audit Interface Limitations	163
Runtime Linker Debugger Interface	164
Interaction Between Controlling and Target Process	165
Debugger Interface Agents	166
Debugger Exported Interface	167
Debugger Import Interface	175
7 Object File Format	177
File Format	177
Data Representation	179
ELF Header	180
ELF Identification	184
Data Encoding	186
Sections	187
Special Sections	199
COMDAT Section	203
Group Section	204
Hash Table Section	205
Move Section	206
Note Section	208
Relocation Sections	210
String Table Section	221
Symbol Table Section	222
Syminfo Table Section	229
Versioning Sections	231
Dynamic Linking	236
Program Header	236
Program Loading (Processor-Specific)	241
Runtime Linker	247
Dynamic Section	248
Global Offset Table (Processor-Specific)	260
Procedure Linkage Table (Processor-Specific)	261

8	Thread-Local Storage	271
	C/C++ Programming Interface	271
	Thread-Local Storage Section	272
	Runtime Allocation of Thread-Local Storage	274
	Program Startup	274
	Thread Creation	275
	Post-Startup Dynamic Loading	276
	Deferred Allocation of Thread-Local Storage Blocks	276
	Thread-Local Storage Access Models	277
	SPARC: Thread-Local Variable Access	278
	SPARC: Thread-Local Storage Relocation Types	283
	x86: Thread-Local Variable Access	286
	x86: Thread-Local Storage Relocation Types	290
9	Mapfile Option	293
	Mapfile Structure and Syntax	293
	Segment Declarations	294
	Mapping Directives	298
	Section-Within-Segment Ordering	299
	Size-Symbol Declarations	300
	File Control Directives	300
	Mapping Example	300
	Mapfile Option Defaults	302
	Internal Map Structure	303
A	Link-Editor Quick Reference	307
	Static Mode	307
	Creating a Relocatable Object	308
	Creating a Static Executable	308
	Dynamic Mode	308
	Creating a Shared Object	308
	Creating a Dynamic Executable	310
B	Versioning Quick Reference	311
	Naming Conventions	311
	Defining a Shared Object's Interface	313
	Versioning a Shared Object	313

Versioning an Existing (Non-versioned) Shared Object	314
Updating a Versioned Shared Object	315
Adding New Symbols	315
Internal Implementation Changes	316
New Symbols and Internal Implementation Changes	316
Migrating Symbols to a Standard Interface	317
C Establishing Dependencies with Dynamic String Tokens	321
Instruction Set Specific Shared Objects	321
Reducing Filtee Searches	322
System Specific Shared Objects	323
Locating Associated Dependencies	323
Dependencies Between Unbundled Products	325
Security	326
D New Linker and Libraries Features and Updates	329
Solaris 9 4/04 Release	329
Solaris 9 12/03 Release	329
Solaris 9 8/03 Release	330
Solaris 9 12/02 Release	330
Solaris 9 Release	330
Solaris 8 07/01 Release	331
Solaris 8 01/01 Release	331
Solaris 8 10/00 Release	332
Solaris 8 Release	332
Solaris 7 Release	333
Solaris 2.6 Release	333
Index	335

Tables

TABLE 5-1	Interface Compatibility Examples	130
TABLE 7-1	ELF 32-Bit Data Types	179
TABLE 7-2	ELF 64-Bit Data Types	179
TABLE 7-3	ELF File Identifiers	181
TABLE 7-4	ELF Machines	181
TABLE 7-5	ELF Versions	182
TABLE 7-6	SPARC: ELF Flags	182
TABLE 7-7	ELF Identification Index	184
TABLE 7-8	ELF Magic Number	184
TABLE 7-9	ELF File Class	185
TABLE 7-10	ELF Data Encoding	185
TABLE 7-11	ELF Special Section Indexes	187
TABLE 7-12	ELF Section Types, <i>sh_type</i>	191
TABLE 7-13	ELF Section Header Table Entry: Index 0	195
TABLE 7-14	ELF Section Attribute Flags	195
TABLE 7-15	ELF <i>sh_link</i> and <i>sh_info</i> Interpretation	198
TABLE 7-16	ELF Special Sections	199
TABLE 7-17	ELF Group Section Flag	204
TABLE 7-18	SPARC: ELF Relocation Types	215
TABLE 7-19	64-bit SPARC: ELF Relocation Types	219
TABLE 7-20	x86: ELF Relocation Types	220
TABLE 7-21	ELF String Table Indexes	222
TABLE 7-22	ELF Symbol Binding, <i>ELF32_ST_BIND</i> and <i>ELF64_ST_BIND</i>	224
TABLE 7-23	ELF Symbol Types, <i>ELF32_ST_TYPE</i> and <i>ELF64_ST_TYPE</i>	225
TABLE 7-24	ELF Symbol Visibility	226
TABLE 7-25	ELF Symbol Table Entry: Index 0	228

TABLE 7-26	SPARC: ELF Symbol Table Entry: Register Symbol	229
TABLE 7-27	SPARC: ELF Register Numbers	229
TABLE 7-28	ELF <code>si_boundto</code> Reserved Values	230
TABLE 7-29	ELF Syminfo Flags	230
TABLE 7-30	ELF Version Definition Structure Versions	232
TABLE 7-31	ELF Version Definition Section Flags	232
TABLE 7-32	ELF Version Dependency Indexes	233
TABLE 7-33	ELF Version Dependency Structure Versions	235
TABLE 7-34	ELF Version Dependency Structure Flags	235
TABLE 7-35	ELF Segment Types	238
TABLE 7-36	ELF Segment Flags	240
TABLE 7-37	ELF Segment Permissions	240
TABLE 7-38	SPARC: ELF Program Header Segments (64K alignment)	243
TABLE 7-39	x86: ELF Program Header Segments (64K alignment)	244
TABLE 7-40	SPARC: ELF Example Shared Object Segment Addresses	247
TABLE 7-41	x86: ELF Example Shared Object Segment Addresses	247
TABLE 7-42	ELF Dynamic Array Tags	249
TABLE 7-43	ELF Dynamic Flags, <code>DT_FLAGS</code>	256
TABLE 7-44	ELF Dynamic Flags, <code>DT_FLAGS_1</code>	257
TABLE 7-45	ELF Dynamic Position Flags, <code>DT_POSFLAG_1</code>	259
TABLE 7-46	ELF Dynamic Feature Flags, <code>DT_FEATURE_1</code>	260
TABLE 7-47	SPARC: Procedure Linkage Table Example	262
TABLE 7-48	64-bit SPARC: Procedure Linkage Table Example	265
TABLE 7-49	x86: Absolute Procedure Linkage Table Example	268
TABLE 7-50	x86: Position-Independent Procedure Linkage Table Example	268
TABLE 8-1	ELF <code>PT_TLS</code> Program Header Entry	273
TABLE 8-2	SPARC: 32-bit and 64-bit General Dynamic Thread-Local Variable Access Codes	279
TABLE 8-3	SPARC: 32-bit and 64-bit Local Dynamic Thread-Local Variable Access Codes	280
TABLE 8-4	SPARC: 32-bit Initial Executable Thread-Local Variable Access Codes	282
TABLE 8-5	SPARC: 64-bit Initial Executable Thread-Local Variable Access Codes	283
TABLE 8-6	SPARC: 32-bit and 64-bit Local Executable Thread-Local Variable Access Codes	283
TABLE 8-7	SPARC: Thread-Local Storage Relocation Types	284
TABLE 8-8	x86: General Dynamic Thread-Local Variable Access Codes	286
TABLE 8-9	x86: Local Dynamic Thread-Local Variable Access Codes	287

TABLE 8-10	x86: Initial Executable, Position Independent, Thread-Local Variable Access Codes	288
TABLE 8-11	x86: Initial Executable, Position Dependent, Thread-Local Variable Access Codes	288
TABLE 8-12	x86: Initial Executable, Position Independent, Dynamic Thread-Local Variable Access Codes	289
TABLE 8-13	x86: Initial Executable, Position Independent, Thread-Local Variable Access Codes	289
TABLE 8-14	x86: Local Executable Thread-Local Variable Access Codes	290
TABLE 8-15	x86: Local Executable Thread-Local Variable Access Codes	290
TABLE 8-16	x86: Local Executable Thread-Local Variable Access Codes	290
TABLE 8-17	x86: Thread-Local Storage Relocation Types	291
TABLE 9-1	Mapfile Segment Attributes	295
TABLE 9-2	Section Attributes	298

Figures

FIGURE 1-1	Static or Dynamic Link-Editing	20
FIGURE 3-1	A Single <code>dlopen()</code> Request	85
FIGURE 3-2	Multiple <code>dlopen()</code> Requests	86
FIGURE 3-3	Multiple <code>dlopen()</code> Requests With A Common Dependency	87
FIGURE 6-1	<code>rtld-debugger</code> Information Flow	165
FIGURE 7-1	Object File Format	177
FIGURE 7-2	Data Encoding <code>ELFDATA2LSB</code>	186
FIGURE 7-3	Data Encoding <code>ELFDATA2MSB</code>	186
FIGURE 7-4	Symbol Hash Table	205
FIGURE 7-5	Note Information	208
FIGURE 7-6	Example Note Segment	209
FIGURE 7-7	ELF String Table	221
FIGURE 7-8	SPARC: Executable File (64K alignment)	242
FIGURE 7-9	x86: Executable File (64K alignment)	243
FIGURE 7-10	SPARC: Process Image Segments	244
FIGURE 7-11	x86: Process Image Segments	245
FIGURE 8-1	Runtime Storage Layout of Thread-Local Storage	274
FIGURE 8-2	Thread-Local Storage Access Models and Transitions	277
FIGURE 9-1	Simple Map Structure	303
FIGURE C-1	Unbundled Dependencies	323
FIGURE C-2	Unbundled Co-Dependencies	325

Preface

In the Solaris™ operating environment, application developers can create applications and libraries using the link-editor `ld(1)`, and execute these objects with the aid of the runtime linker `ld.so.1(1)`. This manual is for those who want to understand more fully the concepts involved in using the Solaris linkers.

About This Manual

This manual describes the operations of the Solaris link-editor and runtime linker. Special emphasis is placed on the generation and use of dynamic executables and shared objects because of their importance in a dynamic runtime environment.

Intended Audience

This manual is intended for a range of programmers who are interested in the Solaris linkers, from the curious beginner to the advanced user.

- Beginners learn the principle operations of the link-editor and runtime linker.
- Intermediate programmers learn to create, and use, efficient custom libraries.
- Advanced programmers, such as language-tools developers, learn how to interpret and generate object files.

Not many programmers should need to read this manual from cover to cover.

Organization

Chapter 1 gives an overview of the linking processes under the Solaris operating environment, together with an introduction of new features added with this release. This chapter is intended for all programmers.

Chapter 2 describes the functions of the link-editor, its two modes of linking (*static* and *dynamic*), scope and forms of input, and forms of output. This chapter is intended for all programmers.

Chapter 3 describes the execution environment and program-controlled runtime binding of code and data. This chapter is intended for all programmers.

Chapter 4 provides definitions of shared objects, describes their mechanisms, and explains how to create and use them. This chapter is intended for all programmers.

Chapter 5 describes how to manage the evolution of an interface provided by a dynamic object. This chapter is intended for all programmers.

Chapter 6 describes interfaces for monitoring, and in some cases modifying, link-editor and runtime linker processing. This chapter is intended for advanced programmers.

Chapter 7 is a reference chapter on ELF files. This chapter is intended for advanced programmers.

Chapter 8 describes Thread-Local Storage. This chapter is intended for advanced programmers.

Chapter 9 describes the `mapfile` directives to the link-editor, which specify the layout of the output file. This chapter is intended for advanced programmers.

Appendix A provides an overview of the most commonly used link-editor options, and is intended for all programmers.

Appendix B provides naming conventions and guidelines for versioning shared objects, and is intended for all programmers.

Appendix C provides examples of how to use reserved dynamic string tokens to define dynamic dependencies, and is intended for all programmers.

Appendix D provides an overview of new features and updates that have been added to the link-editors and indicates to which release they were added.

Throughout this document, all command-line examples use `sh(1)` syntax, and all programming examples are written in the C language.

Note – In this document the term “x86” refers to the Intel 32-bit family of microprocessor chips compatible microprocessor chips made by AMD.

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is <http://docs.sun.com>.

Ordering Sun Documentation

Sun Microsystems offers select product documentation in print. For a list of documents and how to order them, see “Buy printed documentation” at <http://docs.sun.com>.

Typographic Conventions

The following table describes the typographic changes used in this book.

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name% su</code> Password:

TABLE P-1 Typographic Conventions (Continued)

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type rm <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You must be <i>root</i> to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Introduction to the Solaris Linkers

This manual describes the operations of the Solaris link-editor and runtime linker, together with the objects on which they operate. The basic operation of the Solaris linkers involves the combination of objects. The symbolic references from one object are connected to the symbolic definitions within another object. This operation is often referred to as *binding*.

This manual expands the following areas:

Link-Editor

The link-editor, `ld(1)`, concatenates and interprets data from one or more input files. These files can be relocatable objects, shared objects, or archive libraries. From these input files, one output file is created. This file is either a relocatable object, an executable application, or a shared object. The link-editor is most commonly invoked as part of the compilation environment.

Runtime Linker

The runtime linker, `ld.so.1(1)`, processes dynamic executables and shared objects at runtime, and binds them to create a runnable process.

Shared Objects

Shared objects are one form of output from the link-edit phase. Shared objects are sometimes referred to as *Shared Libraries*. Shared objects are important in creating a powerful, flexible runtime environment.

Object Files

The Solaris linkers work with files that conform to the executable and linking format (ELF).

These areas, although separable into individual topics, have a great deal of overlap. While explaining each area, this document brings together the connecting principles.

Link-Editing

Link-editing takes a variety of input files, typically generated from compilers, assemblers, or `ld(1)`. The link-editor concatenates and interprets the data within these input files to form a single output file. Although the link-editor provides numerous options, the output file that it produces is one of four basic types:

- *Relocatable object* – A concatenation of input relocatable objects that can be used in subsequent link-edit phases.
- *Static executable* – A concatenation of input relocatable objects that has all symbolic references bound to the executable, and thus represents a ready-to-run process.
- *Dynamic executable* – A concatenation of input relocatable objects that requires intervention by the runtime linker to produce a runnable process. A dynamic executable might still need symbolic references bound at runtime, and can have one or more dependencies in the form of shared objects.
- *Shared object* – A concatenation of input relocatable objects that provides services that might be bound to a dynamic executable at runtime. The shared object can have dependencies on other shared objects.

These output files, and the key link-editor options used to create them, are shown in Figure 1-1.

Dynamic executables and *shared objects* are often referred to jointly as *dynamic objects*. Dynamic objects are the main focus of this document.




FIGURE 1-1 Static or Dynamic Link-Editing

Runtime Linking

Runtime linking involves the binding of objects, usually generated from one or more previous link-edits, to generate a runnable process. During the generation of these objects by the link-editor, the binding requirements are verified and appropriate bookkeeping information is added to each object to enable the runtime linker to load, relocate, and complete the binding process.

During process execution the facilities of the runtime linker are made available, and can be used to extend the process' address space by adding additional shared objects on demand. The two most common components involved in runtime linking are *dynamic executables* and *shared objects*.

Dynamic executables are applications that are executed under the control of a runtime linker. These applications usually have dependencies in the form of shared objects, which are located and bound by the runtime linker to create a runnable process. Dynamic executables are the default output file generated by the link-editor.

Shared objects provide the key building block to a dynamically linked system. A shared object is similar to a dynamic executable; however, shared objects have not yet been assigned a virtual address.

Dynamic executables usually have dependencies on one or more shared objects. That is, the shared object(s) must be bound to the dynamic executable to produce a runnable process. Because shared objects can be used by many applications, aspects of their construction directly affect shareability, versioning, and performance.

You can distinguish the processing of shared objects by either the link-editor or the runtime linker by referring to the *environments* in which the shared objects are being used:

compilation environment

Shared objects are processed by the link-editor to generate dynamic executables or other shared objects. The shared objects become dependencies of the output file being generated.

runtime environment

Shared objects are processed by the runtime linker, together with a dynamic executable, to produce a runnable process.

Related Topics

Dynamic Linking

Dynamic linking is a term often used to embrace those portions of the link-editing process that generate dynamic executables and shared objects, together with the runtime linking of these objects to generate a runnable process. Dynamic linking enables multiple applications to use the code provided by a shared object by enabling the application to bind to the shared object at runtime.

By separating an application from the services of standard libraries, dynamic linking also increases the portability and extensibility of an application. This separation between the interface of a service and its implementation enables the system to evolve while maintaining application stability. Dynamic linking is a crucial factor in providing an *application binary interface* (ABI), and is the preferred compilation method for Solaris applications.

Application Binary Interfaces

Binary interfaces between system and application components are defined to enable the asynchronous evolution of these facilities. The Solaris linkers operate upon these interfaces to assemble applications for execution. Although all components handled by the Solaris linkers have binary interfaces, the whole set of interfaces provided by the system is referred to as the *Solaris ABI*.

The Solaris ABI is a technological descendent for work on ABIs that started with the *System V Application Binary Interface* and the successor work performed by SPARC International, Inc.® for SPARC® processors called the *SPARC Compliance Definition* (SCD).

32–Bit and 64–Bit Environments

The link-editors operate on 32-bit objects, and on SPARCV9 systems are also capable of operating on 64-bit objects. On SPARC systems, the 64-bit link-editor (ld(1)) is capable of generating 32-bit objects and the 32-bit link-editor is capable of generating 64-bit objects. In the latter case, the size of the generated object, not including the `.bss`, is restricted to 2 Gbytes.

No command-line option is required to distinguish a 32-bit or 64-bit link-edit. The link-editor uses the ELF class of the first input relocatable object file it sees on the command-line to govern the mode in which it will operate. Specialized link-edits, such

as linking solely from a `mapfile` or an archive library, are uninfluenced by their input files, and will default to a 32-bit mode. In these cases a 64-bit link-edit can be enforced with the `-64` option. Intermixing of 32-bit and 64-bit objects is not permitted.

The operations of the link-editors on 32-bit and 64-bit objects is identical. This document typically uses 32-bit examples. Cases where 64-bit processing differs from the 32-bit processing are highlighted.

For more information regarding 64-bit applications, refer to the *Solaris 64-bit Developer's Guide*.

Environment Variables

The link-editors support a number of environment variables that begin with the characters `LD_`, for example `LD_LIBRARY_PATH`. Each environment variable can exist in its generic form, or can be specified with a `_32` or `_64` suffix, for example `LD_LIBRARY_PATH_64`. This suffix makes the environment variable specific, respectively, to 32-bit or 64-bit processes. This suffix also overrides any generic, non-suffixed, version of the environment variable that may be in effect.

Throughout this document, any reference to link-editor environment variables uses the generic, non-suffixed, variant. For a list of all supported environment variables refer to the `ld(1)` and `ld.so.1(1)` man pages.

Support Tools

The Solaris operating environment also provides several support tools and libraries. These tools provide for the analysis and inspection of these objects and the linking processes. These tools include `elfdump(1)`, `nm(1)`, `dump(1)`, `ldd(1)`, `pvs(1)`, `elf(3ELF)`, and a linker debugging support library. Throughout this document, many discussions are augmented with examples of these tools.

Link-Editor

The link-editing process creates an output file from one or more input files. The creation of the output file is directed by the options supplied to the link-editor together with the input sections provided by the input files.

All files are represented in the *executable and linking format* (ELF). For a complete description of the ELF format see Chapter 7. For this introduction, however, it is first necessary to introduce two ELF structures, *sections* and *segments*.

Sections are the smallest indivisible units that can be processed within an ELF file. Segments are a collection of sections that represent the smallest individual units that can be mapped to a memory image by `exec(2)` or by the runtime linker `ld.so.1(1)`.

Although there are many types of ELF sections, they all fall into two categories with respect to the link-editing phase:

- Sections that contain *program data*, whose interpretation is meaningful only to the application itself, such as the program instructions `.text` and the associated data `.data` and `.bss`.
- Sections that contain *link-editing information*, such as the symbol table information found from `.symtab` and `.strtab`, and relocation information such as `.rela.text`.

Basically, the link-editor concatenates the *program data* sections into the output file. The *link-editing information* sections are interpreted by the link-editor to modify other sections or to generate new output information sections used in later processing of the output file.

The following simple breakdown of link-editor functionality introduces the topics covered in this chapter:

- It verifies and checks for consistency all the options passed to it.
- It concatenates sections of the same characteristics (for example, type, attributes, and name) from the input relocatable objects to form new sections within the output file. These concatenated sections can in turn be associated to output

segments.

- It reads symbol table information from both relocatable objects and shared objects to verify and unite references with definitions, and usually generates a new symbol table, or tables, within the output file.
- It reads relocation information from the input relocatable objects and applies this information to the output file by updating other input sections. In addition, output relocation sections might be generated for use by the runtime linker.
- It generates *program headers* that describe all segments created.
- It generates dynamic linking information sections if necessary, which provide information such as shared object dependencies and symbol bindings to the runtime linker.

The process of concatenating like *sections* and associating *sections* to *segments* is carried out using default information within the link-editor. The default *section* and *segment* handling provided by the link-editor is usually sufficient for most link-edits. However, these defaults can be manipulated using the `-M` option with an associated `mapfile`. See Chapter 9.

Invoking the Link-Editor

You can either run the link-editor directly from the command line or have a compiler driver invoke it for you. In the following two sections the description of both methods are expanded. However, using the compiler driver is the preferred choice. The compilation environment is often the consequence of a complex and occasionally changing series of operations known only to compiler drivers.

Direct Invocation

When you invoke the link-editor directly, you have to supply every object file and library required to create the intended output. The link-editor makes no assumptions about the object modules or libraries that you meant to use in creating the output. For example, when you issue the command:

```
$ ld test.o
```

the link-editor creates a dynamic executable named `a.out` using only the input file `test.o`. For the `a.out` to be a useful executable, it should include startup and exit processing code. This code can be language or operating system specific, and is usually provided through files supplied by the compiler drivers.

Additionally, you can also supply your own initialization and termination code. This code must be encapsulated and labeled correctly for it to be correctly recognized and made available to the runtime linker. This encapsulation and labeling can also be provided through files supplied by the compiler drivers.

When creating runtime objects such as executables and shared objects, you should use a compiler driver to invoke the link-editor. Invoking the link-editor directly is recommended only when creating intermediate relocatable objects when using the `-r` option.

Using a Compiler Driver

The conventional way to use the link-editor is through a language-specific compiler driver. You supply the compiler driver, `cc(1)`, `CC(1)`, and so forth, with the input files that make up your application. The compiler driver adds additional files and default libraries to complete the link-edit. These additional files can be seen by expanding the compilation invocation, for example:

```
$ cc -# -o prog main.o
/usr/ccs/bin/ld -dy /opt/COMPILER/crti.o /opt/COMPILER/crt1.o \
/usr/ccs/lib/values-Xt.o -o prog main.o \
-YP,/opt/COMPILER/lib:/usr/ccs/lib:/usr/lib -Qy -lc \
/opt/COMPILER/crtn.o
```

Note – The actual files included by your compiler driver and the mechanism used to display the link-editor invocation might differ.

Specifying the Link-Editor Options

Most options to the link-editor can be passed through the compiler driver command line. For the most part the compiler and the link-editor options do not conflict. Where a conflict arises, the compiler drivers usually provide a command-line syntax you can use to pass specific options to the link-editor. You can also provide options to the link-editor by setting the `LD_OPTIONS` environment variable. For example:

```
$ LD_OPTIONS="-R /home/me/libs -L /home/me/libs" cc -o prog main.c -lfoo
```

The `-R` and `-L` options are interpreted by the link-editor and prepended to any command-line options received from the compiler driver.

The link-editor parses the entire option list for any invalid options or any options with invalid associated arguments. When either of these cases is found, a suitable error message is generated. If the error is deemed fatal, the link-edit terminates. In the following example, the illegal option `-X` is identified, and the illegal argument to the `-z` option is caught by the link-editor's checking.

```
$ ld -X -z sillydefs main.o
ld: illegal option -- X
ld: fatal: option -z has illegal argument `sillydefs'
```

If an option requiring an associated argument is mistakenly specified twice, the link-editor will provide a suitable warning but will continue with the link-edit. For example:

```
$ ld -e foo ..... -e bar main.o
ld: warning: option -e appears more than once, first setting taken
```

The link-editor also checks the option list for any fatal inconsistencies. For example:

```
$ ld -dy -a main.o
ld: fatal: option -dy and -a are incompatible
```

After processing all options, if no fatal error conditions have been detected, the link-editor proceeds to process the input files.

See Appendix A for the most commonly used link-editor options, and the `ld(1)` man page for a complete description of all link-editor options.

Input File Processing

The link-editor reads input files in the order in which they appear on the command line. Each file is opened and inspected to determine its ELF file type and therefore determine how it must be processed. The file types that apply as input for the link-edit are determined by the binding mode of the link-edit, either *static* or *dynamic*.

Under *static* mode, the link-editor accepts only relocatable objects or archive libraries as input files. Under *dynamic* mode, the link-editor also accepts shared objects.

Relocatable objects represent the most basic input file type to the link-editing process. The *program data* sections within these files are concatenated into the output file image being generated. The *link-edit information* sections are organized for later use, but do not become part of the output file image, as new sections are generated to take their places. Symbols are gathered into an internal symbol table for verification and resolution. This table is then used to create one or more symbol tables in the output image.

Although any input file can be specified directly on the link-edit command-line, archive libraries and shared objects are commonly specified using the `-l` option. See “Linking With Additional Libraries” on page 31 for coverage of this mechanism and how it relates to the two different linking modes. However, even though shared objects are often referred to as shared *libraries*, and both of these objects can be specified using the same option, the interpretation of shared objects and archive libraries is quite different. The next two sections expand upon these differences.

Archive Processing

Archives are built using `ar(1)`, and usually consist of a collection of relocatable objects with an archive symbol table. This symbol table provides an association of symbol definitions with the objects that supply these definitions. By default, the link-editor provides *selective* extraction of archive members. When the link-editor reads an archive, it uses information within the internal symbol table it is creating to select only the objects from the archive it requires to complete the binding process. You can also explicitly extract all members of an archive.

The link-editor extracts a relocatable object from an archive if:

- The archive member contains a symbol definition that satisfies a symbol reference, sometimes referred to as an *undefined* symbol, presently held in the link-editor’s internal symbol table.
- The archive member contains a data symbol definition that satisfies a tentative symbol definition presently held in the link-editor’s internal symbol table. An example of this is a FORTRAN COMMON block definition, which causes the extraction of a relocatable object that defines the same DATA symbol.
- The link-editor’s `-z allextract` is in effect. This option suspends selective archive extraction and causes all archive members to be extracted from the archive being processed.

Under selective archive extraction, a weak symbol reference does not extract an object from an archive unless the `-z weakextract` option is in effect. See “Simple Resolutions” on page 39 for more information.

Note – The options `-z weakextract`, `-z allextract`, and `-z defaultextract` enable you to toggle the archive extraction mechanism among multiple archives.

With selective archive extraction, the link-editor makes multiple passes through an archive to extract relocatable objects as needed to satisfy the symbol information being accumulated in the link-editor internal symbol table. After the link-editor has made a complete pass through the archive without extracting any relocatable objects, it moves on to process the next input file.

By extracting from the archive only the relocatable objects needed at the time the archive was encountered, the position of the archive within the input file list can be significant. See “Position of an Archive on the Command Line” on page 32.

Note – Although the link-editor makes multiple passes through an archive to resolve symbols, this mechanism can be quite costly for large archives containing random organizations of relocatable objects. In these cases, you should use tools like `lorder(1)` and `tsort(1)` to order the relocatable objects within the archive and so reduce the number of passes the link-editor must carry out.

Shared Object Processing

Shared objects are indivisible whole units that have been generated by a previous link-edit of one or more input files. When the link-editor processes a shared object, the entire contents of the shared object become a logical part of the resulting output file image. This logical inclusion means that all symbol entries defined in the shared object are made available to the link-editing process. The shared object is actually copied during process execution.

The shared object’s program data sections and most of the link-editing information sections are unused by the link-editor. These sections are interpreted by the runtime linker when the shared object is bound to generate a runnable process. However, the occurrence of a shared object is remembered, and information is stored in the output file image to indicate that this object is a dependency and must be made available at runtime.

By default, all shared objects specified as part of a link-edit are recorded as dependencies in the object being built. This recording is made regardless of whether the object being built actually references symbols offered by the shared object. To minimize runtime linking overhead, specify only those dependencies required to resolve symbol references from the object being built as part of the link-edit. The link-editor’s debugging capabilities, and `ldd(1)` with the `-u` option, can be used to determine unused dependencies. Alternatively, the link-editor’s `-z ignore` option can suppress the dependency recording of unused shared objects.

If a shared object has dependencies on other shared objects, these dependencies are also processed. This processing occurs after all command-line input files have been processed. These shared objects will be used to complete the symbol resolution process; however, their names will not be recorded as dependencies in the output file image being generated.

Although the position of a shared object on the link-edit command-line has less significance than it does for archive processing, the position can have a global effect. Multiple symbols of the same name are allowed to occur between relocatable objects and shared objects, and between multiple shared objects. See “Symbol Resolution” on page 38.

The order of shared objects processed by the link-editor is maintained in the dependency information stored in the output file image. As the runtime linker reads this information, it loads the specified shared objects in the same order. Therefore, the link-editor and the runtime linker select the first occurrence of a symbol of a multiply-defined series of symbols.

Note – Multiple symbol definitions, and thus the information to describe the interposing of one definition of a symbol for another, are reported in the load map output generated using the `-m` option.

Linking With Additional Libraries

Although the compiler drivers often ensure that appropriate libraries are specified to the link-editor, frequently you must supply your own. Shared objects and archives can be specified by explicitly naming the input files required to the link-editor, but a more common and more flexible method involves using the link-editor's `-l` option.

Library Naming Conventions

By convention, shared objects are usually designated by the prefix `lib` and the suffix `.so`, and archives are designated by the prefix `lib` and the suffix `.a`. For example, `libc.so` is the shared object version of the standard C library made available to the compilation environment, and `libc.a` is the library's archive version.

These conventions are recognized by the `-l` option of the link-editor. This option is commonly used to supply additional libraries to a link-edit. The following example directs the link-editor to search for `libfoo.so`. If the link-editor does not find `libfoo.so`, it searches for `libfoo.a` before moving on to the next directory to be searched.

```
$ cc -o prog file1.c file2.c -lfoo
```

Note – There is a naming convention regarding the compilation environment and the runtime environment use of shared objects. The compilation environment uses the simple `.so` suffix, whereas the runtime environment commonly uses the suffix with an additional version number. See “Naming Conventions” on page 102 and “Coordination of Versioned Filenames” on page 146.

When link-editing in dynamic mode, you can choose to link with a mix of shared objects and archives. When link-editing in static mode, only archive libraries are acceptable for input.

When in dynamic mode and using the `-l` option to enable a library search, the link-editor will first search in a given directory for a shared object that matches the specified name. If no match is found, the link-editor looks for an archive library in the same directory. When in static mode and using the `-l` option, only archive libraries are sought.

Linking With a Mix of Shared Objects and Archives

The library search mechanism in dynamic mode searches a given directory for a shared object, and then searches an archive library. Finer control of the type of search required is possible through the `-B` option.

By specifying the `-B dynamic` and `-B static` options on the command line as many times as required, you can toggle the library search between shared objects or archives respectively. For example, to link an application with the archive `libfoo.a` and the shared object `libbar.so`, issue the following command:

```
$ cc -o prog main.o file1.c -Bstatic -lfoo -Bdynamic -lbar
```

The `-B static` and `-B dynamic` keywords are not exactly symmetrical. When you specify `-B static`, the link-editor does not accept shared objects as input until the next occurrence of `-B dynamic`. However, when you specify `-B dynamic`, the link-editor first looks for shared objects and then archive library's in any given directory.

The precise description of the previous example is that the link-editor first searches for `libfoo.a`, and then for `libbar.so`, and if that search fails, for `libbar.a`. Finally, it searches for `libc.so`, and if that search fails, `libc.a`.

Position of an Archive on the Command Line

The position of an archive on the command line can affect the output file being produced. The link-editor searches an archive only to resolve undefined or tentative external references it has previously seen. After this search is completed and any required members have been extracted, the link-editor moves onto the next input file on the command line.

Therefore by default, the archive is not available to resolve any new references from the input files that follow the archive on the command line. For example, the following command directs the link-editor to search `libfoo.a` only to resolve symbol references that have been obtained from `file1.c`. The `libfoo.a` archive is not available to resolve symbol references from `file2.c` or `file3.c`.

```
$ cc -o prog file1.c -Bstatic -lfoo file2.c file3.c -Bdynamic
```

Note – You should specify any archives at the end of the command line unless multiple-definition conflicts require you to do otherwise.

In some instances users have interdependencies between archives such that the extraction of members from one archive is resolved by extracting members from another archive. If these dependencies are cyclic, the archives must be specified repeatedly on the command line to satisfy previous references. For example:

```
$ cc -o prog .... -lA -lB -lC -lA -lB -lC -lA
```

The determination, and maintenance, of repeated archive specifications can be tedious. The `-z rescan` option makes this process simpler. Following all input file processing, this option causes the entire archive list to be reprocessed in an attempt to locate additional archive members that resolve symbol references. This archive rescanning continues until a pass over the archive list occurs in which no new members are extracted. The previous example could therefore be simplified to:

```
$ cc -o prog -z rescan .... -lA -lB -lC
```

Directories Searched by the Link-Editor

All previous examples assume the link-editor knows where to search for the libraries listed on the command line. By default, when linking 32-bit objects, the link-editor knows of only two standard directories in which to look for libraries, `/usr/ccs/lib` and `/usr/lib`. When linking 64-bit objects, only one standard directory is used, `/usr/lib/64`. All other directories to be searched must be added to the link-editor's search path explicitly.

You can change the link-editor search path in two ways: using a command-line option, or using an environment variable.

Using a Command-Line Option

You can use the `-L` option to add a new path name to the library search path. This option affects the search path at the point it is encountered on the command line. For example, the following command searches `path1`, then `/usr/ccs/lib` and `/usr/lib`, to find `libfoo`. It searches `path1` and then `path2`, and then `/usr/ccs/lib` and `/usr/lib`, to find `libbar`.

```
$ cc -o prog main.o -Lpath1 file1.c -lfoo file2.c -Lpath2 -lbar
```

Path names defined using the `-L` option are used only by the link-editor. These path names are not recorded in the output file image created for use by the runtime linker.

Note – You must specify `-L` if you want the link-editor to search for libraries in your current directory. You can use a period (`.`) to represent the current directory.

You can use the `-Y` option to change the default directories searched by the link-editor. The argument supplied with this option takes the form of a colon separated list of directories. For example, the following command searches for `libfoo` only in the directories `/opt/COMPILER/lib` and `/home/me/lib`.

```
$ cc -o prog main.c -YP,/opt/COMPILER/lib:/home/me/lib -lfoo
```

The directories specified using the `-Y` option can be supplemented by using the `-L` option.

Using an Environment Variable

You can also use the environment variable `LD_LIBRARY_PATH`, which takes a colon-separated list of directories, to add to the link-editor's library search path. In its most general form, `LD_LIBRARY_PATH` takes two directory lists separated by a semicolon. The first list is searched before the lists supplied on the command line, and the second list is searched after.

The following example shows the combined effect of setting `LD_LIBRARY_PATH` and calling the link-editor with several `-L` occurrences:

```
$ LD_LIBRARY_PATH=dir1:dir2:dir3
$ export LD_LIBRARY_PATH
$ cc -o prog main.c -Lpath1 ... -Lpath2 ... -Lpathn -lfoo
```

The effective search path is `dir1:dir2:path1:path2...pathn:dir3:/usr/ccs/lib:/usr/lib`.

If no semicolon is specified as part of the `LD_LIBRARY_PATH` definition, the specified directory list is interpreted *after* any `-L` options. In the following example, the effective search path is `path1:path2...pathn:dir1:dir2:/usr/ccs/lib:/usr/lib`.

```
$ LD_LIBRARY_PATH=dir1:dir2
$ export LD_LIBRARY_PATH
$ cc -o prog main.c -Lpath1 ... -Lpath2 ... -Lpathn -lfoo
```

Note – This environment variable can also be used to augment the search path of the runtime linker. See “Directories Searched by the Runtime Linker” on page 64. To prevent this environment variable from influencing the link-editor, use the `-i` option.

Directories Searched by the Runtime Linker

The runtime linker only looks in one default location for dependencies. This location is `/usr/lib` when processing 32-bit objects, and `/usr/lib/64` when processing 64-bit objects. All other directories to be searched must be added to the runtime linker's search path explicitly.

When a dynamic executable or shared object is linked with additional shared objects, these shared objects are recorded as dependencies. These dependencies must be located during process execution by the runtime linker. During the link-edit, one or more search paths can be recorded in the output file. These search paths are used by the runtime linker to locate any dependencies. These recorded search paths are referred to as a *runpath*.

Specialized objects may be built with the `-z nodefaultlib` option to suppress any search of the default location at runtime. Use of this option implies that all the dependencies of an object can be located using its runpaths. Without this option, no matter how you augment the runtime linker's search path, its last element is always the default location. `/usr/lib` for 32-bit objects and `/usr/lib/64` for 64-bit objects.

Note – The default search path can be administrated using a runtime configuration file. See “Configuring the Default Search Paths” on page 67. However, the creator of an object should not rely on the existence of this file. You should always ensure that an object can locate its dependencies with only its runpaths or the default location.

You can use the `-R` option, which takes a colon-separated list of directories, to record a runpath in a dynamic executable or shared object. The following example records the runpath `/home/me/lib:/home/you/lib` in the dynamic executable `prog`.

```
$ cc -o prog main.c -R/home/me/lib:/home/you/lib -Lpath1 \  
-Lpath2 file1.c file2.c -lfoo -lbar
```

The runtime linker uses these paths, followed by the default location, to obtain any shared object dependencies. In this case, this runpath is used to locate `libfoo.so.1` and `libbar.so.1`.

The link-editor accepts multiple `-R` options. These multiple specifications are concatenate together, separated by a colon. Thus, the previous example can also be expressed as follows.

```
$ cc -o prog main.c -R/home/me/lib -Lpath1 -R/home/you/lib \  
-Lpath2 file1.c file2.c -lfoo -lbar
```

For objects that may be installed in various locations, the `$ORIGIN` dynamic string token provides a flexible means of recording a runpath. See “Locating Associated Dependencies” on page 323.

Note – A historic alternative to specifying the `-R` option is to set the environment variable `LD_RUN_PATH`, and make this available to the link-editor. The scope and function of `LD_RUN_PATH` and `-R` are identical, but when both are specified, `-R` supersedes `LD_RUN_PATH`.

Initialization and Termination Sections

Dynamic objects may supply code that provides for runtime initialization and termination processing. This code can be encapsulated in one of two section types, either an array of function pointers or a single code block. Each of these section types is built from a concatenation of like sections from the input relocatable objects.

The sections `.preinit_array`, `.init_array` and `.fini_array` provide arrays of runtime pre-initialization, initialization, and termination functions, respectively. When creating a dynamic object, the link-editor identifies these arrays with the `.dynamic` tag pairs `DT_PREINIT_ARRAY/ARRAYSZ`, `DT_INIT_ARRAY/ARRAYSZ`, and `DT_FINI_ARRAY/ARRAYSZ` accordingly. These tags identify the associated sections so they may be called by the runtime linker. A pre-initialization array is applicable to dynamic executables only.

The sections `.init` and `.fini` provide a runtime initialization and termination code block, respectively. However, the compiler drivers typically supply `.init` and `.fini` sections with files they add to the beginning and end of your input file list. These files have the effect of encapsulating the `.init` and `.fini` code into individual functions. These functions are identified by the reserved symbol names `_init` and `_fini` respectively. When creating a dynamic object, the link-editor identifies these symbols with the `.dynamic` tags `DT_INIT` and `DT_FINI` accordingly. These tags identify the associated sections so they may be called by the runtime linker.

For more information regarding the execution of initialization and termination code at runtime see “Initialization and Termination Routines” on page 76.

The registration of initialization and termination functions can be carried out directly by the link-editor using the `-z initarray` and `-z finiarray` options. For example, the following command places the address of `foo()` in an `.initarray` element, and the address of `bar()` in a `.finiarray` element.

```
$ cat main.c
#include <stdio.h>

void foo()
{
    (void) printf("initializing: foo()\n");
}

void bar()
```

```

{
    (void) printf("finalizing: bar()\n");
}

main()
{
    (void) printf("main()\n");
    return (0);
}

$ cc -o main -zinitarray=foo -zfiniarray=bar main.c
$ main
initializing: foo()
main()
finalizing: bar()

```

The creation of initialization and termination sections can be carried out directly using an assembler. However, most compilers offer special primitives to simplify their declaration. For example, the previous code example can be rewritten using the following `#pragma` definitions. These definitions result in a call to `foo()` being placed in an `.init` section, and a call to `bar()` being placed in a `.fini` section.

```

$ cat main.c
#include <stdio.h>

#pragma init (foo)
#pragma fini (bar)

.....
$ cc -o main main.c
$ main
initializing: foo()
main()
finalizing: bar()

```

Initialization and termination code, spread throughout several relocatable objects, can result in different behavior when included in an archive library or shared object. The link-edit of an application using this archive might extract only a fraction of the objects contained in the archive. These objects might provide only a portion of the initialization and termination code spread throughout the members of the archive. At runtime, only this portion of code is executed. The same application built against the shared object will have all the accumulated initialization and termination code executed when the dependency is loaded at runtime.

To determine the order of executing initialization and termination code within a process at runtime is a complex issue involving dependency analysis. Limiting the content of initialization and termination code can simplify this analysis, while providing both flexible, and predictable runtime behavior. See “Initialization and Termination Order” on page 77 for more details.

Data initialization should be independent if the initialization code is involved with a dynamic object whose memory can be dumped using `dldump(3DL)`.

Symbol Processing

During input file processing, all *local* symbols from the input relocatable objects are passed through to the output file image. All global symbols are accumulated internally within the link-editor. Each *global* symbol supplied by a relocatable object is searched for within this internal symbol table. If a symbol with the same name has already been encountered from a previous input file, a symbol resolution process is called. This symbol resolution process determines which of the two entries is kept.

On completing input file processing, and providing no fatal error conditions have been encountered during symbol resolution, the link-editor determines if any unresolved symbol references remain. Unresolved symbol references can cause the link-edit to terminate.

Finally, the link-editor's internal symbol table is added to the symbol tables of the image being created.

The following sections expand upon symbol resolution and undefined symbol processing.

Symbol Resolution

Symbol resolution runs the entire spectrum, from simple and intuitive to complex and perplexing. Resolutions can be carried out silently by the link-editor, can be accompanied by warning diagnostics, or can result in a fatal error condition.

The resolution of two symbols depends on their attributes, the type of file providing the symbol, and the type of file being generated. For a complete description of symbol attributes, see "Symbol Table Section" on page 222. For the following discussions, however, it is worth identifying three basic symbol types:

- *Undefined* – Symbols that have been referenced in a file but have not been assigned a storage address.
- *Tentative* – Symbols that have been created within a file but have not yet been sized or allocated in storage. They appear as uninitialized C symbols, or FORTRAN COMMON blocks within the file.
- *Defined* – Symbols that have been created and assigned storage addresses and space within the file.

In its simplest form, symbol resolution involves the use of a precedence relationship that has *defined* symbols dominating *tentative* symbols, which in turn dominate *undefined* symbols.

The following C code example shows how these symbol types can be generated. Undefined symbols are prefixed with `u_`, tentative symbols are prefixed with `t_`, and defined symbols are prefixed with `d_`.

```
$ cat main.c
extern int      u_bar;
extern int      u_foo();

int             t_bar;
int             d_bar = 1;

d_foo()
{
    return (u_foo(u_bar, t_bar, d_bar));
}
$ cc -o main.o -c main.c
$ nm -x main.o
```

[Index]	Value	Size	Type	Bind	Other	Shndx	Name
[8]	0x00000000	0x00000000	NOTY	GLOB	0x0	UNDEF	u_foo
[9]	0x00000000	0x00000040	FUNC	GLOB	0x0	2	d_foo
[10]	0x00000004	0x00000004	OBJT	GLOB	0x0	COMMON	t_bar
[11]	0x00000000	0x00000000	NOTY	GLOB	0x0	UNDEF	u_bar
[12]	0x00000000	0x00000004	OBJT	GLOB	0x0	3	d_bar

Simple Resolutions

Simple symbol resolutions are by far the most common, and result when two symbols with similar characteristics are detected and one symbol takes precedence over the other. This symbol resolution is carried out silently by the link-editor. For example, for symbols with the same binding, a reference to an undefined symbol from one file is bound to, or satisfied by, a defined or tentative symbol definition from another file. Or, a tentative symbol definition from one file is bound to a defined symbol definition from another file.

Symbols that undergo resolution can have either a global or weak binding. Weak bindings have lower precedence than global binding, so symbols with different bindings are resolved according to a slight alteration of the basic rules.

Weak symbols can usually be defined via the compiler, either individually or as aliases to global symbols. One mechanism uses a `#pragma` definition:

```
$ cat main.c
#pragma weak   bar
#pragma weak   foo = _foo

int            bar = 1;

_foo()
{
    return (bar);
}
```

```

}
$ cc -o main.o -c main.c
$ nm -x main.o
[Index]  Value          Size          Type Bind  Other Shndx  Name
.....
[7]      |0x00000000|0x00000004|OBJT |WEAK |0x0  |3      |bar
[8]      |0x00000000|0x00000028|FUNC |WEAK |0x0  |2      |foo
[9]      |0x00000000|0x00000028|FUNC |GLOB |0x0  |2      |_foo

```

Notice that the weak alias `foo` is assigned the same attributes as the global symbol `_foo`. This relationship is maintained by the link-editor and results in the symbols being assigned the same value in the output image. In symbol resolution, weak defined symbols are silently overridden by any global definition of the same name.

Another form of simple symbol resolution, interposition, occurs between relocatable objects and shared objects, or between multiple shared objects. In these cases, when a symbol is multiply-defined, the relocatable object, or the first definition between multiple shared objects, is silently taken by the link-editor. The relocatable object's definition, or the first shared object's definition, is said to *interpose* on all other definitions. This interposition can be used to override the functionality provided by one shared object, by a dynamic executable, or by another shared object.

The combination of weak symbols and interposition provides a useful programming technique. For example, the standard C library provides several services that you are allowed to redefine. However, ANSI C defines a set of standard services that must be present on the system and cannot be replaced in a strictly conforming program.

The function `fread(3C)`, for example, is an ANSI C library function, whereas the system function `read(2)` is not. A conforming ANSI C program must be able to redefine `read(2)` and still use `fread(3C)` in a predictable way.

The problem here is that `read(2)` underlies the `fread(3C)` implementation in the standard C library. Therefore, a program that redefines `read(2)` might confuse the `fread(3C)` implementation. To guard against this occurrence, ANSI C states that an implementation cannot use a name that is not reserved for it. Using the following `#pragma` directive you can define just such a reserved name, and from it generate an alias for the function `read(2)`.

```
#pragma weak read = _read
```

Thus, you can quite freely define your own `read()` function without compromising the `fread(3C)` implementation, which in turn is implemented to use the `_read()` function.

The link-editor will not have difficulty with your redefinition of `read()`, either when linking against the shared object or archive version of the standard C library. In the former case, interposition takes its course. In the latter case, the fact that the C library's definition of `read(2)` is weak allows that definition to be quietly overridden.

You can use the link-editor's `-m` option to write a list of all interposed symbol references, along with section load address information, to the standard output.

Complex Resolutions

Complex resolutions occur when two symbols of the same name are found with differing attributes. In these cases, the link-editor selects the most appropriate symbol and generates a warning message indicating the symbol, the attributes that conflict, and the identity of the file from which the symbol definition is taken. In the following example two files with a definition of the data item `array` have different size requirements.

```
$ cat foo.c
int array[1];

$ cat bar.c
int array[2] = { 1, 2 };

$ cc -dn -r -o temp.o foo.c bar.c
ld: warning: symbol `array' has differing sizes:
      (file foo.o value=0x4; file bar.o value=0x8);
      bar.o definition taken
```

A similar diagnostic is produced if the symbol's alignment requirements differ. In both of these cases, the diagnostic can be suppressed by using the link-editor's `-t` option.

Another form of attribute difference is the symbol's type. In the following example the symbol `bar()` has been defined as both a data item and a function.

```
$ cat foo.c
bar()
{
    return (0);
}
$ cc -o libfoo.so -G -K pic foo.c
$ cat main.c
int    bar = 1;

main()
{
    return (bar);
}
$ cc -o main main.c -L. -lfoo
ld: warning: symbol `bar' has differing types:
      (file main.o type=OBJT; file ./libfoo.so type=FUNC);
      main.o definition taken
```

Note – Symbol types in this context are classifications that can be expressed in ELF. They are not related to the data types as employed by the programming language, except in the crudest fashion.

In cases like the previous example, the relocatable object definition is taken when the resolution occurs between a relocatable object and a shared object, or the first definition is taken when the resolution occurs between two shared objects. When such resolutions occur between symbols of different bindings (weak or global), a warning is also produced.

Inconsistencies between symbol types are not suppressed by the link-editor's `-t` option.

Fatal Resolutions

Symbol conflicts that cannot be resolved result in a fatal error condition. In this case, an appropriate error message is provided indicating the symbol name together with the names of the files that provided the symbols, and no output file is generated. Although the fatal condition is sufficient to terminate the link-edit, all input file processing is first completed. In this manner, all fatal resolution errors can be identified.

The most common fatal error condition exists when two relocatable objects both define symbols of the same name, and neither symbol is a weak definition:

```
$ cat foo.c
int bar = 1;

$ cat bar.c
bar()
{
    return (0);
}

$ cc -dn -r -o temp.o foo.c bar.c
ld: fatal: symbol 'bar' is multiply-defined:
      (file foo.o and file bar.o);
ld: fatal: File processing errors. No output written to int.o
```

`foo.c` and `bar.c` have conflicting definitions for the symbol `bar`. Because the link-editor cannot determine which should dominate, the link-edit usually terminates with an error message. You can use the link-editor's `-z muldefs` option to suppress this error condition, and allow the first symbol definition to be taken.

Undefined Symbols

After all of the input files have been read and all symbol resolution is complete, the link-editor searches the internal symbol table for any symbol references that have not been bound to symbol definitions. These symbol references are referred to as *undefined* symbols. The effect of these undefined symbols on the link-edit process can vary according to the type of output file being generated, and possibly the type of symbol.

Generating an Executable Output File

When the link-editor is generating an executable output file, the link-editor's default behavior is to terminate with an appropriate error message should any symbols remain undefined. A symbol remains undefined when a symbol reference in a relocatable object is never matched to a symbol definition:

```
$ cat main.c
extern int foo();

main()
{
    return (foo());
}
$ cc -o prog main.c
Undefined      first referenced
 symbol              in file
foo                main.o
ld: fatal: Symbol referencing errors. No output written to prog
```

In a similar manner, a symbol reference within a shared object that is never matched to a symbol definition when the shared object is being used to create a dynamic executable will also result in an undefined symbol:

```
$ cat foo.c
extern int bar;
foo()
{
    return (bar);
}

$ cc -o libfoo.so -G -K pic foo.c
$ cc -o prog main.c -L. -lfoo
Undefined      first referenced
 symbol              in file
bar                ./libfoo.so
ld: fatal: Symbol referencing errors. No output written to prog
```

If you want to allow undefined symbols, as in cases like the previous example, then the default fatal error condition can be suppressed by using the link-editor's `-z nodefs` option.

Note – Take care when using the `-z nodefs` option. If an unavailable symbol reference is required during the execution of a process, a fatal runtime relocation error occurs. It may be possible to detect this error during the initial execution and testing of an application. However, more complex execution paths can result in this error condition taking much longer to detect, which can be time consuming and costly.

Symbols can also remain undefined when a symbol reference in a relocatable object is bound to a symbol definition in an implicitly defined shared object. For example, continuing with the files `main.c` and `foo.c` used in the previous example:

```
$ cat bar.c
int bar = 1;

$ cc -o libbar.so -R. -G -K pic bar.c -L. -lfoo
$ ldd libbar.so
        libfoo.so =>      ./libfoo.so

$ cc -o prog main.c -L. -lbar
Undefined          first referenced
 symbol            in file
foo                main.o (symbol belongs to implicit \
                  dependency ./libfoo.so)
ld: fatal: Symbol referencing errors. No output written to prog
```

`prog` is built with an *explicit* reference to `libbar.so`. `libbar.so` has a dependency on `libfoo.so`, and therefore an implicit reference to `libfoo.so` from `prog` is established.

Because `main.c` made a specific reference to the interface provided by `libfoo.so`, `prog` really has a dependency on `libfoo.so`. However, only explicit shared object dependencies are recorded in the output file being generated. Thus, `prog` fails to run if a new version of `libbar.so` is developed that no longer has a dependency on `libfoo.so`.

For this reason, bindings of this type are deemed fatal, and the implicit reference must be made explicit by referencing the library directly during the link-edit of `prog`. The required reference is hinted at in the fatal error message shown in the preceding example.

Generating a Shared Object Output File

When the link-editor is generating a shared object output file, it allows undefined symbols to remain at the end of the link-edit. This default behavior allows the shared object to import symbols from either relocatable objects or from other shared objects when the object is used to create a dynamic executable.

The link-editor's `-z defs` option can be used to force a fatal error if any undefined symbols remain. This option is recommended when creating any shared objects. Shared objects that reference symbols from an application can use the `-z defs` option and define the applications symbols using the `extern mapfile` directive, as described in "Defining Additional Symbols" on page 46.

A self-contained shared object, in which all references to external symbols are satisfied by named dependencies, provides maximum flexibility. The shared object can be employed by many users without those users having to determine and establish dependencies to satisfy the shared object's requirements.

Weak Symbols

Weak symbol references that are not bound during a link-edit do not result in a fatal error condition, no matter what output file type is being generated.

If a static executable is being generated, the symbol is converted to an absolute symbol and assigned a value of zero.

If a dynamic executable or shared object is being produced, the symbol is left as an undefined weak reference and assigned the value zero. During process execution, the runtime linker searches for this symbol. If the runtime linker does not find a match, it binds the reference to an address of zero instead of generating a fatal runtime relocation error.

Historically, these undefined weak referenced symbols have been employed as a mechanism to test for the existence of functionality. For example, the following C code fragment might have been used in the shared object `libfoo.so.1`:

```
#pragma weak    foo

extern void    foo(char *);

void bar(char * path)
{
    void (* fptr)(char *);

    if ((fptr = foo) != 0)
        (* fptr)(path);
}
```

When an application is built that references `libfoo.so.1`, the link-edit will complete successfully regardless of whether a definition for the symbol `foo` is found. If during execution of the application the function address tests nonzero, the function is called. However, if the symbol definition is not found, the function address tests zero and so it is not called.

Compilation systems view this address comparison technique as having undefined semantics, which can result in the test statement being removed under optimization. In addition, the runtime symbol binding mechanism places other restrictions on the use of this technique, which prevents a consistent model from being available for all dynamic objects.

Note – Undefined weak references in this manner are discouraged. Instead, you should use `dlsym(3DL)` with the `RTLD_DEFAULT` flag as a means of testing for a symbol's existence. See "Testing for Functionality" on page 91.

Tentative Symbol Order Within the Output File

Contributions from input files usually appear in the output file in the order of their contribution. An exception occurs when processing tentative symbols and their associated storage. These symbols are not fully defined until their resolution is complete. If the resolution occurs as a result of encountering a *defined* symbol from a relocatable object, then the order of appearance is that which would have occurred for the definition.

If you need to control the ordering of a group of symbols, then any tentative definition should be redefined to a zero-initialized data item. For example, the following tentative definitions result in a reordering of the data items within the output file, compared to the original order described in the source file `foo.c`:

```
$ cat foo.c
char A_array[0x10];
char B_array[0x20];
char C_array[0x30];

$ cc -o prog main.c foo.c
$ nm -vx prog | grep array
[32] | 0x00020754|0x00000010|OBJT |GLOB |0x0 |15 |A_array
[34] | 0x00020764|0x00000030|OBJT |GLOB |0x0 |15 |C_array
[42] | 0x00020794|0x00000020|OBJT |GLOB |0x0 |15 |B_array
```

By defining these symbols as initialized data items, the relative ordering of these symbols within the input file is carried over to the output file:

```
$ cat foo.c
char A_array[0x10] = { 0 };
char B_array[0x20] = { 0 };
char C_array[0x30] = { 0 };

$ cc -o prog main.c foo.c
$ nm -vx prog | grep array
[32] | 0x000206bc|0x00000010|OBJT |GLOB |0x0 |12 |A_array
[42] | 0x000206cc|0x00000020|OBJT |GLOB |0x0 |12 |B_array
[34] | 0x000206ec|0x00000030|OBJT |GLOB |0x0 |12 |C_array
```

Defining Additional Symbols

Besides the symbols provided from input files, you can supply additional symbol references or definitions to a link-edit. In the simplest form, symbol references can be generated using the link-editor's `-u` option. Greater flexibility is provided with the link-editor's `-M` option and an associated `mapfile` that enables you to define symbol references and a variety of symbol definitions.

The `-u` option provides a mechanism for generating a symbol reference from the link-edit command line. This option can be used to perform a link-edit entirely from archives, or to provide additional flexibility in selecting the objects to extract from multiple archives. See section “Archive Processing” on page 29 for an overview of archive extraction.

For example, perhaps you want to generate a dynamic executable from the relocatable object `main.o`, which refers to the symbols `foo` and `bar`. You want to obtain the symbol definition `foo` from the relocatable object `foo.o` contained in `lib1.a`, and the symbol definition `bar` from the relocatable object `bar.o`, contained in `lib2.a`.

However, the archive `lib1.a` also contains a relocatable object defining the symbol `bar`. This relocatable object is presumably of differing functionality to the relocatable object provided in `lib2.a`. To specify the required archive extraction, you can use the following link-edit:

```
$ cc -o prog -L. -u foo -l1 main.o -l2
```

The `-u` option generates a reference to the symbol `foo`. This reference causes extraction of the relocatable object `foo.o` from the archive `lib1.a`. The first reference to the symbol `bar` occurs in `main.o`, which is encountered after `lib1.a` has been processed. Therefore, the relocatable object `bar.o` is obtained from the archive `lib2.a`.

Note – This simple example assumes that the relocatable object `foo.o` from `lib1.a` does not directly or indirectly reference the symbol `bar`. If it does then the relocatable object `bar.o` is also extracted from `lib1.a` during its processing. See “Archive Processing” on page 29 for a discussion of the link-editor’s multi-pass processing of an archive.

A more extensive set of symbol definitions can be provided using the link-editor’s `-M` option and an associated `mapfile`. The syntax for these `mapfile` entries is:

```
[ name ] {  
    scope:  
        symbol [ = [ type ] [ value ] [ size ] [ attribute ] ];  
} [ dependency ];
```

name

A label for this set of symbol definitions, if present, identifies a *version definition* within the image. See Chapter 5.

scope

Indicates the visibility of the symbols’ binding within the output file being generated. All symbols defined with a `mapfile` are treated as global in scope during the link-edit process. That is, they are resolved against any other symbols of the same name obtained from any of the input files. The following definitions, and aliases, define a symbols’ visibility in the object being created:

default / global

Symbols of this scope remain visible to other external objects. References to such symbols from within the object are bound at runtime, thus allowing interposition to take place.

protected / symbolic

Symbols of this scope remain visible to other external objects. References to these symbols from within the object are bound at link-edit, thus preventing runtime interposition. This scope definition has the same affect as a symbol with STV_PROTECTED visibility. See Table 7-24.

hidden / local

Symbols of this scope are reduced to symbols with a local binding. Symbols of this scope are not visible to other external objects. This scope definition has the same affect as a symbol with STV_HIDDEN visibility. See Table 7-24.

eliminate

Symbols of this scope are hidden. Their symbol table entries are eliminated.

symbol

The name of the symbol required. If the name is not followed by one of the symbol attributes, *type*, *value*, *size* or *extern*, a symbol reference is created. This reference is exactly the same as would be generated using the -u option discussed earlier in this section. If the symbol name is followed by any symbol attributes, then a symbol definition is generated using the associated attributes.

When in *local* scope, this symbol name can be defined as the special *auto-reduction* directive ****. This directive results in all global symbols, not explicitly defined to be *global* in the *mapfile*, receiving a local binding within any dynamic object file being generated.

type

Indicates the symbol type attribute. This attribute can be either *data*, *function*, or *COMMON*. The former two type attributes result in an absolute symbol definition. See "Symbol Table Section" on page 222. The latter type attribute results in a tentative symbol definition.

value

Indicates the value attribute and takes the form of *Vnumber*.

size

Indicates the size attribute and takes the form of *Snumber*.

attribute

This keyword provides additional attributes for the symbol:

EXTERN

Indicates the symbol is defined externally to the object being created. This attribute can be used with the *DIRECT* or *NODIRECT* attributes to establish individual direct, or no-direct references. Undefined symbols that would be flagged with the -z *defs* option can also be suppressed with this option.

DIRECT

Indicates that this symbol should be directly bound to. This attribute can be used with the `EXTERN` attribute to control binding to an external symbol. See “Direct Binding” on page 69.

NODIRECT

Indicates that this symbol should not be directly bound to. This state applies to references from within the object being created and from external references. This attribute can be used with the `EXTERN` attribute to control binding to an external symbol. See “Direct Binding” on page 69.

dependency

Represents a *version definition* that is inherited by this definition. See Chapter 5.

If either a version definition or the auto-reduction directive is specified, then versioning information is recorded in the image created. If this image is an executable or shared object, then any symbol reduction is also applied.

If the image being created is a relocatable object, then by default, no symbol reduction is applied. In this case, any symbol reductions are recorded as part of the versioning information. These reductions are applied when the relocatable object is finally used to generate an executable or shared object. The link-editor’s `-B reduce` option can be used to force symbol reduction when generating a relocatable object.

A more detailed description of the versioning information is provided in Chapter 5.

Note – To ensure interface definition stability, no wildcard expansion is provided for defining symbol names.

This section presents several examples of using the `mapfile` syntax.

The following example shows how three symbol references can be defined. These references are then used to extract members of an archive. Although this archive extraction can be achieved by specifying multiple `-u` options to the `link-edit`, this example also shows how the eventual scope of a symbol can be reduced to *local*.

```
$ cat foo.c
foo()
{
    (void) printf("foo: called from lib.a\n");
}
$ cat bar.c
bar()
{
    (void) printf("bar: called from lib.a\n");
}
$ cat main.c
extern void    foo(), bar();

main()
```

```

{
    foo();
    bar();
}
$ ar -rc lib.a foo.o bar.o main.o
$ cat mapfile
{
    local:
        foo;
        bar;
    global:
        main;
};
$ cc -o prog -M mapfile lib.a
$ prog
foo: called from lib.a
bar: called from lib.a
$ nm -x prog | egrep "main$|foo$|bar$"
[28] |0x00010604|0x00000024|FUNC |LOCL |0x0 |7 |foo
[30] |0x00010628|0x00000024|FUNC |LOCL |0x0 |7 |bar
[49] |0x0001064c|0x00000024|FUNC |GLOB |0x0 |7 |main

```

The significance of reducing symbol scope from global to local is covered in more detail in the section “Reducing Symbol Scope” on page 52.

The following example shows how two absolute symbol definitions can be defined. These definitions are then used to resolve the references from the input file `main.c`.

```

$ cat main.c
extern int    foo();
extern int    bar;

main()
{
    (void) printf("&foo = %x\n", &foo);
    (void) printf("&bar = %x\n", &bar);
}
$ cat mapfile
{
    global:
        foo = FUNCTION V0x400;
        bar = DATA V0x800;
};
$ cc -o prog -M mapfile main.c
$ prog
&foo = 400 &bar = 800
$ nm -x prog | egrep "foo$|bar$"
[37] |0x00000800|0x00000000|OBJT |GLOB |0x0 |ABS |bar
[42] |0x00000400|0x00000000|FUNC |GLOB |0x0 |ABS |foo

```

When obtained from an input file, symbol definitions for functions or data items are usually associated with elements of data storage. A `mapfile` definition is insufficient to be able to construct this data storage, so these symbols must remain as absolute values.

However, a `mapfile` can also be used to define a `COMMON`, or tentative, symbol. Unlike other types of symbol definition, tentative symbols do not occupy storage within a file, but define storage that must be allocated at runtime. Therefore, symbol definitions of this kind can contribute to the storage allocation of the output file being generated.

A feature of tentative symbols that differs from other symbol types is that their *value* attribute indicates their alignment requirement. A `mapfile` definition can therefore be used to realign tentative definitions obtained from the input files of a link-edit.

The following example shows the definition of two tentative symbols. The symbol `foo` defines a new storage region whereas the symbol `bar` is actually used to change the alignment of the same tentative definition within the file `main.c`.

```
$ cat main.c
extern int    foo;
int          bar[0x10];

main()
{
    (void) printf("&foo = %x\n", &foo);
    (void) printf("&bar = %x\n", &bar);
}
$ cat mapfile
{
    global:
        foo = COMMON V0x4 S0x200;
        bar = COMMON V0x100 S0x40;
};
$ cc -o prog -M mapfile main.c
ld: warning: symbol 'bar' has differing alignments:
      (file mapfile value=0x100; file main.o value=0x4);
      largest value applied
$ prog
&foo = 20940
&bar = 20900
$ nm -x prog | egrep "foo$|bar$"
[37]  |0x00020900|0x00000040|OBJT |GLOB |0x0 |16 |bar
[42]  |0x00020940|0x00000200|OBJT |GLOB |0x0 |16 |foo
```

Note – This symbol resolution diagnostic can be suppressed by using the link-editor's `-t` option.

Reducing Symbol Scope

Symbol definitions defined to have local scope within a mapfile can be used to reduce the symbol's eventual binding. This mechanism can play an important role in reducing the symbol's visibility to future link-edits that use the generated file as part of their input. In fact, this mechanism can provide for the precise definition of a file's interface, and so restrict the functionality made available to others.

For example, say you want to generate a simple shared object from the files `foo.c` and `bar.c`. The file `foo.c` contains the global symbol `foo`, which provides the service that you want to make available to others. The file `bar.c` contains the symbols `bar` and `str`, which provide the underlying implementation of the shared object. The creation of a simple shared object usually results in all three of these symbols having global scope.

```
$ cat foo.c
extern const char * bar();

const char * foo()
{
    return (bar());
}
$ cat bar.c
const char * str = "returned from bar.c";

const char * bar()
{
    return (str);
}
$ cc -o lib.so.1 -G foo.c bar.c
$ nm -x lib.so.1 | egrep "foo$|bar$|str$"
[29] |0x000104d0|0x00000004|OBJT |GLOB |0x0 |12 |str
[32] |0x00000418|0x00000028|FUNC |GLOB |0x0 |6 |bar
[33] |0x000003f0|0x00000028|FUNC |GLOB |0x0 |6 |foo
```

You can now use the functionality offered by this shared object as part of the link-edit of another application. References to the symbol `foo` are bound to the implementation provided by the shared object.

Because of their global binding, direct reference to the symbols `bar` and `str` is also possible. This can have dangerous consequences, as you might later change the implementation underlying the function `foo`. In so doing, you could unintentionally cause an existing application that had bound to `bar` or `str` to fail or misbehave.

Another consequence of the global binding of the symbols `bar` and `str` is that they can be interposed upon by symbols of the same name. The interposition of symbols within shared objects is covered in section "Simple Resolutions" on page 39. This interposition can be intentional and be used as a means of circumventing the intended functionality offered by the shared object. On the other hand, this interposition can be unintentional, the result of the same common symbol name used for both the application and the shared object.

When developing the shared object, you can protect against this scenario by reducing the scope of the symbols `bar` and `str` to a local binding. In the following example, the symbols `bar` and `str` are no longer available as part of the shared objects interface. Thus, these symbols cannot be referenced, or interposed upon, by an external object. You have effectively defined an interface for the shared object. This interface can be managed while hiding the details of the underlying implementation.

```
$ cat mapfile
{
    local:
        bar;
        str;
};
$ cc -o lib.so.1 -M mapfile -G foo.c bar.c
$ nm -x lib.so.1 | egrep "foo$|bar$|str$"
[27] | 0x000003dc | 0x00000028 | FUNC | LOCL | 0x0 | 6 | bar
[28] | 0x00010494 | 0x00000004 | OBJT | LOCL | 0x0 | 12 | str
[33] | 0x000003b4 | 0x00000028 | FUNC | GLOB | 0x0 | 6 | foo
```

This symbol scope reduction has an additional performance advantage. The symbolic relocations against the symbols `bar` and `str` that would have been necessary at runtime are now reduced to relative relocations. This reduces the runtime overhead of initializing and processing the shared object. See “When Relocations are Performed” on page 121 for details of symbolic relocation overhead.

As the number of symbols processed during a link-edit increases, the ability to define each local scope reduction within a `mapfile` becomes harder to maintain. An alternative and more flexible mechanism enables you to define the shared objects interface in terms of the global symbols that should be maintained, and instructs the link-editor to reduce all other symbols to local binding. This mechanism is achieved using the special *auto-reduction* directive `“*”`. For example, the previous `mapfile` definition can be rewritten to define `foo` as the only global symbol required in the output file generated:

```
$ cat mapfile
lib.so.1.1
{
    global:
        foo;
    local:
        *;
};
$ cc -o lib.so.1 -M mapfile -G foo.c bar.c
$ nm -x lib.so.1 | egrep "foo$|bar$|str$"
[30] | 0x00000370 | 0x00000028 | FUNC | LOCL | 0x0 | 6 | bar
[31] | 0x00010428 | 0x00000004 | OBJT | LOCL | 0x0 | 12 | str
[35] | 0x00000348 | 0x00000028 | FUNC | GLOB | 0x0 | 6 | foo
```

This example also defines a version name, `lib.so.1.1`, as part of the `mapfile` directive. This version name establishes an internal version definition that defines the file's symbolic interface. The creation of a version definition is recommended. The definition forms the foundation of an internal versioning mechanism that can be used throughout the evolution of the file. See Chapter 5.

Note – If a version name is not supplied, the output file name is used to label the version definition. The versioning information created within the output file can be suppressed using the link-editor's `-z noversion` option.

Whenever a version name is specified, *all* global symbols must be assigned to a version definition. If any global symbols remain unassigned to a version definition, the link-editor generates a fatal error condition:

```
$ cat mapfile
lib.so.1.1 {
    global:
        foo;
};
$ cc -o lib.so.1 -M mapfile -G foo.c bar.c
Undefined          first referenced
 symbol            in file
str                 bar.o (symbol has no version assigned)
bar                 bar.o (symbol has no version assigned)
ld: fatal: Symbol referencing errors. No output written to lib.so.1
```

The `-B local` option can be used to assert the *auto-reduction* directive `""` from the command line. The previous example could be compiled successfully with:

```
$ cc -o lib.so.1 -M mapfile -B local -G foo.c bar.c
```

When generating an executable or shared object, any symbol reduction results in the recording of version definitions within the output image, together with the reduction of the appropriate symbols. When generating a relocatable object, the version definitions are created but the symbol reductions are not processed. The result is that the symbol entries for any symbol reductions still remain global. For example, using the previous `mapfile` with the *auto-reduction* directive and associated relocatable objects, an intermediate relocatable object is created that shows no symbol reduction.

```
$ cat mapfile
lib.so.1.1 {
    global:
        foo;
    local:
        *;
};
$ ld -o lib.o -M mapfile -r foo.o bar.o
$ nm -x lib.o | egrep "foo$|bar$|str$"
[17] 0x00000000|0x00000004|OBJT |GLOB |0x0 |3 |str
[19] 0x00000028|0x00000028|FUNC |GLOB |0x0 |1 |bar
```

```
[20] |0x00000000|0x00000028|FUNC |GLOB |0x0 |1 |foo
```

The version definitions created within this image show that symbol reductions are required. When the relocatable object is used eventually to generate an executable or shared object, the symbol reductions occur. In other words, the link-editor reads and interprets symbol reduction information contained in relocatable objects in the same manner as it processes the data from a `mapfile`.

Thus, the intermediate relocatable object produced in the previous example can now be used to generate a shared object:

```
$ ld -o lib.so.1 -G lib.o
$ nm -x lib.so.1 | egrep "foo$|bar$|str$"
[22] |0x000104a4|0x00000004|OBJT |LOCL |0x0 |14 |str
[24] |0x000003dc|0x00000028|FUNC |LOCL |0x0 |8 |bar
[36] |0x000003b4|0x00000028|FUNC |GLOB |0x0 |8 |foo
```

Symbol reduction at the point at which an executable or shared object is created is typically the most common requirement. However, symbol reductions can be forced to occur when creating a relocatable object by using the link-editor's `-B reduce` option.

```
$ ld -o lib.o -M mapfile -B reduce -r foo.o bar.o
$ nm -x lib.o | egrep "foo$|bar$|str$"
[15] |0x00000000|0x00000004|OBJT |LOCL |0x0 |3 |str
[16] |0x00000028|0x00000028|FUNC |LOCL |0x0 |1 |bar
[20] |0x00000000|0x00000028|FUNC |GLOB |0x0 |1 |foo
```

Symbol Elimination

An extension to symbol reduction is the elimination of a symbol entry from an object's symbol table. Local symbols are only maintained in an object's `.symtab` symbol table. This entire table can be removed from the object using the link-editor's `-s` option, or `strip(1)`. On occasion, you might want to maintain the `.symtab` symbol table but remove selected local symbol definitions from it.

Symbol elimination can be carried out using the `mapfile` directive `eliminate`. As with the `local` directive, symbols can be individually defined, or the symbol name can be defined as the special *auto-elimination* directive `"*"`. The following example shows the elimination of the symbol `bar` for the previous symbol reduction example.

```
$ cat mapfile
lib.so.1.1
{
    global:
        foo;
    local:
        str;
    eliminate:
        *;
};
$ cc -o lib.so.1 -M mapfile -G foo.c bar.c
```

```
$ nm -x lib.so.1 | egrep "foo$|bar$|str$"
[31]      |0x00010428|0x00000004|OBJT |LOCL |0x0   |12      |str
[35]      |0x00000348|0x00000028|FUNC |GLOB |0x0   |6       |foo
```

The `-B eliminate` option can be used to assert the *auto-elimination* directive `**` from the command line.

External Bindings

When a symbol reference from the object being created is satisfied by a definition within a shared object, the symbol remains undefined. The relocation information associated with the symbol provides for its lookup at runtime. The shared object that provided the definition typically becomes a dependency.

The runtime linker employs a default search model to locate this definition at runtime. It typically searches each object, starting with the dynamic executable, and progressing through each dependency in the same order in which the objects were loaded.

Objects can also be created to use direct bindings. With this technique, the relationship between the symbol reference and the object that provides the symbol definition is maintained within the object being created. The runtime linker uses this information to directly bind the reference to the object that defines the symbol, thus bypassing the default symbol search model. See “Direct Binding” on page 69.

String Table Compression

String tables are compressed by the link-editor by removing duplicate entries and tail substrings. This compression can significantly reduce the size of any string tables. A compressed `.dynstr` table can produce a smaller text segment and hence reduce runtime paging activity. Because of these benefits, string table compression is enabled by default.

Linking objects that contribute a very large number of symbols may increase the link-edit time due to the string table compression. To avoid this cost during development use the link-editors `-z nocompstrtab` option. Any string table compression performed during a link-edit can be displayed using the link-editors debugging tokens `-D strtab,detail`.

Generating the Output File

After all input file processing and symbol resolution is completed with no fatal errors, the link-editor can start generating the output file. The link-editor establishes the additional sections that must be generated to complete the output file. These sections include the symbol tables that contain local symbol definitions from the input files, together with the global and weak symbol information that has been collected in the link-editor's internal symbol table.

Also included are any output relocation and dynamic information sections required by the runtime linker. After all the output section information has been established, the total output file size is calculated and the output file image is created accordingly.

When creating a dynamic executable or shared object, two symbol tables are usually generated. The `.dynsym` table and its associated string table `.dynstr` contain register (even if these are local), global, weak, and section symbols. These sections become part of the `text` segment that is mapped as part of the process image at runtime (see the `mmap(2)` man page). This enables the runtime linker to read these sections and perform any necessary relocations.

The `.symtab` table, and its associated string table `.strtab` contain all the symbols collected from the input file processing. These sections are not mapped as part of the process image. They can even be stripped from the image using the link-editor's `-s` option, or after the link-edit using `strip(1)`.

During the generation of the symbol tables, reserved symbols are created. These symbols have special meaning to the linking process and should not be defined in your code.

`__etext`
The first location after the text segment.

`__edata`
The first location after initialized data.

`__end`
The first location after all data.

`__DYNAMIC`
The address of the dynamic information section (the `.dynamic` section).

`__END__`
The same as `__end`. The symbol has local scope and, together with `__START__`, provides a means of establishing an object's address range.

`__GLOBAL_OFFSET_TABLE__`
The position-independent reference to a link-editor supplied table of addresses, the `.got` section. This table is constructed from position-independent data references

occurring in objects that have been compiled with the `-K pic` option. See “Position-Independent Code” on page 115.

`_PROCEDURE_LINKAGE_TABLE_`

The position-independent reference to a link-editor supplied table of addresses, the `.plt` section. This table is constructed from position-independent function references occurring in objects that have been compiled with the `-K pic` option. See “Position-Independent Code” on page 115.

`_START_`

The first location within the text segment. The symbol has local scope and, together with `_END_`, provides a means of establishing an object’s address range.

When generating an executable, the link-editor looks for additional symbols to define the executable’s entry point. If a symbol was specified using the link-editor’s `-e` option, that symbol is used. Otherwise the link-editor looks for the reserved symbol names `_start`, and then `main`. If none of these symbols exists, the first address of the text segment is used.

Relocation Processing

After you have created the output file, all data sections from the input files are copied to the new image. Any relocations specified by the input files are applied to the output image. Any additional relocation information that must be generated is also written to the new image.

Relocation processing is normally uneventful, although error conditions might arise that are accompanied by specific error messages. Two conditions are worth more discussion. The first condition involves text relocations that result from position-dependent code. This condition is covered in more detail in “Position-Independent Code” on page 115. The second condition can arise from displacement relocations, which is described more fully in the next section.

Displacement Relocations

Error conditions might occur if displacement relocations are applied to a data item, which itself can be used in a copy relocation. The details of copy relocations are covered in “Copy Relocations” on page 122.

A displacement relocation remains valid when both the relocated offset and the target to which it is relocated remain separated by the same displacement. A copy relocation is one where a global data item within a shared object is copied to the `.bss` of an

executable, to preserve the executable's read-only text segment. If the copied data has a displacement relocation applied to it, or an external relocation is a displacement into the copied data, the displacement relocation becomes invalidated.

The areas to address in trying to catch these sorts of errors are:

- When generating a shared object, flag any potential copy relocatable data items that can be problematic if they are involved in a displacement relocation. During construction of a shared object, the link-editor has no knowledge of what references might be made to it. Thus, all that can be flagged are *potential* problems.
- When generating an executable, flag the creation of a copy relocation whose data is involved in a displacement relocation.

However, displacement relocations applied to a shared object might be completed during its creation at link-edit time. Therefore, a link-edit of an application that references this shared object has no knowledge of a displacement being in effect in any copy-relocated data.

To help diagnose these problem areas, the link-editor indicates the displacement relocation use of a dynamic object with one or more dynamic `DT_FLAGS_1` flags, as shown in Table 7-44. In addition, the link-editor's `-z verbose` option can be used to display suspicious relocations.

For example, say you create a shared object with a global data item, `bar []`, which has a displacement relocation applied to it. This item could be copy-relocated if referenced from a dynamic executable. The link-editor warns of this condition with:

```
$ cc -G -o libfoo.so.1 -z verbose -K pic foo.o
ld: warning: relocation warning: R_SPARC_DISP32: file foo.o: symbol foo: \
displacement relocation to be applied to the symbol bar: at 0x194: \
displacement relocation will be visible in output image
```

If you now create an application that references the data item `bar []`, a copy relocation will be created which results in the displacement relocation being invalidated. Because the link-editor can explicitly discover this situation, an error message is generated regardless of the use of the `-z verbose` option.

```
$ cc -o prog prog.o -L. -lfoo
ld: warning: relocation error: R_SPARC_DISP32: file foo.so: symbol foo: \
displacement relocation applied to the symbol bar at: 0x194: \
the symbol bar is a copy relocated symbol
```

Note – `ldd(1)`, when used with either the `-d` or `-r` options, uses the displacement dynamic flags to generate similar relocation warnings.

These error conditions can be avoided by ensuring that the symbol definition being relocated (offset) and the symbol target of the relocation are both local. Use static definitions or the link-editor's scoping technology. See "Reducing Symbol Scope" on page 52. Relocation problems such as these can be avoided by accessing data within shared objects using functional interfaces.

Debugging Aids

A debugging library is provided with the Solaris linkers. This library enables you to trace the link-editing process in more detail. This library helps you understand, or debug, the link-edit of your own applications or libraries. Although the type of information displayed using this library is expected to remain constant, the exact format of the information might change slightly from release to release.

Some of the debugging output might be unfamiliar if you do not have an intimate knowledge of the ELF format. However, many aspects might be of general interest to you.

Debugging is enabled by using the `-D` option, and all output produced is directed to the standard error. This option must be augmented with one or more tokens to indicate the type of debugging required. The tokens available can be displayed by typing `-D help` at the command line.

```
$ ld -Dhelp
debug:
debug:          For debugging the link-editing of an application:
debug:          LD_OPTIONS=-Dtoken1,token2 cc -o prog ...
debug:          or,
debug:          ld -Dtoken1,token2 -o prog ...
debug:          where placement of -D on the command line is significant
debug:          and options can be switched off by prepending with '!'.
debug:
debug: args      display input argument processing
debug: basic     provide basic trace information/warnings
debug: detail    provide more information in conjunction with other options
debug: entry     display entrance criteria descriptors
debug: files     display input file processing (files and libraries)
debug: got       display GOT symbol information
debug: help      display this help message
debug: libs      display library search paths; detail flag shows actual
debug:           library lookup (-l) processing
debug: map       display map file processing
debug: move      display move section processing
debug: reloc     display relocation processing
debug: sections  display input section processing
debug: segments display available output segments and address/offset
```

```

debug:                processing; detail flag shows associated sections
debug: statistics    display processing statistics
debug: strtabs       display information about string table compression; detail
debug:                shows layout of string tables
debug: support       display support library processing
debug: symbols       display symbol table processing; detail flag shows
debug:                internal symbol table addition and resolution
debug: tls           display TLS processing info
debug: unused        display unused/unreferenced files; detail flag shows
debug:                unused sections
debug: versions      display version processing

```

Note – This listing is an example, and shows the options meaningful to the link-editor. The exact options might differ from release to release.

Most compiler drivers interpret the `-D` option during their preprocessing phase. Therefore, the `LD_OPTIONS` environment variable is a suitable mechanism for passing this option to the link-editor.

The following example shows how input files can be traced. This syntax can be especially useful in determining what libraries have been located, or what relocatable objects have been extracted from an archive during a link-edit.

```

$ LD_OPTIONS=-Dfiles cc -o prog main.o -L. -lfoo
.....
debug: file=main.o [ ET_REL ]
debug: file=./libfoo.a [ archive ]
debug: file=./libfoo.a(foo.o) [ ET_REL ]
debug: file=./libfoo.a [ archive ] (again)
.....

```

Here the member `foo.o` is extracted from the archive library `libfoo.a` to satisfy the link-edit of `prog`. Notice that the archive is searched twice to verify that the extraction of `foo.o` did not warrant the extraction of additional relocatable objects. More than one “(again)” display indicates that the archive is a candidate for ordering using `lorder(1)` and `tsort(1)`.

By using the `symbols` token, you can determine which symbol caused an archive member to be extracted, and which object made the initial symbol reference.

```

$ LD_OPTIONS=-Dsymbols cc -o prog main.o -L. -lfoo
.....
debug: symbol table processing; input file=main.o [ ET_REL ]
.....
debug: symbol[7]=foo (global); adding
debug:
debug: symbol table processing; input file=./libfoo.a [ archive ]
debug: archive[0]=bar
debug: archive[1]=foo (foo.o) resolves undefined or tentative symbol
debug:
debug: symbol table processing; input file=./libfoo(foo.o) [ ET_REL ]

```

.....

The symbol `foo` is referenced by `main.o` and is added to the link-editor's internal symbol table. This symbol reference causes the extraction of the relocatable object `foo.o` from the archive `libfoo.a`.

Note – This output has been simplified for this document.

By using the `detail` token together with the `symbols` token, the details of symbol resolution during input file processing can be observed.

```
$ LD_OPTIONS=-Dsymbols,detail cc -o prog main.o -L. -lfoo
.....
debug: symbol table processing; input file=main.o [ ET_REL ]
.....
debug: symbol[7]=foo (global); adding
debug:   entered 0x000000 0x000000 NOTY GLOB UNDEF REF_REL_NEED
debug:
debug: symbol table processing; input file=./libfoo.a [ archive ]
debug: archive[0]=bar
debug: archive[1]=foo (foo.o) resolves undefined or tentative symbol
debug:
debug: symbol table processing; input file=./libfoo.a(foo.o) [ ET_REL ]
debug: symbol[1]=foo.c
.....
debug: symbol[7]=bar (global); adding
debug:   entered 0x000000 0x000004 OBJT GLOB 3      REF_REL_NEED
debug: symbol[8]=foo (global); resolving [7][0]
debug:   old 0x000000 0x000000 NOTY GLOB UNDEF main.o
debug:   new 0x000000 0x000024 FUNC GLOB 2      ./libfoo.a(foo.o)
debug: resolved 0x000000 0x000024 FUNC GLOB 2      REF_REL_NEED
.....
```

The original undefined symbol `foo` from `main.o` has been overridden with the symbol definition from the extracted archive member `foo.o`. The detailed symbol information reflects the attributes of each symbol.

In the previous example, you can see that using some of the debugging tokens can produce a wealth of output. In cases where you are interested only in the activity around a subset of the input files, the `-D` option can be placed directly in the link-edit command-line, and toggled on and off. In the following example the display of symbol processing is switched on only during the processing of the library `libbar`.

```
$ ld .... -o prog main.o -L. -Dsymbols -lbar -D!symbols .....
```

Note – To obtain the link-edit command line you might have to expand the compilation line from any driver being used. See “Using a Compiler Driver” on page 27.

Runtime Linker

As part of the initialization and execution of a *dynamic executable*, an *interpreter* is called to complete the binding of the application to its dependencies. In the Solaris operating environment, this interpreter is referred to as the runtime linker.

During the link-editing of a dynamic executable, a special `.interp` section, together with an associated program header, are created. This section contains a path name specifying the program's interpreter. The default name supplied by the link-editor is that of the runtime linker: `/usr/lib/ld.so.1` for a 32-bit executable and `/usr/lib/64/ld.so.1` for a 64-bit executable.

Note – `ld.so.1` is a special case of a shared object. Here, a version number of 1 is used. However, later Solaris releases might provide higher version numbers.

During the process of executing a dynamic object the kernel loads the file and reads the program header information. See “Program Header” on page 236. From this information the kernel locates the name of the required interpreter. The kernel loads this interpreter and transfers control to it, passing sufficient information to enable the interpreter to continue binding the application and run it.

In addition to initializing an application, the runtime linker provides services that enable the application to extend its address space. This process involves loading additional objects and binding to symbols within them.

The runtime linker:

- Analyzes the executable's dynamic information section (`.dynamic`) and determines what dependencies are required.
- Locates and loads in these dependencies, and analyzes their dynamic information sections to determine if any additional dependencies are required.
- Performs any necessary relocations to bind these objects in preparation for process execution.

- Calls any initialization functions provided by the dependencies.
- Passes control to the application.
- Can be called upon during the application's execution, to perform any delayed function binding.
- Can be called upon by the application to acquire additional objects with `dlopen(3DL)`, and bind to symbols within these objects with `dlsym(3DL)`.

Shared Object Dependencies

When the runtime linker creates the memory segments for a program, the dependencies tell what shared objects are needed to supply the program's services. By repeatedly connecting referenced shared objects and their dependencies, the runtime linker generates a complete process image.

Note – Even when a shared object is referenced multiple times in the dependency list, the runtime linker connects the object only once to the process.

Locating Shared Object Dependencies

During the link-edit of a dynamic executable, one or more shared objects are explicitly referenced. These objects are recorded as dependencies within the dynamic executable.

The runtime linker uses this dependency information to locate, and load, the associated objects. These dependencies are processed in the same order as they were referenced during the link-edit of the executable.

Once all the dynamic executable's dependencies are loaded, they too are inspected, in the order they are loaded, to locate any additional dependencies. This process continues until all dependencies are located and loaded. This technique results in a breadth-first ordering of all dependencies.

Directories Searched by the Runtime Linker

The runtime linker looks in one default location for dependencies. This location is `/usr/lib` when processing 32-bit objects, and `/usr/lib/64` when processing 64-bit objects. Any dependency specified as a simple file name is prefixed with this default directory name. The resulting path name is used to locate the actual file.

The dependencies of a dynamic executable or shared object can be displayed using `ldd(1)`. For example, the file `/usr/bin/cat` has the following dependencies:

```
$ ldd /usr/bin/cat
      libc.so.1 =>      /usr/lib/libc.so.1
      libdl.so.1 =>    /usr/lib/libdl.so.1
```

The file `/usr/bin/cat` has a dependency, or *needs*, the files `libc.so.1` and `libdl.so.1`.

The dependencies recorded in an object can be inspected using `dump(1)`. Use this command to display the file's `.dynamic` section, and look for entries that have a `NEEDED` tag. In the following example, the dependency `libdl.so.1`, displayed in the previous `ldd(1)` example, is not recorded in the file `/usr/bin/cat`. `ldd(1)` shows the *total* dependencies of the specified file, and `libdl.so.1` is actually a dependency of `/usr/lib/libc.so.1`.

```
$ dump -Lvp /usr/bin/cat
```

```
/usr/bin/cat:
[INDEX] Tag      Value
[1]      NEEDED   libc.so.1
.....
```

In the previous `dump(1)` example, the dependencies are expressed as simple file names. In other words, there is no `'/'` in the name. The use of a simple file name requires the runtime linker to generate the path name from a set of rules. File names that contain an embedded `'/'`, are used as provided.

The simple file name recording is the standard, most flexible mechanism of recording dependencies. The `-h` option of the link-editor records a simple name within the dependency. See “Naming Conventions” on page 102 and “Recording a Shared Object Name” on page 103.

Frequently, dependencies are distributed in directories other than `/usr/lib` or `/usr/lib/64`. If a dynamic executable or shared object needs to locate dependencies in another directory, the runtime linker must explicitly be told to search this directory.

You can specify additional search path, on a per-object basis, by recording a *runpath* during the link-edit of an object. See “Directories Searched by the Runtime Linker” on page 35 for details on recording this information.

Any *runpath* recording can be displayed using `dump(1)`. Reference the `.dynamic` entry that has the `RUNPATH` tag. In the following example, `prog` has a dependency on `libfoo.so.1`. The runtime linker must search directories `/home/me/lib` and `/home/you/lib` before it looks in the default location.

```
$ dump -Lvp prog
```

```
prog:
[INDEX] Tag      Value
[1]      NEEDED   libfoo.so.1
```

```
[2]    NEEDED    libc.so.1
[3]    RUNPATH  /home/me/lib:/home/you/lib
.....
```

Another way to add to the runtime linker's search path is to set the environment variable `LD_LIBRARY_PATH`. This environment variable, which is analyzed once at process startup, can be set to a colon-separated list of directories. These directories are searched by the runtime linker before any `runpath` specification or default directory.

These environment variables are well suited to debugging purposes, such as forcing an application to bind to a local dependency. In the following example, the file `prog` from the previous example is bound to `libfoo.so.1`, found in the present working directory.

```
$ LD_LIBRARY_PATH=. prog
```

Although useful as a temporary mechanism of influencing the runtime linker's search path, the use of `LD_LIBRARY_PATH` is strongly discouraged in production software. Any dynamic executables that can reference this environment variable will have their search paths augmented. This augmentation can result in an overall degradation in performance. Also, as pointed out in "Using an Environment Variable" on page 34 and "Directories Searched by the Runtime Linker" on page 35, `LD_LIBRARY_PATH` affects the link-editor.

Environmental search paths can result in a 64-bit executable searching a path that contains a 32-bit library matching the name being looked for. Or, vice versa. The runtime linker rejects the mismatched 32-bit library and continues down its search path looking for a valid 64-bit match. If no match is found, an error message is generated. This can be observed in detail by setting the `LD_DEBUG` environment variable to include the `files` token. See "Debugging Library" on page 93.

```
$ LD_LIBRARY_PATH=/usr/bin/64 LD_DEBUG=files /usr/bin/ls
...
00283: file=libc.so.1;  needed by /usr/bin/ls
00283:
00283: file=/usr/lib/64/libc.so.1  rejected: ELF class mismatch: \
00283:                               32-bit/64-bit
00283:
00283: file=/usr/lib/libc.so.1  [ ELF ]; generating link map
00283:   dynamic: 0xef631180  base: 0xef580000  size:      0xb8000
00283:   entry:   0xef5a1240  phdr: 0xef580034  phnum:      3
00283:   lmid:    0x0
00283:
00283: file=/usr/lib/libc.so.1;  analyzing [ RTLD_GLOBAL RTLD_LAZY ]
...
```

If a dependency cannot be located, `ldd(1)` indicates that the object cannot be found. Any attempt to execute the application results in an appropriate error message from the runtime linker:

```
$ ldd prog
      libfoo.so.1 =>      (file not found)
      libc.so.1  =>      /usr/lib/libc.so.1
```

```
libdl.so.1 => /usr/lib/libdl.so.1
$ prog
ld.so.1: prog: fatal: libfoo.so.1: open failed: No such file or directory
```

Configuring the Default Search Paths

The default search paths used by the runtime linker are `/usr/lib` or `/usr/lib/64`. These search paths can be administered using a runtime configuration file created by the `crle(1)` utility. This file is often a useful aid for establishing search paths for applications that have not been built with the correct runpaths.

A configuration file can be constructed in the default location `/var/ld/ld.config`, for 32-bit applications, or `/var/ld/64/ld.config`, for 64-bit applications. This file affects all applications of the respective type on a system. Configuration files can also be created in other locations, and the runtime linker's `LD_CONFIG` environment variable used to select these files. This latter method is useful for testing a configuration file before installing it in the default location.

Dynamic String Tokens

The runtime linker allows for the expansion of various dynamic string tokens. These tokens may be applicable for filter, runpath and dependency definitions.

- `$ISALIST` – Expands to the native instruction sets executable on this platform. See “Instruction Set Specific Shared Objects” on page 321.
- `$ORIGIN` – Provides the directory the object was loaded from. See “Locating Associated Dependencies” on page 323.
- `$OSNAME` – Expands to the name of the operating system. See “System Specific Shared Objects” on page 323.
- `$OSREL` – Expands to the operating system release level. See “System Specific Shared Objects” on page 323.
- `$PLATFORM` – Expands to the processor type of the current machine. See “System Specific Shared Objects” on page 323.

Relocation Processing

After the runtime linker has loaded all the dependencies required by an application, the linker processes each object and performs all necessary relocations.

During the link-editing of an object, any relocation information supplied with the input relocatable objects is applied to the output file. However, when creating a dynamic executable or shared object, many of the relocations cannot be completed at

link-edit time because they require logical addresses that are known only when the objects are loaded into memory. In these cases the link-editor generates new relocation records as part of the output file image. The runtime linker must then process these new relocation records.

For a more detailed description of the many relocation types, see “Relocation Types (Processor-Specific)” on page 212. There are two basic types of relocations:

- Non-symbolic relocations
- Symbolic relocations

The relocation records for an object can be displayed by using `dump(1)`. In the following example, the file `libbar.so.1` contains two relocation records that indicate that the *global offset table* (the `.got` section) must be updated.

```
$ dump -rvp libbar.so.1

libbar.so.1:

.rela.got:
Offset      Symndx          Type              Addend
0x10438     0                R_SPARC_RELATIVE  0
0x1043c     foo              R_SPARC_GLOB_DAT  0
```

The first relocation is a simple relative relocation that can be seen from its relocation type and the symbol index (`Symndx`) field being zero. This relocation needs to use the base address at which the object was loaded into memory to update the associated `.got` offset.

The second relocation requires the address of the symbol `foo`. To complete this relocation, the runtime linker must locate this symbol from either the dynamic executable or one of its dependencies.

Symbol Lookup

When an object requires a symbol, the runtime linker searches for that symbol based upon the requesting object’s *symbol search scope*, and the symbol *visibility* offered by each object within the process. These attributes are applied as defaults to an object at the time the object is loaded, as specific modes to `dlopen(3DL)`, and in some cases can be recorded within the object at the time it is built.

Typically, an average user becomes familiar with the default symbol search models that are applied to a dynamic executable and its dependencies, and to objects obtained through `dlopen(3DL)`. The former is outlined in the next section “Default Lookup” on page 69, and the latter, which is also able to exploit the various symbol lookup attributes, is discussed in “Symbol Lookup” on page 84.

An alternative model for symbol lookup is provided when a dynamic object employs direct bindings. This model directs the runtime linker to search for a symbol directly in the object that provided the symbol at link-edit time. See “Direct Binding” on page 69.

Default Lookup

A dynamic executable and all the dependencies loaded with it are assigned *world* search scope, and *global* symbol visibility. See “Symbol Lookup” on page 84. When the runtime linker looks up a symbol for a dynamic executable or for any of the dependencies loaded with the executable, it does so by searching each object. The runtime linker starts with the dynamic executable, and progresses through each dependency in the same order in which the objects were loaded.

As discussed in previous sections, `ldd(1)` lists the dependencies of a dynamic executable in the order in which they are loaded. Therefore, if the shared object `libbar.so.1` requires the address of symbol `foo` to complete its relocation, and this shared object is a dependency of the dynamic executable `prog`:

```
$ ldd prog
    libfoo.so.1 => /home/me/lib/libfoo.so.1
    libbar.so.1 => /home/me/lib/libbar.so.1
```

The runtime linker first looks for `foo` in the dynamic executable `prog`, then in the shared object `/home/me/lib/libfoo.so.1`, and finally in the shared object `/home/me/lib/libbar.so.1`.

Note – Symbol lookup can be an expensive operation, especially when the size of symbol names increases and the number of dependencies increases. This aspect of performance is discussed in more detail in “Performance Considerations” on page 112. See “Direct Binding” on page 69 for an alternative lookup model.

Interposition

The runtime linker’s default mechanism of searching for a symbol first in the dynamic executable and then in each of the dependencies means that the first occurrence of the required symbol will satisfy the search. Therefore, if more than one instance of the same symbol exists, the first instance interposes on all others. See also “Shared Object Processing” on page 30.

Direct Binding

When creating an object to use direct bindings, the relationship between the referenced symbol and the dependency that provided the definition is recorded in the object. The runtime linker uses this information to search directly for the symbol in the

associated object, rather than carry out the default symbol search model. Direct binding information can only be established to dependencies specified with the `link-edit`. Therefore, use of the `-z defs` option is recommended.

Direct bindings can be established with one of the following mechanisms.

- With the `-B direct` option. This option establishes direct bindings between the object being built and all its dependencies. The use of `-B direct` also enables lazy loading, which is equivalent to adding the option `-z lazyload` to the front of the `link-edit` command line. See “Lazy Loading of Dynamic Dependencies” on page 74.
- With the `-z direct` option. This option establishes direct bindings between the object being built and any dependencies that follow the option on the command line. This option can be used together with the `-z nodirect` option, to toggle the use of direct bindings between dependencies.
- With the `DIRECT mapfile` attribute. This attribute provides for directly binding individual symbols. The alternative attribute `NODIRECT`, can be used to prevent direct binding to individual symbols. See “Defining Additional Symbols” on page 46.

The direct binding model can significantly reduce the symbol lookup overhead within a dynamic process that has many symbolic relocations and many dependencies. This model also enables multiple symbols of the same name to be located from different objects that have been bound to directly.

Direct binding can circumvent the traditional use of interposition symbols because it bypasses the default search model. The default model ensures that all references to a symbol bind to one definition.

Interposition can still be achieved in a direct binding environment, on a per-object basis, if an object is identified as an interposer. Any object loaded using the environment variable `LD_PRELOAD` or created with the link-editor’s `-z interpose` option, is identified as an interposer. When the runtime linker searches for a directly bound symbol, it first looks in any object identified as an interposer before it looks in the object that supplies the symbol definition.

Note – Direct bindings can be disabled at runtime by setting the environment variable `LD_NODIRECT` to a non-null value.

Some interfaces exist that offer alternative implementations of a default technology. These implementations also assume they are the only instance of that technology within a process. An example of this is the `malloc(3C)` family. Directly binding to interfaces within such a family should be avoided, as it is possible for more than one instance of the technology to be referenced by the same process. For example, one dependency within a process may directly bind against `libc.so.1`, while another dependency directly binds against `libmapmalloc.so.1`.

Objects that provide a single implementation for a process, should define the interfaces to that implementation using the `mapfile` directive `NODIRECT`. This directive insures no users directly bind to an implementation, but use the default symbol search model.

Note – `NODIRECT` `mapfile` directives can be combined with the command line options `-B direct` or `-z direct`. Symbols that are not explicitly defined `NODIRECT` will follow the command line directive.

When Relocations Are Performed

Relocations can be distinguish by when they are performed. This distinction arises due to the type of *reference* being made to the relocated offset, and is either:

- An immediate reference
- A lazy reference

An *immediate reference* refers to a relocation that must be determined immediately when an object is loaded. These references are typically to data items used by the object code, pointers to functions, and even calls to functions made from position-dependent shared objects. These relocations cannot provide the runtime linker with knowledge of when the relocated item is referenced. Therefore, all immediate relocations must be carried out when an object is loaded, and before the application gains, or regains, control.

A *lazy reference* refers to a relocation that can be determined as an object executes. These references are typically calls to global functions made from position-independent shared objects, or calls to external functions made from a dynamic executable. During the compilation and link-editing of any dynamic module that provide these references, the associated function calls become calls to a procedure linkage table entry. These entries make up the `.plt` section. Each procedure linkage table entry becomes a lazy reference with a relocation associated with it.

Procedure linkage table entries are constructed so that when they are first called, control is passed to the runtime linker. The runtime linker looks up the required symbol and rewrites information in the associated object so that any future calls to this procedure linkage table entry go directly to the function. This mechanism enables relocations of this type to be deferred until the first instance of a function is called. This process is sometimes referred to as *lazy binding*.

The runtime linker's default mode is to perform lazy binding whenever procedure linkage table relocations are provided. This default can be overridden by setting the environment variable `LD_BIND_NOW` to any non-null value. This environment variable setting causes the runtime linker to perform both immediate and lazy reference

relocations when an object is loaded, and before the application gains, or regains, control. For example, setting the environment variable as follows means that all relocations within the file `prog` and within its dependencies, will be processed before control is transferred to the application.

```
$ LD_BIND_NOW=1 prog
```

Objects can also be accessed with `dlopen(3DL)` with the mode defined as `RTLD_NOW`. Objects can also be built using the link-editor's `-z now` option to indicate that they require complete relocation processing at the time they are loaded. This relocation requirement is also propagated to any dependencies of the marked object at runtime.

Note – Although the preceding examples of immediate and lazy references are typical, the creation of procedure linkage table entries is ultimately controlled by the relocation information provided by the relocatable object files used as input to a link-edit. Relocation records such as `R_SPARC_WPLT30` and `R_386_PLT32` instruct the link-editor to create a procedure linkage table entry are common for position-independent code. However, as a dynamic executable has a fixed location, external function references that can be determined at link-edit time can be converted to procedure linkage table entries regardless of the original relocation records.

Relocation Errors

The most common relocation error occurs when a symbol cannot be found. This condition results in an appropriate runtime linker error message and the termination of the application. For example:

```
$ ldd prog
  libfoo.so.1 => ./libfoo.so.1
  libc.so.1 => /usr/lib/libc.so.1
  libbar.so.1 => ./libbar.so.1
  libdl.so.1 => /usr/lib/libdl.so.1

$ prog
ld.so.1: prog: fatal: relocation error: file ./libfoo.so.1: \
symbol bar: referenced symbol not found
```

The symbol `bar`, which is referenced in the file `libfoo.so.1`, cannot be located.

During the link-edit of a dynamic executable, any potential relocation errors of this sort are flagged as fatal undefined symbols. See “Generating an Executable Output File” on page 43 for examples. This runtime relocation error can occur if the link-edit of `main` used a different version of the shared object `libbar.so.1` that contained a symbol definition for `bar`, or if the `-z nodefs` option was used as part of the link-edit.

If a relocation error of this type occurs because a symbol used as an immediate reference cannot be located, the error condition will occur immediately during process initialization. Because of the default mode of lazy binding, if a symbol used as a lazy

reference cannot be found, the error condition will occur after the application has gained control. This latter case can take minutes or months, or might never occur, depending on the execution paths exercised throughout the code.

To guard against errors of this kind, the relocation requirements of any dynamic executable or shared object can be validated using `ldd(1)`.

When the `-d` option is specified with `ldd(1)`, all dependencies will be printed and all immediate reference relocations will be processed. If a reference cannot be resolved, a diagnostic message is produced. From the previous example this option would result in:

```
$ ldd -d prog
libfoo.so.1 => ./libfoo.so.1
libc.so.1 => /usr/lib/libc.so.1
libbar.so.1 => ./libbar.so.1
libdl.so.1 => /usr/lib/libdl.so.1
symbol not found: bar          (./libfoo.so.1)
```

When the `-r` option is specified with `ldd(1)`, all immediate *and* lazy reference relocations are processed. If either type of relocation cannot be resolved, a diagnostic message is produced.

Loading Additional Objects

The runtime linker provides an additional level of flexibility by enabling you to introduce new objects during process initialization.

The environment variable `LD_PRELOAD` can be initialized to a shared object or relocatable object file name, or a string of file names separated by white space. These objects are loaded after the dynamic executable and before any dependencies. These objects are assigned *world* search scope, and *global* symbol visibility.

```
$ LD_PRELOAD=./newstuff.so.1 prog
```

The dynamic executable `prog` is loaded, followed by the shared object `newstuff.so.1`, and then by the dependencies defined within `prog`.

The order in which these objects are processed can be displayed using `ldd(1)`:

```
$ LD_PRELOAD=./newstuff.so.1 ldd prog
./newstuff.so.1 => ./newstuff.so
libc.so.1 => /usr/lib/libc.so.1
```

In another example the preloading is a little more complex and time consuming.

```
$ LD_PRELOAD="./foo.o ./bar.o" prog
```

The runtime linker first link-edits the relocatable objects `foo.o` and `bar.o` to generate a shared object that is maintained in memory. This memory image is then inserted between the dynamic executable and its dependencies in the same manner as the shared object `newstuff.so.1` was preloaded in the previous example. Again, the order in which these objects are processed can be displayed with `ldd(1)`:

```
$ LD_PRELOAD="./foo.o ./bar.o" ldd prog
./foo.o => ./foo.o
./bar.o => ./bar.o
libc.so.1 => /usr/lib/libc.so.1
```

These mechanisms of inserting an object after a dynamic executable take the concept of interposition to another level. You can use these mechanisms to experiment with a new implementation of a function that resides in a standard shared object. If you preload an object containing this function, the object interposes on the original. Thus the old functionality can be completely hidden with the new preloaded version.

Another use of preloading is to augment a function that resides in a standard shared object. The intention is to interpose the new symbol on the original, enabling the new function to carry out some additional processing while calling through to the original function. This mechanism requires either a symbol alias that is to be associated with the original function or the ability to look up the original symbol's address.

Lazy Loading of Dynamic Dependencies

When a dynamic object is loaded into memory, it is examined for any additional dependencies. By default, if any dependencies exist they are immediately loaded. This cycle continues until the full dependency tree is exhausted. At which point all inter-object references, specified by relocations, are resolved.

Under this default model, all the dependencies of an application are loaded into memory, and all data relocations are performed. These operations are performed regardless of whether the code in these dependencies is referenced by the application during its execution.

Under a lazy loading model, any dependencies that are labeled for lazy loading are loaded only when explicitly referenced. By taking advantage of a function call's lazy binding, the loading of a dependency is delayed until it is first referenced. In fact, objects that are never referenced are never loaded.

A relocation reference can be immediate or lazy. Because immediate references must be resolved when an object is initialized, any dependency that satisfies this reference must be immediately loaded. Therefore, identifying such a dependency as lazy loadable has little effect. See "When Relocations Are Performed" on page 71. Immediate references between dynamic objects are generally discouraged.

Lazy loading is used by the link-editor itself, which references a debugging library, `liblddbg`. Because debugging is only called upon infrequently, loading this library every time the link-editor is invoked is unnecessary and expensive. By indicating that this library can be lazily loaded, the expense of processing it can be moved to those invocations that ask for debugging output.

The alternate method of achieving a lazy loading model is to use `dlopen()` and `dlsym()` to load and bind to a dependency when needed. This is ideal if the number of `dlsym()` references is small, or the dependency name or location is not known at link-edit time. For more complex interactions with known dependencies, coding to normal symbol references and designating the dependency to be lazily loaded is simpler.

An object is designated as lazily or normally loaded through the link-editor options `-z lazyload` and `-z nolazyload` respectfully. These options are position-dependent on the link-edit command line. Any dependency found following the option takes on the loading attribute specified by the option. By default, the `-z nolazyload` option is in effect.

The following simple program has a dependency on `libdebug.so.1`. The dynamic section (`.dynamic`), shows `libdebug.so.1` is marked for lazy loading. The symbol information section (`.SUNW_syminfo`), shows the symbol reference that triggers `libdebug.so.1` loading.

```
$ cc -o prog prog.c -L. -zlazyload -ldebug -znolazyload -R'$ORIGIN'
$ elfdump -d prog
```

```
Dynamic Section: .dynamic
  index  tag          value
  [0]    POSFLAG_1    0x1          [ LAZY ]
  [1]    NEEDED       0x123        libdebug.so.1
  [2]    NEEDED       0x131        libc.so.1
  [3]    RUNPATH      0x13b        $ORIGIN
  ...
```

```
$ elfdump -y prog
```

```
Syminfo section: .SUNW_syminfo
  index flgs boundto          symbol
  ...
  [52] DL      [1] libdebug.so.1  debug
```

The `POSFLAG_1` with the value of `LAZY` designates that the following `NEEDED` entry, `libdebug.so.1`, should be lazily loaded. Because `libc.so.1` has no preceding `LAZY` flag it is loaded at the initial startup of the program.

The use of lazy loading can require a precise declaration of dependencies and runpaths through out the objects used by an application. For example, suppose two objects, `libA.so` and `libB.so`, both make reference to symbols in `libX.so`. `libA.so` declares it has a dependency on `libX.so`, but `libB.so` does not. Typically, when `libA.so` and `libB.so` are used together, `libB.so` can reference `libX.so` because `libA.so` made it available. But, if `libA.so` declares `libX.so` to be lazy

loaded, it is possible that `libX.so` may not be loaded when `libB.so` makes reference to it. A similar failure can occur if `libB.so` declares `libX.so` as a dependency but fails to provide a runpath necessary to locate it.

Regardless of lazy loading, it is recommended that dynamic objects declare all their dependencies and how to locate them. With lazy loading, this dependency information becomes even more important.

Note – Lazy loading can be disabled at runtime by setting the environment variable `LD_NOLAZYLOAD` to a non-null value.

Initialization and Termination Routines

Before transferring control to an application, the runtime linker processes any initialization sections found in the application and any loaded dependencies. The initialization sections `.preinit_array`, `.init_array`, and `.init` are created by the link-editor when a dynamic object is built.

The runtime linker executes functions whose addresses are contained in the `.preinit_array` and `.init_array` sections. These functions are executed in the same order in which their addresses appear in the array. The runtime linker executes an `.init` section as an individual function. If an object contains both `.init` and `.init_array` sections, the `.init` section is processed before the functions defined by the `.init_array` section for that object.

A dynamic executable may provide pre-initialization functions in a `.preinit_array` section. These functions are executed after the runtime linker has built the process image and performed relocations but before any other initialization functions. Pre-initialization functions are not permitted in shared objects.

Note – Any `.init` section within the dynamic executable is called from the application itself by the process startup mechanism supplied by the compiler driver. The `.init` section within the dynamic executable is called last, after all dependency initialization sections are executed.

Dynamic objects can also provide termination sections. The termination sections `.fini_array` and `.fini` are created by the link-editor when a dynamic object is built.

Any termination sections are organized such that they can be recorded by `atexit(3C)`. These routines are called when the process calls `exit(2)`, or when objects are removed from the running process with `dlclose(3DL)`.

The runtime linker executes functions whose addresses are contained in the `.fini_array` section. These functions are executed in the reverse order in which their addresses appear in the array. The runtime linker executes a `.fini` section as an individual function. If an object contains both `.fini` and `.fini_array` sections, the functions defined by the `.fini_array` section are processed before the `.fini` section for that object.

Note – Any `.fini` section within the dynamic executable is called from the application itself by the process termination mechanism supplied by the compiler driver. The `.fini` section of the dynamic executable is called first, before all dependency termination sections are executed.

For more information regarding the creation of initialization and termination sections by the link-editor see “Initialization and Termination Sections” on page 36.

Initialization and Termination Order

To determine the order of executing initialization and termination code within a process at runtime is a complex issue involving dependency analysis. This process has evolved substantially from the original inception of initialization and termination sections. This process attempts to fulfill the expectations of modern languages and current programming techniques. However, scenarios can exist, where user expectations are hard to meet. Understanding these scenarios, and limiting the content of initialization and termination code can provide both flexible and predictable runtime behavior.

Prior to the Solaris 2.6 release, dependency initialization routines were called in *reverse* load order, which is the reverse order of the dependencies displayed with `ldd(1)`. Similarly, dependency termination routines were called in load order. However, as dependency hierarchies became more complex, this simple ordering approach became inadequate.

Starting with the Solaris 2.6 release, the runtime linker constructs a topologically sorted list of objects that have been loaded. This list is built from the dependency relationship expressed by each object, together with any symbol bindings that occur outside of the expressed dependencies.

Initialization sections are executed in the reverse topological order of the dependencies. If cyclic dependencies are found, the objects that form the cycle cannot be topologically sorted. The initialization sections of any cyclic dependencies are executed in their reverse load order. Similarly, termination routines are called in the topological order of dependencies and any cyclic dependencies are executed in their load order.

Use `ldd(1)` with the `-i` option to display the initialization order of an object's dependencies. For example, the following dynamic executable and its dependencies exhibit a cyclic dependency:

```
$ dump -Lv B.so.1 | grep NEEDED
[1]  NEEDED      C.so.1
$ dump -Lv C.so.1 | grep NEEDED
[1]  NEEDED      B.so.1
$ dump -Lv main | grep NEEDED
[1]  NEEDED      A.so.1
[2]  NEEDED      B.so.1
[3]  NEEDED      libc.so.1
$ ldd -i main
    A.so.1 =>      ./A.so.1
    B.so.1 =>      ./B.so.1
    libc.so.1 =>   /usr/lib/libc.so.1
    C.so.1 =>      ./C.so.1
    libdl.so.1 =>  /usr/lib/libdl.so.1

cyclic dependencies detected, group[1]:
./libC.so.1
./libB.so.1

init object=/usr/lib/libc.so.1
init object=./A.so.1
init object=./C.so.1 - cyclic group [1], referenced by:
./B.so.1
init object=./B.so.1 - cyclic group [1], referenced by:
./C.so.1
```



Caution – Prior to Solaris 8 10/00, the environment variable `LD_BREADTH` could be set to a non-null value to force the runtime linker to execute initialization and termination sections in pre-Solaris 2.6 order. This functionality has since been disabled, as the initialization dependencies of many applications have become complex and mandate topological sorting. Any `LD_BREADTH` setting is now silently ignored.

Initialization processing is repeated for any objects added to the running process with `dlopen(3DL)`. Termination processing is also carried out for any objects unloaded from the process as a result of a call to `dldclose(3DL)`.

Symbol bindings are incorporated as part of dependency analysis because many shared objects exist that do not express their dependencies accurately. Incorporating symbol bindings can therefore help produce a more accurate dependency relationship.

However, the addition of symbol binding information to objects that do not express all their dependencies, may still be insufficient to determine an objects complete dependencies. The most common model of loading objects uses lazy binding. With this model, only *immediate reference* symbol bindings are processed before initialization processing. Symbol bindings from *lazy references* may still be pending, and may extend the dependency relationships so far established.

As the dependency analysis of an object may be incomplete, and as cyclic dependencies often exist, the runtime linker also provides for dynamic initialization. This initialization attempts to execute any initialization sections before any functions in the same object are called. During lazy symbol binding, the runtime linker determines whether the initialization sections of the object being bound to have been called. If not, the runtime linker calls them before returning from the symbol binding procedure.

Dynamic initialization can not be revealed with `ldd(1)`. However, the exact sequence of initialization calls can be observed at runtime by setting the `LD_DEBUG` environment variable to include the token *basic*. See “Debugging Library” on page 93.

Dynamic initialization is only available when processing lazy references. Use of the environment variable `LD_BIND_NOW`, objects built with the `-z now` option, or objects referenced by `dlopen(3DL)` with mode `RTLD_NOW`, circumvent any dynamic initialization.

Note – Objects that are pending initialization, and are referenced through `dlopen(3DL)`, will be initialized prior to returning control from this function.

The preceding sections describe the various techniques employed to execute initialization and termination sections in a manner that attempts to meet user expectations. However, coding style and link-editing practices should also be employed to simplify the initialization and termination relationships between dependencies. This simplification, helps keep initialization and termination processing predictable, and less prone to any side affects of unexpected dependency ordering.

Keep the content of initialization and termination sections to a minimum. Avoid global constructors by initializing objects at runtime. Reduce the dependency of initialization and termination code on other dependencies. Define the dependency requirements of all dynamic objects. See “Generating a Shared Object Output File” on page 44. Do not express dependencies that are not required. See “Shared Object Processing” on page 30. Avoid cyclic dependencies. Do not depend on the order of an initialization or termination sequence. The ordering of objects can be affected by both shared object and application development. See “Dependency Ordering” on page 106.

Security

Secure processes have some restrictions applied to the evaluation of their dependencies and runpaths to prevent malicious dependency substitution or symbol interposition.

The runtime linker categorizes a process as secure if the user is not a super-user, and the real user and effective user identifiers are not equal. Similarly, if the user is not a super-user and the real group and effective group identifiers are not equal, the process is deemed secure. See the `getuid(2)`, `geteuid(2)`, `getgid(2)` and `getegid(2)` man pages.

The default trusted directory known to the runtime linker is `/usr/lib/secure` for 32-bit objects or `/usr/lib/secure/64` for 64-bit objects. The utility `crle(1)` may be used to specify additional trusted directories applicable for secure applications. Administrators who use this technique should ensure that the target directories are suitably protected from malicious intrusion.

If an `LD_LIBRARY_PATH` family environment variable is in effect for a secure process, only the trusted directories specified by this variable are used to augment the runtime linker's search rules. See "Directories Searched by the Runtime Linker" on page 64.

In a secure process, any runpath specifications provided by the application or any of its dependencies is used, provided it is a full pathname, that is, the pathname starts with a `('/')`.

In a secure process, the expansion of the `$ORIGIN` string is allowed only if it expands to a trusted directory. See "Security" on page 326.

In a secure process, `LD_CONFIG` is ignored. A secure process uses the default configuration file, if it exists. See `crle(1)`.

In a secure process, `LD_SIGNAL` is ignored.

Additional objects can be loaded with a secure process using the `LD_PRELOAD` or `LD_AUDIT` environment variables. These objects must be specified as full path names or simple file names. Full path names are restricted to known trusted directories. Simple file names, in which no `('/')` appears in the name, are located subject to the search path restrictions previously described. Simple file names resolve only to known trusted directories.

In a secure process, any dependencies that consist of simple file names are processed using the path name restrictions previously described. Dependencies expressed as full or relative path names are used as is. Therefore, the developer of a secure process should ensure that the target directory referenced as a full or relative path name dependency is suitably protected from malicious intrusion.

When creating a secure process, do not use relative path names to express dependencies or to construct `dlopen(3DL)` path names. This restriction should be applied to the application and to all dependencies.

Runtime Linking Programming Interface

Dependencies specified during the link-edit of an application are processed by the runtime linker during process initialization. In addition to this mechanism, the application can extend its address space during its execution by binding to additional objects. The application can request the same services of the runtime linker that are used to process the dependencies specified during the link-edit of the application.

This delayed object binding has several advantages:

- By processing an object when it is required rather than during the initialization of an application, startup time can be greatly reduced. In fact, the object might not be required if its services are not needed during a particular run of the application, such as for help or debugging information.
- The application can choose between several different objects, depending on the exact services required, such as for a networking protocol.
- Any objects added to the process address space during execution can be freed after use.

An application can use the following typical scenario to access an additional shared object.

- A shared object is located and added to the address space of a running application using `dlopen(3DL)`. Any dependencies that this shared object has are located and added at this time.
- The added shared object and its dependencies are relocated. Any initialization sections within these objects are called.
- The application locates symbols within the added objects using `dlsym(3DL)`. The application can then reference the data or call the functions defined by these new symbols.
- After the application has finished with the objects, the address space can be freed using `dlclose(3DL)`. Any termination sections within the objects being freed is called at this time.
- Any error conditions that occur as a result of using these runtime linker interface routines can be displayed using `dLError(3DL)`.

The services of the runtime linker are defined in the header file `dlfcn.h` and are made available to an application by the shared object `libdl.so.1`. In the following example, the file `main.c` can make reference to any of the `dlopen(3DL)` family of routines, and the application `prog` can bind to these routines at runtime.

```
$ cc -o prog main.c -ldl
```

Loading Additional Objects

Additional objects can be added to a running process's address space using `dlopen(3DL)`. This function takes a path name and a binding mode as arguments, and returns a handle to the application. This handle can be used to locate symbols for use by the application using `dlsym(3DL)`.

If the path name is specified as a *simple* file name, one with no `/'` in the name, then the runtime linker will use a set of rules to generate an appropriate path name. Path names that contain a `/'` will be used as provided.

These search path rules are exactly the same as are used to locate any initial dependencies. See "Directories Searched by the Runtime Linker" on page 64. For example, if the file `main.c` contains the following code fragment:

```
#include      <stdio.h>
#include      <dlfcn.h>

main(int argc, char ** argv)
{
    void *   handle;
    .....

    if ((handle = dlopen("foo.so.1", RTLD_LAZY)) == NULL) {
        (void) printf("dlopen: %s\n", dlerror());
        exit (1);
    }
    .....
}
```

then to locate the shared object `foo.so.1`, the runtime linker uses any `LD_LIBRARY_PATH` definition present at process initialization, followed by any `runpath` specified during the link-edit of `prog`. Finally, the runtime linker uses the default location `/usr/lib` for 32-bit objects, and `/usr/lib/64` for 64-bit objects.

If the path name is specified as:

```
if ((handle = dlopen("./foo.so.1", RTLD_LAZY)) == NULL) {
```

then the runtime linker searches for the file only in the current working directory of the process.

Note – Any shared object specified using `dlopen(3DL)` should be referenced by its *versioned* file name. For more information on versioning, see "Coordination of Versioned Filenames" on page 146.

If the required object cannot be located, `dlopen(3DL)` returns a `NULL` handle. In this case `dLError(3DL)` can be used to display the true reason for the failure. For example:

```
$ cc -o prog main.c -ldl
$ prog
dlopen: ld.so.1: prog: fatal: foo.so.1: open failed: No such \
file or directory
```

If the object being added by `dlopen(3DL)` has dependencies on other objects, they too are brought into the process's address space. This process continues until all the dependencies of the specified object are loaded. This dependency tree is referred to as a *group*.

If the object specified by `dlopen(3DL)`, or any of its dependencies, are already part of the process image, then the objects are not processed any further. A valid handle is returned to the application. This mechanism prevents the same object from being loaded more than once, and enables an application to obtain a handle to itself. For example, if the previous `main.c` example contained the following `dlopen()` call:

```
if ((handle = dlopen((const char *)0, RTLD_LAZY)) == NULL) {
```

then the handle returned from `dlopen(3DL)` can be used to locate symbols within the application itself, within any of the dependencies loaded as part of the process's initialization, or within any objects added to the process's address space, using a `dlopen(3DL)` that specified the `RTLD_GLOBAL` flag.

Relocation Processing

As described in Chapter 3, after locating and loading any objects, the runtime linker must process each object and perform any necessary relocations. Any objects brought into the process's address space with `dlopen(3DL)` must also be relocated in the same manner.

For simple applications this process is straightforward. However, for users who have more complex applications with many `dlopen(3DL)` calls involving many objects, possibly with common dependencies, this process can be quite important.

Relocations can be categorized according to when they occur. The default behavior of the runtime linker is to process all immediate reference relocations at initialization and all lazy references during process execution, a mechanism commonly referred to as lazy binding.

This same mechanism is applied to any objects added with `dlopen(3DL)` when the mode is defined as `RTLD_LAZY`. An alternative is to require all relocations of an object to be performed immediately when the object is added. You can use a mode of `RTLD_NOW`, or record this requirement in the object when it is built using the `link-editor's -z now` option. This relocation requirement is propagated to any dependencies of the object being opened.

Relocations can also be categorized into non-symbolic and symbolic. The remainder of this section covers issues regarding symbolic relocations, regardless of when these relocations occur, with a focus on some of the subtleties of symbol lookup.

Symbol Lookup

If an object acquired by `dlopen(3DL)` refers to a global symbol, the runtime linker must locate this symbol from the pool of objects that make up the process. In the absence of direct binding, a default symbol search model is applied to objects obtained by `dlopen(3DL)`. However, the mode of a `dlopen(3DL)`, combined with the attributes of the objects that make up the process, provide for alternative symbol search models.

Objects that required direct binding, although maintaining all the attributes described later, search for symbols directly in the associated dependency. See “Direct Binding” on page 69.

Two attributes of an object affect symbol lookup. The first is the requesting object’s symbol *search scope*, and the second is the symbol *visibility* offered by each object within the process. An object’s search scope can be:

`world`

The object can look in any other global object within the process.

`group`

The object can look only in an object of the same *group*. The dependency tree created from an object obtained with `dlopen(3DL)`, or from an object built using the link-editor’s `-B group` option, forms a unique group.

The visibility of a symbol from an object can be:

`global`

The object’s symbols can be referenced from any object having *world* search scope.

`local`

The object’s symbols can be referenced only from other objects that make up the same group.

By default, objects obtained with `dlopen(3DL)` are assigned *world* symbol search scope, and *local* symbol visibility. The section, “Default Symbol Lookup Model” on page 85, uses this default model to illustrate typical object group interactions. The sections “Defining a Global Object” on page 88, “Isolating a Group” on page 88, and “Object Hierarchies” on page 88 show examples of using `dlopen(3DL)` modes and file attributes to extend the default symbol lookup model.

Default Symbol Lookup Model

For each object added by `dlopen(3DL)` the runtime linker first looks for the symbol in the dynamic executable. The runtime linker then looks in each of the objects provided during the initialization of the process. If the symbol is still not found, the runtime linker continues the search, looking in the object acquired through the `dlopen(3DL)` and in any of its dependencies.

In the following example, the dynamic executable `prog` and the shared object `B.so.1` each have the following (simplified) dependencies:

```
$ ldd prog
    A.so.1 =>          ./A.so.1
$ ldd B.so.1
    C.so.1 =>          ./C.so.1
```

If `prog` acquires the shared object `B.so.1` by `dlopen(3DL)`, then any symbol required to relocate the shared objects `B.so.1` and `C.so.1` will first be looked for in `prog`, followed by `A.so.1`, followed by `B.so.1`, and finally in `C.so.1`. In this simple case, think of the shared objects acquired through the `dlopen(3DL)` as if they had been added to the end of the original link-edit of the application. For example, the objects referenced in the previous listing can be expressed diagrammatically as shown in the following figure.




FIGURE 3-1 A Single `dlopen()` Request

Any symbol lookup required by the objects acquired from the `dlopen(3DL)`, shown as shaded blocks, proceeds from the dynamic executable `prog` through to the final shared object `C.so.1`.

This symbol lookup is established by the attributes assigned to the objects as they were loaded. Recall that the dynamic executable and all the dependencies loaded with it are assigned global symbol visibility, and that the new objects are assigned world symbol search scope. Therefore, the new objects are able to look for symbols in the original objects. The new objects also form a unique group in which each object has local symbol visibility. Therefore, each object within the group can look for symbols within the other group members.

These new objects do not affect the normal symbol lookup required by either the application or its initial object dependencies. For example, if `A.so.1` requires a function relocation after the above `dlopen(3DL)` has occurred, the runtime linker's normal search for the relocation symbol is to look in `prog` and then `A.so.1`. The runtime linker does not follow through and look in `B.so.1` or `C.so.1`.

This symbol lookup is again a result of the attributes assigned to the objects as they were loaded. The world symbol search scope is assigned to the dynamic executable and all the dependencies loaded with it. This scope does not allow them to look for symbols in the new objects that only offer local symbol visibility.

These symbol search and symbol visibility attributes maintain associations between objects based on their introduction into the process address space, and on any dependency relationship between the objects. Assigning the objects associated with a given `dlopen(3DL)` to a unique group ensures that only objects associated with the same `dlopen(3DL)` are allowed to look up symbols within themselves and their related dependencies.

This concept of defining associations between objects becomes more clear in applications that carry out more than one `dlopen(3DL)`. For example, suppose the shared object `D.so.1` has the following dependency:

```
$ ldd D.so.1
      E.so.1 =>          ../E.so.1
```

and the `prog` application used `dlopen(3DL)` to load this shared object in addition to the shared object `B.so.1`. The following figure illustrates the symbol lookup relationship between the objects.




FIGURE 3-2 Multiple `dlopen()` Requests

Suppose that both `B.so.1` and `D.so.1` contain a definition for the symbol `foo`, and both `C.so.1` and `E.so.1` contain a relocation that requires this symbol. Because of the association of objects to a unique group, `C.so.1` is bound to the definition in `B.so.1`, and `E.so.1` is bound to the definition in `D.so.1`. This mechanism is intended to provide the most intuitive binding of objects obtained from multiple calls to `dlopen(3DL)`.

When objects are used in the scenarios that have so far been described, the order in which each `dlopen(3DL)` occurs has no effect on the resulting symbol binding. However, when objects have common dependencies, the resultant bindings can be affected by the order in which the `dlopen(3DL)` calls are made.

In the following example, the shared objects `O.so.1` and `P.so.1` have the same common dependency.

```
$ ldd O.so.1
    Z.so.1 =>          ./Z.so.1
$ ldd P.so.1
    Z.so.1 =>          ./Z.so.1
```

In this example, the `prog` application will `dlopen(3DL)` each of these shared objects. Because the shared object `Z.so.1` is a common dependency of both `O.so.1` and `P.so.1`, `Z.so.1` is assigned to both of the groups that are associated with the two `dlopen(3DL)` calls. This relationship is shown in the following figure.




FIGURE 3-3 Multiple `dlopen()` Requests With A Common Dependency

`Z.so.1` is available for both `O.so.1` and `P.so.1` to look up symbols. More importantly, as far as `dlopen(3DL)` ordering is concerned, `Z.so.1` is also able to look up symbols in both `O.so.1` and `P.so.1`.

Therefore, if both `O.so.1` and `P.so.1` contain a definition for the symbol `foo`, which is required for a `Z.so.1` relocation, the actual binding that occurs is unpredictable because it is affected by the order of the `dlopen(3DL)` calls. If the functionality of symbol `foo` differs between the two shared objects in which it is defined, the overall outcome of executing code within `Z.so.1` might vary depending on the application's `dlopen(3DL)` ordering.

Defining a Global Object

The default assignment of local symbol visibility to the objects obtained by a `dlopen(3DL)` can be promoted to global by augmenting the mode argument with the `RTLD_GLOBAL` flag. Under this mode, any objects obtained through a `dlopen(3DL)` can be used by any other objects with world symbol search scope to locate symbols.

In addition, any object obtained by `dlopen(3DL)` with the `RTLD_GLOBAL` flag is available for symbol lookup using `dlopen()` with a path name whose value is 0.

Note – If a member of a group having local symbol visibility is referenced by another group requiring global symbol visibility, the object’s visibility will become a concatenation of both local and global. This promotion of attributes remains even if the global group reference is later removed.

Isolating a Group

The default assignment of world symbol search scope to the objects obtained by a `dlopen(3DL)` can be reduced to group by augmenting the mode argument with the `RTLD_GROUP` flag. Under this mode, any objects obtained through a `dlopen(3DL)` will only be allowed to look for symbols within their own group.

Using the link-editor’s `-B group` option, you can assign the group symbol search scope to objects when they are built.

Note – If a member of a group, having group search capability, is referenced by another group requiring world search capability, the object’s search capability will become a concatenation of both group and world. This promotion of attributes remains even if the world group reference is later removed.

Object Hierarchies

If an initial object, obtained from a `dlopen(3DL)`, was to use `dlopen(3DL)` to open a secondary object, both objects would be assigned to a unique group. This situation can prevent either object from locating symbols from one another.

In some implementations the initial object has to export symbols for the relocation of the secondary object. This requirement can be satisfied by one of two mechanisms:

- Making the initial object an explicit dependency of the second object
- Use the `RTLD_PARENT` mode flag to `dlopen(3DL)` the secondary object

If the initial object is an explicit dependency of the secondary object, the initial object is assigned to the secondary objects’ group. The initial object is therefore able to provide symbols for the secondary objects’ relocation.

If many objects can use `dlopen(3DL)` to open the secondary object, and each of these initial objects must export the same symbols to satisfy the secondary objects' relocation, then the secondary object cannot be assigned an explicit dependency. In this case, the `dlopen(3DL)` mode of the secondary object can be augmented with the `RTLD_PARENT` flag. This flag causes the propagation of the secondary objects' group to the initial object in the same manner as an explicit dependency would do.

There is one small difference between these two techniques. If you specify an explicit dependency, the dependency itself becomes part of the secondary objects' `dlopen(3DL)` dependency tree, and thus becomes available for symbol lookup with `dlsym(3DL)`. If you obtain the secondary object with `RTLD_PARENT`, the initial object does not become available for symbol lookup with `dlsym(3DL)`.

When a secondary object is obtained by `dlopen(3DL)` from an initial object with global symbol visibility, the `RTLD_PARENT` mode is both redundant and harmless. This case commonly occurs when `dlopen(3DL)` is called from an application or from one of the dependencies of the application.

Obtaining New Symbols

A process can obtain the address of a specific symbol using `dlsym(3DL)`. This function takes a *handle* and a *symbol name*, and returns the address of the symbol to the caller. The handle directs the search for the symbol in the following manner:

- A handle can be returned from a `dlopen(3DL)` of a named object. This handle enables symbols to be obtained from the named object and the objects that define its dependency tree. A handle returned using the mode `RTLD_FIRST`, enables symbols to be obtained only from the named object.
- A handle can be returned from a `dlopen(3DL)` of a path name whose value is 0. This handle enables symbols to be obtained from the initiating object of the associated link-map and the objects that define its dependency tree. Typically, the initiating object is the dynamic executable. This handle also enables symbols to be obtained from any object obtained by a `dlopen(3DL)` with the `RTLD_GLOBAL` mode, on the associated link-map. A handle returned using the mode `RTLD_FIRST`, enables symbols to be obtained only from the initiating object of the associated link-map.
- The special handle `RTLD_DEFAULT` enables symbols to be obtained from the initiating object of the associated link-map and objects that define its dependency tree. This handle also enables symbols to be obtained from any object obtained by a `dlopen(3DL)` that belongs to the same group as the caller. Use of `RTLD_DEFAULT` follows the same model as used to resolve a symbolic relocation from the calling object.
- The special handle `RTLD_NEXT` enables symbols to be obtained from the next associated object on the callers link-map list.

In the following example, which is probably the most common, an application adds additional objects to its address space. The application then uses `dlsym(3DL)` to locate function or data symbols. The application then uses these symbols to call upon services provided in these new objects. The file `main.c` contains the following code:

```
#include <stdio.h>
#include <dlfcn.h>

main()
{
    void * handle;
    int * dptr, (* fptr)();

    if ((handle = dlopen("foo.so.1", RTLD_LAZY)) == NULL) {
        (void) printf("dlopen: %s\n", dlerror());
        exit (1);
    }

    if (((fptr = (int (*)())dlsym(handle, "foo")) == NULL) ||
        ((dptr = (int *)dlsym(handle, "bar")) == NULL)) {
        (void) printf("dlsym: %s\n", dlerror());
        exit (1);
    }

    return ((*fptr)(*dptr));
}
```

The symbols `foo` and `bar` are searched for in the file `foo.so.1`, followed by any dependencies that are associated with this file. The function `foo` is then called with the single argument `bar` as part of the `return()` statement.

The application `prog`, built using the above file `main.c`, contains the following dependencies:

```
$ ldd prog
    libdl.so.1 => /usr/lib/libdl.so.1
    libc.so.1 => /usr/lib/libc.so.1
```

If the file name specified in the `dlopen(3DL)` had the value `0`, the symbols `foo` and `bar` are searched for in `prog`, followed by `/usr/lib/libdl.so.1`, and finally `/usr/lib/libc.so.1`.

Once the `handle` has indicated the root at which to start a symbol search, the search mechanism follows the same model as described in "Symbol Lookup" on page 68.

If the required symbol cannot be located, `dlsym(3DL)` returns a `NULL` value. In this case, `dlerror(3DL)` can be used to indicate the true reason for the failure. In the following example, the application `prog` is unable to locate the symbol `bar`.

```
$ prog
dlsym: ld.so.1: main: fatal: bar: can't find symbol
```

Testing for Functionality

The special handle `RTLD_DEFAULT` enables an application to test for the existence of another symbol. The symbol search follows the same model as used to relocate the calling object. See “Default Symbol Lookup Model” on page 85. For example, if the application `prog` contained the following code fragment:

```
if ((fptr = (int (*)())dlsym(RTLD_DEFAULT, "foo")) != NULL)
    (*fptr)();
```

then `foo` is searched for in `prog`, followed by `/usr/lib/libdl.so.1`, and then `/usr/lib/libc.so.1`. If this code fragment was contained in the file `B.so.1` from the example shown in Figure 3-1, then the search for `foo` continues into `B.so.1` and then `C.so.1`.

This mechanism provides a robust and flexible alternative to the use of undefined weak references, discussed in “Weak Symbols” on page 45.

Using Interposition

The special handle `RTLD_NEXT` enables an application to locate the next symbol in a symbol scope. For example, if the application `prog` contained the following code fragment:

```
if ((fptr = (int (*)())dlsym(RTLD_NEXT, "foo")) == NULL) {
    (void) printf("dlsym: %s\n", dlerror());
    exit (1);
}

return ((*fptr)());
```

then `foo` is searched for in the shared objects associated with `prog`, which in this case are `/usr/lib/libdl.so.1` and then `/usr/lib/libc.so.1`. If this code fragment was contained in the file `B.so.1` from the example shown in Figure 3-1, then `foo` is searched for in the associated shared object `C.so.1` only.

Use of `RTLD_NEXT` provides a means to exploit symbol interposition. For example, a function within an object can be interposed upon by a preceding object, which can then augment the processing of the original function. For example, the following code fragment is placed in the shared object `malloc.so.1`:

```
#include <sys/types.h>
#include <dlfcn.h>
#include <stdio.h>

void *
malloc(size_t size)
{
    static void * (* fptr)() = 0;
    char          buffer[50];
```

```

    if (fptr == 0) {
        fptr = (void * (*)())dlsym(RTLD_NEXT, "malloc");
        if (fptr == NULL) {
            (void) printf("dlopen: %s\n", dlerror());
            return (0);
        }
    }

    (void) sprintf(buffer, "malloc: %#x bytes\n", size);
    (void) write(1, buffer, strlen(buffer));
    return ((*fptr)(size));
}

```

This shared object can be interposed before the system library `/usr/lib/libc.so.1` where `malloc(3C)` usually resides. Any calls to `malloc()` are now interposed upon before the original function is called to complete the allocation:

```

$ cc -o malloc.so.1 -G -K pic malloc.c
$ cc -o prog file1.o file2.o ..... -R. malloc.so.1
$ prog
malloc: 0x32 bytes
malloc: 0x14 bytes
.....

```

Alternatively, this same interposition can be achieved using the following:

```

$ cc -o malloc.so.1 -G -K pic malloc.c
$ cc -o prog main.c
$ LD_PRELOAD=./malloc.so.1 prog
malloc: 0x32 bytes
malloc: 0x14 bytes
.....

```

Note – Users of any interposition technique must be careful to handle any possibility of recursion. The previous example formats the diagnostic message using `sprintf(3C)`, instead of using `printf(3C)` directly, to avoid any recursion caused by `printf(3C)`'s possible use of `malloc(3C)`.

The use of `RTLD_NEXT` within a dynamic executable or preloaded object, provides a predictable and useful interposition technique. Be careful when using this technique in a generic object dependency, as the actual load order of objects is not always predictable.

Feature Checking

Dynamic objects built by the link-editor sometimes require new runtime linker features. The function `_check_rtd_feature()` can be used to check if the runtime features required for execution are supported by the running runtime linker. The runtime features currently identified are listed in Table 7-46.

Debugging Aids

A debugging library and `mdb(1)` module are provided with the Solaris linkers. The debugging library enables you to trace the runtime linking process in more detail. The `mdb(1)` module enables interactive process debugging.

Debugging Library

This debugging library helps you understand, or debug, the execution of applications and dependencies. Although the type of information displayed using this library is expected to remain constant, the exact format of the information might change slightly from release to release.

Some of the debugging output might be unfamiliar to those who do not have an intimate knowledge of the runtime linker. However, many aspects may be of general interest to you.

Debugging is enabled by using the environment variable `LD_DEBUG`. All debugging output is prefixed with the process identifier and by default is directed to the standard error. This environment variable must be augmented with one or more tokens to indicate the type of debugging required.

The tokens available with this debugging option can be displayed by using `LD_DEBUG=help`. Any dynamic executable can be used to solicit this information, as the process itself terminates following the display of the information. For example:

```
$ LD_DEBUG=help prog
11693:
11693:          For debugging the runtime linking of an application:
11693:          LD_DEBUG=token1,token2 prog
11693:          enables diagnostics to the stderr.  The additional
11693:          option:
11693:          LD_DEBUG_OUTPUT=file
11693:          redirects the diagnostics to an output file created
11593:          using the specified name and the process id as a
11693:          suffix.  All diagnostics are prepended with the
11693:          process id.
11693:
11693:
11693: audit      display runtime link-audit processing
11693: basic      provide basic trace information/warnings
11693: bindings   display symbol binding; detail flag shows
11693:            absolute:relative addresses
11693: detail     provide more information in conjunction with other
11693:            options
11693: files      display input file processing (files and libraries)
11693: help       display this help message
```

```

11693: libs      display library search paths
11693: move      display move section processing
11693: reloc     display relocation processing
11693: symbols   display symbol table processing
11693: tls       display TLS processing info
11693: unused    display unused/unreferenced files
11693: versions  display version processing

```

This example shows the options meaningful to the runtime linker. The exact options might differ from release to release.

The environment variable `LD_DEBUG_OUTPUT` can be used to specify an output file for use instead of the standard error. The process identifier is added as a suffix to the output file.

Debugging of secure applications is not allowed.

One of the most useful debugging options is to display the symbol bindings that occur at runtime. The following example uses a very trivial dynamic executable that has a dependency on two local shared objects.

```

$ cat bar.c
int bar = 10;
$ cc -o bar.so.1 -K pic -G bar.c

$ cat foo.c
foo(int data)
{
    return (data);
}
$ cc -o foo.so.1 -K pic -G foo.c

$ cat main.c
extern int  foo();
extern int  bar;

main()
{
    return (foo(bar));
}
$ cc -o prog main.c -R/tmp:. foo.so.1 bar.so.1

```

The runtime symbol bindings can be displayed by setting `LD_DEBUG=bindings`:

```

$ LD_DEBUG=bindings prog
11753: .....
11753: binding file=prog to file=./bar.so.1: symbol bar
11753: .....
11753: transferring control: prog
11753: .....
11753: binding file=prog to file=./foo.so.1: symbol foo
11753: .....

```

The symbol `bar`, which is required by an immediate relocation, is bound *before* the application gains control. Whereas the symbol `foo`, which is required by a lazy relocation, is bound *after* the application gains control when the function is first called. This demonstrates the default mode of lazy binding. If the environment variable `LD_BIND_NOW` is set, all symbol bindings occur before the application gains control.

Setting `LD_DEBUG=bindings,detail`, provides additional information regarding the real and relative addresses of the actual binding locations.

When the runtime linker performs a function relocation, it rewrites data associated with the functions `.plt` so that any subsequent calls will go directly to the function. The environment variable `LD_BIND_NOT` can be set to any value to prevent this data update. By using this variable together with the debugging request for detailed bindings, you can get a complete runtime account of all function binding. The output from this combination can be excessive, in which case the performance of the application is degraded.

You can use `LD_DEBUG` to display the various search paths used. For example, the search path mechanism used to locate any dependencies can be displayed by setting `LD_DEBUG=libs`.

```
$ LD_DEBUG=libs prog
11775:
11775: find object=foo.so.1; searching
11775:  search path=/tmp:. (RPATH from file prog)
11775:  trying path=/tmp/foo.so.1
11775:  trying path=./foo.so.1
11775:
11775: find object=bar.so.1; searching
11775:  search path=/tmp:. (RPATH from file prog)
11775:  trying path=/tmp/bar.so.1
11775:  trying path=./bar.so.1
11775: .....
```

The runpath recorded in the application `prog` affects the search for the two dependencies `foo.so.1` and `bar.so.1`.

In a similar manner, the search paths of each symbol lookup can be displayed by setting `LD_DEBUG=symbols`. If this is combined with a `bindings` request, you can obtain a complete picture of the symbol relocation process.

```
$ LD_DEBUG=bindings,symbols
11782: .....
11782: symbol=bar; lookup in file=./foo.so.1 [ ELF ]
11782: symbol=bar; lookup in file=./bar.so.1 [ ELF ]
11782: binding file=prog to file=./bar.so.1: symbol bar
11782: .....
11782: transferring control: prog
11782: .....
11782: symbol=foo; lookup in file=prog [ ELF ]
11782: symbol=foo; lookup in file=./foo.so.1 [ ELF ]
11782: binding file=prog to file=./foo.so.1: symbol foo
11782: .....
```

In the previous example, the symbol bar is not searched for in the application prog. This is due to an optimization used when processing copy relocations. See “Copy Relocations” on page 122 for more details of this relocation type.

Debugger Module

The debugger module provides a set of dcmds and walkers that can be loaded under mdb(1). This module can be used to inspect various internal data structures of the runtime linker. Much of this information requires familiarity with the internals of the runtime linker, and may change from release to release. However, some elements of these data structures reveal the basic components of a dynamically linked process and may aid general debugging.

The following examples show some simple scenarios of using mdb(1) with the runtime linker debugger module.

```
$ cat main.c
#include <dlfcn.h>

int main()
{
    void * handle;
    void (* fptr)();

    if ((handle = dlopen("foo.so.1", RTLD_LAZY)) == NULL)
        return (1);

    if ((fptr = (void (*)())dlsym(handle, "foo")) == NULL)
        return (1);

    (*fptr)();
    return (0);
}
$ cc -o main main.c -R. -ldl
```

If mdb(1) has not automatically loaded the debugger module, ld.so, explicitly do so. The capabilities of the debugger module can then be inspected.

```
$ mdb main
> ::load ld.so
> ::dmods -l ld.so

ld.so
-----
dcmd Bind                - Display a Binding descriptor
dcmd Callers             - Display Rt_map CALLERS binding descriptors
dcmd Depends            - Display Rt_map DEPENDS binding descriptors
dcmd ElfDyn             - Display Elf_Dyn entry
dcmd ElfEhdr            - Display Elf_Ehdr entry
dcmd ElfPhdr           - Display Elf_Phdr entry
dcmd Groups             - Display Rt_map GROUPS group handles
```



```

dcmd GrpDesc           - Display a Group Descriptor
dcmd GrpHdl           - Display a Group Handle
dcmd Handles          - Display Rt_map HANDLES group descriptors
dcmd List             - Display entries in a List
dcmd ListRtmap        - Display a List of Rt_Map's
dcmd Lm_list          - Display ld.so.1 Lm_list structure
dcmd Rt_map           - Display ld.so.1 Rt_map structure
dcmd Rt_maps          - Display list of Rt_map structures
walk List             - Walk a List
walk Rt_maps          - Walk a List of Rt_maps
> ::bp main
> :r

```

Each dynamic object within a process is expressed as a link-map, `Rt_map`, which is maintained on a link-map list. All link-maps for the process can be displayed with `Rt_maps`.

```

> ::Rt_maps
Link-map lists (dynlm_list): 0xffbfe0d0
-----
Lm_list: 0xff3f6f60 (LM_ID_BASE)
-----
Link-map*  ADDR()      NAME()
-----
0xff3f9040 0x00010000 main
0xff3f9460 0xff3b0000 /lib/libdl.so.1
0xff3f977c 0xff280000 /lib/libc.so.1
-----
Lm_list: 0xff3f6f88 (LM_ID_LDSO)
-----
0xff3f8cc0 0xff3c0000 /lib/ld.so.1

```

An individual link-map can be displayed with `Rt_map`.

```

> 0xff3f9040::Rt_map
Rt_map located at: 0xff3f9040
NAME: main
ADDR: 0x00010000      DYN: 0x000207bc
NEXT: 0xff3f9460     PREV: 0x00000000
FCT: 0xff3f6f18     TLSMODID: 0
INIT: 0x00010710     FINI: 0x0001071c
GROUPS: 0x00000000   HANDLES: 0x00000000
DEPENDS: 0xff3f96e8  CALLERS: 0x00000000
.....

```

The object's `.dynamic` section can be displayed with the `ElfDyn` `dcmd`. The following example shows the first 4 entries.

```

> 0x000207bc,4::ElfDyn
Elf_Dyn located at: 0x207bc
0x207bc NEEDED 0x0000010f
Elf_Dyn located at: 0x207c4
0x207c4 NEEDED 0x00000124
Elf_Dyn located at: 0x207cc
0x207cc INIT 0x00010710

```

```
Elf_Dyn located at: 0x207d4
0x207d4 FINI 0x0001071c
```

mdb(1) is also very useful for setting deferred break points. In this example, it might be useful to put a break point on the function `foo()`. However, until the `dlopen(3DL)` of `foo.so.1` occurs, this symbol isn't known to the debugger. A deferred break point instructs the debugger to set a real breakpoint when the dynamic object is loaded.

```
> ::bp foo.so.1`foo
> :c
> mdb: You've got symbols!
> mdb: stop at foo.so.1`foo
mdb: target stopped at:
foo.so.1`foo: save %sp, -0x68, %sp
```

At this point, new objects have been loaded:

```
> *ld.so`lml_main::Rt_maps
Link-map* ADDR() NAME()
-----
0xff3f9040 0x00010000 main
0xff3f9460 0xff3b0000 /lib/libdl.so.1
0xff3f977c 0xff280000 /lib/libc.so.1
0xff3f9ca4 0xff380000 ./foo.so.1
0xff37006c 0xff260000 ./bar.so.1
```

The link-map for `foo.so.1` shows the handle returned by `dlopen(3DL)`. You can expand this structure using `Handles`.

```
> 0xff3f9ca4::Handles -v
HANDLES for ./foo.so.1
-----
HANDLE: 0xff3f9f60 Alist[used 1: total 1]
-----
Group Handle located at: 0xff3f9f28
-----
owner: ./foo.so.1
flags: 0x00000000 [ 0 ]
refcnt: 1 depends: 0xff3f9fa0 Alist[used 2: total 4]
-----
Group Descriptor located at: 0xff3f9fac
depend: 0xff3f9ca4 ./foo.so.1
flags: 0x00000003 [ AVAIL-TO-DLSYM,ADD-DEPENDENCIES ]
-----
Group Descriptor located at: 0xff3f9fd8
depend: 0xff37006c ./bar.so.1
flags: 0x00000003 [ AVAIL-TO-DLSYM,ADD-DEPENDENCIES ]
```

The dependencies of a handle are a list of link-maps that represent the objects of the handle that can satisfy a `dlsym(3DL)` request. In this case, the dependencies are `foo.so.1` and `bar.so.1`.

Note – The above examples provide a basic guide to the debugger module capabilities, but the exact commands, usage, and output may change from release to release. Refer to usage and help information for the exact capabilities available on your system.

Shared Objects

Shared objects are one form of output created by the link-editor and are generated by specifying the `-G` option. In the following example, the shared object `libfoo.so.1` is generated from the input file `foo.c`.

```
$ cc -o libfoo.so.1 -G -K pic foo.c
```

A shared object is an *indivisible* unit that is generated from one or more relocatable objects. Shared objects can be bound with dynamic executables to form a runnable process. As their name implies, shared objects can be shared by more than one application. Because of this potentially far-reaching effect, this chapter describes this form of link-editor output in greater depth than has been covered in previous chapters.

For a shared object to be bound to a dynamic executable or another shared object, it must first be available to the link-edit of the required output file. During this link-edit, any input shared objects are interpreted as if they had been added to the logical address space of the output file being produced. All the functionality of the shared object is made available to the output file.

Any input shared objects become dependencies of this output file. A small amount of bookkeeping information is maintained within the output file to describe these dependencies. The runtime linker interprets this information and completes the processing of these shared objects as part of creating a runnable process.

The following sections expand upon the use of shared objects within the compilation and runtime environments. These environments are introduced in “Runtime Linking” on page 21.

Naming Conventions

Neither the link-editor nor the runtime linker interprets any file by virtue of its file name. All files are inspected to determine their ELF type (see “ELF Header” on page 180). This information enables the link-editor to deduce the processing requirements of the file. However, shared objects usually follow one of two naming conventions, depending on whether they are being used as part of the compilation environment or the runtime environment.

When used as part of the compilation environment, shared objects are read and processed by the link-editor. Although these shared objects can be specified by explicit file names as part of the command passed to the link-editor, the `-l` option is usually used to take advantage of the link-editor’s library search capabilities. See “Shared Object Processing” on page 30.

A shared object that is applicable to this link-editor processing, should be designated with the prefix `lib` and the suffix `.so`. For example, `/usr/lib/libc.so` is the shared object representation of the standard C library made available to the compilation environment. By convention, 64-bit shared objects are placed in a subdirectory of the `lib` directory called `64`. For example, the 64-bit counterpart of `/usr/lib/libc.so.1`, is `/usr/lib/64/libc.so.1`.

When used as part of the runtime environment, shared objects are read and processed by the runtime linker. To allow for change in the exported interface of the shared object over a series of software releases, provide the shared object as a *versioned* file name.

A versioned file name commonly takes the form of a `.so` suffix followed by a version number. For example, `/usr/lib/libc.so.1` is the shared object representation of version *one* of the standard C library made available to the runtime environment.

If a shared object is never intended for use within a compilation environment, its name might drop the conventional `lib` prefix. Examples of shared objects that fall into this category are those used solely with `dlopen(3DL)`. A suffix of `.so` is still recommended to indicate the actual file type. In additions, a version number is strongly recommended to provide for the correct binding of the shared object across a series of software releases. Chapter 5 describes versioning in more detail.

Note – The shared object name used in a `dlopen(3DL)` is usually represented as a *simple* file name, that has no `’/’` in the name. The runtime linker can then use a set of rules to locate the actual file. See “Loading Additional Objects” on page 73 for more details.

Recording a Shared Object Name

The recording of a dependency in a dynamic executable or shared object will, by default, be the file name of the associated shared object as it is referenced by the link-editor. For example, the following dynamic executables, that are built against the same shared object `libfoo.so`, result in different interpretations of the same dependency.

```
$ cc -o ../tmp/libfoo.so -G foo.o
$ cc -o prog main.o -L../tmp -lfoo
$ dump -Lv prog | grep NEEDED
[1]      NEEDED   libfoo.so

$ cc -o prog main.o ../tmp/libfoo.so
$ dump -Lv prog | grep NEEDED
[1]      NEEDED   ../tmp/libfoo.so

$ cc -o prog main.o /usr/tmp/libfoo.so
$ dump -Lv prog | grep NEEDED
[1]      NEEDED   /usr/tmp/libfoo.so
```

As these examples show, this mechanism of recording dependencies can result in inconsistencies due to different compilation techniques. Also, the location of a shared object as referenced during the link-edit might differ from the eventual location of the shared object on an installed system. To provide a more consistent means of specifying dependencies, shared objects can record within themselves the file name by which they should be referenced at runtime.

During the link-edit of a shared object, its runtime name can be recorded within the shared object itself by using the `-h` option. In the following example, the shared object's runtime name `libfoo.so.1`, is recorded within the file itself. This identification is known as an *soname*.

```
$ cc -o ../tmp/libfoo.so -G -K pic -h libfoo.so.1 foo.c
```

The following example shows how the soname recording can be displayed using `dump(1)` and referring to the entry that has the `SONAME` tag.

```
$ dump -Lvp ../tmp/libfoo.so

../tmp/libfoo.so:
[INDEX] Tag      Value
[1]      SONAME   libfoo.so.1
.....
```

When the link-editor processes a shared object that contains an soname, this is the name that is recorded as a dependency within the output file being generated.

If this new version of `libfoo.so` is used during the creation of the dynamic executable `prog` from the previous example, all three methods of creating the executable result in the same dependency recording.

```

$ cc -o prog main.o -L./tmp -lfoo
$ dump -Lv prog | grep NEEDED
[1]      NEEDED  libfoo.so.1

$ cc -o prog main.o ../tmp/libfoo.so
$ dump -Lv prog | grep NEEDED
[1]      NEEDED  libfoo.so.1

$ cc -o prog main.o /usr/tmp/libfoo.so
$ dump -Lv prog | grep NEEDED
[1]      NEEDED  libfoo.so.1

```

In the previous examples, the `-h` option is used to specify a simple file name, that has no `'/'` in the name. This convention enables the runtime linker to use a set of rules to locate the actual file. See “Locating Shared Object Dependencies” on page 64 for more details.

Inclusion of Shared Objects in Archives

The mechanism of recording an soname within a shared object is essential if the shared object is ever processed from an archive library.

An archive can be built from one or more shared objects and then used to generate a dynamic executable or shared object. Shared objects can be extracted from the archive to satisfy the requirements of the link-edit. Unlike the processing of relocatable objects, which are concatenated to the output file being created, any shared objects extracted from the archive will be recorded as dependencies. See “Archive Processing” on page 29 for more details on the criteria for archive extraction.

The name of an archive member is constructed by the link-editor and is a concatenation of the archive name and the object within the archive. For example:

```

$ cc -o libfoo.so.1 -G -K pic foo.c
$ ar -r libfoo.a libfoo.so.1
$ cc -o main main.o libfoo.a
$ dump -Lv main | grep NEEDED
[1]      NEEDED  libfoo.a(libfoo.so.1)

```

Because a file with this concatenated name is unlikely to exist at runtime, providing an soname within the shared object is the only means of generating a meaningful runtime file name for the dependency.

Note – The runtime linker does not extract objects from archives. Therefore, in the previous example, the required shared object dependencies must be extracted from the archive and made available to the runtime environment.

Recorded Name Conflicts

When shared objects are used to create a dynamic executable or another shared object, the link-editor performs several consistency checks. These checks ensure that any dependency names recorded in the output file are unique.

Conflicts in dependency names can occur if two shared objects used as input files to a link-edit both contain the same soname. For example:

```
$ cc -o libfoo.so -G -K pic -h libsame.so.1 foo.c
$ cc -o libbar.so -G -K pic -h libsame.so.1 bar.c
$ cc -o prog main.o -L. -lfoo -lbar
ld: fatal: recording name conflict: file './libfoo.so' and \
      file './libbar.so' provide identical dependency names: libsame.so.1
ld: fatal: File processing errors. No output written to prog
```

A similar error condition occurs if the file name of a shared object that does not have a recorded soname matches the soname of another shared object used during the same link-edit.

If the runtime name of a shared object being generated matches one of its dependencies, the link-editor also reports a name conflict

```
$ cc -o libbar.so -G -K pic -h libsame.so.1 bar.c -L. -lfoo
ld: fatal: recording name conflict: file './libfoo.so' and \
      -h option provide identical dependency names: libsame.so.1
ld: fatal: File processing errors. No output written to libbar.so
```

Shared Objects With Dependencies

Shared objects can have their own dependencies. The search rules used by the runtime linker to locate shared object dependencies are covered in “Directories Searched by the Runtime Linker” on page 64. If a shared object does not reside in the default directory `/usr/lib` (for 32-bit objects), or `/usr/lib/64` (for 64-bit objects), then the runtime linker must explicitly be told where to look. The preferred mechanism of indicating any requirement of this kind is to record a runpath in the object that has the dependencies by using the link-editor’s `-R` option.

In the following example, the shared object `libfoo.so` has a dependency on `libbar.so`, which is expected to reside in the directory `/home/me/lib` at runtime or, failing that, in the default location.

```
$ cc -o libbar.so -G -K pic bar.c
$ cc -o libfoo.so -G -K pic foo.c -R/home/me/lib -L. -lbar
$ dump -Lv libfoo.so
```

```
libfoo.so:
```

```

**** DYNAMIC SECTION INFORMATION ****
.dynamic:
[INDEX] Tag      Value
[1]      NEEDED  libbar.so
[2]      RUNPATH /home/me/lib
.....

```

The shared object is responsible for specifying any runpath required to locate its dependencies. Any runpath specified in the dynamic executable is only used to locate the dependencies of the dynamic executable. These runpaths are not used to locate any dependencies of the shared objects.

The environment variable `LD_LIBRARY_PATH` has a more global scope. Any path names specified using this variable are used by the runtime linker to search for any shared object dependencies. Although useful as a temporary mechanism that influences the runtime linker's search path, the use of this environment variable is strongly discouraged in production software. See "Directories Searched by the Runtime Linker" on page 64 for a more extensive discussion.

Dependency Ordering

When dynamic executables and shared objects have dependencies on the same common shared objects, the order in which the objects are processed can become less predictable.

For example, assume a shared object developer generates `libfoo.so.1` with the following dependencies:

```

$ ldd libfoo.so.1
    libA.so.1 =>      ./libA.so.1
    libB.so.1 =>      ./libB.so.1
    libC.so.1 =>      ./libC.so.1

```

If you create a dynamic executable `prog`, using this shared object, and define an explicit dependency on `libC.so.1`, the resulting shared object order will be:

```

$ cc -o prog main.c -R. -L. -lC -lfoo
$ ldd prog
    libC.so.1 =>      ./libC.so.1
    libfoo.so.1 =>    ./libfoo.so.1
    libA.so.1 =>      ./libA.so.1
    libB.so.1 =>      ./libB.so.1

```

Any requirement on the order of processing the shared object `libfoo.so.1` dependencies would be compromised by the construction of the dynamic executable `prog`.

Developers who place special emphasis on symbol interposition and `.init` section processing should be aware of this potential change in shared object processing order.

Shared Objects as Filters

Shared objects can be defined to act as *filters* for the interfaces they provide. This technique involves associating filter interfaces with an alternative shared object. The alternative shared object supplies symbol definitions at runtime, and is referred to as a *filtee*. A filtee is built in the same manner as any shared object is built.

Filtering provides a mechanism of abstracting the compilation environment from the runtime environment. During a link-edit, a symbol reference that binds to a filter interface, is resolved to the filters symbol definition. At runtime, a symbol reference that binds to a filter interface, can be redirected to an alternative shared object. A shared object can define all the interfaces it offers as filters.

Two forms of filtering exist.

Standard filtering

This filtering requires only a symbol table entry for the interface being filtered. At runtime, the implementation of a filter symbol definition must be provided from a filtee.

Interfaces are defined to act as standard filters by using the link-editor's `-F` flag. This flag is qualified with the name of one or more filtees that must supply the symbol definition at runtime.

If a filtee cannot be processed at runtime, or a filter symbol cannot be located within the filtee, the filtee is skipped for this symbol lookup.

Auxiliary filtering

This filtering provides a similar mechanism to standard filtering, except the filter provides a fall back implementation corresponding to the auxiliary filter interfaces. At runtime, the implementation of the symbol definition can be provided from a filtee.

Interfaces are defined to act as auxiliary filters by using the link-editor's `-f` flag. This flag is qualified with the name of one or more filtees that may supply the symbol definition at runtime.

If a filtee cannot be processed at runtime, or an auxiliary filter symbol cannot be located within the filtee, the implementation of the symbol within the filter is used.

Generating Standard Filters

To generate a standard filter, you first define a filtee on which the filtering is applied. The following example builds a filtee `filtee.so.1`, that supplies the symbols `foo` and `bar`.

```
$ cat filtee.c
char * bar = "defined in filtee";

char * foo()
{
    return("defined in filtee");
}
$ cc -o filtee.so.1 -G -K pic filtee.c
```

To declare all the interfaces offered by a shared object to be filters, the shared object is defined a filter by using the link-editor's `-F` flag.

In the following example, the shared object `filter.so.1`, is defined to be a filter. This shared object offers the symbols `foo` and `bar`, and indicates it is a filter on the filtee `filtee.so.1`. In this example, the environment variable `LD_OPTIONS` is used to circumvent the compiler driver from interpreting the `-F` option.

```
$ cat filter.c
char * bar = 0;

char * foo()
{
    return (0);
}
$ LD_OPTIONS='-F filtee.so.1' \
cc -o filter.so.1 -G -K pic -h filter.so.1 -R. filter.c
$ elfdump -d filter.so.1 | egrep "SONAME|FILTER"
    [2] SONAME          0xee   filter.so.1
    [3] FILTER         0xfb   filtee.so.1
```

If the link-editor references the standard filter `filtee.so.1`, to create a dynamic executable or shared object, it uses information from the symbol table of the filter to satisfy any symbol resolution. However, at runtime, any reference to the symbols of the filter result in the additional loading of the filtee `filtee.so.1`. The runtime linker uses the filtee to resolve any symbols defined by `filter.so.1`. If the filtee is not found, or a filter symbol is not found in the filtee, the filter is skipped for this symbol lookup.

For example, the following dynamic executable `prog`, references the symbols `foo` and `bar`, which are resolved during link-edit from the filter `filter.so.1`. The execution of `prog` results in `foo` and `bar` being obtained from the filtee `filtee.so.1`, *not* from the filter `filter.so.1`.

```
$ cat main.c
extern char * bar, * foo();

main()
```

```

{
    (void) printf("foo is %s: bar is %s\n", foo(), bar);
}
$ cc -o prog main.c -R. filter.so.1
$ prog
foo is defined in filtee: bar is defined in filtee

```

In these examples, the filtee `filtee.so.1` is uniquely associated to the filter, and is not available to satisfy symbol lookup from any other objects that might be loaded as a consequence of executing `prog`.

Standard filters provide a mechanism for defining a subset interface of an existing shared object, or an interface group spanning a number of existing shared objects. Standard filters also provide a means of redirecting an interface to its implementation. Several standard filters are used in the Solaris operating environment.

The `/usr/lib/libsys.so.1` filter provides a subset of the standard C library `/usr/lib/libc.so.1`. This subset represents the ABI-conforming functions and data items that reside in the C library that must be imported by a conforming application.

The `/usr/lib/libxnet.so.1` filter uses multiple filtees. This library provides socket and XTI interfaces from `/usr/lib/libsocket.so.1`, `/usr/lib/libnsl.so.1`, and `/usr/lib/libc.so.1`.

The `/usr/lib/libdl.so.1` filter defines the user interface to the runtime linker. This interface provides an abstraction between the symbols referenced in a compilation environment from `libdl.so.1` and the actual implementation binding produced within the runtime environment from `ld.so.1(1)`.

Because the code in a standard filter is never referenced at runtime, adding content to any functions defined as filters is redundant. Any filter code might require relocation, which would result in an unnecessary overhead when processing the filter at runtime. Functions are best defined as empty routines, or directly from a `mapfile`. See “Defining Additional Symbols” on page 46.

Note – The link-editor uses the ELF class of the first relocatable file that is processed to govern the class of object that is created. Use the link-editor’s `-64` option to create a 64-bit filter solely from a `mapfile`.

When generating data symbols within a filter, always initialize the data items. The resulting data definition insures that references are established correctly from dynamic executables.

Some of the more complex symbol resolutions carried out by the link-editor require knowledge of a symbol's attributes, including the symbol's size. Therefore, you should generate the symbols in the filter so that their attributes match those of the symbols in the filtee. This ensures that the link-editing process analyzes the filter in a manner compatible with the symbol definitions used at runtime. See "Symbol Resolution" on page 38.

Generating Auxiliary Filters

To generate an auxiliary filter, you first define a filtee on which the filtering is applied. The following example builds a filtee `filtee.so.1`, that supplies the symbol `foo`.

```
$ cat filtee.c
char * foo()
{
    return("defined in filtee");
}
$ cc -o filtee.so.1 -G -K pic filtee.c
```

To declare all of the interfaces offered by a shared object to be auxiliary filters, the shared object is defined to be an auxiliary filter using the link-editor's `-f` flag.

In the following example, the shared object `filter.so.1` is defined to be an auxiliary filter. This shared object offers the symbols `foo` and `bar`, and indicates it is an auxiliary filter on the filtee `filtee.so.1`. In this example, the environment variable `LD_OPTIONS` is used to circumvent the compiler driver from interpreting the `-f` option.

```
$ cat filter.c
char * bar = "defined in filter";

char * foo()
{
    return ("defined in filter");
}
$ LD_OPTIONS='-f libbar.so.1' \
cc -o filter.so.1 -G -K pic -h filter.so.1 -R. filter.c
$ elfdump -d filter.so.1 | egrep "SONAME|AUXILIARY"
[2] SONAME          0xee  filter.so.1
[3] AUXILIARY       0xfb  filtee.so.1
```

If the link-editor references the auxiliary filter `filter.so.1`, to create a dynamic executable or shared object, it uses information from the symbol table of the filter to satisfy any symbol resolution. However, at runtime, any reference to the symbols of the filter result in a search for the filtee `filtee.so.1`. If this filtee is found, the runtime linker uses the filtee to resolve any symbols defined by `filter.so.1`. If the filtee is not found, or a symbol from the filter is not found in the filtee, then the original symbol within the filter is used.

For example, the following dynamic executable `prog`, references the symbols `foo` and `bar`, which are resolved during link-edit from the filter `filter.so.1`. The execution of `prog` results in `foo` being obtained from the filtee `filtee.so.1`, *not* from the filter `filter.so.1`. However, `bar` is obtained from the filter `filter.so.1`, as this symbol has no alternative definition in the filtee `filtee.so.1`.

```
$ cat main.c
extern char * bar, * foo();

main()
{
    (void) printf("foo is %s: bar is %s\n", foo(), bar);
}
$ cc -o prog main.c -R. filter.so.1
$ prog
foo is defined in filtee: bar is defined in filter
```

In these examples, the filtee `filtee.so.1` is uniquely associated to the filter, and is not available to satisfy symbol lookup from any other objects that might be loaded as a consequence of executing `prog`.

Auxiliary filters provide a mechanism for defining an alternative interface of an existing shared object. This mechanism is used in the Solaris operating environment to provide optimized functionality within platform specific shared objects. See “Instruction Set Specific Shared Objects” on page 321, and “System Specific Shared Objects” on page 323 for examples.

Note – The environment variable `LD_NOAUXFLTR` can be set to disable the runtime linker's auxiliary filter processing. Because auxiliary filters are frequently employed to provide platform specific optimizations, this option can be useful in evaluating filtee use and their performance impact.

Filtee Processing

The runtime linker's processing of a filter defers the loading of a filtee until a filter symbol is referenced. This implementation is analogous to the filter performing a `dlopen(3DL)`, using mode `RTLD_LOCAL`, on each of its filtees as they are required. This implementation accounts for differences in dependency reporting that can be produced by tools such as `ldd(1)`.

The link-editor's `-z loadfltr` option can be used when creating a filter to cause the immediate processing of its filtees at runtime. In addition, the immediate processing of any filtees within a process can be triggered by setting the `LD_LOADFLTR` environment variable to any value.

Performance Considerations

A shared object can be used by multiple applications within the same system. The performance of a shared object affects the applications that use it and the system as a whole.

Although the actual code within a shared object will directly affect the performance of a running process, the performance issues focused upon here target the runtime processing of the shared object itself. The following sections investigate this processing in more detail by looking at aspects such as text size and purity, together with relocation overhead.

Analyzing Files

Various tools are available to analyze the contents of an ELF file. To display the size of a file use the `size(1)` command. For example:

```
$ size -x libfoo.so.1
59c + 10c + 20 = 0x6c8

$ size -xf libfoo.so.1
..... + 1c(.init) + ac(.text) + c(.fini) + 4(.rodata) + \
..... + 18(.data) + 20(.bss) .....
```

The first example indicates the size of the shared objects *text*, *data*, and *bss*, a categorization used in previous releases of the SunOS operating system.

The ELF format provides a finer granularity for expressing data within a file by organizing the data into *sections*. The second example displays the size of each of the file's loadable sections.

Sections are allocated to units known as *segments*, some of which describe how portions of a file are mapped into memory (see the `mmap(2)` man page). These loadable segments can be displayed by using the `dump(1)` command and examining the `LOAD` entries. For example:

```
$ dump -ov libfoo.so.1

libfoo.so.1:
***** PROGRAM EXECUTION HEADER *****
Type      Offset      Vaddr      Paddr
Filesz    Memsz      Flags      Align

LOAD      0x94        0x94        0x0
0x59c     0x59c      r-x        0x10000

LOAD      0x630      0x10630     0x0
```



```
0x10c      0x12c      rwx      0x1000
```

There are two loadable segments in the shared object `libfoo.so.1`, commonly referred to as the *text* and *data* segments. The text segment is mapped to allow reading and execution of its contents (*r-x*), whereas the data segment is mapped to also allow its contents to be modified (*rwX*). The memory size (`Memsz`) of the data segment differs from the file size (`Filesz`). This difference accounts for the `.bss` section, which is part of the data segment, and is dynamically created when the segment is loaded.

Programmers usually think of a file in terms of the symbols that define the functions and data elements within their code. These symbols can be displayed using `nm(1)`. For example:

```
$ nm -x libfoo.so.1
```

[Index]	Value	Size	Type	Bind	Other	Shndx	Name
[39]	0x00000538	0x00000000	FUNC	GLOB	0x0	7	__init
[40]	0x00000588	0x00000034	FUNC	GLOB	0x0	8	foo
[41]	0x00000600	0x00000000	FUNC	GLOB	0x0	9	__fini
[42]	0x00010688	0x00000010	OBJT	GLOB	0x0	13	data
[43]	0x0001073c	0x00000020	OBJT	GLOB	0x0	16	bss

The section that contains a symbol can be determined by referencing the section index (`Shndx`) field from the symbol table and by using `dump(1)` to display the sections within the file. For example:

```
$ dump -hv libfoo.so.1
```

```
libfoo.so.1:
**** SECTION HEADER TABLE ****
[No]  Type   Flags  Addr      Offset    Size     Name
.....
[7]   PBIT   -AI    0x538     0x538     0x1c    .init
[8]   PBIT   -AI    0x554     0x554     0xac    .text
[9]   PBIT   -AI    0x600     0x600     0xc     .fini
.....
[13]  PBIT   WA-    0x10688   0x688     0x18    .data
[16]  NOBI   WA-    0x1073c   0x73c     0x20    .bss
.....
```

The output from both the previous `nm(1)` and `dump(1)` examples shows the association of the functions `__init`, `foo`, and `__fini` to the sections `.init`, `.text` and `.fini`. These sections, because of their read-only nature, are part of the *text* segment.

Similarly, the data arrays `data`, and `bss` are associated with the sections `.data` and `.bss` respectively. These sections, because of their writable nature, are part of the *data* segment.

Note – The previous dump(1) display has been simplified for this example.

Underlying System

When an application is built using a shared object, the entire loadable contents of the object are mapped into the virtual address space of that process at runtime. Each process that uses a shared object starts by referencing a single copy of the shared object in memory.

Relocations within the shared object are processed to bind symbolic references to their appropriate definitions. This results in the calculation of true virtual addresses that could not be derived at the time the shared object was generated by the link-editor. These relocations usually result in updates to entries within the process's data segments.

The memory management scheme underlying the dynamic linking of shared objects shares memory among processes at the granularity of a page. Memory pages can be shared as long as they are not modified at runtime. If a process writes to a page of a shared object when writing a data item, or relocating a reference to a shared object, it generates a private copy of that page. This private copy will have no effect on other users of the shared object. However, this page has lost any benefit of sharing between other processes. Text pages that become modified in this manner are referred to as *impure*.

The segments of a shared object that are mapped into memory fall into two basic categories; the *text* segment, which is read-only, and the *data* segment, which is read-write. See "Analyzing Files" on page 112 on how to obtain this information from an ELF file. An overriding goal when developing a shared object is to maximize the text segment and minimize the data segment. This optimizes the amount of code sharing while reducing the amount of processing needed to initialize and use a shared object. The following sections present mechanisms that can help achieve this goal.

Lazy Loading of Dynamic Dependencies

You can defer the loading of a shared object dependency until the dependency is first referenced by establishing the object as lazy loadable. See "Lazy Loading of Dynamic Dependencies" on page 74.

For small applications a typical thread of execution may reference all the applications dependencies. The application loads all of its dependencies whether they are defined lazy loadable or not. However, under lazy loading, dependency processing may be deferred from process startup and spread throughout the process's execution.

For applications with many dependencies, lazy loading often results in some dependencies not being loaded at all. These dependencies are those not referenced for the particular thread of execution.

Position-Independent Code

The compiler can generate position-independent code under the `-Kpic` option. Whereas the code within a dynamic executable is usually tied to a fixed address in memory, position-independent code can be loaded anywhere in the address space of a process. Because the code is not tied to a specific address, it will execute correctly without page modification at a different address in each process that uses it. This code creates programs that require the smallest amount of page modification at runtime.

When you use position-independent code, relocatable references are generated as an indirection that use data in the shared object's data segment. The text segment code remains read-only, and all relocation updates are applied to corresponding entries within the data segment. See "Global Offset Table (Processor-Specific)" on page 260 and "Procedure Linkage Table (Processor-Specific)" on page 261 for more details on the use of these two sections.

If a shared object is built from code that is not position-independent, the text segment will usually require a large number of relocations to be performed at runtime. Although the runtime linker is equipped to handle this, the system overhead this creates can cause serious performance degradation.

You can identify a shared object that requires relocations against its text segment. Use `dump(1)` and inspect the output for any `TEXTREL` entry. For example:

```
$ cc -o libfoo.so.1 -G -R. foo.c
$ dump -Lv libfoo.so.1 | grep TEXTREL
[9]      TEXTREL  0
```

Note – The value of the `TEXTREL` entry is irrelevant. Its presence in a shared object indicates that text relocations exist.

To prevent the creation of a shared object that contains text relocations use the link-editor's `-z text` flag. This flag causes the link-editor to generate diagnostics indicating the source of any non-position-independent code used as input. Such code results in a failure to generate the intended shared object. For example:

```
$ cc -o libfoo.so.1 -z text -G -R. foo.c
Text relocation remains          referenced
      against symbol              offset    in file
foo                                0x0      foo.o
bar                                0x8      foo.o
ld: fatal: relocations remain against allocatable but \
```

non-writable sections

Two relocations are generated against the text segment because of the non-position-independent code generated from the file `foo.o`. Where possible, these diagnostics indicate any symbolic references that are required to carry out the relocations. In this case, the relocations are against the symbols `foo` and `bar`.

Another common cause of creating text relocations when generating a shared object is by including hand-written assembler code that has not been coded with the appropriate position-independent prototypes.

Note – You may want to experiment with some simple source files to determine coding sequences that enable position-independence. Use the compiler's ability to generate intermediate assembler output.

SPARC: `-K pic` and `-K PIC` Options

For SPARC binaries, a subtle difference between the `-K pic` option and an alternative `-K PIC` option affects references to global offset table entries. See “Global Offset Table (Processor-Specific)” on page 260.

The global offset table is an array of pointers, the size of whose entries are constant for 32-bit (4-bytes) and 64-bit (8-bytes). The following code sequence makes reference to an entry under `-K pic`:

```
ld    [%l7 + j], %o0    ! load &j into %o0
```

Where `%l7` is the precomputed value of the symbol `_GLOBAL_OFFSET_TABLE_` of the object making the reference.

This code sequence provides a 13-bit displacement constant for the global offset table entry. This displacement therefore provides for 2048 unique entries for 32-bit objects, and 1024 unique entries for 64-bit objects. If an object is built that requires more than the available number of entries, the link-editor produces a fatal error:

```
$ cc -K pic -G -o lobfoo.so.1 a.o b.o ... z.o
ld: fatal: too many symbols require 'small' PIC references:
      have 2050, maximum 2048 -- recompile some modules -K PIC.
```

To overcome this error condition, compile some of the input relocatable objects with the `-K PIC` option. This option provides a 32-bit constant for the global offset table entry:

```
sethi %hi(j), %g1
or    %g1, %lo(j), %g1    ! get 32-bit constant GOT offset
ld    [%l7 + %g1], %o0    ! load &j into %o0
```

You can investigate the global offset table requirements of an object using `elfdump(1)` with the `-G` option. You can also examine the processing of these entries during a link-edit using the link-editors debugging tokens `-D got, detail`.

Ideally, frequently accessed data items benefit from using the `-K pic` model. You can reference a single entry using both models. However, determining which relocatable objects should be compiled with either option can be time consuming, and the performance improvement realized small. A recompilation of all relocatable objects with the `-K PIC` option is typically easier.

Remove Unused Material

The inclusion of functions and data that are not used by the object being built, is wasteful. This material bloats the object, which can result in unnecessary relocation overhead and associated paging activity. References to unused dependencies are also wasteful. These references result in the unnecessary loading and processing of other shared objects.

Unused sections are displayed during a link-edit when using the link-editors debugging token `-D unused`. Sections identified as unused should be removed from the link-edit, or eliminated using the link-editors `-z ignore` option.

The link-editor identifies a section from a relocatable object as unused if:

- The section is allocatable
- No other sections bind to (relocate) to this section
- The section provides no global symbols

You can improve the link-editors ability to eliminate sections by defining the shared objects external interfaces. By defining an interface, global symbols that are not defined as part of the interface are reduced to locals. These reduced symbols, if unreferenced from other objects, are now clearly identified as candidates for elimination.

Individual functions and data variables can be eliminated by the link-editor if these items are assigned to their own sections. This section refinement is achieved using compiler options such as `-xF`. Earlier compilers only provided for the assignment of functions to their own sections. Newer compilers have extended the `-xF` syntax to assign data variables to their own sections. Earlier compilers required C++ exception handling to be disabled when using `-xF`. This restriction has been dropped with later compilers.

If all allocatable sections from a relocatable object can be eliminated, the entire file is discarded from the link-edit.

In addition to input file elimination, the link-editor also identifies unused dependencies. A dependency is deemed unused if it is not bound to by the object being produced. An object built with the `-z ignore` option, will have no unused dependencies recorded in it.

The `-z ignore` option applies only to the files that follow the option on the `link-edit` command line. The `-z ignore` option is cancelled with `-z record`.

Maximizing Shareability

As mentioned in “Underlying System” on page 114, only a shared object’s text segment is shared by all processes that use it. The object’s data segment typically is not shared. Each process using a shared object, generates a private memory copy of its entire data segment as data items within the segment are written to. Reduce the data segment, either by moving data elements that are never written to the text segment, or by removing the data items completely.

The following sections describe several mechanisms that can be used to reduce the size of the data segment.

Move Read-Only Data to Text

Data elements that are read-only should be moved into the text segment using `const` declarations. For example, the following character string resides in the `.data` section, which is part of the writable data segment:

```
char * rdstr = "this is a read-only string";
```

In contrast, the following character string resides in the `.rodata` section, which is the read-only data section contained within the text segment:

```
const char * rdstr = "this is a read-only string";
```

Reducing the data segment by moving read-only elements into the text segment is admirable. However, moving data elements that require relocations can be counterproductive. For example, examine the following array of strings:

```
char * rdstrs[] = { "this is a read-only string",  
                  "this is another read-only string" };
```

A better definition might seem to be:

```
const char * const rdstrs[] = { ..... };
```

This definition ensures that the strings and the array of pointers to these strings are placed in a `.rodata` section. Unfortunately, although the user perceives the array of addresses as read-only, these addresses must be relocated at runtime. This definition therefore results in the creation of text relocations. Representing it as:

```
const char * rdstrs[] = { ..... };
```

insures the array pointers are maintained in the writable data segment where they can be relocated. The array strings are maintained in the read-only text segment.

Note – Some compilers, when generating position-independent code, can detect read-only assignments that result in runtime relocations. These compilers arrange for placing such items in writable segments. For example, `.picdata`.

Collapse Multiply-Defined Data

Data can be reduced by collapsing multiply-defined data. A program with multiple occurrences of the same error messages can be better off by defining one global datum, and have all other instances reference this. For example:

```
const char * Errmsg = "prog: error encountered: %d";

foo()
{
    .....
    (void) fprintf(stderr, Errmsg, error);
    .....
}
```

The main candidates for this sort of data reduction are strings. String usage in a shared object can be investigated using `strings(1)`. The following example generates a sorted list of the data strings within the file `libfoo.so.1`. Each entry in the list is prefixed with the number of occurrences of the string.

```
$ strings -10 libfoo.so.1 | sort | uniq -c | sort -rn
```

Use Automatic Variables

Permanent storage for data items can be removed entirely if the associated functionality can be designed to use automatic (stack) variables. Any removal of permanent storage usually results in a corresponding reduction in the number of runtime relocations required.

Allocate Buffers Dynamically

Large data buffers should usually be allocated dynamically rather than being defined using permanent storage. Often this results in an overall saving in memory, as only those buffers needed by the present invocation of an application are allocated. Dynamic allocation also provides greater flexibility by enabling the buffer's size to change without affecting compatibility.

Minimizing Paging Activity

Any process that accesses a new page causes a page fault, which is an expensive operation. Because shared objects can be used by many processes, any reduction in the number of page faults generated by accessing a shared object will benefit the process and the system as a whole.

Organizing frequently used routines and their data to an adjacent set of pages frequently improves performance because it improves the locality of reference. When a process calls one of these functions, the function might already be in memory because of its proximity to the other frequently used functions. Similarly, grouping interrelated functions improves locality of references. For example, if every call to the function `foo()` results in a call to the function `bar()`, place these functions on the same page. Tools like `cflow(1)`, `tcov(1)`, `prof(1)` and `gprof(1)` are useful in determining code coverage and profiling.

Isolate related functionality to its own shared object. The standard C library has historically been built containing many unrelated functions. Only rarely, for example, will any single executable use everything in this library. Because of widespread use, determining what set of functions are really the most frequently used is also somewhat difficult. In contrast, when designing a shared object from scratch, maintain only related functions within the shared object. This will improve locality of reference and has the side effect of reducing the object's overall size.

Relocations

In "Relocation Processing" on page 67, the mechanisms by which the runtime linker relocates dynamic executables and shared objects to create a runnable process was covered. "Symbol Lookup" on page 68 and "When Relocations Are Performed" on page 71 categorized this relocation processing into two areas to simplify and help illustrate the mechanisms involved. These same two categorizations are also ideally suited for considering the performance impact of relocations.

Symbol Lookup

When the runtime linker needs to look up a symbol, by default it does so by searching in each object. The runtime linker starts with the dynamic executable, and progresses through each shared object in the same order that the objects are loaded. In many instances, the shared object that requires a symbolic relocation turns out to be the provider of the symbol definition.

In this situation, if the symbol used for this relocation is not required as part of the shared object's interface, then this symbol is a strong candidate for conversion to a *static* or *automatic* variable. A symbol reduction can also be applied to removed

symbols from a shared objects interface. See “Reducing Symbol Scope” on page 52 for more details. By making these conversions, the link-editor incurs the expense of processing any symbolic relocation against these symbols during the shared object’s creation.

The only global data items that should be visible from a shared object are those that contribute to its user interface. Historically this has been a hard goal to accomplish, because global data are often defined to allow reference from two or more functions located in different source files. By applying symbol reduction, unnecessary global symbols can be removed. See “Reducing Symbol Scope” on page 52. Any reduction in the number of global symbols exported from a shared object results in lower relocation costs and an overall performance improvement.

The use of direct bindings can also significantly reduce the symbol lookup overhead within a dynamic process that has many symbolic relocations and many dependencies. See “Direct Binding” on page 69.

When Relocations are Performed

All immediate reference relocations must be carried out during process initialization before the application gains control. However, any lazy reference relocations can be deferred until the first instance of a function being called. Immediate relocations typically result from data references. Therefore, reducing the number of data references also reduces the runtime initialization of a process.

Initialization relocation costs can also be deferred by converting data references into function references. For example, you can return data items by a functional interface. This conversion usually results in a perceived performance improvement because the initialization relocation costs are effectively spread throughout the process’s execution. Some of the functional interfaces might never be called by a particular invocation of a process, thus removing their relocation overhead altogether.

The advantage of using a functional interface can be seen in the section, “Copy Relocations” on page 122. This section examines a special, and somewhat expensive, relocation mechanism employed between dynamic executables and shared objects. It also provides an example of how this relocation overhead can be avoided.

Combined Relocation Sections

Relocations by default are grouped by the sections against which they are to be applied. However, when an object is built with the `-z combrelloc` option, all but the procedure linkage table relocations are placed into a single common section named `.SUNW_reloc`. See “Procedure Linkage Table (Processor-Specific)” on page 261.

Combining relocation records in this manner enables all `RELATIVE` relocations to be grouped together. All symbolic relocations are sorted by symbol name. The grouping of `RELATIVE` relocations permits optimized runtime processing using the `DT_RELACOUNT/DT_RELCOUNT` `.dynamic` entries. Sorted symbolic entries help reduce runtime symbol lookup.

Copy Relocations

Shared objects are usually built with position-independent code. References to external data items from code of this type employs indirect addressing through a set of tables. See “Position-Independent Code” on page 115 for more details. These tables are updated at runtime with the real address of the data items. These updated tables enable access to the data without the code itself being modified.

Dynamic executables, however, are generally not created from position-independent code. Any references to external data they make can seemingly only be achieved at runtime by modifying the code that makes the reference. Modifying a read-only text segment is to be avoided. The *copy* relocation technique can solve this reference.

Suppose the link-editor is used to create a dynamic executable, and a reference to a data item is found to reside in one of the dependent shared objects. Space is allocated in the dynamic executable’s `.bss`, equivalent in size to the data item found in the shared object. This space is also assigned the same symbolic name as defined in the shared object. Along with this data allocation, the link-editor generates a special copy relocation record that will instruct the runtime linker to copy the data from the shared object to this allocated space within the dynamic executable.

Because the symbol assigned to this space is global, it is used to satisfy any references from any shared objects. The dynamic executable inherits the data item. Any other objects within the process that make reference to this item are bound to this copy. The original data from which the copy is made effectively becomes unused.

The following example of this mechanism uses an array of system error messages that is maintained within the standard C library. In previous SunOS operating system releases, the interface to this information was provided by two global variables, `sys_errlist []`, and `sys_nerr`. The first variable provided the array of error message strings, while the second conveyed the size of the array itself. These variables were commonly used within an application in the following manner:

```
$ cat foo.c
extern int      sys_nerr;
extern char *  sys_errlist[];

char *
error(int errnum)
{
    if ((errnum < 0) || (errnum >= sys_nerr))
        return (0);
    return (sys_errlist[errnum]);
}
```

```
}
```

The application uses the function `error` to provide a focal point to obtain the system error message associated with the number `errno`.

Examining a dynamic executable built using this code shows the implementation of the copy relocation in more detail:

```
$ cc -o prog main.c foo.c
$ nm -x prog | grep sys_
[36] |0x00020910|0x00000260|OBJT |WEAK |0x0 |16 |sys_errlist
[37] |0x0002090c|0x00000004|OBJT |WEAK |0x0 |16 |sys_nerr
$ dump -hv prog | grep bss
[16] NOBI WA- 0x20908 0x908 0x268 .bss
$ dump -rv prog
```

```
**** RELOCATION INFORMATION ****
```

```
.rela.bss:
Offset      Symndx          Type            Addend

0x2090c     sys_nerr        R_SPARC_COPY    0
0x20910     sys_errlist     R_SPARC_COPY    0
.....
```

The link-editor has allocated space in the dynamic executable's `.bss` to receive the data represented by `sys_errlist` and `sys_nerr`. These data are copied from the C library by the runtime linker at process initialization. Thus, each application that uses these data gets a private copy of the data in its own data segment.

There are two drawbacks to this technique. First, each application pays a performance penalty for the overhead of copying the data at runtime. Second, the size of the data array `sys_errlist` has now become part of the C library's interface. Suppose the size of this array were to change, perhaps as new error messages are added. Any dynamic executables that reference this array have to undergo a new link-edit to be able to access any of the new error messages. Without this new link-edit, the allocated space within the dynamic executable is insufficient to hold the new data.

These drawbacks can be eliminated if the data required by a dynamic executable are provided by a functional interface. The ANSI C function `strerror(3C)` returns a pointer to the appropriate error string, based on the error number supplied to it. One implementation of this function might be:

```
$ cat strerror.c
static const char * sys_errlist[] = {
    "Error 0",
    "Not owner",
    "No such file or directory",
    .....
};
static const int sys_nerr =
    sizeof (sys_errlist) / sizeof (char *);
```

```

char *
strerror(int errnum)
{
    if ((errnum < 0) || (errnum >= sys_nerr))
        return (0);
    return ((char *)sys_errlist[errnum]);
}

```

The error routine in `foo.c` can now be simplified to use this functional interface. This simplification in turn removes any need to perform the original copy relocations at process initialization.

Additionally, because the data are now local to the shared object, the data are no longer part of its interface. The shared object therefore has the flexibility of changing the data without adversely effecting any dynamic executables that use it. Eliminating data items from a shared object's interface generally improves performance while making the shared object's interface and code easier to maintain.

`ldd(1)`, when used with either the `-d` or `-r` options, can verify any copy relocations that exist within a dynamic executable.

For example, suppose the dynamic executable `prog` had originally been built against the shared object `libfoo.so.1` and the following two copy relocations had been recorded:

```

$ nm -x prog | grep _size_
[36] | 0x000207d8|0x40|OBJT |GLOB |15 | _size_gets_smaller
[39] | 0x00020818|0x40|OBJT |GLOB |15 | _size_gets_larger
$ dump -rv size | grep _size_
0x207d8      _size_gets_smaller      R_SPARC_COPY      0
0x20818      _size_gets_larger       R_SPARC_COPY      0

```

A new version of this shared object is supplied that contains different data sizes for these symbols:

```

$ nm -x libfoo.so.1 | grep _size_
[26] | 0x00010378|0x10|OBJT |GLOB |8 | _size_gets_smaller
[28] | 0x00010388|0x80|OBJT |GLOB |8 | _size_gets_larger

```

Running `ldd(1)` against the dynamic executable reveals:

```

$ ldd -d prog
libfoo.so.1 => ./libfoo.so.1
.....
copy relocation sizes differ: _size_gets_smaller
(file prog size=40; file ./libfoo.so.1 size=10);
./libfoo.so.1 size used; possible insufficient data copied
copy relocation sizes differ: _size_gets_larger
(file prog size=40; file ./libfoo.so.1 size=80);
./prog size used; possible data truncation

```

`ldd(1)` shows that the dynamic executable will copy as much data as the shared object has to offer, but only accepts as much as its allocated space allows.

Copy relocations can be eliminated by building the application from position-independent code. See “Position-Independent Code” on page 115.

Using `-B symbolic`

The link-editor’s `-B symbolic` option enables you to bind symbol references to their global definitions within a shared object. This option is historic, in that it was designed for use in creating the runtime linker itself.

Defining an object’s interface and reducing non-public symbols to local is preferable to using the `-B symbolic` option. See “Reducing Symbol Scope” on page 52. Using `-B symbolic` can often result in some non-intuitive side effects.

If a symbolically bound symbol is interposed upon, then references to the symbol from outside of the symbolically bound object bind to the interposer. The object itself is already bound internally. Essentially, two symbols with the same name are now being referenced from within the process. A symbolically bound data symbol that results in a copy relocation creates the same interposition situation. See “Copy Relocations” on page 122.

Note – Symbolically bound shared objects are identified by the `.dynamic` flag `DF_SYMBOLIC`. This flag is informational only. The runtime linker processes symbol lookups from these objects in the same manner as any other object. Any symbolic binding is assumed to have been created at the link-edit phase.

Profiling Shared Objects

The runtime linker can generate profiling information for any shared objects that are processed during the running of an application. The runtime linker is responsible for binding shared objects to an application and is therefore able to intercept any *global* function bindings. These bindings take place through `.plt` entries. See “When Relocations Are Performed” on page 71 for details of this mechanism.

The `LD_PROFILE` environment variable specifies the name of a shared object to profile. You can analyze one shared object at a time using this environment variable. The setting of the environment variable can be used to analyze the use of the shared object by one or more applications. In the following example, the use of `libc` by the single invocation of the command `ls(1)` is analyzed:

```
$ LD_PROFILE=libc.so.1 ls -l
```

In the following example, the environment variable setting is recorded in a configuration file. This setting causes any application’s use of `libc` to accumulate the analyzed information:

```
# crle -e LD_PROFILE=libc.so.1
$ ls -l
$ make
$ ...
```

When profiling is enabled, a profile data file is created, if it does not already exist. The file is mapped by the runtime linker. In the above examples, this data file is `/var/tmp/libc.so.1.profile`. 64-bit libraries require an extended profile format and are written using the `.profilex` suffix. You can also specify an alternative directory to store the profile data using the `LD_PROFILE_OUTPUT` environment variable.

This profile data file is used to deposit `profil(2)` data and call count information related to the use of the specified shared object. This profiled data can be directly examined with `gprof(1)`.

Note – `gprof(1)` is most commonly used to analyze the `gmon.out` profile data created by an executable that has been compiled with the `-xpg` option of `cc(1)`. The runtime linker's profile analysis does not require any code to be compiled with this option. Applications whose dependent shared objects are being profiled should not make calls to `profil(2)`, because this system call does not provide for multiple invocations within the same process. For the same reason, these applications must not be compiled with the `-xpg` option of `cc(1)`. This compiler-generated mechanism of profiling is also built on top of `profil(2)`.

One of the most powerful features of this profiling mechanism is to enable the analysis of a shared object as used by multiple applications. Frequently, profiling analysis is carried out using one or two applications. However, a shared object, by its very nature, can be used by a multitude of applications. Analyzing how these applications use the shared object can offer insights into where energy might be spent to improvement the overall performance of the shared object.

The following example shows a performance analysis of `libc` over a creation of several applications within a source hierarchy.

```
$ LD_PROFILE=libc.so.1 ; export LD_PROFILE
$ make
$ gprof -b /usr/lib/libc.so.1 /var/tmp/libc.so.1.profile
.....
```

granularity: each sample hit covers 4 byte(s)

index	%time	self	descendents	called/total	parents		
				called+self	name	index	
				called/total	children		
.....							
				52/29381	<code>_gettext</code>	[96]	
				174/29381	<code>_tzload</code>	[54]	

```

          10.50      0.00   1634/29381   <external>
          16.14      0.00   2512/29381   _opendir [15]
          160.65     0.00  25009/29381  _endopen [3]
[2]      35.0   188.74      0.00   29381      _open [2]
-----

```

```

.....
granularity: each sample hit covers 4 byte(s) ....

```

	% cumulative	self	self	self	total	
time	seconds	seconds	calls	ms/call	ms/call	name
35.0	188.74	188.74	29381	6.42	6.42	_open [2]
13.0	258.80	70.06	12094	5.79	5.79	_write [4]
9.9	312.32	53.52	34303	1.56	1.56	_read [6]
7.1	350.53	38.21	1177	32.46	32.46	_fork [9]

```

.....

```

The special name *<external>* indicates a reference from outside of the address range of the shared object being profiled. Thus, in the above example, 1634 calls to the function `open(2)` within `libc` occurred from the dynamic executables, or from other shared objects, bound with `libc` while the profiling analysis was in progress.

Note – The profiling of shared objects is multithread safe, except in the case where one thread calls `fork(2)` while another thread is updating the profile data information. The use of `fork1(2)` removes this restriction.

Application Binary Interfaces and Versioning

ELF objects processed by the link-editors provide many global symbols to which other objects can bind. These symbols describe the object's application binary interface (ABI). During the evolution of an object, this interface can change due to the addition or deletion of global symbols. In addition, the object's evolution can involve internal implementation changes.

Versioning refers to several techniques that can be applied to an object to indicate interface and implementation changes. These techniques provide for the object's controlled evolution while maintaining backward compatibility.

This chapter describes how an object's ABI can be defined and classifies how changes to this interface can affect backward compatibility. It also presents models by which interface and implementation changes can be incorporated into new releases of the object.

The focus of this chapter is on the runtime interfaces of dynamic executables and shared objects. The techniques used to describe and manage changes within these dynamic objects are presented in generic terms. A common set of naming conventions and versioning scenarios as applied to shared objects can be found in Appendix B.

Developers of dynamic objects must be aware of the ramifications of an interface change and understand how such changes can be managed, especially in regards to maintaining backward compatibility with previously shipped objects.

The global symbols made available by any dynamic object represent the object's public interface. Frequently, the number of global symbols remaining in an object at the end of a link-edit are more than you would like to make public. These global symbols result from the relationship required between relocatable objects used to create the object. They represent private interfaces within the object itself.

Before defining an object's binary interface, you should first determine those global symbols you wish to make publicly available from the object being created. These public symbols can be established using the link-editor's `-M` option and an associated `mapfile` as part of the final link-edit. This technique is introduced in "Reducing

Symbol Scope” on page 52. This public interface establishes one or more version definitions within the object being created. These definitions form the foundation for the addition of new interfaces as the object evolves.

The following sections build upon this initial public interface. First though, you should understand how various changes to an interface can be categorized so that they can be managed appropriately.

Interface Compatibility

Many types of change can be made to an object. In their simplest terms, these changes can be categorized into one of two groups:

- *Compatible* updates. These updates are additive, in that all previously available interfaces remain intact.
- *Incompatible* updates. These updates have changed the existing interface in such a way that existing users of the interface can fail or behave incorrectly.

The following table categorizes some common object changes.

TABLE 5-1 Interface Compatibility Examples

Object Change	Update Type
The addition of a symbol	Compatible
The removal of a symbol	Incompatible
The addition of an argument to a non- <code>varargs(3HEAD)</code> function	Incompatible
The removal of an argument from a function	Incompatible
The change of size, or content, of a data item to a function or as an external definition	Incompatible
A bug fix, or internal enhancement to a function, providing the semantic properties of the object remain unchanged	Compatible
A bug fix, or internal enhancement to a function when the semantic properties of the object change	Incompatible

Because of interposition, the addition of a symbol can constitute an incompatible update, such that the new symbol might conflict with an applications use of that symbol. However, this does seem rare in practice as source-level name space management is commonly used.

Compatible updates can be accommodated by maintaining version definitions internal to the object being generated. Incompatible updates can be accommodated by producing a new object with a new external versioned name. Both of these versioning techniques enable the selective binding of applications. They also enable verification of correct version binding at runtime. These two techniques are explored in more detail in the following sections.

Internal Versioning

A dynamic object can have one or more internal version definitions associated with it. Each version definition is commonly associated with one or more symbol names. A symbol name can only be associated with *one* version definition. However, a version definition can inherit the symbols from other version definitions. Thus, a structure exists to define one or more independent, or related, version definitions within the object being created. As new changes are made to the object, new version definitions can be added to express these changes.

There are two consequences of providing version definitions within a shared object:

- Dynamic objects that are built against a versioned shared object can record their dependency on the version definitions bound to. These version dependencies are verified at runtime to ensure that the appropriate interfaces, or functionality, are available for the correct execution of an application.
- Dynamic objects can select the version definitions of a shared object to bind to during their link-edit. This mechanism enables developers to control their dependency on a shared object to the interfaces, or functionality, that provide the most flexibility.

Creating a Version Definition

Version definitions commonly consist of an association of symbol names to a unique version name. These associations are established within a `mapfile` and supplied to the final link-edit of an object using the link-editor's `-M` option. This technique is introduced in the section "Reducing Symbol Scope" on page 52.

A version definition is established whenever a version name is specified as part of the `mapfile` directive. In the following example, two source files are combined, together with `mapfile` directives, to produce an object with a defined public interface:

```
$ cat foo.c
extern const char * _foo1;

void foo1()
```

```

{
    (void) printf(_foo1);
}

$ cat data.c
const char * _foo1 = "string used by foo1()\n";

$ cat mapfile
SUNW_1.1 {
    global:
        foo1;
    local:
        *;
};
$ cc -o libfoo.so.1 -M mapfile -G foo.o data.o
$ nm -x libfoo.so.1 | grep "foo.$"
[33] |0x0001058c|0x00000004|OBJT |LOCL |0x0 |17 |_foo1
[35] |0x00000454|0x00000034|FUNC |GLOB |0x0 |9 |foo1

```

The symbol `foo1` is the only global symbol defined to provide the shared object's public interface. The special auto-reduction directive `"*"` causes the reduction of all other global symbols to have local binding within the object being generated. This directive is introduced in "Defining Additional Symbols" on page 46. The associated version name, `SUNW_1.1`, causes the generation of a version definition. Thus, the shared object's public interface consists of the internal version definition `SUNW_1.1`, associated with the global symbol `foo1`.

Whenever a version definition, or the auto-reduction directive, are used to generate an object, a base version definition is also created. This base version is defined using the name of the file itself, and is used to associate any reserved symbols generated by the link-editor. See "Generating the Output File" on page 57 for a list of these reserved symbols.

The version definitions contained within an object can be displayed using `pvs(1)` with the `-d` option:

```

$ pvs -d libfoo.so.1
libfoo.so.1;
SUNW_1.1;

```

The object `libfoo.so.1` has an internal version definition named `SUNW_1.1`, together with a base version definition `libfoo.so.1`.

Note – The link-editor's `-z noversion` option allows symbol reduction to be directed by a `mapfile` but suppresses the creation of version definitions.

Starting with this initial version definition, the object can evolve by adding new interfaces and updated functionality. For example, a new function, `foo2`, together with its supporting data structures, can be added to the object by updating the source files `foo.c` and `data.c`:

```

$ cat foo.c
extern const char * _foo1;
extern const char * _foo2;

void foo1()
{
    (void) printf(_foo1);
}

void foo2()
{
    (void) printf(_foo2);
}

$ cat data.c
const char * _foo1 = "string used by foo1()\n";
const char * _foo2 = "string used by foo2()\n";

```

A new version definition, `SUNW_1.2`, can be created to define a new interface representing the symbol `foo2`. In addition, this new interface can be defined to inherit the original version definition `SUNW_1.1`.

The creation of this new interface is important as it identifies the evolution of the object and enables users to verify and select the interfaces to which they bind. These concepts are covered in more detail in “Binding to a Version Definition” on page 136 and in “Specifying a Version Binding” on page 140.

The following example shows the `mapfile` directives that create these two interfaces.

```

$ cat mapfile
SUNW_1.1 {
    global:
        foo1;
    local:
        *;
};

SUNW_1.2 {
    global:
        foo2;
} SUNW_1.1;

$ cc -o libfoo.so.1 -M mapfile -G foo.o data.o
$ nm -x libfoo.so.1 | grep "foo.$"
[33] |0x00010644|0x00000004|OBJT |LOCL |0x0 |17 |_foo1
[34] |0x00010648|0x00000004|OBJT |LOCL |0x0 |17 |_foo2
[36] |0x000004bc|0x00000034|FUNC |GLOB |0x0 |9 |foo1
[37] |0x000004f0|0x00000034|FUNC |GLOB |0x0 |9 |foo2

```

The symbols `foo1` and `foo2` are both defined to be part of the shared object’s public interface. However, each of these symbols is assigned to a different version definition; `foo1` is assigned to `SUNW_1.1`, and `foo2` is assigned to `SUNW_1.2`.

These version definitions, their inheritance, and their symbol association can be displayed using `pvs(1)` together with the `-d`, `-v` and `-s` options:

```
$ pvs -dsv libfoo.so.1
libfoo.so.1:
    _end;
    _GLOBAL_OFFSET_TABLE_;
    _DYNAMIC;
    _edata;
    _PROCEDURE_LINKAGE_TABLE_;
    _etext;
SUNW_1.1:
    foo1;
    SUNW_1.1;
SUNW_1.2:                {SUNW_1.1}:
    foo2;
    SUNW_1.2
```

The version definition `SUNW_1.2` has a dependency on the version definition `SUNW_1.1`.

The inheritance of one version definition by another is a useful technique that reduces the version information that will eventually be recorded by any object that binds to a version dependency. Version inheritance is covered in more detail in the section “Binding to a Version Definition” on page 136.

Any internal version definition has an associated *version definition symbol* created. As shown in the previous `pvs(1)` example, these symbols are displayed when using the `-v` option.

Creating a Weak Version Definition

Internal changes to an object that do not require the introduction of a new interface definition can be defined by creating a *weak* version definition. Examples of such changes are bug fixes or performance improvements.

Such a version definition is empty, in that it has no global interface symbols associated with it.

For example, suppose the data file `data.c`, used in the previous examples, is updated to provide more detailed string definitions:

```
$ cat data.c
const char * _foo1 = "string used by function foo1()\n";
const char * _foo2 = "string used by function foo2()\n";
```

A weak version definition can be introduced to identify this change:

```
$ cat mapfile
SUNW_1.1 {                # Release X
    global:
        foo1;
```

```

        local:
            *;
};

SUNW_1.2 {
    global:
        foo2;
} SUNW_1.1;

SUNW_1.2.1 { } SUNW_1.2;    # Release X+2

$ cc -o libfoo.so.1 -M mapfile -G foo.o data.o
$ pvs -dv libfoo.so.1
    libfoo.so.1;
    SUNW_1.1;
    SUNW_1.2:          {SUNW_1.1};
    SUNW_1.2.1 [WEAK]: {SUNW_1.2};

```

The empty version definition is signified by the weak label. These weak version definitions enable applications to verify the existence of a particular implementation by binding to the version definition associated with that functionality. The section “Binding to a Version Definition” on page 136 illustrates how these definitions can be used in more detail.

Defining Unrelated Interfaces

The previous examples show how new version definitions added to an object inherit any existing version definitions. You can also create version definitions that are unique and independent. In the following example, two new files, `bar1.c` and `bar2.c`, are added to the object `libfoo.so.1`. These files contribute two new symbols, `bar1` and `bar2`, respectively:

```

$ cat bar1.c
extern void foo1();

void bar1()
{
    foo1();
}
$ cat bar2.c
extern void foo2();

void bar2()
{
    foo2();
}

```

These two symbols are intended to define two new public interfaces. Neither of these new interfaces are related to each other. However, each expresses a dependency on the original `SUNW_1.2` interface.

The following `mapfile` definition creates this required association:

```

$ cat mapfile
SUNW_1.1 {                                # Release X
    global:
        foo1;
    local:
        *;
};

SUNW_1.2 {                                # Release X+1
    global:
        foo2;
} SUNW_1.1;

SUNW_1.2.1 { } SUNW_1.2;                 # Release X+2

SUNW_1.3a {                               # Release X+3
    global:
        bar1;
} SUNW_1.2;

SUNW_1.3b {                               # Release X+3
    global:
        bar2;
} SUNW_1.2;

```

Again, the version definitions created in `libfoo.so.1` when using this mapfile, and their related dependencies, can be inspected using `pvs(1)`:

```

$ cc -o libfoo.so.1 -M mapfile -G foo.o bar1.o bar2.o data.o
$ pvs -dv libfoo.so.1
    libfoo.so.1;
    SUNW_1.1;
    SUNW_1.2:           {SUNW_1.1};
    SUNW_1.2.1 [WEAK]: {SUNW_1.2};
    SUNW_1.3a:         {SUNW_1.2};
    SUNW_1.3b:         {SUNW_1.2};

```

The following sections explore how these version definition recordings can be used to verify runtime binding requirements and control the binding of an object during its creation.

Binding to a Version Definition

When a dynamic executable or shared object is built against other shared objects, these dependencies are recorded in the resulting object. See “Shared Object Processing” on page 30 and “Recording a Shared Object Name” on page 103 for more details. If these shared object dependencies also contain version definitions, then an associated version dependency is recorded in the object being built.

The following example takes the data files from the previous section and generates a shared object suitable for a compile time environment. This shared object, `libfoo.so.1`, is used in the succeeding binding examples.


```

$ cc -o libfoo.so.1 -h libfoo.so.1 -M mapfile -G foo.o bar.o \
data.o
$ ln -s libfoo.so.1 libfoo.so
$ pvs -dsv libfoo.so.1
libfoo.so.1:
    _end;
    _GLOBAL_OFFSET_TABLE_;
    _DYNAMIC;
    _edata;
    _PROCEDURE_LINKAGE_TABLE_;
    _etext;
SUNW_1.1:
    foo1;
    SUNW_1.1;
SUNW_1.2:                {SUNW_1.1}:
    foo2;
    SUNW_1.2;
SUNW_1.2.1 [WEAK]:      {SUNW_1.2}:
    SUNW_1.2.1;
SUNW_1.3a:              {SUNW_1.2}:
    bar1;
    SUNW_1.3a;
SUNW_1.3b:              {SUNW_1.2}:
    bar2;
    SUNW_1.3b

```

In effect, there are six public interfaces being offered by the shared object. Four of these interfaces, `SUNW_1.1`, `SUNW_1.2`, `SUNW_1.3a`, and `SUNW_1.3b`, define exported symbol names. One interface, `SUNW_1.2.1`, describes an internal implementation change to the shared object, and one interface, `libfoo.so.1`, defines several reserved labels. Dynamic objects created with this shared object as a dependency, record the version names of the interfaces the dynamic object binds to.

The following example creates an application that references symbols `foo1` and `foo2`. The versioning dependency information recorded in the application can be examined using `pvs(1)` with the `-r` option.

```

$ cat prog.c
extern void foo1();
extern void foo2();

main()
{
    foo1();
    foo2();
}
$ cc -o prog prog.c -L. -R. -lfoo
$ pvs -r prog
libfoo.so.1 (SUNW_1.2, SUNW_1.2.1);

```

In this example, the application `prog` has bound to the two interfaces `SUNW_1.1` and `SUNW_1.2`. These interfaces provided the global symbols `foo1` and `foo2` respectively.

Because version definition `SUNW_1.1` is defined within `libfoo.so.1` as being inherited by the version definition `SUNW_1.2`, you also need to record the latter version dependency. This normalization of version definition dependencies reduces the amount of version information maintained within an object, and reduces the processing required at runtime.

Because the application `prog` was built against the shared object's implementation containing the weak version definition `SUNW_1.2.1`, this dependency is also recorded. Even though this version definition is defined to inherit the version definition `SUNW_1.2`, the version's weak nature precludes its normalization with `SUNW_1.1`, and results in a separate dependency recording.

Had there been multiple weak version definitions that inherited from each other, then these definitions will be normalized in the same manner as non-weak version definitions are.

Note – The recording of a version dependency can be suppressed by the link-editor's `-z noversion` option.

Having recorded these version definition dependencies, the runtime linker validates the existence of the required version definitions in the objects that are bound to when the application is executed. This validation can be displayed using `ldd(1)` with the `-v` option. For example, by running `ldd(1)` on the application `prog`, the version definition dependencies are shown to be found correctly in the shared object `libfoo.so.1`:

```
$ ldd -v prog

find object=libfoo.so.1; required by prog
  libfoo.so.1 =>    ./libfoo.so.1
find version=libfoo.so.1;
  libfoo.so.1 (SUNW_1.2) =>          ./libfoo.so.1
  libfoo.so.1 (SUNW_1.2.1) =>      ./libfoo.so.1
....
```

Note – `ldd(1)` with the `-v` option implies *verbose* output. A recursive list of all dependencies, together with all versioning requirements, is generated.

If a non-weak version definition dependency cannot be found, a fatal error occurs during application initialization. Any weak version definition dependency that cannot be found is silently ignored. For example, if the application `prog` is run in an environment in which `libfoo.so.1` only contains the version definition `SUNW_1.1`, then the following fatal error occurs:

```
$ pvs -dv libfoo.so.1
  libfoo.so.1;
```

```

        SUNW_1.1;
$ prog
ld.so.1: prog: fatal: libfoo.so.1: version 'SUNW_1.2' not \
found (required by file prog)

```

Had the application `prog` not recorded any version definition dependencies, the nonexistence of the required interface symbol `foo2` would have manifested itself some time during the execution of the application as a fatal relocation error. This relocation error might occur at process initialization, during process execution, or might not occur at all if the execution path of the application did not call the function `foo2`. See “Relocation Errors” on page 72.

Recording version definition dependencies provides an alternative and immediate indication of the availability of the interfaces required by the application.

If the application `prog` is run in an environment in which `libfoo.so.1` only contains the version definitions `SUNW_1.1` and `SUNW_1.2`, then all non-weak version definition requirements will be satisfied. The absence of the weak version definition `SUNW_1.2.1` is deemed nonfatal, and so no runtime error condition is generated. However, `ldd(1)` can be used to display all version definitions that cannot be found:

```

$ pvs -dv libfoo.so.1
    libfoo.so.1;
    SUNW_1.1;
    SUNW_1.2:                {SUNW_1.1};
$ prog
string used by foo1()
string used by foo2()
$ ldd prog
    libfoo.so.1 =>    ./libfoo.so.1
    libfoo.so.1 (SUNW_1.2.1) =>    (version not found)
    .....

```

Note – If an object requires a version definition from a given dependency, and at runtime an implementation of that dependency is found that contains no version definition information, the version verification of the dependency will be silently ignored. This policy provides a level of backward compatibility as a transition from non-versioned to versioned shared objects occurs. `ldd(1)`, however, can still be used to display any version requirement discrepancies. The environment variable `LD_NOVERSION` can be used to suppress all runtime versioning verification.

Verifying Versions in Additional Objects

Version definition symbols also provide a mechanism for verifying the version requirements of an object obtained by `dlopen(3DL)`. Any object added to the process’s address space using this function will have no automatic version dependency verification carried out by the runtime linker. Thus, the caller of this function is responsible for verifying that any versioning requirements are met.

The presence of a required version definition can be verified by looking up the associated version definition symbol using `dlsym(3DL)`. The following example adds the shared object `libfoo.so.1` to a process using `dlopen(3DL)`, and verifies the interface `SUNW_1.2` is available.

```
#include      <stdio.h>
#include      <dlfcn.h>

main()
{
    void *      handle;
    const char * file = "libfoo.so.1";
    const char * vers = "SUNW_1.2";
    ....

    if ((handle = dlopen(file, (RTLD_LAZY | RTLD_FIRST))) == NULL) {
        (void) printf("dlopen: %s\n", dlerror());
        exit (1);
    }

    if (dlsym(handle, vers) == NULL) {
        (void) printf("fatal: %s: version '%s' not found\n",
            file, vers);
        exit (1);
    }
    ....
}
```

Specifying a Version Binding

When creating a dynamic object against a shared object containing version definitions, you can instruct the link-editor to limit the binding to specific version definitions. Effectively, the link-editor enables you to control an object's binding to specific interfaces.

An object's binding requirements can be controlled using a *file control directive*. This directive is supplied using the link-editor's `-M` option and an associated `mapfile`. The following syntax for file control directives is available:

```
name - version [ version ... ] [ $ADDVERS=version ] ;
```

- *name* – Represents the name of the shared object dependency. This name should match the shared object's compilation environment name as used by the link-editor. See "Library Naming Conventions" on page 31.
- *version* – Represents the version definition name within the shared object that should be made available for binding. Multiple version definitions can be specified.
- `$ADDVERS` – Allows for additional version definitions to be recorded.

This binding control can be useful in the following scenarios:

- When a shared object defines independent, unique versions. This versioning is possible when defining different standards interfaces. An object can be built with binding controls to ensure it only binds to a specific interface.
- When a shared object has been versioned over several software releases. An object can be built with binding controls to restrict its binding to the interfaces available in a previous software release. Thus, an object can run with an old release of the shared object dependency, after being built using the latest release of the shared object.

The following example illustrates the use of the version control mechanism. This example uses the shared object `libfoo.so.1` containing the following version interface definitions:

```
$ pvs -dsv libfoo.so.1
libfoo.so.1:
    _end;
    _GLOBAL_OFFSET_TABLE_;
    _DYNAMIC;
    _edata;
    _PROCEDURE_LINKAGE_TABLE_;
    _etext;
SUNW_1.1:
    fool;
    foo2;
    SUNW_1.1;
SUNW_1.2:          {SUNW_1.1}:
    bar;
```

The version definitions `SUNW_1.1` and `SUNW_1.2` represent interfaces within `libfoo.so.1` that were made available in software Release `X` and Release `X+1` respectively.

An application can be built to bind only to the interfaces available in Release `X` by using the following version control `mapfile` directive:

```
$ cat mapfile
libfoo.so - SUNW_1.1;
```

For example, suppose you develop an application, `prog`, and want to ensure that the application can run on Release `X`. The application can then only use the interfaces available in that release. If the application mistakenly references the symbol `bar`, then the application is not compliant with the required interface. This condition is signalled by the link-editor as an undefined symbol error:

```
$ cat prog.c
extern void fool();
extern void bar();

main()
{
    fool();
    bar();
}
```

```

$ cc -o prog prog.c -M mapfile -L. -R. -lfoo
Undefined      first referenced
 symbol        in file
bar            prog.o (symbol belongs to unavailable \
              version ./libfoo.so (SUNW_1.2))
ld: fatal: Symbol referencing errors. No output written to prog

```

To be compliant with the SUNW_1.1 interface, you must remove the reference to bar. You can either rework the application to remove the requirement on bar, or add an implementation of bar to the creation of the application.

Note – By default, shared object dependencies encountered as part of a link-edit, are also verified against any file control directives. Use the environment variable LD_NOVERSION to suppress the version verification of any shared object dependencies.

Binding to Additional Version Definitions

To record more version dependencies than would be produced from the normal symbol binding of an object, use the \$ADDVERS file control directive. This section describes scenarios where this additional binding might be useful.

From the previous libfoo.so.1 example, assume that in Release X+2, the version definition SUNW_1.1 is subdivided into two standard releases, STAND_A and STAND_B. To preserve compatibility, the SUNW_1.1 version definition must be maintained. In this example, this version definition is expressed as inheriting the two standard definitions:

```

$ pvs -dsv libfoo.so.1
libfoo.so.1:
    _end;
    _GLOBAL_OFFSET_TABLE_;
    _DYNAMIC;
    _edata;
    _PROCEDURE_LINKAGE_TABLE_;
    _etext;
SUNW_1.1:    {STAND_A, STAND_B}:
    SUNW_1.1;
SUNW_1.2:    {SUNW_1.1}:
    bar;
STAND_A:
    foo1;
    STAND_A;
STAND_B:
    foo2;
    STAND_B;

```

If the only requirement of application `prog` is the interface symbol `foo1`, the application will have a single dependency on the version definition `STAND_A`. This precludes running `prog` on a system where `libfoo.so.1` is less than Release `X+2`. The version definition `STAND_A` did not exist in previous releases, even though the interface `foo1` did.

The application `prog` can be built to align its requirement with previous releases by creating a dependency on `SUNW_1.1`:

```
$ cat mapfile
libfoo.so - SUNW_1.1 $ADDVERS=SUNW_1.1;
$ cat prog
extern void foo1();

main()
{
    foo1();
}
$ cc -M mapfile -o prog prog.c -L. -R. -lfoo
$ pvs -r prog
    libfoo.so.1 (SUNW_1.1);
```

This explicit dependency is sufficient to encapsulate the true dependency requirements. This dependency satisfies compatibility with older releases.

“Creating a Weak Version Definition” on page 134 described how weak version definitions can be used to mark an internal implementation change. These version definitions are well suited to indicate bug fixes and performance improvements made to an object. If the existence of a weak version is required, an explicit dependency on this version definition can be generated. The creation of such a dependency can be important when a bug fix, or performance improvement, is critical for the object to function correctly.

From the previous `libfoo.so.1` example, assume a bug fix is incorporated as the weak version definition `SUNW_1.2.1` in software Release `X+3`:

```
$ pvs -dsv libfoo.so.1
libfoo.so.1:
    _end;
    _GLOBAL_OFFSET_TABLE_;
    _DYNAMIC;
    _edata;
    _PROCEDURE_LINKAGE_TABLE_;
    _etext;
SUNW_1.1:      {STAND_A, STAND_B}:
    SUNW_1.1;
SUNW_1.2:      {SUNW_1.1}:
    bar;
STAND_A:
    foo1;
    STAND_A;
STAND_B:
    foo2;
```

```

        STAND_B;
SUNW_1.2.1 [WEAK]: {SUNW_1.2}:
        SUNW_1.2.1;

```

Normally, if an application is built against this shared object, the application records a weak dependency on the version definition `SUNW_1.2.1`. This dependency is informational only. This dependency does not cause termination of the application should the version definition not exist in the `libfoo.so.1` used at runtime.

The file control directive `$ADDVERS` can be used to generate an explicit dependency on a version definition. If this definition is weak, then this explicit reference also causes the version definition to be promoted to a strong dependency.

The application `prog` can be built to enforce the requirement that the `SUNW_1.2.1` interface be available at runtime by using the following file control directive:

```

$ cat mapfile
libfoo.so - SUNW_1.1 $ADDVERS=SUNW_1.2.1;
$ cat prog
extern void fool();

main()
{
    fool();
}
$ cc -M mapfile -o prog prog.c -L. -R. -lfoo
$ pvs -r prog
    libfoo.so.1 (SUNW_1.2.1);

```

`prog` has been built with an explicit dependency on the interface `STAND_A`. Because the version definition `SUNW_1.2.1` is promoted to a strong version, it is also normalized with the dependency `STAND_A`. At runtime, if the version definition `SUNW_1.2.1` cannot be found, a fatal error is generated.

Note – When working with a small number of dependencies, you can use the link-editor’s `-u` option to explicitly bind to a version definition. Use this option to reference the version definition symbol. However, a symbol reference is nonselective. When working with multiple dependencies, that might contain similarly named version definitions, this technique may be insufficient to create explicit bindings.

Version Stability

The various models for binding to versions within an object only remain intact if the individual version definitions remain constant over the life time of the object.

Once a version definition for an object has been created and made public, it must exist in subsequent releases of that object unchanged. Both the version name and the symbols associated with it must remain constant. For this reason, wildcard expansion of the symbol names defined within a version definition is not supported. The number of symbols matching the wildcard might differ over the course of an objects evolution.

Relocatable Objects

Version information can be recorded and used within dynamic objects. Relocatable objects can maintain versioning information in a similar manner. However, there are some subtle differences in how this information is used.

Any version definitions supplied to the link-edit of a relocatable object are recorded in the same format as they are when building dynamic executables or shared objects. However, by default, symbol reduction is not carried out on the object being created. Instead, when the relocatable object is finally used as input to the generation of a dynamic object, the version recording itself will be used to determine the symbol reductions to apply.

In addition, any version definitions found in relocatable objects are propagated to the dynamic object. For an example of version processing in relocatable objects, see “Reducing Symbol Scope” on page 52.

External Versioning

Runtime references to a shared object should always refer to the file’s version file name. This is usually expressed as a file name with a version number suffix. When a shared object’s interface changes in an incompatible manner, such that it will break old applications, a new shared object should be distributed with a new versioned file name. In addition, the original versioned file name must still be distributed to provide the interfaces required by the old applications.

You should provide shared objects as separate versioned file names within the runtime environment when building applications over a series of software releases. You can then guarantee that the interface against which the applications were built is available for them to bind during their execution.

The following section describes how to coordinate the binding of an interface between the compilation and runtime environments.

Coordination of Versioned Filenames

During a link-edit, the most common method to input shared objects is to use the `-l` option. This option uses the link-editor's library search mechanism to locate shared objects that are prefixed with `lib` and suffixed with `.so`.

However, at runtime, any shared object dependencies should exist in their *versioned* name form. Instead of maintaining two distinct shared objects that follow these naming conventions, create file system links between the two file names.

To make the runtime shared object `libfoo.so.1` available to the compilation environment, provide a symbolic link from the compilation file name to the runtime file name. For example:

```
$ cc -o libfoo.so.1 -G -K pic foo.c
$ ln -s libfoo.so.1 libfoo.so
$ ls -l libfoo*
lrwxrwxrwx 1 usr grp          11 1991 libfoo.so -> libfoo.so.1
-rwxrwxr-x 1 usr grp       3136 1991 libfoo.so.1
```

Either a symbolic link or hard link can be used. However, as a documentation and diagnostic aid, symbolic links are more useful.

The shared object `libfoo.so.1` has been generated for the runtime environment. Generating a symbolic link `libfoo.so`, has also enabled this file's use in a compilation environment. For example:

```
$ cc -o prog main.o -L. -lfoo
```

The link-editor processes the relocatable object `main.o` with the interface described by the shared object `libfoo.so.1`, which is found by following the symbolic link `libfoo.so`.

Over a series of software releases, new versions of this shared object may be distributed with changed interfaces. The compilation environment can be constructed to use the interface that is applicable by changing the symbolic link. For example:

```
$ ls -l libfoo*
lrwxrwxrwx 1 usr grp          11 1993 libfoo.so -> libfoo.so.3
-rwxrwxr-x 1 usr grp       3136 1991 libfoo.so.1
-rwxrwxr-x 1 usr grp       3237 1992 libfoo.so.2
-rwxrwxr-x 1 usr grp       3554 1993 libfoo.so.3
```

Three major versions of the shared object are available. Two of these shared objects, `libfoo.so.1` and `libfoo.so.2`, provide the dependencies for existing applications. `libfoo.so.3` offers the latest major release for creating and running new applications.

Using this symbolic link mechanism itself is insufficient to coordinate the correct binding of a shared object from its use in the compilation environment to its requirement in the runtime environment. As the example presently stands, the link-editor records in the dynamic executable `prog` the file name of the shared object it has processed. In this case, that file name is the compilation environment file name.

```
$ dump -Lv prog

prog:
**** DYNAMIC SECTION INFORMATION ****
.dynamic:
[INDEX] Tag      Value
[1]      NEEDED   libfoo.so
.....
```

When the application `prog` is executed, the runtime linker searches for the dependency `libfoo.so`. `prog` binds to the file to which this symbolic link is pointing.

To provide the correct runtime name to be recorded as a dependency, the shared object `libfoo.so.1` should be built with an `soname` definition. This definition identifies the shared object's runtime name. This name is used as the dependency name by any object that links against this shared object. This definition can be provided using the `-h` option during the link-edit of the shared object itself. For example:

```
$ cc -o libfoo.so.1 -G -K pic -h libfoo.so.1 foo.c
$ ln -s libfoo.so.1 libfoo.so
$ cc -o prog main.o -L. -lfoo
$ dump -Lv prog
```

```
prog:
**** DYNAMIC SECTION INFORMATION ****
.dynamic:
[INDEX] Tag      Value
[1]      NEEDED   libfoo.so.1
.....
```

This symbolic link and the `soname` mechanism have established a robust coordination between the shared-object naming conventions of the compilation and runtime environment. The interface processed during the link-edit is accurately recorded in the output file generated. This recording ensures that the intended interface are furnished at runtime.



Caution – Creating a new externally versioned shared object is a major change. Be sure you understand the complete dependencies of any processes that use this shared object.

For example, an application might have dependencies on `libfoo.so.1` and an externally delivered object `libISV.so.1`. This latter object might also have a dependency on `libfoo.so.1`. If the application is redesigned to use the new interfaces in `libfoo.so.2` without any change to its use of the external object `libISV.so.1`, then both major versions of `libfoo.so` will be brought into the running process. Because the only reason to change the version of `libfoo.so` is to mark an incompatible change, having both versions of the object within a process can lead to incorrect symbol binding and hence undesirable interactions.

Support Interfaces

The link-editors provide a number of support interfaces that enable the monitoring, and in some cases modification, of link-editor and runtime linker processing. These interfaces typically require a more advanced understanding of link-editing concepts than has been described in previous chapters. The following interfaces are described in this chapter:

- *ld-support* – “Link-Editor Support Interface” on page 149
- *rtld-audit* – “Runtime Linker Auditing Interface” on page 155
- *rtld-debugger* – “Runtime Linker Debugger Interface” on page 164

Link-Editor Support Interface

The link-editor performs many operations including the opening of files and the concatenation of sections from these files. Monitoring, and sometimes modifying, these operations can often be beneficial to components of a compilation system.

This section describes the *ld-support* interface for input file inspection, and to some degree, input file data modification of those files that compose a link-edit. Two applications that employ this interface are the link-editor itself, which uses it to process debugging information within relocatable objects, and the *make(1S)* utility, which uses it to save state information.

The *ld-support* interface is composed of a support library that offers one or more support interface routines. This library is loaded as part of the link-edit process, and any support routines found are called at various stages of link-editing.

You should be familiar with the *elf(3ELF)* structures and file format when using this interface.

Invoking the Support Interface

The link-editor accepts one or more support libraries provided by either the `SGS_SUPPORT` environment variable or with the link-editor's `-S` option. The environment variable consists of a colon separated list of support libraries:

```
$ SGS_SUPPORT=./support.so.1:libldstab.so.1 cc ...
```

The `-S` option specifies a single support library. Multiple `-S` options can be specified:

```
$ LD_OPTIONS="-S./support.so.1 -Slibldstab.so.1" cc ...
```

A support library is a shared object. The link-editor opens each support library, in the order they are specified, using `dlopen(3DL)`. If both the environment variable and `-S` option are encountered, then the support libraries specified with the environment variable are processed first. Each support library is then searched, using `dlsym(3DL)`, for any support interface routines. These support routines are then called at various stages of link-editing.

A support library must be consistent with the ELF class of the link-editor being invoked, either 32-bit or 64-bit. See "32-Bit and 64-Bit Environments" on page 150 for more details.

Note – By default, the Solaris support library `libldstab.so.1` is used by the link-editor to process, and compact, compiler-generated debugging information supplied within input relocatable objects. This default processing is suppressed if you invoke the link-editor with any support libraries specified using the `-S` option. If the default processing of `libldstab.so.1` is required in addition to your support library services, add `libldstab.so.1` explicitly to the list of support libraries supplied to the link-editor.

32-Bit and 64-Bit Environments

As described in "32-Bit and 64-Bit Environments" on page 22, the 64-bit link-editor (`ld(1)`) is capable of generating 32-bit objects and the 32-bit link-editor is capable of generating 64-bit objects. Each of these objects has an associated support interface defined.

The support interface for 64-bit objects is similar to that of 32-bit objects, but ends in a `64` suffix, for example `ld_start()` and `ld_start64()`. This convention allows both implementations of the support interface to reside in a single shared object `libldstab.so.1` of each class, 32-bit and 64-bit.

The `SGS_SUPPORT` environment variable can be specified with a `_32` or `_64` suffix, and the link-editor options `-z ld32` and `-z ld64` can be used to define `-S` option requirements. These definitions will only be interpreted, respectively, by the 32-bit or 64-bit class of the link-editor. This enables both classes of support library to be specified when the class of the link-editor may not be known.

Support Interface Functions

All ld-support interfaces are defined in the header file `link.h`. All interface arguments are basic C types or ELF types. The ELF data types can be examined with the ELF access library `libelf`. See `elf(3ELF)` for a description of `libelf` contents. The following interface functions are provided by the ld-support interface, and are described in their expected order of use.

`ld_version()`

This function provides the initial handshake between the link-editor and the support library.

```
uint_t ld_version(uint_t version);
```

The link-editor calls this interface with the highest version of the ld-support interface it is capable of supporting. The support library can verify that this version is sufficient for its use, and return the version it expects to use. This version is normally `LD_SUP_VCURRENT`.

If the support library does not provide this interface, the initial support level `LD_SUP_VERSION1` is assumed.

If the support library returns a version of zero, or a value greater than the ld-support interface the link-editor supports, the support library will not be used.

`ld_start()`

This function is called after initial validation of the link-editor command line, and indicates the start of input file processing.

```
void ld_start(const char * name, const Elf32_Half type,  
             const char * caller);
```

```
void ld_start64(const char * name, const Elf64_Half type,  
               const char * caller);
```

name is the output file name being created. *type* is the output file type, which is either `ET_DYN`, `ET_REL`, or `ET_EXEC`, as defined in `sys/elf.h`. *caller* is the application calling the interface, which is normally `/usr/ccs/bin/ld`.

`ld_file()`

This function is called for each input file before any processing of the files data is carried out.

```
void ld_file(const char * name, const Elf_Kind kind, int flags,  
            Elf * elf);
```

```
void ld_file64(const char * name, const Elf_Kind kind, int flags,  
              Elf * elf);
```

name is the input file about to be processed. *kind* indicates the input file type, which is either `ELF_K_AR`, or `ELF_K_ELF`, as defined in `libelf.h`. *flags* indicates how the link-editor obtained the file, and can be one or more of the following definitions:

- LD_SUP_DERIVED – The file name was not explicitly named on the command line. It was either derived from a -l expansion, or it identifies an extracted archive member.
- LD_SUP_EXTRACTED – The file was extracted from an archive.
- LD_SUP_INHERITED – The file was obtained as a dependency of a command-line shared object.

If no *flags* values are specified then the input file has been explicitly named on the command line. *elf* is a pointer to the file's ELF descriptor.

`ld_input_section()`

This function is called for each section of the input file. This function is called before the link-editor has determined whether the section should be propagated to the output file. This function differs from `ld_section()` processing, which is only called for sections that contribute to the output file.

```
void ld_input_section(const char * name, Elf32_Shdr ** shdr,
                    Elf32_Word sndx, Elf_Data * data, Elf * elf, uint_t flags);
```

```
void ld_input_section64(const char * name, Elf64_Shdr ** shdr,
                      Elf64_Word sndx, Elf_Data * data, Elf * elf, uint_t flags);
```

name is the input section name. *shdr* is a pointer to the associated section header. *sndx* is the section index within the input file. *data* is a pointer to the associated data buffer. *elf* is a pointer to the file's ELF descriptor. *flags* is reserved for future use.

Modification of the section header is permitted by reallocating a section header and reassigning the **shdr* to the new header. The link-editor uses the section header information that **shdr* points to upon return from `ld_input_section()` to process the section.

You can modify the data by reallocating the data and reassigning the `Elf_Data` buffer's `d_buf` pointer. Any modification to the data should ensure the correct setting of the `Elf_Data` buffer's `d_size` element. For input sections that become part of the output image, setting the `d_size` element to zero effectively removes the data from the output image.

The *flags* field points to a `uint_t` data field that is initially zero filled. No flags are currently assigned, although the ability to assign flags in future updates, by the link-editor or the support library, is provided.

`ld_section()`

This function is called for each section of the input file that will be propagated to the output file, but before any processing of the section data is carried out.

```
void ld_section(const char * name, Elf32_Shdr * shdr,
               Elf32_Word sndx, Elf_Data * data, Elf * elf);
```

```
void ld_section64(const char * name, Elf64_Shdr * shdr,
                 Elf64_Word sndx, Elf_Data * data, Elf * elf);
```


name is the input section name. *shdr* is a pointer to the associated section header. *sndx* is the section index within the input file. *data* is a pointer to the associated data buffer. *elf* is a pointer to the files ELF descriptor.

You can modify the data by reallocating the data itself and reassigning the `Elf_Data` buffer's `d_buf` pointer. Any modification to the data should ensure the correct setting of the `Elf_Data` buffer's `d_size` element. For input sections that will become part of the output image, setting the `d_size` element to zero will effectively remove the data from the output image.

Note – Sections that are stripped using the link-editor's `-s` option, or discarded due to `SHT_SUNW_COMDAT` processing or `SHF_EXCLUDE` identification, are not reported to `ld_section()`. See “COMDAT Section” on page 203, and Table 7–14.

`ld_input_done()`

This function is called when input file processing is complete but before the output file is laid out.

```
void ld_input_done(uint_t flags);
```

The *flags* field points to a `uint_t` data field that is initially zero filled. No flags are currently assigned, although the ability to assign flags in future updates, by the link-editor or the support library, is provided.

`ld_atexit()`

This function is called when the link-edit is complete.

```
void ld_atexit(int status);
```

```
void ld_atexit64(int status);
```

status is the `exit(2)` code that will be returned by the link-editor and is either `EXIT_FAILURE` or `EXIT_SUCCESS`, as defined in `stdlib.h`.

Support Interface Example

The following example creates a support library that prints the section name of any relocatable object file processed as part of a 32-bit link-edit.

```
$ cat support.c
#include <link.h>
#include <stdio.h>

static int indent = 0;

void
ld_start(const char * name, const Elf32_Half type,
         const char * caller)
```

```

    {
        (void) printf("output image: %s\n", name);
    }

void
ld_file(const char * name, const Elf_Kind kind, int flags,
        Elf * elf)
{
    if (flags & LD_SUP_EXTRACTED)
        indent = 4;
    else
        indent = 2;

    (void) printf("%*sfile: %s\n", indent, "", name);
}

void
ld_section(const char * name, Elf32_Shdr * shdr, Elf32_Word sndx,
           Elf_Data * data, Elf * elf)
{
    Elf32_Ehdr * ehdr = elf32_getehdr(elf);

    if (ehdr->e_type == ET_REL)
        (void) printf("%*s section [%ld]: %s\n", indent,
                    "", (long)sndx, name);
}

```

This support library is dependent upon libelf to provide the ELF access function `elf32_getehdr(3ELF)` that is used to determine the input file type. The support library is built using:

```
$ cc -o support.so.1 -G -K pic support.c -lelf -lc
```

The following example shows the section diagnostics resulting from the construction of a trivial application from a relocatable object and a local archive library. The invocation of the support library, in addition to default debugging information processing, is brought about by the `-S` option usage.

```
$ LD_OPTIONS="-S./support.so.1 -Slibldstab.so.1" \
cc -o prog main.c -L. -lfoo
```

```

output image: prog
  file: /opt/COMPILER/crti.o
    section [1]: .shstrtab
    section [2]: .text
    .....
  file: /opt/COMPILER/crt1.o
    section [1]: .shstrtab
    section [2]: .text
    .....
  file: /opt/COMPILER/values-xt.o
    section [1]: .shstrtab
    section [2]: .text
    .....
  file: main.o

```

```
    section [1]: .shstrtab
    section [2]: .text
    .....
file: ./libfoo.a
    file: ./libfoo.a(foo.o)
    section [1]: .shstrtab
    section [2]: .text
    .....
file: /usr/lib/libc.so
file: /opt/COMPILER/crt.o
    section [1]: .shstrtab
    section [2]: .text
    .....
file: /usr/lib/libdl.so.1
```

Note – The number of sections displayed in this example have been reduced to simplify the output. Also, the files included by the compiler driver can vary.

Runtime Linker Auditing Interface

The `rtld-audit` interface enables a process to access runtime linking information regarding itself. An example of using this mechanism is the runtime profiling of shared objects described in “Profiling Shared Objects” on page 125.

The `rtld-audit` interface is implemented as an audit library that offers one or more auditing interface routines. If this library is loaded as part of a process, the audit routines are called by the runtime linker at various stages of process execution. These interfaces enable the audit library to access:

- The search for dependencies. Search paths may be substituted by the audit library.
- Information regarding loaded objects.
- Symbol bindings that occur between loaded objects. These bindings can be altered by the audit library.
- Exploitation of the lazy binding mechanism provided by procedure linkage table entries to allow auditing of function calls and their return values. The arguments to a function and its return value can be modified by the audit library. See “Procedure Linkage Table (Processor-Specific)” on page 261.

Some of these facilities can be achieved by preloading specialized shared objects. However, a preloaded object exists within the same namespace as the objects of a process. This often restricts or complicates the implementation of the preloaded shared object. The `rtld-audit` interface offers the user a unique namespace in which to execute their audit libraries. This namespace ensures that the audit library does not intrude upon the normal bindings that occur within the process.

Establishing a Namespace

When the runtime linker binds a dynamic executable with its dependencies, it generates a linked list of *link-maps* to describe the process. The link-map structure, defined in `/usr/include/sys/link.h`, describes each object within the process. The symbol search mechanism required to bind objects of an application traverses this list of link-maps. This link-map list is said to provide the *namespace* for process symbol resolution.

The runtime linker itself is also described by a link-map. This link-map is maintained on a different list from that of the application objects. The runtime linker therefore resides in its own unique name space, which prevents any direct binding of the application to services within the runtime linker. An application can only call upon the public services of the runtime linker by the filter `libdl.so.1`.

The `rtld-audit` interface employs its own link-map list on which it maintains any audit libraries. The audit libraries are thus isolated from the symbol binding requirements of the application. Inspection of the application link-map list is possible with `dlmopen(3DL)`. When used with the `RTLTD_NOLOAD` flag, `dlmopen(3DL)` allows the audit library to query an object's existence without causing its loading.

Two identifiers are defined in `/usr/include/link.h` to define the application and runtime linker link-map lists:

```
#define LM_ID_BASE      0      /* application link-map list */
#define LM_ID_LDSO     1      /* runtime linker link-map list */
```

Each `rtld-audit` support library is assigned a unique new link-map identifier.

Creating an Audit Library

An audit library is built like any other shared object. Its unique namespace within a process requires some additional care.

- The library must provide all dependency requirements.
- The library should not use system interfaces that do not provide for multiple instances of the interface within a process.

If the audit library calls `printf(3C)`, then the audit library must define a dependency on `libc`. See "Generating a Shared Object Output File" on page 44. Because the audit library has a unique namespace, symbol references cannot be satisfied by the `libc` present in the application being audited. If an audit library has a dependency on `libc`, then two versions of `libc.so.1` are loaded into the process. One version satisfies the binding requirements of the application link-map list. The other version satisfies the binding requirements of the audit link-map list.

To ensure that audit libraries are built with all dependencies recorded, use the link-editor `-z defs` option.

Some system interfaces assume that they are the only instance of their implementation within a process, for example, signals and `malloc(3C)`. Audit libraries should avoid using such interfaces, as doing so can inadvertently alter the behavior of the application.

Note – An audit library can allocate memory using `mapmalloc(3MALLOC)`, as this allocation method can exist with any allocation scheme normally employed by the application.

Invoking the Auditing Interface

The `rtld-audit` interface is enabled by one of two means. Each method implies a scope to the objects that are audited.

- *Global* auditing is enabled using the environment variable `LD_AUDIT`. The audit libraries made available by this method are provided with information regarding all dynamic objects used by the process.
- *Local* auditing is enabled through dynamic entries recorded within an object at the time it was built. The audit libraries made available by this method are provided with information regarding those dynamic objects identified for auditing.

Either method of invocation consists of a string that contains a colon-separated list of shared objects that are loaded by `dlopen(3DL)`. Each object is loaded onto its own audit link-map list. Each object is searched for audit routines using `dlsym(3DL)`. Audit routines that are found are called at various stages during the applications execution.

The `rtld-audit` interface enables multiple audit libraries to be supplied. Audit libraries that expect to be employed in this fashion should not alter the bindings that would normally be returned by the runtime linker. Alteration of these bindings can produce unexpected results from audit libraries that follow.

Secure applications can only obtain audit libraries from trusted directories. By default, the only trusted directory known to the runtime linker for 32-bit objects is `/usr/lib/secure`. For 64-bit objects, the trusted directory is `/usr/lib/secure/64`.

Recording Local Auditors

Local auditing requirements can be established when an object is built using the link-editor options `-p` or `-P`. For example, to audit `libfoo.so.1`, with the audit library `audit.so.1`, record the requirement at link-edit time using the `-p` option:

```
$ cc -G -o libfoo.so.1 -Wl,-paudit.so.1 -K pic foo.c
$ dump -Lv libfoo.so.1 | fgrep AUDIT
[3]   AUDIT      audit.so.1
```

At runtime, the existence of this audit identifier results in the audit library being loaded and information being passed to it regarding the identifying object.

With this mechanism alone, information such as searching for the identifying object has occurred prior to the audit library being loaded. To provide as much auditing information as possible, the existence of an object requiring local auditing is propagated to users of that object. For example, if an application is built that depends on `libfoo.so.1`, then the application is identified to indicate its dependencies require auditing:

```
$ cc -o main main.c libfoo.so.1
$ dump -Lv main | fgrep AUDIT
[5]  DEPAUDIT  audit.so.1
```

The auditing enabled via this mechanism results in the audit library being passed information regarding *all* of the applications explicit dependencies. This dependency auditing can also be recorded directly when creating an object by using the link-editor's `-P` option:

```
$ cc -o main main.c -Wl,-Paudit.so.1
$ dump -Lv main | fgrep AUDIT
[5]  DEPAUDIT  audit.so.1
```

Note – Auditing can be disabled at runtime by setting the environment variable `LD_NOAUDIT` to a non-null value.

Audit Interface Functions

The following functions are provided by the `rtld-audit` interface and are described in their expected order of use.

Note – References to architecture, or object class specific interfaces are reduced to their generic name to simplify the discussions. For example, a reference to `la_symbind32()` and `la_symbind64()` is specified as `la_symbind()`.

`la_version()`

This function provides the initial handshake between the runtime linker and the audit library. This interface must be provided for the audit library to be loaded.

```
uint_t la_version(uint_t version);
```

The runtime linker calls this interface with the highest *version* of the `rtld-audit` interface it is capable of supporting. The audit library can verify that this version is sufficient for its use, and return the version it expects to use. This version is normally `LAV_CURRENT`, which is defined in `/usr/include/link.h`.

If the audit library return is zero, or a version greater than the rtdl-audit interface the runtime linker supports, the audit library is discarded.

`la_activity()`

This function informs an auditor that link-map activity is occurring.

```
void la_activity(uintptr_t * cookie, uint_t flags);
```

cookie identifies the object heading the link-map. *flags* indicates the type of activity as defined in `/usr/include/link.h`:

- `LA_ACT_ADD` – Objects are being added to the link-map list.
- `LA_ACT_DELETE` – Objects are being deleted from the link-map list.
- `LA_ACT_CONSISTENT` – Object activity has been completed.

`la_objsearch()`

This function informs an auditor that an object is about to be searched for.

```
char * la_objsearch(const char * name, uintptr_t * cookie, uint_t flags);
```

name indicates the file or path name being searched for. *cookie* identifies the object initiating the search. *flags* identifies the origin and creation of *name* as defined in `/usr/include/link.h`:

- `LA_SER_ORIG` – This is the initial search name. Typically, this indicates the file name that is recorded as a `DT_NEEDED` entry, or the argument supplied to `dlopen(3DL)`.
- `LA_SER_LIBPATH` – The path name has been created from a `LD_LIBRARY_PATH` component.
- `LA_SER_RUNPATH` – The path name has been created from a runpath component.
- `LA_SER_DEFAULT` – The path name has been created from a default search path component.
- `LA_SER_CONFIG` – The path component originated from a configuration file. See the `crle(1)` man page.
- `LA_SER_SECURE` – The path component is specific to secure objects.

The return value indicates the search path name that the runtime linker should continue to process. A value of zero indicates that this path should be ignored. An audit library that simply monitors search paths should return *name*.

`la_objopen()`

This function is called each time a new object is loaded by the runtime linker.

```
uint_t la_objopen(Link_map * lmp, Lmid_t lmid, uintptr_t * cookie);
```

lmp provides the link-map structure that describes the new object. *lmid* identifies the link-map list to which the object has been added. *cookie* provides a pointer to an identifier. This identifier is initialized to the objects *lmp*. This identifier can be modified by the audit library to better identify the object to other rtdl-audit interface routines

The `la_objopen()` function returns a value that indicates the symbol bindings of interest for this object. These values can result in later calls to `la_symbind()`. The return value is a mask of the following values defined in `/usr/include/link.h`:

- `LA_FLG_BINDTO` – Audit symbol bindings *to* this object.
- `LA_FLG_BINDFROM` – Audit symbol bindings *from* this object.

See the `la_symbind()` function for more details on the use of these two flags.

A return value of zero indicates that binding information is of no interest for this object.

`la_preinit()`

This function is called once after all objects have been loaded for the application, but before transfer of control to the application occurs.

```
void la_preinit(uintptr_t * cookie);
```

`cookie` identifies the primary object that started the process, normally the dynamic executable.

`la_symbind()`

This function is called when a binding occurs between two objects that have been tagged for binding notification from `la_objopen()`.

```
uintptr_t la_symbind32(Elf32_Sym * sym, uint_t ndx,  
                    uintptr_t * refcook, uintptr_t * defcook, uint_t * flags);
```

```
uintptr_t la_symbind64(Elf64_Sym * sym, uint_t ndx,  
                    uintptr_t * refcook, uintptr_t * defcook, uint_t * flags,  
                    const char * sym_name);
```

`sym` is a constructed symbol structure, whose `sym->st_value` indicates the address of the symbol definition being bound. See `/usr/include/sys/elf.h`. `la_symbind32()` adjusts the `sym->st_name` to point to the actual symbol name. `la_symbind64()` leaves `sym->st_name` to be the index into the bound objects string table.

`ndx` indicates the symbol index within the bound object's dynamic symbol table. `refcook` describes the object making reference to this symbol. This identifier is the same identifier as passed to the `la_objopen()` function that returned `LA_FLG_BINDFROM`. `defcook` describes the object defining this symbol. This identifier is the same as passed to the `la_objopen()` that returned `LA_FLG_BINDTO`.

`flags` points to a data item that can convey information regarding the binding. This data item can be used to modify the continued auditing of procedure linkage table entries. This value is a mask of the following flags defined in `/usr/include/link.h`:

- `LA_SYMB_NOPLTENTER` – The `la_pltenter()` function is *not* called for this symbol.

- `LA_SYMB_NOPLTEXTIT` – The `la_pltexit()` function is *not* called for this symbol.
- `LA_SYMB_DLSYM` – The symbol binding occurred as a result of calling `dlsym(3DL)`.
- `LA_SYMB_ALTVALUE` (`LAV_VERSION2`) – An alternate value was returned for the symbol value by a previous call to `la_symbind()`.

By default, if `la_pltenter()` or `la_pltexit()` functions exist, they are called after `la_symbind()` for procedure linkage table entries each time the symbol is referenced. See also “Audit Interface Limitations” on page 163.

The return value indicates the address to which control should be passed following this call. An audit library that simply monitors symbol binding should return the value of `sym->st_value` so that control is passed to the bound symbol definition. An audit library can intentionally redirect a symbol binding by returning a different value.

`sym_name`, which is applicable for `la_symbind64()` only, contains the name of the symbol being processed. This name is available in the `sym->st_name` field for the 32-bit interface.

`la_pltenter()`

These functions are called on a SPARC and x86 system respectively. These functions are called when a procedure linkage table entry, between two objects that have been tagged for binding notification, is called.

```
uintptr_t la_sparcv8_pltenter(Elf32_Sym * sym, uint_t ndx,
                             uintptr_t * refcook, uintptr_t * defcook,
                             La_sparcv8_regs * regs, uint_t * flags);
```

```
uintptr_t la_sparcv9_pltenter(Elf64_Sym * sym, uint_t ndx,
                             uintptr_t * refcook, uintptr_t * defcook,
                             La_sparcv9_regs * regs, uint_t * flags,
                             const char * sym_name);
```

```
uintptr_t la_i86_pltenter(Elf32_Sym * sym, uint_t ndx,
                          uintptr_t * refcook, uintptr_t * defcook,
                          La_i86_regs * regs, uint_t * flags);
```

`sym`, `ndx`, `refcook`, `defcook` and `sym_name` provide the same information as passed to `la_symbind()`.

`regs` points to the out registers on a SPARC system, and the stack and frame registers on a x86 system, as defined in `/usr/include/link.h`.

`flags` points to a data item that can convey information regarding the binding. This data can be used to modify the continuing auditing of this procedure linkage table entry. This data item is the same as pointed to by the `flags` from `la_symbind()`. This value is a mask of the following flags defined in `/usr/include/link.h`:

- `LA_SYMB_NOPLTENTER` – `la_pltenter()` is *not* be called again for this symbol.

- `LA_SYMB_NOPLTEXTIT` – `la_pltexit()` is *not* be called for this symbol.

The return value indicates the address to which control should be passed following this call. An audit library that simply monitors symbol binding should return the value of `sym->st_value` so that control is passed to the bound symbol definition. An audit library can intentionally redirect a symbol binding by returning a different value.

`la_pltexit()`

This function is called when a procedure linkage table entry, between two objects that have been tagged for binding notification, returns. This function is called before control reaches the caller.

```
uintptr_t la_pltexit(Elf32_Sym * sym, uint_t ndx, uintptr_t * refcook,
                    uintptr_t * defcook, uintptr_t retval);
```

```
uintptr_t la_pltexit64(Elf64_Sym * sym, uint_t ndx, uintptr_t * refcook,
                      uintptr_t * defcook, uintptr_t retval, const char * sym_name);
```

`sym`, `ndx`, `refcook`, `defcook` and `sym_name` provide the same information as passed to `la_symbind()`. `retval` is the return code from the bound function. An audit library that simply monitors symbol binding should return `retval`. An audit library can intentionally return a different value.

Note – The `la_pltexit()` interface is experimental. See “Audit Interface Limitations” on page 163.

`la_objclose()`

This function is called after any termination code for an object has been executed and prior to the object being unloaded.

```
uint_t la_objclose(uintptr_t * cookie);
```

`cookie` identifies the object, and was obtained from a previous `la_objopen()`. Any return value is presently ignored.

Audit Interface Example

The following simple example creates an audit library that prints the name of each shared object dependency loaded by the dynamic executable `date(1)`.

```
$ cat audit.c
#include <link.h>
#include <stdio.h>

uint_t
la_version(uint_t version)
{
```

```

        return (LAV_CURRENT);
    }

    uint_t
    la_objopen(Link_map * lmp, Lmid_t lmid, uintptr_t * cookie)
    {
        if (lmid == LM_ID_BASE)
            (void) printf("file: %s loaded\n", lmp->l_name);
        return (0);
    }
}
$ cc -o audit.so.1 -G -K pic -z defs audit.c -lmapmalloc -lc
$ LD_AUDIT=./audit.so.1 date
file: date loaded
file: /usr/lib/libc.so.1 loaded
file: /usr/lib/libdl.so.1 loaded
file: /usr/lib/locale/en_US/en_US.so.2 loaded
Thur Aug 10 17:03:55 PST 2000

```

Audit Interface Demonstrations

A number of demonstration applications that use the `rtdl-audit` interface are provided in the `SUNWosdem` package under `/usr/demo/link_audit`:

`sotruss`

This demo provides tracing of procedure calls between the dynamic objects of a named application.

`whocalls`

This demo provides a stack trace for a specified function whenever called by a named application.

`perfcnt`

This demo traces the amount of time spent in each function for a named application.

`symbindrep`

This demo reports all symbol bindings performed to load a named application.

`sotruss(1)` and `whocalls(1)` are included in the `SUNWtoo` package. `perfcnt` and `symbindrep` are example programs. They are not intended for use in a production environment.

Audit Interface Limitations

There are some limitations to the use of the `la_pltexit()` family. These limitations stem from the need to insert an extra stack frame between the caller and callee to provide a `la_pltexit()` return value. This requirement is not a problem when calling just the `la_pltenter()` routines, as. In this case, any intervening stack can be cleaned up prior to transferring control to the destination function.

Because of these limitations, `la_pltexit()` should be considered an experimental interface. When in doubt, avoid the use of the `la_pltexit()` routines.

Functions That Directly Inspect the Stack

A small number of functions exist that directly inspect the stack or make assumptions of its state. Some examples of these functions are the `setjmp(3C)` family, `vfork(2)`, and any function that returns a structure, not a pointer to a structure. These functions are compromised by the extra stack created to support `la_pltexit()`.

The runtime linker cannot detect functions of this type, and thus the audit library creator is responsible for disabling `la_pltexit()` for such routines.

Runtime Linker Debugger Interface

The runtime linker performs many operations including the mapping of objects into memory and the binding of symbols. Debugging programs often need to access information that describes these runtime linker operations as part of analyzing an application. These debugging programs run as a separate process to the application they are analyzing.

This section describes the `rtld-debugger` interface for monitoring and modifying a dynamically linked application from another process. The architecture of this interface follows the model used in `libthread_db(3THR)`.

When using the `rtld-debugger` interface, at least two processes are involved:

- One or more *target* processes. The target processes must be dynamically linked and use the runtime linker `/usr/lib/ld.so.1` for 32-bit processes, or `/usr/lib/64/ld.so.1` for 64-bit processes.
- A *controlling* process links with the `rtld-debugger` interface library and uses it to inspect the dynamic aspects of the target processes. A 64-bit controlling process can debug both 64-bit and 32-bit targets. However, a 32-bit controlling process is limited to 32-bit targets.

The most anticipated use of the `rtld-debugger` interface is when the controlling process is a debugger and its target is a dynamic executable.

The `rtld-debugger` interface enables the following activities with a target process:

- Initial rendezvous with the runtime linker.
- Notification of the loading and unloading of dynamic objects.
- Retrieval of information regarding any loaded objects.
- Stepping over procedure linkage table entries.

- Enabling object padding.

Interaction Between Controlling and Target Process

To be able to inspect and manipulate a target process, the `rtld-debugger` interface employs an *exported* interface, an *imported* interface, and *agents* for communicating between these interfaces.

The controlling process is linked with the `rtld-debugger` interface provided by `librtld_db.so.1`, and makes requests of the interface exported from this library. This interface is defined in `/usr/include/rtld_db.h`. In turn, `librtld_db.so.1` makes requests of the interface imported from the controlling process. This interaction allows the `rtld-debugger` interface to:

- Look up symbols in a target process.
- Read and write memory in the target process.

The imported interface consists of a number of `proc_service` routines that most debuggers already employ to analyze processes. These routines are described in “Debugger Import Interface” on page 175.

The `rtld-debugger` interface assumes that the process being analyzed is stopped when requests are made of the `rtld-debugger` interface. If this halt does not occur, data structures within the runtime linker of the target process might not be in a consistent state for examination.

The flow of information between `librtld_db.so.1`, the controlling process (debugger) and the target process (dynamic executable) is diagrammed in the following figure.

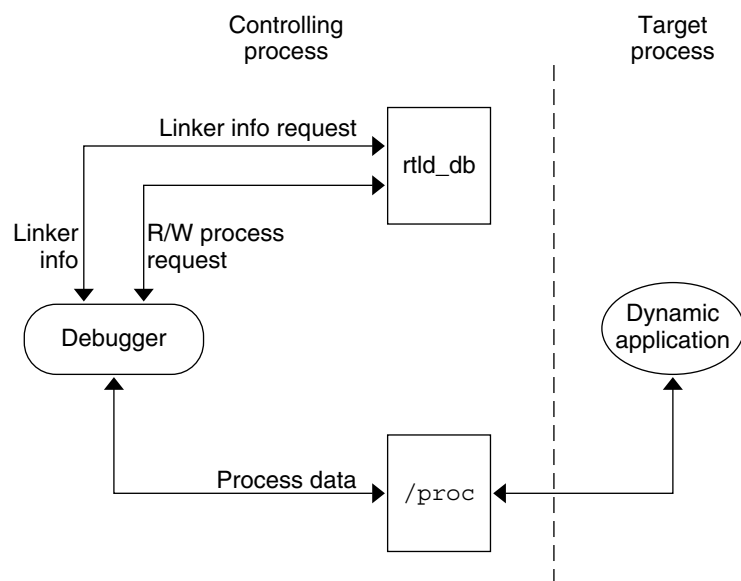


FIGURE 6-1 rtld-debugger Information Flow

Note – The rtld-debugger interface is dependent upon the `proc_service` interface, `/usr/include/proc_service.h`, which is considered experimental. The rtld-debugger interface might have to track changes in the `proc_service` interface as it evolves.

A sample implementation of a controlling process that uses the rtld-debugger interface is provided in the `SUNWosdem` package under `/usr/demo/librtld_db`. This debugger, `rdb`, provides an example of using the `proc_service` imported interface, and shows the required calling sequence for all `librtld_db.so.1` exported interfaces. The following sections describe the rtld-debugger interfaces. More detailed information can be obtained by examining the sample debugger.

Debugger Interface Agents

An agent provides an opaque handle that can describe internal interface structures. The agent also provides a mechanism of communication between the exported and imported interfaces. The rtld-debugger interface is intended to be used by a debugger which can manipulate several processes at the same time, these agents are used to identify the process.

`struct ps_prochandle`

Is an opaque structure that is created by the controlling process to identify the target process that is passed between the exported and imported interface.

`struct rd_agent`

Is an opaque structure created by the rtdl-debugger interface that identifies the target process that is passed between the exported and imported interface.

Debugger Exported Interface

This section describes the various interfaces exported by the `/usr/lib/librtdl_db.so.1` audit library. It is broken down into functional groups.

Agent Manipulation Interfaces

`rd_init()`

This function establishes the rtdl-debugger version requirements. The base *version* is defined as `RD_VERSION1`. The current *version* is always defined by `RD_VERSION`.

```
rd_err_e rd_init(int version);
```

Version `RD_VERSION2`, added in Solaris 8 10/00, extends the `rd_loadobj_t` structure. See the `rl_flags`, `rl_bend` and `rl_dynamic` fields in “Scanning Loadable Objects” on page 168.

Version `RD_VERSION3`, added in Solaris 8 01/01, extends the `rd_plt_info_t` structure. See the `pi_baddr` and `pi_flags` fields in “Procedure Linkage Table Skipping” on page 172.

If the version requirement of the controlling process is greater than the rtdl-debugger interface available, then `RD_NOCAPAB` is returned.

`rd_new()`

This function creates a new exported interface agent.

```
rd_agent_t * rd_new(struct ps_prochandle * php);
```

php is a cookie created by the controlling process to identify the target process. This cookie is used by the imported interface offered by the controlling process to maintain context, and is opaque to the rtdl-debugger interface.

`rd_reset()`

This function resets the information within the agent based off the same *ps_prochandle* structure given to `rd_new()`.

```
rd_err_e rd_reset(struct rd_agent * rdap);
```

This function is called when a target process is restarted.

`rd_delete()`

This function deletes an agent and frees any state associated with it.

```
void rd_delete(struct rd_agent * rdap);
```

Error Handling

The following error states can be returned by the `rtld-debugger` interface (defined in `rtld_db.h`):

```
typedef enum {
    RD_ERR,
    RD_OK,
    RD_NOCAPAB,
    RD_DBERR,
    RD_NOBASE,
    RD_NODYNAM,
    RD_NOMAPS
} rd_err_e;
```

The following interfaces can be used to gather the error information.

`rd_errstr()`

This function returns a descriptive error string describing the error code `rderr`.

```
char * rd_errstr(rd_err_e rderr);
```

`rd_log()`

This function turns logging on (1) or off (0).

```
void rd_log(const int onoff);
```

When logging is turned on, the imported interface function `ps_plog()` provided by the controlling process, is called with more detailed diagnostic information.

Scanning Loadable Objects

You can obtain information for each object maintained on the runtime linkers link-map is achieved by using the following structure, defined in `rtld_db.h`:

```
typedef struct rd_loadobj {
    psaddr_t    rl_nameaddr;
    unsigned    rl_flags;
    psaddr_t    rl_base;
    psaddr_t    rl_data_base;
    unsigned    rl_lmident;
    psaddr_t    rl_refnameaddr;
    psaddr_t    rl_plt_base;
    unsigned    rl_plt_size;
    psaddr_t    rl_bend;
    psaddr_t    rl_padstart;
    psaddr_t    rl_padend;
```



```

        psaddt_t      rl_dynamic;
} rd_loadobj_t;

```

Notice that all addresses given in this structure, including string pointers, are addresses in the target process and not in the address space of the controlling process itself.

`rl_nameaddr`

A pointer to a string that contains the name of the dynamic object.

`rl_flags`

With revision `RD_VERSION2`, dynamically loaded relocatable objects are identified with `RD_FLG_MEM_OBJECT`.

`rl_base`

The base address of the dynamic object.

`rl_data_base`

The base address of the data segment of the dynamic object.

`rl_lmident`

The link-map identifier (see “Establishing a Namespace” on page 156).

`rl_refnameaddr`

If the dynamic object is a standard filter, then this points to the name of the filtees.

`rl_plt_base, rl_plt_size`

These elements are present for backward compatibility and are currently unused.

`rl_bend`

The end address of the object (text + data + bss). With revision `RD_VERSION2`, a dynamically loaded relocatable object will cause this element to point to the end of the created object, which will include its section headers.

`rl_padstart`

The base address of the padding before the dynamic object (refer to “Dynamic Object Padding” on page 174).

`rl_padend`

The base address of the padding after the dynamic object (refer to “Dynamic Object Padding” on page 174).

`rl_dynamic`

This field, added with `RD_VERSION2`, provides the base address of the object’s dynamic section, which allows reference to such entries as `DT_CHECKSUM` (see Table 7-42).

The `rd_loadobj_iter()` routine uses this object data structure to access information from the runtime linker’s link-map lists:

`rd_loadobj_iter()`

This function iterates over all dynamic objects currently loaded in the target process.

```

typedef int rl_iter_f(const rd_loadobj_t *, void *);

rd_err_e rd_loadobj_iter(rd_agent_t * rap, rl_iter_f * cb,
                        void * clnt_data);

```

On each iteration the imported function specified by *cb* is called. *clnt_data* can be used to pass data to the *cb* call. Information about each object is returned via a pointer to a volatile (stack allocated) `rd_loadobj_t` structure.

Return codes from the *cb* routine are examined by `rd_loadobj_iter()` and have the following meaning:

- 1 – continue processing link-maps.
- 0 – stop processing link-maps and return control to the controlling process.

`rd_loadobj_iter()` returns `RD_OK` on success. A return of `RD_NOMAPS` indicates the runtime linker has not yet loaded the initial link-maps.

Event Notification

A controlling process can track certain events that occur within the scope of the runtime linker that. These events are:

`RD_PREINIT`

The runtime linker has loaded and relocated all the dynamic objects and is about to start calling the `.init` sections of each object loaded.

`RD_POSTINIT`

The runtime linker has finished calling all of the `.init` sections and is about to transfer control to the primary executable.

`RD_DLACTIONIVITY`

The runtime linker has been invoked to either load or unload a dynamic object.

These events can be monitored using the following interface, defined in `sys/link.h` and `rtld_db.h`:

```

typedef enum {
    RD_NONE = 0,
    RD_PREINIT,
    RD_POSTINIT,
    RD_DLACTIONIVITY
} rd_event_e;

/*
 * Ways that the event notification can take place:
 */
typedef enum {
    RD_NOTIFY_BPT,
    RD_NOTIFY_AUTOBPT,
    RD_NOTIFY_SYSCALL
} rd_notify_e;

```

```

/*
 * Information on ways that the event notification can take place:
 */
typedef struct rd_notify {
    rd_notify_e    type;
    union {
        psaddr_t    bptaddr;
        long        syscallno;
    } u;
} rd_notify_t;

```

The following functions track events:

`rd_event_enable()`

This function enables (1) or disables (0) event monitoring.

```
rd_err_e rd_event_enable(struct rd_agent * rdap, int onoff);
```

Note – Presently, for performance reasons, the runtime linker ignores event disabling. The controlling process should not assume that a given break-point can not be reached because of the last call to this routine.

`rd_event_addr()`

This function specifies how the controlling program is notified of a given event.

```
rd_err_e rd_event_addr(rd_agent_t * rdap, rd_event_e event,
    rd_notify_t * notify);
```

Depending on the event type, the notification of the controlling process takes place by calling a benign, cheap system call that is identified by `notify->u.syscallno`, or executing a break point at the address specified by `notify->u.bptaddr`. The controlling process is responsible for tracing the system call or place the actual break-point.

When an event has occurred, additional information can be obtained by this interface, defined in `rtld_db.h`:

```

typedef enum {
    RD_NOSTATE = 0,
    RD_CONSISTENT,
    RD_ADD,
    RD_DELETE
} rd_state_e;

typedef struct rd_event_msg {
    rd_event_e    type;
    union {
        rd_state_e    state;
    } u;
} rd_event_msg_t;

```

The `rd_state_e` values are:

`RD_NOSTATE`

There is no additional state information available.

`RD_CONSISTANT`

The link-maps are in a stable state and can be examined.

`RD_ADD`

A dynamic object is in the process of being loaded and the link-maps are not in a stable state. They should not be examined until the `RD_CONSISTANT` state is reached.

`RD_DELETE`

A dynamic object is in the process of being deleted and the link-maps are not in a stable state. They should not be examined until the `RD_CONSISTANT` state is reached.

The `rd_event_getmsg()` function is used to obtain this event state information.

`rd_event_getmsg()`

This function provides additional information concerning an event.

```
rd_err_e rd_event_getmsg(struct rd_agent * rdap, rd_event_msg_t * msg);
```

The following table shows the possible state for each of the different event types.

<code>RD_PREINIT</code>	<code>RD_POSTINIT</code>	<code>RD_DLACTIVITY</code>
<code>RD_NOSTATE</code>	<code>RD_NOSTATE</code>	<code>RD_CONSISTANT</code>
		<code>RD_ADD</code>
		<code>RD_DELETE</code>

Procedure Linkage Table Skipping

The `rtld-debugger` interface enables a controlling process to skip over procedure linkage table entries. When a controlling process, such as a debugger, is asked to step into a function for the first time, the procedure linkage table processing, causes control to be passed to the runtime linker to search for the function definition.

The following interface enables a controlling process to step over the runtime linker's procedure linkage table processing. The controlling process can determine when a procedure linkage table entry is encountered based on external information provided in the ELF file.

Once a target process has stepped into a procedure linkage table entry, the process calls the `rd_plt_resolution()` interface:

`rd_plt_resolution()`

This function returns the resolution state of the current procedure linkage table entry and information on how to skip it.

```
rd_err_e rd_plt_resolution(rd_agent_t * rdap, paddr_t pc,
                          lwpid_t lwpid, paddr_t plt_base, rd_plt_info_t * rpi);
```

pc represents the first instruction of the procedure linkage table entry. *lwpid* provides the lwp identifier and *plt_base* provides the base address of the procedure linkage table. These three variables provide information sufficient for various architectures to process the procedure linkage table.

rpi provides detailed information regarding the procedure linkage table entry as defined in the following data structure, defined in `rtld_db.h`:

```
typedef enum {
    RD_RESOLVE_NONE,
    RD_RESOLVE_STEP,
    RD_RESOLVE_TARGET,
    RD_RESOLVE_TARGET_STEP
} rd_skip_e;

typedef struct rd_plt_info {
    rd_skip_e      pi_skip_method;
    long          pi_nstep;
    psaddr_t      pi_target;
    psaddr_t      pi_baddr;
    unsigned int  pi_flags;
} rd_plt_info_t;

#define RD_FLG_PI_PLTBOUND    0x0001
```

The elements of the `rd_plt_info_t` structure are:

`pi_skip_method`

Identifies how the procedure linkage table entry can be traversed. This method is set to one of the `rd_skip_e` values.

`pi_nstep`

Identifies how many instructions to step over when `RD_RESOLVE_STEP` or `RD_RESOLVE_TARGET_STEP` are returned.

`pi_target`

Specifies the address at which to set a breakpoint when `RD_RESOLVE_TARGET_STEP` or `RD_RESOLVE_TARGET` are returned.

`pi_baddr`

The procedure linkage table destination address, added with `RD_VERSION3`. When the `RD_FLG_PI_PLTBOUND` flag of the `pi_flags` field is set, this element identifies the resolved (bound) destination address.

`pi_flags`

A flags field, added with `RD_VERSION3`. The flag `RD_FLG_PI_PLTBOUND` identifies the procedure linkage entry as having been resolved (bound) to its destination

address, which is available in the `pi_baddr` field.

The following scenarios are possible from the `rd_plt_info_t` return values:

- The first call through this procedure linkage table must be resolved by the runtime linker. In this case, the `rd_plt_info_t` contains:

```
{RD_RESOLVE_TARGET_STEP, M, <BREAK>, 0, 0}
```

The controlling process sets a breakpoint at `BREAK` and continues the target process. When the breakpoint is reached, the procedure linkage table entry processing has finished. The controlling process can then step `M` instructions to the destination function. Notice that the bound address (`pi_baddr`) has not been set since this is the first call through a procedure linkage table entry.

- On the `N`th time through this procedure linkage table, `rd_plt_info_t` contains:

```
{RD_RESOLVE_STEP, M, 0, <BoundAddr>, RD_FLG_PI_PLTBOUND}
```

The procedure linkage table entry has already been resolved and the controlling process can step `M` instructions to the destination function. The address that the procedure linkage table entry is bound to is `<BoundAddr>` and the `RD_FLG_PI_PLTBOUND` bit has been set in the flags field.

Dynamic Object Padding

The default behavior of the runtime linker relies on the operating system to load dynamic objects where they can be most efficiently referenced. Some controlling processes benefit from the existence of padding around the objects loaded into memory of the target process. This interface enables a controlling process to request this padding.

```
rd_objpad_enable()
```

This function enables or disables the padding of any subsequently loaded objects with the target process. Padding occurs on both sides of the loaded object.

```
rd_err_e rd_objpad_enable(struct rd_agent * rdap, size_t padsiz);
```

padsiz specifies the size of the padding, in bytes, to be preserved both before and after any objects loaded into memory. This padding is reserved as a memory mapping using `mmap(2)` with `PROT_NONE` permissions and the `MAP_NORESERVE` flag. Effectively, the runtime linker reserves areas of the virtual address space of the target process adjacent to any loaded objects. These areas can later be utilized by the controlling process.

A *padsiz* of 0 disables any object padding for later objects.

Note – Reservations obtained using `mmap(2)` from `/dev/zero` with `MAP_NORESERVE` can be reported using the `proc(1)` facilities and by referring to the link-map information provided in `rd_loadobj_t`.

Debugger Import Interface

The imported interface that a controlling process must provide to `librtld_db.so.1` is defined in `/usr/include/proc_service.h`. A sample implementation of these `proc_service` functions can be found in the `rdb` demonstration debugger. The `rtld-debugger` interface uses only a subset of the `proc_service` interfaces available. Future versions of the `rtld-debugger` interface might take advantage of additional `proc_service` interfaces without creating an incompatible change.

The following interfaces are currently being used by the `rtld-debugger` interface:

`ps_pauxv()`

This function returns a pointer to a copy of the `auxv` vector.

```
ps_err_e ps_pauxv(const struct ps_prochandle * ph, auxv_t ** aux);
```

Because the `auxv` vector information is copied to an allocated structure, the pointer remains as long as the `ps_prochandle` is valid.

`ps_pread()`

This function reads data from the target process.

```
ps_err_e ps_pread(const struct ps_prochandle * ph, paddr_t addr,
                 char * buf, int size);
```

From address `addr` in the target process, `size` bytes are copied to `buf`.

`ps_pwrite()`

This function writes data to the target process.

```
ps_err_e ps_pwrite(const struct ps_prochandle * ph, paddr_t addr,
                  char * buf, int size);
```

`size` bytes from `buf` are copied into the target process at address `addr`.

`ps_plog()`

This function is called with additional diagnostic information from the `rtld-debugger` interface.

```
void ps_plog(const char * fmt, ...);
```

The controlling process determines where, or if, to log this diagnostic information. The arguments to `ps_plog()` follow the `printf(3C)` format.

`ps_pglobal_lookup()`

This function searches for the symbol in the target process.

```
ps_err_e ps_pglobal_lookup(const struct ps_prochandle * ph,
                          const char * obj, const char * name, ulong_t * sym_addr);
```

The symbol named *name* is searched for within the object named *obj* within the target process *ph*. If the symbol is found, the symbol address is stored in *sym_addr*.

```
ps_pglobal_sym()
```

This function searches for the symbol in the target process.

```
ps_err_e ps_pglobal_sym(const struct ps_prochandle * ph,
                       const char * obj, const char * name, ps_sym_t * sym_desc);
```

The symbol named *name* is searched for within the object named *obj* within the target process *ph*. If the symbol is found, the symbol descriptor is stored in *sym_desc*.

In the event that the rtdl-debugger interface needs to find symbols within the application or runtime linker prior to any link-map creation, the following reserved values for *obj* are available:

```
#define PS_OBJ_EXEC ((const char *)0x0) /* application id */
#define PS_OBJ_LDSO ((const char *)0x1) /* runtime linker id */
```

The controlling process can use the procfs file system for these objects, using the following pseudo code:

```
ioctl(.., PIOCNAUXV, ...)      - obtain AUX vectors
ldsoaddr = auxv[AT_BASE];
ldsofd = ioctl(..., PIOCOPENM, &ldsoaddr);

/* process elf information found in ldsofd ... */

execfd = ioctl(.., PIOCOPENM, 0);

/* process elf information found in execfd ... */
```

Once the file descriptors are found, the ELF files can be examined for their symbol information by the controlling program.

Object File Format

This chapter describes the executable and linking format (ELF) of the object files produced by the assembler and link-editor. There are three main types of object files:

- A relocatable file holds sections containing code and data. These files are suitable to be linked with other object files to create executable files, shared object files, or another relocatable object.
- An executable file holds a program that is ready to execute. The file specifies how `exec(2)` creates a program's process image.
- A shared object file holds code and data suitable to be linked in two contexts. First, the link-editor can process this file with other relocatable and shared object files to create other object files. Second, the runtime linker combines this file with a dynamic executable file and other shared objects to create a process image.

The first section in this chapter, "File Format" on page 177, focuses on the format of object files and how that pertains to creating programs. The second section, "Dynamic Linking" on page 236, focuses on how the format pertains to loading programs.

Programs manipulate object files with the functions contained in the ELF access library, `libelf`. Refer to the `elf(3ELF)` man page for a description of `libelf` contents. Sample source code that uses `libelf` is provided in the `SUNWosdem` package under the `/usr/demo/ELF` directory.

File Format

Object files participate in both program linking and program execution. For convenience and efficiency, the object file format provides parallel views of a file's contents, reflecting the differing needs of these activities. The following figure shows an object file's organization.

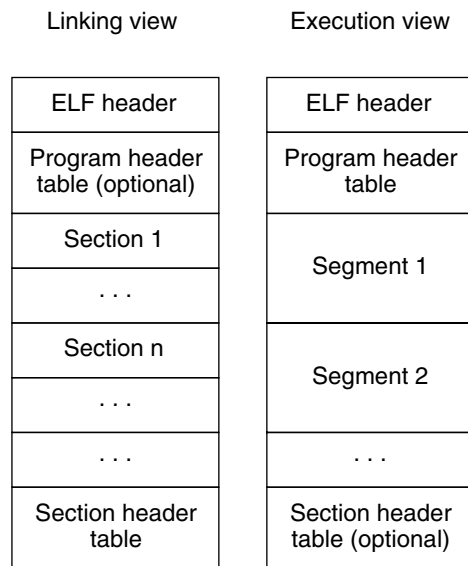


FIGURE 7-1 Object File Format

An ELF header resides at the beginning of an object file and holds a *road map* describing the file's organization.

Note – Only the ELF header has a fixed position in the file. The flexibility of the ELF format requires no specified order for header tables, sections or segments. However, this figure is typical of the layout used in Solaris.

Sections represent the smallest indivisible units that can be processed within an ELF file. *Segments* are a collection of sections that represent the smallest individual units that can be mapped to a memory image by `exec(2)` or by the runtime linker.

Sections hold the bulk of object file information for the linking view: instructions, data, symbol table, relocation information, and so on. Descriptions of sections appear in the first part of this chapter. The second part of this chapter discusses segments and the program execution view of the file.

A program header table, if present, tells the system how to create a process image. Files used to generate a process image, executables and shared objects, must have a program header table; relocatable objects do not need such a table.

A section header table contains information describing the file's sections. Every section has an entry in the table. Each entry gives information such as the section name, the section size, and so forth. Files used in link-editing must have a section header table; other object files might or might not have one.

Data Representation

The object file format supports various processors with 8-bit bytes, 32-bit and 64-bit architectures. Nevertheless, it is intended to be extensible to larger (or smaller) architectures. Table 7-1 and Table 7-2 list the 32-bit and 64-bit data types.

Object files represent some control data with a machine-independent format, making it possible to identify object files and interpret their contents in a common way. The remaining data in an object file use the encoding of the target processor, regardless of the machine on which the file was created.

TABLE 7-1 ELF 32-Bit Data Types

Name	Size	Alignment	Purpose
Elf32_Addr	4	4	Unsigned program address
Elf32_Half	2	2	Unsigned medium integer
Elf32_Off	4	4	Unsigned file offset
Elf32_Sword	4	4	Signed integer
Elf32_Word	4	4	Unsigned integer
unsigned char	1	1	Unsigned small integer

TABLE 7-2 ELF 64-Bit Data Types

Name	Size	Alignment	Purpose
Elf64_Addr	8	8	Unsigned program address
Elf64_Half	2	2	Unsigned medium integer
Elf64_Off	8	8	Unsigned file offset
Elf64_Sword	4	4	Signed integer
Elf64_Word	4	4	Unsigned integer
Elf64_Xword	8	8	Unsigned long integer
Elf64_Sxword	8	8	Signed long integer
unsigned char	1	1	Unsigned small integer

All data structures that the object file format defines follow the natural size and alignment guidelines for the relevant class. If necessary, data structures contain explicit padding to ensure 4-byte alignment for 4-byte objects, to force structure sizes to a multiple of 4, and so forth. Data also have suitable alignment from the beginning of the file. Thus, for example, a structure containing an `Elf32_Addr` member will be aligned on a 4-byte boundary within the file, and a structure containing an `Elf64_Addr` member will be aligned on an 8-byte boundary.

Note – For portability, ELF uses no bit-fields.

ELF Header

Some object file control structures can grow because the ELF header contains their actual sizes. If the object file format changes, a program can encounter control structures that are larger or smaller than expected. Programs might therefore ignore extra information. The treatment of missing information depends on context and will be specified if and when extensions are defined.

The ELF header has the following structure, defined in `sys/elf.h`:

```
#define EI_NIDENT      16

typedef struct {
    unsigned char    e_ident[EI_NIDENT];
    Elf32_Half      e_type;
    Elf32_Half      e_machine;
    Elf32_Word      e_version;
    Elf32_Addr      e_entry;
    Elf32_Off       e_phoff;
    Elf32_Off       e_shoff;
    Elf32_Word      e_flags;
    Elf32_Half      e_ehsize;
    Elf32_Half      e_phentsize;
    Elf32_Half      e_phnum;
    Elf32_Half      e_shentsize;
    Elf32_Half      e_shnum;
    Elf32_Half      e_shstrndx;
} Elf32_Ehdr;

typedef struct {
    unsigned char    e_ident[EI_NIDENT];
    Elf64_Half      e_type;
    Elf64_Half      e_machine;
    Elf64_Word      e_version;
    Elf64_Addr      e_entry;
    Elf64_Off       e_phoff;
    Elf64_Off       e_shoff;
    Elf64_Word      e_flags;
    Elf64_Half      e_ehsize;
```

```

Elf64_Half    e_phentsize;
Elf64_Half    e_phnum;
Elf64_Half    e_shentsize;
Elf64_Half    e_shnum;
Elf64_Half    e_shstrndx;
} Elf64_Ehdr;

```

The elements of this structure are:

`e_ident`

The initial bytes mark the file as an object file and provide machine-independent data with which to decode and interpret the file's contents. Complete descriptions appear in "ELF Identification" on page 184.

`e_type`

Identifies the object file type, as listed in the following table.

TABLE 7-3 ELF File Identifiers

Name	Value	Meaning
ET_NONE	0	No file type
ET_REL	1	Relocatable file
ET_EXEC	2	Executable file
ET_DYN	3	Shared object file
ET_CORE	4	Core file
ET_LOPROC	0xff00	Processor-specific
ET_HIPROC	0xffff	Processor-specific

Although the core file contents are unspecified, type `ET_CORE` is reserved to mark the file. Values from `ET_LOPROC` through `ET_HIPROC` (inclusive) are reserved for processor-specific semantics. Other values are reserved and will be assigned to new object file types as necessary.

`e_machine`

Specifies the required architecture for an individual file. Relevant architectures are listed in the following table.

TABLE 7-4 ELF Machines

Name	Value	Meaning
EM_NONE	0	No machine
EM_SPARC	2	SPARC
EM_386	3	Intel 80386

TABLE 7-4 ELF Machines (Continued)

Name	Value	Meaning
EM_SPARC32PLUS	18	Sun SPARC 32+
EM_SPARCV9	43	SPARC V9

Other values are reserved and will be assigned to new machines as necessary (see `sys/elf.h`). Processor-specific ELF names use the machine name to distinguish them. For example, the flags defined in Table 7-5 use the prefix `EF_`. A flag named `WIDGET` for the `EM_XYZ` machine would be called `EF_XYZ_WIDGET`.

e_version

Identifies the object file version, as listed in the following table.

TABLE 7-5 ELF Versions

Name	Value	Meaning
EV_NONE	0	Invalid version
EV_CURRENT	>=1	Current version

The value 1 signifies the original file format. The value of `EV_CURRENT` changes as necessary to reflect the current version number.

e_entry

The virtual address to which the system first transfers control, thus starting the process. If the file has no associated entry point, this member holds zero.

e_phoff

The program header table's file offset in bytes. If the file has no program header table, this member holds zero.

e_shoff

The section header table's file offset in bytes. If the file has no section header table, this member holds zero.

e_flags

Processor-specific flags associated with the file. Flag names take the form `EF_machine_flag`. This member is presently zero for x86. The SPARC flags are listed in the following table.

TABLE 7-6 SPARC: ELF Flags

Name	Value	Meaning
EF_SPARC_EXT_MASK	0xffff00	Vendor Extension mask
EF_SPARC_32PLUS	0x000100	Generic V8+ features

TABLE 7-6 SPARC: ELF Flags (Continued)

Name	Value	Meaning
EF_SPARC_SUN_US1	0x000200	Sun UltraSPARC™ 1 Extensions
EF_SPARC_HAL_R1	0x000400	HAL R1 Extensions
EF_SPARC_SUN_US3	0x000800	Sun UltraSPARC 3 Extensions
EF_SPARCV9_MM	0x3	Mask for Memory Model
EF_SPARCV9_TSO	0x0	Total Store Ordering
EF_SPARCV9_PSO	0x1	Partial Store Ordering
EF_SPARCV9_RMO	0x2	Relaxed Memory Ordering

e_ehsize

The ELF header's size in bytes.

e_phentsize

The size in bytes of one entry in the file's program header table. All entries are the same size.

e_phnumThe number of entries in the program header table. The product of `e_phentsize` and `e_phnum` gives the table's size in bytes. If a file has no program header table, `e_phnum` holds the value zero.**e_shentsize**

A section header's size in bytes. A section header is one entry in the section header table. All entries are the same size.

e_shnumThe number of entries in the section header table. The product of `e_shentsize` and `e_shnum` gives the section header table's size in bytes. If a file has no section header table, `e_shnum` holds the value zero.

If the number of sections is greater than or equal to `SHN_LORESERVE (0xff00)`, this member has the value zero and the actual number of section header table entries is contained in the `sh_size` field of the section header at index 0. Otherwise, the `sh_size` member of the initial entry contains 0.

e_shstrndxThe section header table index of the entry that is associated with the section name string table. If the file has no section name string table, this member holds the value `SHN_UNDEF`.

If the section name string table section index is greater than or equal to `SHN_LORESERVE (0xff00)`, this member has the value `SHN_XINDEX (0xffff)` and the actual index of the section name string table section is contained in the `sh_link` field of the section header at index 0. Otherwise, the `sh_link` member of the initial entry contains 0.

ELF Identification

ELF provides an object file framework to support multiple processors, multiple data encoding, and multiple classes of machines. To support this object file family, the initial bytes of the file specify how to interpret the file. These bytes are independent of the processor on which the inquiry is made and independent of the file's remaining contents.

The initial bytes of an ELF header and an object file correspond to the `e_ident` member.

TABLE 7-7 ELF Identification Index

Name	Value	Purpose
EI_MAG0	0	File identification
EI_MAG1	1	File identification
EI_MAG2	2	File identification
EI_MAG3	3	File identification
EI_CLASS	4	File class
EI_DATA	5	Data encoding
EI_VERSION	6	File version
EI_OSABI	7	Operating system/ABI identification
EI_ABIVERSION	8	ABI version
EI_PAD	9	Start of padding bytes
EI_NIDENT	16	Size of <code>e_ident</code> []

These indexes access bytes that hold the values described below.

EI_MAG0 - EI_MAG3

A 4-byte *magic number*, identifying the file as an ELF object file, as listed in the following table.

TABLE 7-8 ELF Magic Number

Name	Value	Position
ELFMAG0	0x7f	<code>e_ident[EI_MAG0]</code>
ELFMAG1	'E'	<code>e_ident[EI_MAG1]</code>
ELFMAG2	'L'	<code>e_ident[EI_MAG2]</code>

TABLE 7-8 ELF Magic Number (Continued)

Name	Value	Position
ELFMAG3	'F'	e_ident[EI_MAG3]

EI_CLASS

Byte e_ident[EI_CLASS] identifies the file's class, or capacity, as listed in the following table.

TABLE 7-9 ELF File Class

Name	Value	Meaning
ELFCLASSNONE	0	Invalid class
ELFCLASS32	1	32-bit objects
ELFCLASS64	2	64-bit objects

The file format is designed to be portable among machines of various sizes, without imposing the sizes of the largest machine on the smallest. The class of the file defines the basic types used by the data structures of the object file container itself. The data contained in object file sections may follow a different programming model.

Class ELFCLASS32 supports machines with files and virtual address spaces up to 4 gigabytes. It uses the basic types defined in Table 7-1.

Class ELFCLASS64 is reserved for 64-bit architectures such as SPARC. It uses the basic types defined in Table 7-2.

EI_DATA

Byte e_ident[EI_DATA] specifies the data encoding of the processor-specific data in the object file, as listed in the following table.

TABLE 7-10 ELF Data Encoding

Name	Value	Meaning
ELFDATANONE	0	Invalid data encoding
ELFDATA2LSB	1	See Figure 7-2.
ELFDATA2MSB	2	See Figure 7-3.

More information on these encodings appears in the section "Data Encoding" on page 186. Other values are reserved and will be assigned to new encodings as necessary.

EI_VERSION

Byte `e_ident[EI_VERSION]` specifies the ELF header version number. Currently, this value must be `EV_CURRENT`.

EI_OSABI

Byte `e_ident[EI_OSABI]` identifies the operating system and ABI to which the object is targeted. Some fields in other ELF structures have flags and values that have operating system or ABI specific meanings. The interpretation of those fields is determined by the value of this byte.

EI_ABIVERSION

Byte `e_ident[EI_ABIVERSION]` identifies the version of the ABI to which the object is targeted. This field is used to distinguish among incompatible versions of an ABI. The interpretation of this version number is dependent on the ABI identified by the `EI_OSABI` field. If no values are specified for the `EI_OSABI` field for the processor, or no version values are specified for the ABI determined by a particular value of the `EI_OSABI` byte, the value 0 is used to indicate unspecified.

EI_PAD

This value marks the beginning of the unused bytes in `e_ident`. These bytes are reserved and set to zero. Programs that read object files should ignore them.

Data Encoding

A file's data encoding specifies how to interpret the basic objects in a file. Class `ELFCLASS32` files use objects that occupy 1, 2, and 4 bytes. Class `ELFCLASS64` files use objects that occupy 1, 2, 4, and 8 bytes. Under the defined encodings, objects are represented as shown below. Byte numbers appear in the upper left corners.

Encoding `ELFDATA2LSB` specifies 2's complement values, with the least significant byte occupying the lowest address.




FIGURE 7-2 Data Encoding `ELFDATA2LSB`

Encoding `ELFDATA2MSB` specifies 2's complement values, with the most significant byte occupying the lowest address.




FIGURE 7-3 Data Encoding ELFDATA2MSB

Sections

An object file's section header table helps you locate all of the sections of the file. The section header table is an array of `Elf32_Shdr` or `Elf64_Shdr` structures, as described below. A section header table index is a subscript into this array. The ELF header's `e_shoff` member gives the byte offset from the beginning of the file to the section header table; `e_shnum` tells how many entries the section header table contains; `e_shentsize` gives the size in bytes of each entry.

If the number of sections is greater than or equal to `SHN_LORESERVE` (`0xff00`), `e_shnum` has the value `SHN_UNDEF` (0) and the actual number of section header table entries is contained in the `sh_size` field of the section header at index 0. Otherwise, the `sh_size` member of the initial entry contains 0.

Some section header table indexes are reserved in contexts where index size is restricted. For example, the `st_shndx` member of a symbol table entry and the `e_shnum` and `e_shstrndx` members of the ELF header. In such contexts, the reserved values do not represent actual sections in the object file. Also in such contexts, an escape value indicates that the actual section index is to be found elsewhere, in a larger field.

TABLE 7-11 ELF Special Section Indexes

Name	Value
<code>SHN_UNDEF</code>	0
<code>SHN_LORESERVE</code>	<code>0xff00</code>
<code>SHN_LOPROC</code>	<code>0xff00</code>
<code>SHN_BEFORE</code>	<code>0xff00</code>
<code>SHN_AFTER</code>	<code>0xff01</code>
<code>SHN_HIPROC</code>	<code>0xff1f</code>

TABLE 7-11 ELF Special Section Indexes (Continued)

Name	Value
SHN_LOOS	0xff20
SHN_HIOS	0xff3f
SHN_ABS	0xffff1
SHN_COMMON	0xffff2
SHN_XINDEX	0xffff
SHN_HIRESERVE	0xffff

Note – Although index 0 is reserved as the undefined value, the section header table contains an entry for index 0. That is, if the `e_shnum` member of the ELF header says a file has 6 entries in the section header table, they have the indexes 0 through 5. The contents of the initial entry are specified later in this section.

SHN_UNDEF

An undefined, missing, irrelevant, or otherwise meaningless section reference. For example, a symbol *defined* relative to section number `SHN_UNDEF` is an undefined symbol.

SHN_LORESERVE

The lower boundary of the range of reserved indexes.

SHN_LOPROC - SHN_HIPROC

Values in this inclusive range are reserved for processor-specific semantics.

SHN_LOOS - SHN_HIOS

Values in this inclusive range are reserved for operating system-specific semantics.

SHN_BEFORE, SHN_AFTER

Provide for initial and final section ordering in conjunction with the `SHF_LINK_ORDER` and `SHF_ORDERED` section flags, listed in Table 7-14.

SHN_ABS

Absolute values for the corresponding reference. For example, symbols defined relative to section number `SHN_ABS` have absolute values and are not affected by relocation.

SHN_COMMON

Symbols defined relative to this section are common symbols, such as FORTRAN COMMON or unallocated C external variables. These symbols are sometimes referred to as tentative.

SHN_XINDEX

An escape value indicating that the actual section header index is too large to fit in the containing field. The header section index is found in another location specific to the structure where it appears.

SHN_HIRESERVE

The upper boundary of the range of reserved indexes. The system reserves indexes between SHN_LORESERVE and SHN_HIRESERVE, inclusive. The values do not reference the section header table. The section header table does not contain entries for the reserved indexes.

Sections contain all information in an object file except the ELF header, the program header table, and the section header table. Moreover, the sections in object files satisfy several conditions:

- Every section in an object file has exactly one section header describing it. Section headers can exist that do not have a section.
- Each section occupies one contiguous, possibly empty, sequence of bytes within a file.
- Sections in a file cannot overlap. No byte in a file resides in more than one section.
- An object file can have inactive space. The various headers and the sections might not cover every byte in an object file. The contents of the inactive data are unspecified.

A section header has the following structure, defined in `sys/elf.h`:

```
typedef struct {
    Elf32_Word    sh_name;
    Elf32_Word    sh_type;
    Elf32_Word    sh_flags;
    Elf32_Addr    sh_addr;
    Elf32_Off     sh_offset;
    Elf32_Word    sh_size;
    Elf32_Word    sh_link;
    Elf32_Word    sh_info;
    Elf32_Word    sh_addralign;
    Elf32_Word    sh_entsize;
} Elf32_Shdr;

typedef struct {
    Elf64_Word    sh_name;
    Elf64_Word    sh_type;
    Elf64_Xword   sh_flags;
    Elf64_Addr    sh_addr;
    Elf64_Off     sh_offset;
    Elf64_Xword   sh_size;
    Elf64_Word    sh_link;
    Elf64_Word    sh_info;
    Elf64_Xword   sh_addralign;
    Elf64_Xword   sh_entsize;
} Elf64_Shdr;
```

The elements of this structure are:

`sh_name`

The name of the section. Its value is an index into the section header string table section giving the location of a null-terminated string. Section names and their descriptions are listed in Table 7-16.

`sh_type`

Categorizes the section's contents and semantics. Section types and their descriptions are listed in Table 7-12.

`sh_flags`

Sections support 1-bit flags that describe miscellaneous attributes. Flag definitions are listed in Table 7-14.

`sh_addr`

If the section is to appear in the memory image of a process, this member gives the address at which the section's first byte should reside. Otherwise, the member contains 0.

`sh_offset`

The byte offset from the beginning of the file to the first byte in the section. For a `SHT_NOBITS` section, this member indicates the conceptual offset in the file, as the section occupies no space in the file.

`sh_size`

The section's size in bytes. Unless the section type is `SHT_NOBITS`, the section occupies `sh_size` bytes in the file. A section of type `SHT_NOBITS` can have a nonzero size, but it occupies no space in the file.

`sh_link`

A section header table index link, whose interpretation depends on the section type. Table 7-15 describes the values.

`sh_info`

Extra information, whose interpretation depends on the section type. Table 7-15 describes the values.

`sh_addralign`

Some sections have address alignment constraints. For example, if a section holds a double-word, the system must ensure double-word alignment for the entire section. That is, the value of `sh_addr` must be congruent to 0, modulo the value of `sh_addralign`. Currently, only 0 and positive integral powers of two are allowed. Values 0 and 1 mean the section has no alignment constraints.

`sh_entsize`

Some sections hold a table of fixed-size entries, such as a symbol table. For such a section, this member gives the size in bytes of each entry. The member contains 0 if the section does not hold a table of fixed-size entries.

A section header's `sh_type` member specifies the section's semantics, as shown in the following table.

TABLE 7-12 ELF Section Types, *sh_type*

Name	Value
SHT_NULL	0
SHT_PROGBITS	1
SHT_SYMTAB	2
SHT_STRTAB	3
SHT_RELA	4
SHT_HASH	5
SHT_DYNAMIC	6
SHT_NOTE	7
SHT_NOBITS	8
SHT_REL	9
SHT_SHLIB	10
SHT_DYNSYM	11
SHT_INIT_ARRAY	14
SHT_FINI_ARRAY	15
SHT_PREINIT_ARRAY	16
SHT_GROUP	17
SHT_SYMTAB_SHNDX	18
SHT_LOOS	0x60000000
SHT_LOSUNW	0x6fffffff7
SHT_SUNW_ANNOTATE	0x6fffffff7
SHT_SUNW_DEBUGSTR	0x6fffffff8
SHT_SUNW_DEBUG	0x6fffffff9
SHT_SUNW_move	0x6fffffffa
SHT_SUNW_COMDAT	0x6fffffffb
SHT_SUNW_syminfo	0x6fffffffc
SHT_SUNW_verdef	0x6fffffffd
SHT_SUNW_verneed	0x6fffffffe
SHT_SUNW_versym	0x6fffffff

TABLE 7-12 ELF Section Types, *sh_type* (Continued)

Name	Value
SHT_HISUNW	0x6fffffff
SHT_HIOS	0x6fffffff
SHT_LOPROC	0x70000000
SHT_SPARC_GOTDATA	0x70000000
SHT_HIPROC	0x7fffffff
SHT_LOUSER	0x80000000
SHT_HIUSER	0xffffffff

SHT_NULL

Identifies the section header as inactive. This section header does not have an associated section. Other members of the section header have undefined values.

SHT_PROGBITS

Identifies information defined by the program, whose format and meaning are determined solely by the program.

SHT_SYMTAB, SHT_DYNSYM

Identifies a symbol table. Typically, a **SHT_SYMTAB** section provides symbols for link-editing. As a complete symbol table, it can contain many symbols unnecessary for dynamic linking. Consequently, an object file can also contain a **SHT_DYNSYM** section, which holds a minimal set of dynamic linking symbols, to save space. See “Symbol Table Section” on page 222 for details.

SHT_STRTAB, SHT_DYNSTR

Identifies a string table. An object file can have multiple string table sections. See “String Table Section” on page 221 for details.

SHT_RELA

Identifies relocation entries with explicit addends, such as type **Elf32_Rela** for the 32-bit class of object files. An object file can have multiple relocation sections. See “Relocation Sections” on page 210 for details.

SHT_HASH

Identifies a symbol hash table. A dynamically linked object file must contain a symbol hash table. Currently, an object file can have only one hash table, but this restriction might be relaxed in the future. See “Hash Table Section” on page 205 for details.

SHT_DYNAMIC

Identifies information for dynamic linking. Currently, an object file can have only one dynamic section. See “Dynamic Section” on page 248 for details.

SHT_NOTE

Identifies information that marks the file in some way. See “Note Section” on page 208 for details.

SHT_NOBITS

Identifies a section that occupies no space in the file but otherwise resembles SHT_PROGBITS. Although this section contains no bytes, the `sh_offset` member contains the conceptual file offset.

SHT_REL

Identifies relocation entries without explicit addends, such as type `Elf32_Rel` for the 32-bit class of object files. An object file can have multiple relocation sections. See “Relocation Sections” on page 210 for details.

SHT_SHLIB

Identifies a reserved section which has unspecified semantics. Programs that contain a section of this type do not conform to the ABI.

SHT_INIT_ARRAY

Identifies a section containing an array of pointers to initialization functions. Each pointer in the array is taken as a parameterless procedure with a void return. See “Initialization and Termination Sections” on page 36 for details.

SHT_FINI_ARRAY

Identifies a section containing an array of pointers to termination functions. Each pointer in the array is taken as a parameterless procedure with a void return. See “Initialization and Termination Sections” on page 36 for details.

SHT_PREINIT_ARRAY

Identifies a section containing an array of pointers to functions that are invoked before all other initialization functions. Each pointer in the array is taken as a parameterless procedure with a void return. See “Initialization and Termination Sections” on page 36 for details.

SHT_GROUP

Identifies a section group. A section group identifies a set of related sections that must be treated as a unit by the link-editor. Sections of type SHT_GROUP may appear only in relocatable objects. See “Group Section” on page 204 for details.

SHT_SYMTAB_SHNDX

Identifies a section containing extended section indexes, that is associated with a symbol table. If any section header indexes referenced by a symbol table, contain the escape value `SHN_XINDEX`, an associated SHT_SYMTAB_SHNDX is required.

The SHT_SYMTAB_SHNDX section is an array of `Elf32_Word` values. There is one value for every entry in the associated symbol table entry. The values represent the section header indexes against which the symbol table entries are defined. Only if corresponding symbol table entry’s `st_shndx` field contains the escape value `SHN_XINDEX` will the matching `Elf32_Word` hold the actual section header index. Otherwise, the entry must be `SHN_UNDEF` (0).

SHT_LOOS – SHT_HIOS
Values in this inclusive range are reserved for operating system-specific semantics.

SHT_LOSUNW – SHT_HISUNW
Values in this inclusive range are reserved for Solaris semantics.

SHT_SUNW_ANNOTATE
The processing of an annotate section follows all of the default rules for processing a section. The only exception occurs if the annotate section is in non-allocatable memory. If the section header flag SHF_ALLOC is not set, the link-editor silently ignores any relocations against this section that cannot be satisfied during the link-edit.

SHT_SUNW_DEBUGSTR, SHT_SUNW_DEBUG
Identifies debugging information. Sections of this type are stripped from the object using the link-editor's `-s` option, or after the link-edit using `strip(1)`.

SHT_SUNW_move
Identifies data to handle partially initialized symbols. See "Move Section" on page 206 for details.

SHT_SUNW_COMDAT
Identifies a section that allows multiple copies of the same data to be reduced to a single copy. See "COMDAT Section" on page 203 for details.

SHT_SUNW_syminfo
Identifies additional symbol information. See "Syminfo Table Section" on page 229 for details.

SHT_SUNW_verdef
Identifies fine-grained versions defined by this file. See "Version Definition Section" on page 231 for details.

SHT_SUNW_verneed
Identifies fine-grained dependencies required by this file. See "Version Dependency Section" on page 234 for details.

SHT_SUNW_versym
Identifies a table describing the relationship of symbols to the version definitions offered by the file. See "Version Symbol Section" on page 233 for details.

SHT_LOPROC - SHT_HIPROC
Values in this inclusive range are reserved for processor-specific semantics.

SHT_SPARC_GOTDATA
Identifies SPARC specific data, referenced using GOT-relative addressing. That is, offsets relative to the address assigned to the symbol `_GLOBAL_OFFSET_TABLE_`. For 64-bit SPARC, data in this section must be bound at link-edit time to locations within $\{+-\} 2^{32}$ bytes of the GOT address.

SHT_LOUSER
Specifies the lower boundary of the range of indexes reserved for application programs.

SHT_HIUSER

Specifies the upper boundary of the range of indexes reserved for application programs. Section types between SHT_LOUSER and SHT_HIUSER can be used by the application without conflicting with current or future system-defined section types.

Other section-type values are reserved. As mentioned before, the section header for index 0 (SHN_UNDEF) exists, even though the index marks undefined section references. The following table shows the values.

TABLE 7-13 ELF Section Header Table Entry: Index 0

Name	Value	Note
sh_name	0	No name
sh_type	SHT_NULL	Inactive
sh_flags	0	No flags
sh_addr	0	No address
sh_offset	0	No file offset
sh_size	0	No size
sh_link	SHN_UNDEF	No link information
sh_info	0	No auxiliary information
sh_addralign	0	No alignment
sh_entsize	0	No entries

A section header's sh_flags member holds 1-bit flags that describe the section's attributes:

TABLE 7-14 ELF Section Attribute Flags

Name	Value
SHF_WRITE	0x1
SHF_ALLOC	0x2
SHF_EXECINSTR	0x4
SHF_MERGE	0x10
SHF_STRINGS	0x20
SHF_INFO_LINK	0x40
SHF_LINK_ORDER	0x80

TABLE 7-14 ELF Section Attribute Flags (Continued)

Name	Value
SHF_OS_NONCONFORMING	0x100
SHF_GROUP	0x200
SHF_TLS	0x400
SHF_MASKOS	0x0ff00000
SHF_ORDERED	0x40000000
SHF_EXCLUDE	0x80000000
SHF_MASKPROC	0xf0000000

If a flag bit is set in `sh_flags`, the attribute is *on* for the section. Otherwise, the attribute is *off* or does not apply. Undefined attributes are reserved and set to zero.

SHF_WRITE

Identifies a section that should be writable during process execution.

SHF_ALLOC

Identifies a section that occupies memory during process execution. Some control sections do not reside in the memory image of an object file. This attribute is off for those sections.

SHF_EXECINSTR

Identifies a section that contains executable machine instructions.

SHF_MERGE

Identifies a section containing data that may be merged to eliminate duplication. Unless the `SHF_STRINGS` flag is also set, the data elements in the section are of a uniform size. The size of each element is specified in the section header's `sh_entsize` field. If the `SHF_STRINGS` flag is also set, the data elements consist of null-terminated character strings. The size of each character is specified in the section header's `sh_entsize` field.

SHF_STRINGS

Identifies a section that consists of null-terminated character strings. The size of each character is specified in the section header's `sh_entsize` field.

SHF_INFO_LINK

This section headers `sh_info` field holds a section header table index.

SHF_LINK_ORDER

This section adds special ordering requirements to the link-editor. The requirements apply if the `sh_link` field of this section's header references another section, the linked-to section. If this section is combined with other sections in the output file, the section appears in the same relative order with respect to those sections. Similarly the linked-to section appears with respect to sections the linked-to section is combined with.

The special `sh_link` values `SHN_BEFORE` and `SHN_AFTER` (see Table 7–11) imply that the sorted section is to precede or follow, respectively, all other sections in the set being ordered. Input file link-line order is preserved if multiple sections in an ordered set have one of these special values.

A typical use of this flag is to build a table that references text or data sections in address order.

In the absence of the `sh_link` ordering information, sections from a single input file combined within one section of the output file will be contiguous and have the same relative ordering as they did in the input file. The contributions from multiple input files appear in link-line order.

`SHF_OS_NONCONFORMING`

This section requires special OS-specific processing beyond the standard linking rules to avoid incorrect behavior. If this section has either an `sh_type` value or contains `sh_flags` bits in the OS-specific ranges for those fields, and the link-editor does not recognize these values, then the link-editor will reject the object file containing this section with an error.

`SHF_GROUP`

This section is a member, perhaps the only one, of a section group. The section must be referenced by a section of type `SHT_GROUP`. The `SHF_GROUP` flag can be set only for sections contained in relocatable objects. See “Group Section” on page 204 for details.

`SHF_TLS`

This section holds thread-local storage, meaning that each separate execution flow has its own distinct instance of this data. See Chapter 8 for details.

`SHF_MASKOS`

All bits included in this mask are reserved for operating system-specific semantics.

`SHF_ORDERED`

This section requires ordering in relation to other sections of the same type. Ordered sections are combined within the section pointed to by the `sh_link` entry. The `sh_link` entry of an ordered section can point to itself.

If the `sh_info` entry of the ordered section is a valid section within the same input file, the ordered section will be sorted based on the relative ordering within the output file of the section pointed to by the `sh_info` entry.

The special `sh_info` values `SHN_BEFORE` and `SHN_AFTER` (see Table 7–11) imply that the sorted section is to precede or follow, respectively, all other sections in the set being ordered. Input file link-line order is preserved if multiple sections in an ordered set have one of these special values.

In the absence of the `sh_info` ordering information, sections from a single input file combined within one section of the output file will be contiguous and have the same relative ordering as they did in the input file. The contributions from multiple input files appear in link-line order.

SHF_EXCLUDE

This section is excluded from input to the link-edit of an executable or shared object. This flag is ignored if the **SHF_ALLOC** flag is also set, or if relocations exist against the section.

SHF_MASKPROC

All bits included in this mask are reserved for processor-specific semantics.

Two members in the section header, **sh_link** and **sh_info**, hold special information, depending on section type.

TABLE 7-15 ELF **sh_link** and **sh_info** Interpretation

sh_type	sh_link	sh_info
SHT_DYNAMIC	The section header index of the associated string table.	0
SHT_HASH	The section header index of the associated symbol table.	0
SHT_REL SHT_RELA	The section header index of the associated symbol table.	The section header index of the section to which the relocation applies. See also Table 7-16 and "Relocation Sections" on page 210.
SHT_SYMTAB SHT_DYNSYM	The section header index of the associated string table.	One greater than the symbol table index of the last local symbol (binding STB_LOCAL).
SHT_GROUP	The section header index of the associated symbol table.	The symbol table index of an entry in the associated symbol table. The name of the specified symbol table entry provides a signature for the section group.
SHT_SYMTAB_SHNDX	The section header index of the associated symbol table.	0
SHT_SUNW_move	The section header index of the associated symbol table.	0
SHT_SUNW_COMDAT	0	0
SHT_SUNW_syminfo	The section header index of the associated symbol table.	The section header index of the associated <code>.dynamic</code> section.
SHT_SUNW_verdef	The section header index of the associated string table.	The number of version definitions within the section.
SHT_SUNW_verneed	The section header index of the associated string table.	The number of version dependencies within the section.

TABLE 7-15 ELF `sh_link` and `sh_info` Interpretation (Continued)

<code>sh_type</code>	<code>sh_link</code>	<code>sh_info</code>
<code>SHT_SUNW_versym</code>	The section header index of the associated symbol table.	0

Special Sections

Various sections hold program and control information. Sections in the following table are used by the system and have the indicated types and attributes.

TABLE 7-16 ELF Special Sections

Name	Type	Attribute
<code>.bss</code>	<code>SHT_NOBITS</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.comment</code>	<code>SHT_PROGBITS</code>	None
<code>.data</code>	<code>SHT_PROGBITS</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.data1</code>	<code>SHT_PROGBITS</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.dynamic</code>	<code>SHT_DYNAMIC</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.dynstr</code>	<code>SHT_STRTAB</code>	<code>SHF_ALLOC</code>
<code>.dynsym</code>	<code>SHT_DYNSYM</code>	<code>SHF_ALLOC</code>
<code>.fini</code>	<code>SHT_PROGBITS</code>	<code>SHF_ALLOC + SHF_EXECINSTR</code>
<code>.fini_array</code>	<code>SHT_FINI_ARRAY</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.got</code>	<code>SHT_PROGBITS</code>	See “Global Offset Table (Processor-Specific)” on page 260
<code>.hash</code>	<code>SHT_HASH</code>	<code>SHF_ALLOC</code>
<code>.init</code>	<code>SHT_PROGBITS</code>	<code>SHF_ALLOC + SHF_EXECINSTR</code>
<code>.init_array</code>	<code>SHT_INIT_ARRAY</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.interp</code>	<code>SHT_PROGBITS</code>	See “Program Interpreter” on page 247
<code>.note</code>	<code>SHT_NOTE</code>	None
<code>.plt</code>	<code>SHT_PROGBITS</code>	See “Procedure Linkage Table (Processor-Specific)” on page 261
<code>.preinit_array</code>	<code>SHT_PREINIT_ARRAY</code>	<code>SHF_ALLOC + SHF_WRITE</code>
<code>.rela</code>	<code>SHT_RELA</code>	None

TABLE 7-16 ELF Special Sections (Continued)

Name	Type	Attribute
.relname	SHT_REL	See "Relocation Sections" on page 210
.relnamex	SHT_RELA	See "Relocation Sections" on page 210
.rodata	SHT_PROGBITS	SHF_ALLOC
.rodata1	SHT_PROGBITS	SHF_ALLOC
.shstrtab	SHT_STRTAB	None
.strtab	SHT_STRTAB	See description below
.symtab	SHT_SYMTAB	See "Symbol Table Section" on page 222
.symtab_shndx	SHT_SYMTAB_SHNDX	See "Symbol Table Section" on page 222
.tbss	SHT_NOBITS	SHF_ALLOC + SHF_WRITE + SHF_TLS
.tdata	SHT_PROGBITS	SHF_ALLOC + SHF_WRITE + SHF_TLS
.tdata1	SHT_PROGBITS	SHF_ALLOC + SHF_WRITE + SHF_TLS
.text	SHT_PROGBITS	SHF_ALLOC + SHF_EXECINSTR
.SUNW_bss	SHT_NOBITS	SHF_ALLOC + SHF_WRITE
.SUNW_heap	SHT_PROGBITS	SHF_ALLOC + SHF_WRITE
.SUNW_move	SHT_SUNW_move	SHF_ALLOC
.SUNW_reloc	SHT_REL SHT_RELA	SHF_ALLOC
.SUNW_syminfo	SHT_SUNW_syminfo	SHF_ALLOC
.SUNW_version	SHT_SUNW_verdef SHT_SUNW_verneed SHT_SUNW_versym	SHF_ALLOC

.bss

Uninitialized data that contribute to the program's memory image. By definition, the system initializes the data with zeros when the program begins to run. The section occupies no file space, as indicated by the section type SHT_NOBITS.

- `.comment`
Comment information, typically contributed by the components of the compilation system. This section can be manipulated by `mcs(1)`.
- `.data, .data1`
Initialized data that contribute to the program's memory image.
- `.dynamic`
Dynamic linking information. See "Dynamic Section" on page 248 for details.
- `.dynstr`
Strings needed for dynamic linking, most commonly the strings that represent the names associated with symbol table entries.
- `.dynsym`
Dynamic linking symbol table. See "Symbol Table Section" on page 222 for details.
- `.fini`
Executable instructions that contribute to a single termination function for the executable or shared object containing the section. See "Initialization and Termination Routines" on page 76 for details.
- `.fini_array`
An array of function pointers that contribute to a single termination array for the executable or shared object containing the section. See "Initialization and Termination Routines" on page 76 for details.
- `.got`
The global offset table. See "Global Offset Table (Processor-Specific)" on page 260 for details.
- `.hash`
Symbol hash table. See "Hash Table Section" on page 205 for details.
- `.init`
Executable instructions that contribute to a single initialization function for the executable or shared object containing the section. See "Initialization and Termination Routines" on page 76 for details.
- `.init_array`
An array of function pointers that contributes to a single initialization array for the executable or shared object containing the section. See "Initialization and Termination Routines" on page 76 for details.
- `.interp`
The path name of a program interpreter. See "Program Interpreter" on page 247 for details.
- `.note`
Information in the format described in "Note Section" on page 208.
- `.plt`
The procedure linkage table. See "Procedure Linkage Table (Processor-Specific)" on page 261 for details.

- `.preinit_array`
An array of function pointers that contribute to a single pre-initialization array for the executable or shared object containing the section. See “Initialization and Termination Routines” on page 76 for details.
- `.rela`
Relocations that do not apply to a particular section. One use of this section is for register relocations. See “Register Symbols” on page 229 for details.
- `.relname, .relaname`
Relocation information, as “Relocation Sections” on page 210 describes. If the file has a loadable segment that includes relocation, the sections’ attributes include the SHF_ALLOC bit. Otherwise, that bit is off. Conventionally, *name* is supplied by the section to which the relocations apply. Thus, a relocation section for `.text` normally will have the name `.rel.text` or `.rela.text`.
- `.rodata, .rodata1`
Read-only data that typically contribute to a non-writable segment in the process image. See “Program Header” on page 236 for details.
- `.shstrtab`
Section names.
- `.strtab`
Strings, most commonly the strings that represent the names associated with symbol table entries. If the file has a loadable segment that includes the symbol string table, the section’s attributes include the SHF_ALLOC bit. Otherwise, that bit is turned off.
- `.symtab`
Symbol table, as “Symbol Table Section” on page 222 describes. If the file has a loadable segment that includes the symbol table, the section’s attributes include the SHF_ALLOC bit. Otherwise, that bit is turned off.
- `.symtab_shndx`
This section holds the special symbol table section index array, as described by `.symtab`. The section’s attributes will include the SHF_ALLOC bit if the associated symbol table section does. Otherwise, that bit is turned off.
- `.tbss`
This section holds uninitialized thread-local data that contribute to the program’s memory image. By definition, the system initializes the data with zeros when the data is instantiated for each new execution flow. The section occupies no file space, as indicated by the section type, SHT_NOBITS. See Chapter 8 for details.
- `.tdata, .tdata1`
These sections hold initialized thread-local data that contribute to the program’s memory image. A copy of its contents is instantiated by the system for each new execution flow. See Chapter 8 for details.
- `.text`
The *text* or executable instructions of a program.

- .SUNW_bss
Partially initialized data for shared objects that contribute to the program's memory image. The data is initialized at runtime. The section occupies no file space, as indicated by the section type SHT_NOBITS.
- .SUNW_heap
The *heap* of a dynamic executable created from `dldump(3DL)`.
- .SUNW_move
Additional information for partially initialized data. See "Move Section" on page 206 for details.
- .SUNW_reloc
Relocation information, as "Relocation Sections" on page 210 describes. This section is a concatenation of relocation sections that provides better locality of reference of the individual relocation records. Only the offset of the relocation record itself is meaningful, thus the section `sh_info` value is zero.
- .SUNW_syminfo
Additional symbol table information. See "Syminfo Table Section" on page 229 for details.
- .SUNW_version
Versioning information. See "Versioning Sections" on page 231 for details.

Section names with a dot (.) prefix are reserved for the system, although applications can use these sections if their existing meanings are satisfactory. Applications can use names without the prefix to avoid conflicts with system sections. The object file format enables you to define sections not in the list above. An object file can have more than one section with the same name.

Section names reserved for a processor architecture are formed by placing an abbreviation of the architecture name ahead of the section name. The name should be taken from the architecture names used for `e_machine`. For example, `.Foo.psect` is the `psect` section defined by the `FOO` architecture.

Existing extensions use their historical names

COMDAT Section

COMDAT sections are uniquely identified by their section name (`sh_name`). If the link-editor encounters multiple sections of type `SHT_SUNW_COMDAT`, with the same section name, the first section is retained and the rest discarded. Any relocations applied to a discarded `SHT_SUNW_COMDAT` section are ignored. Any symbols defined in a discarded section are removed.

Additionally, the link-editor supports the section naming convention used for section reordering when the compiler is invoked with the `-xF` option. If a function is placed in a section with the name `.sectname%funcname`, the final `SHT_SUNW_COMDAT` sections that are retained are coalesced into a section identified by `.sectname`. Using this method, the `SHT_SUNW_COMDAT` sections are placed into the `.text`, `.data`, or any other section as their final destination.

Group Section

Some sections occur in interrelated groups. For example, an out-of-line definition of an inline function might require additional information besides the section containing executable instructions. This additional information can be a read-only data section containing literals referenced, one or more debugging information sections, or other informational sections. Furthermore, there may be internal references among these sections. These references make no sense if one of the sections were removed, or one of the sections were replaced, by a duplicate from another object. Therefore, these groups are included, or these groups are omitted, from the linked object as a unit.

A section of type `SHT_GROUP` defines such a grouping of sections. The name of a symbol from one of the containing object's symbol tables provides a signature for the section group. The section header of the `SHT_GROUP` section specifies the identifying symbol entry. The `sh_link` member contains the section header index of the symbol table section that contains the entry. The `sh_info` member contains the symbol table index of the identifying entry. The `sh_flags` member of the section header contains 0. The name of the section (`sh_name`) is not specified.

The section data of a `SHT_GROUP` section is an array of `Elf32_Word` entries. The first entry is a flag word. The remaining entries are a sequence of section header indices.

The following flag is currently defined:

TABLE 7-17 ELF Group Section Flag

Name	Value
GRP_COMDAT	0x1

GRP_COMDAT

`GRP_COMDAT` is a `COMDAT` group. It may duplicate another `COMDAT` group in another object file, where duplication is defined as having the same group signature. In such cases, only one of the duplicate groups is retained by the link-editor. The members of the remaining groups are discarded.

The section header indices in the `SHT_GROUP` section, identify the sections that make up the group. These sections must have the `SHF_GROUP` flag set in their `sh_flags` section header member. If the link-editor decides to remove the section group, the link-editor removes all members of the group.

To facilitate removing a group without leaving dangling references and with only minimal processing of the symbol table, the following rules are followed:

- References to the sections comprising a group from sections outside of the group must be made through symbol table entries with `STB_GLOBAL` or `STB_WEAK` binding and section index `SHN_UNDEF`. If there is a definition of the same symbol in the object containing the references, it must have a separate symbol table entry from the references. Sections outside of the group may not reference symbols with `STB_LOCAL` binding for addresses contained in the group's sections, including symbols with type `STT_SECTION`.
- Non-symbol references to the sections comprising a group are not allowed from outside the group. For example, you cannot use a group member's section header index in an `sh_link` or `sh_info` member.
- A symbol table entry defined relative to one of the group's sections may be removed if the group members are discarded. This removal occurs if the symbol table entry is contained in a symbol table section that is not part of the group.

Hash Table Section

A hash table consists of `Elf32_Word` or `Elf64_Word` objects that provide for symbol table access. The `SHT_HASH` section provides this hash table. The symbol table to which the hashing is associated is specified in the `sh_link` entry of the hash table's section header. Labels appear below to help explain the hash table organization, but they are not part of the specification.




FIGURE 7-4 Symbol Hash Table

The bucket array contains `nbucket` entries, and the chain array contains `nchain` entries. Indexes start at 0. Both `bucket` and `chain` hold symbol table indexes. Chain table entries parallel the symbol table. The number of symbol table entries should equal `nchain`, so symbol table indexes also select chain table entries.

A hashing function that accepts a symbol name, returns a value to compute a bucket index. Consequently, if the hashing function returns the value x for some name, bucket `[x%nbucket]` gives an index y . This is an index into both the symbol table and the chain table. If the symbol table entry is not the name desired, `chain[y]` gives the next symbol table entry with the same hash value.

The chain links can be followed until the selected symbol table entry holds the desired name, or the chain entry contains the value `STN_UNDEF`.

The hash function is as follows:

```
unsigned long
elf_hash(const unsigned char *name)
{
    unsigned long h = 0, g;

    while (*name)
    {
        h = (h << 4) + *name++;
        if (g = h & 0xf0000000)
            h ^= g >> 24;
        h &= ~g;
    }
    return h;
}
```

Move Section

Typically, within ELF files, initialized data variables are maintained within the object file. If a data variable is very large, and contains only a small number of initialized (nonzero) elements, the entire variable is still maintained in the object file.

Objects that contain large partially initialized data variables, such as FORTRAN COMMON blocks, can result in a significant disk space overhead. The `SHT_SUNW_move` section provides a mechanism of compressing these data variables. This compression reduces the disk size of the associated object.

The `SHT_SUNW_move` section contains multiple entries of the type `Elf32_Move` or `Elf64_Move`. These entries allow data variables to be defined as tentative items (`.bss`), thus occupying no space in the object file but contributing to the object's memory image at runtime. The move records establish how the memory image is initialized with data to construct the complete data variable.

`Elf32_Move` and `Elf64_Move` entries are defined as follows:

```
typedef struct {
    Elf32_Lword    m_value;
    Elf32_Word     m_info;
    Elf32_Word     m_poffset;
    Elf32_Half     m_repeat;
}
```

```

        Elf32_Half      m_stride;
    } Elf32_Move;

#define ELF32_M_SYM(info)      ((info)>>8)
#define ELF32_M_SIZE(info)    ((unsigned char)(info))
#define ELF32_M_INFO(sym, size) (((sym)<<8)+(unsigned char)(size))

typedef struct {
    Elf64_Lword      m_value;
    Elf64_Xword      m_info;
    Elf64_Xword      m_poffset;
    Elf64_Half       m_repeat;
    Elf64_Half       m_stride;
} Elf64_Move;

#define ELF64_M_SYM(info)      ((info)>>8)
#define ELF64_M_SIZE(info)    ((unsigned char)(info))
#define ELF64_M_INFO(sym, size) (((sym)<<8)+(unsigned char)(size))

```

The elements of these structures are as follows:

`m_value`

The initialization value, which is the value that is moved into the memory image.

`m_info`

The symbol table index, with respect to which the initialization is applied, together with the size, in bytes, of the offset being initialized. The lower 8 bits of the member define the size, which can be 1, 2, 4 or 8. The upper bytes define the symbol index.

`m_poffset`

The offset relative to the associated symbol to which the initialization is applied.

`m_repeat`

A repetition count.

`m_stride`

The stride count. This value indicates the number of units that should be skipped when performing a repetitive initialization. A unit is the size of an initialization object as defined by `m_info`. An `m_stride` value of 0 indicates that the initialization be performed contiguously for `m_repeat` units.

The following data definition would traditionally consume 0x8000 bytes within an object file:

```

typedef struct {
    int    one;
    char   two;
} Data

Data move[0x1000] = {
    {0, 0},      {1, '1'},      {0, 0},
    {0xf, 'F'}, {0xf, 'F'},    {0, 0},
    {0xe, 'E'}, {0, 0},      {0xe, 'E'}
};

```

Using an `SHT_SUNW_move` section the data item can be moved to the `.bss` section and initialized with the associated move entries:

```
$ elfdump -s data | fgrep move
[17] 0x00020868 0x00008000 OBJT GLOB 0 .bss      move
$ elfdump -m data
```

```
Move Section: .SUNW_move
      offset  ndx    size  repeat  stride  value  with respect to
      0x8     0x17  4     1       0       0x1    move
      0xc     0x17  1     1       0       0x31   move
      0x18    0x17  4     2       2       0xf    move
      0x1c    0x17  1     2       8       0x46   move
      0x28    0x17  4     2       4       0xe    move
      0x2c    0x17  1     2       16      0x45   move
```

Move sections supplied from relocatable objects are concatenated and output in the object being created by the link-editor. However, the following conditions cause the link-editor to process the move entries and expand their contents into a traditional data item:

- The output file is a static executable.
- The size of the move entries is greater than the size of the symbol into which the move data would be expanded.
- The `-z nopartial` option is in effect.

Note Section

Sometimes a vendor or system engineer needs to mark an object file with special information that other programs will check for conformance, compatibility, and so forth. Sections of type `SHT_NOTE` and program header elements of type `PT_NOTE` can be used for this purpose.

The note information in sections and program header elements holds any number of entries, as shown in the following figure. For 64- and 32-bit objects, each entry is an array of 4-byte words in the format of the target processor. Labels are shown in Figure 7-6 to help explain note information organization, but they are not part of the specification.

namesz
descsz
type
name ...
desc ...

FIGURE 7-5 Note Information

The elements of the structure are:

`namesz` and `name`

The first `namesz` bytes in `name` contain a null-terminated character representation of the entry's owner or originator. There is no formal mechanism for avoiding name conflicts. By convention, vendors use their own name, such as "XYZ Computer Company," as the identifier. If no `name` is present, `namesz` contains 0. Padding is present, if necessary, to ensure 4-byte alignment for the descriptor. Such padding is not included in `namesz`.

`descsz` and `desc`

The first `descsz` bytes in `desc` hold the note descriptor. If no descriptor is present, `descsz` contains 0. Padding is present, if necessary, to ensure 4-byte alignment for the next note entry. Such padding is not included in `descsz`.

`type`

Provides the interpretation of the descriptor. Each originator controls its own types. Multiple interpretations of a single `type` value can exist. A program must recognize both the `name` and the `type` to understand a descriptor. Types currently must be nonnegative.

The note segment shown in the following figure holds two entries.

	+0	+1	+2	+3	
namesz	7				
descsz	0				No descriptor
type	1				
name	X	Y	Z		
	C	o	\0	pad	
namesz	7				
descsz	8				
type	3				
name	X	Y	Z		
	C	o	\0	pad	
desc	word0				
	word1				

FIGURE 7-6 Example Note Segment

Note – The system reserves note information with no name (`namesz == 0`) and with a zero-length name (`name[0] == '\0'`) but currently defines no types. All other names must have at least one non-null character.

Relocation Sections

Relocation is the process of connecting symbolic references with symbolic definitions. For example, when a program calls a function, the associated call instruction must transfer control to the proper destination address at execution. Relocatable files must have information that describes how to modify their section contents, thus allowing executable and shared object files to hold the right information for a process's program image. Relocation entries are these data.

Relocation entries can have the following structure, defined in `sys/elf.h`:

```
typedef struct {
    Elf32_Addr    r_offset;
    Elf32_Word    r_info;
} Elf32_Rel;

typedef struct {
    Elf32_Addr    r_offset;
```

```

        Elf32_Word      r_info;
        Elf32_Sword     r_addend;
} Elf32_Rela;

typedef struct {
        Elf64_Addr      r_offset;
        Elf64_Xword     r_info;
} Elf64_Rel;

typedef struct {
        Elf64_Addr      r_offset;
        Elf64_Xword     r_info;
        Elf64_Sxword    r_addend;
} Elf64_Rela;

```

The elements of this structure are:

`r_offset`

This member gives the location at which to apply the relocation action. Different object files have slightly different interpretations for this member.

For a relocatable file, the value indicates a section offset. The relocation section itself describes how to modify another section in the file. Relocation offsets designate a storage unit within the second section.

For an executable or shared object, the value indicates the virtual address of the storage unit affected by the relocation. This information makes the relocation entries more useful for the runtime linker.

Although the interpretation of the member changes for different object files to allow efficient access by the relevant programs, the meanings of the relocation types stay the same.

`r_info`

This member gives both the symbol table index, with respect to which the relocation must be made, and the type of relocation to apply. For example, a call instruction's relocation entry holds the symbol table index of the function being called. If the index is `STN_UNDEF`, the undefined symbol index, the relocation uses 0 as the symbol value.

Relocation types are processor-specific. A relocation entry's relocation type or symbol table index is the result of applying `ELF32_R_TYPE` or `ELF32_R_SYM`, respectively, to the entry's `r_info` member:

```

#define ELF32_R_SYM(info)          ((info)>>8)
#define ELF32_R_TYPE(info)        ((unsigned char)(info))
#define ELF32_R_INFO(sym, type)    (((sym)<<8)+(unsigned char)(type))

#define ELF64_R_SYM(info)          ((info)>>32)
#define ELF64_R_TYPE(info)        ((Elf64_Word)(info))
#define ELF64_R_INFO(sym, type)    (((Elf64_Xword)(sym)<<32)+ \
                                     (Elf64_Xword)(type))

```

For `Elf64_Rel` and `Elf64_Rela` structures, the `r_info` field is further broken down into an 8-bit type identifier and a 24-bit type dependent data field:

```
#define ELF64_R_TYPE_DATA(info)      (((Elf64_Xword) (info) << 32) >> 40)
#define ELF64_R_TYPE_ID(info)       (((Elf64_Xword) (info) << 56) >> 56)
#define ELF64_R_TYPE_INFO(data, type) (((Elf64_Xword) (data) << 8) + \
                                       (Elf64_Xword) (type))
```

`r_addend`

This member specifies a constant addend used to compute the value to be stored into the relocatable field.


`Rela` entries contain an explicit addend. Entries of type `Rel` store an implicit addend in the location to be modified. 32-bit and 64-bit SPARC use only `Elf32_Rela` and `Elf64_Rela` relocation entries respectively. Thus, the `r_addend` member serves as the relocation addend. x86 uses only `Elf32_Rel` relocation entries. The field to be relocated holds the addend. In all cases, the addend and the computed result use the same byte order.

A relocation section can reference two other sections: a symbol table, identified by the `sh_info` section header entry, and a section to modify, identified by the `sh_link` section header entry. "Sections" on page 187 specifies these relationships. An `sh_link` entry is required when a relocation section exists in a relocatable object, but is optional for executables and shared objects. The relocation offset is sufficient to perform the relocation.


Relocation Types (Processor-Specific)

Relocation entries describe how to alter instruction and data fields in the following figures. Bit numbers appear in the lower box corners.

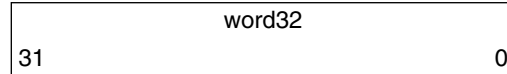
On the SPARC platform, relocation entries apply to bytes (`byte8`), half-words (`half16`), or words (the others).




On 64-bit SPARC, relocations also apply to extended-words (xword64):



On x86, relocation entries apply to words (word32):



word32 specifies a 32-bit field occupying 4 bytes with an arbitrary byte alignment. These values use the same byte order as other word values in the x86 architecture:



In all cases, the `r_offset` value designates the offset or virtual address of the first byte of the affected storage unit. The relocation type specifies which bits to change and how to calculate their values.

Calculations for the following relocation types assume the actions are transforming a relocatable file into either an executable or a shared object file. Conceptually, the link-editor merges one or more relocatable files to form the output. The link-editor first decides how to combine and locate the input files. Then it updates the symbol values. Finally the link-editor performs the relocation. Relocations applied to executable or shared object files are similar and accomplish the same result. Descriptions in the tables in this section use the following notation:

- A The addend used to compute the value of the relocatable field.
- B The base address at which a shared object is loaded into memory during execution. Generally, a shared object file is built with a 0 base virtual address, but the execution address is different. See "Program Header" on page 236.
- G The offset into the global offset table at which the address of the relocation entry's symbol resides during execution. See "Global Offset Table (Processor-Specific)" on page 260.
- GOT The address of the global offset table. See "Global Offset Table (Processor-Specific)" on page 260.
- L The section offset or address of the procedure linkage table entry for a symbol. See "Procedure Linkage Table (Processor-Specific)" on page 261.
- P The section offset or address of the storage unit being relocated, computed using `r_offset`.
- S The value of the symbol whose index resides in the relocation entry.

SPARC: Relocation Types

Field names in the following table tell whether the relocation type checks for overflow. A calculated relocation value can be larger than the intended field, and a relocation type can verify (V) the value fits or truncate (T) the result. As an example, V-simm13 means that the computed value can not have significant, nonzero bits outside the `simm13` field.

TABLE 7-18 SPARC: ELF Relocation Types

Name	Value	Field	Calculation
R_SPARC_NONE	0	None	None
R_SPARC_8	1	V-byte8	S + A
R_SPARC_16	2	V-half16	S + A
R_SPARC_32	3	V-word32	S + A
R_SPARC_DISP8	4	V-byte8	S + A - P
R_SPARC_DISP16	5	V-half16	S + A - P
R_SPARC_DISP32	6	V-disp32	S + A - P
R_SPARC_WDISP30	7	V-disp30	(S + A - P) >> 2
R_SPARC_WDISP22	8	V-disp22	(S + A - P) >> 2
R_SPARC_HI22	9	T-imm22	(S + A) >> 10
R_SPARC_22	10	V-imm22	S + A
R_SPARC_13	11	V-simm13	S + A
R_SPARC_LO10	12	T-simm13	(S + A) & 0x3ff
R_SPARC_GOT10	13	T-simm13	G & 0x3ff
R_SPARC_GOT13	14	V-simm13	G
R_SPARC_GOT22	15	T-simm22	G >> 10
R_SPARC_PC10	16	T-simm13	(S + A - P) & 0x3ff
R_SPARC_PC22	17	V-disp22	(S + A - P) >> 10
R_SPARC_WPLT30	18	V-disp30	(L + A - P) >> 2
R_SPARC_COPY	19	None	None
R_SPARC_GLOB_DAT	20	V-word32	S + A
R_SPARC_JMP_SLOT	21	None	See R_SPARC_JMP_SLOT
R_SPARC_RELATIVE	22	V-word32	B + A

TABLE 7-18 SPARC: ELF Relocation Types (Continued)

Name	Value	Field	Calculation
R_SPARC_UA32	23	V-word32	S + A
R_SPARC_PLT32	24	V-word32	L + A
R_SPARC_HIPLT22	25	T-imm22	(L + A) >> 10
R_SPARC_LOPLT10	26	T-simm13	(L + A) & 0x3ff
R_SPARC_PCPLT32	27	V-word32	L + A - P
R_SPARC_PCPLT22	28	V-disp22	(L + A - P) >> 10
R_SPARC_PCPLT10	29	V-simm13	(L + A - P) & 0x3ff
R_SPARC_10	30	V-simm10	S + A
R_SPARC_11	31	V-simm11	S + A
R_SPARC_HH22	34	V-imm22	(S + A) >> 42
R_SPARC_HM10	35	T-simm13	((S + A) >> 32) & 0x3ff
R_SPARC_LM22	36	T-imm22	(S + A) >> 10
R_SPARC_PC_HH22	37	V-imm22	(S + A - P) >> 42
R_SPARC_PC_HM10	38	T-simm13	((S + A - P) >> 32) & 0x3ff
R_SPARC_PC_LM22	39	T-imm22	(S + A - P) >> 10
R_SPARC_WDISP16	40	V-d2/disp14	(S + A - P) >> 2
R_SPARC_WDISP19	41	V-disp19	(S + A - P) >> 2
R_SPARC_7	43	V-imm7	S + A
R_SPARC_5	44	V-imm5	S + A
R_SPARC_6	45	V-imm6	S + A
R_SPARC_HIX22	48	V-imm22	((S + A) ^ 0xffffffffffffffff) >> 10
R_SPARC_LOX10	49	T-simm13	((S + A) & 0x3ff) 0x1c00
R_SPARC_H44	50	V-imm22	(S + A) >> 22
R_SPARC_M44	51	T-imm10	((S + A) >> 12) & 0x3ff
R_SPARC_L44	52	T-imm13	(S + A) & 0xfff
R_SPARC_REGISTER	53	V-word32	S + A
R_SPARC_UA16	55	V-half16	S + A

TABLE 7-18 SPARC: ELF Relocation Types (Continued)

Name	Value	Field	Calculation
R_SPARC_GOTDATA_HIX22	80	T-imm22	$((S + A - GOT) \gg 10) \wedge ((S + A - GOT) \gg 42)$
R_SPARC_GOTDATA_LOX22	81	T-imm13	$((S + A - GOT) \& 0x3ff) (0x1c00 \wedge \sim ((S + A - GOT) \gg 50) \& 0x1c00)$
R_SPARC_GOTDATA_OP_HIX22	82	T-imm22	$((S + A - GOT) \gg 10) \wedge ((S + A - GOT) \gg 42)$
R_SPARC_GOTDATA_OP_LOX22	83	T-imm13	$((S + A - GOT) \& 0x3ff) (0x1c00 \wedge \sim ((S + A - GOT) \gg 50) \& 0x1c00)$
R_SPARC_GOTDATA_OP	84	Word32	special

Note – Additional relocations are available for thread-local storage references. These relocations are covered in Chapter 8.

Some relocation types have semantics beyond simple calculation:

R_SPARC_GOT10

Resembles R_SPARC_LO10, except that it refers to the address of the symbol's global offset table entry. Additionally, R_SPARC_GOT10 instructs the link-editor to create a global offset table.

R_SPARC_GOT13

Resembles R_SPARC_13, except that it refers to the address of the symbol's global offset table entry. Additionally, R_SPARC_GOT13 instructs the link-editor to create a global offset table.

R_SPARC_GOT22

Resembles R_SPARC_22, except that it refers to the address of the symbol's global offset table entry. Additionally, R_SPARC_GOT22 instructs the link-editor to create a global offset table.

R_SPARC_WPLT30

Resembles R_SPARC_WDISP30, except that it refers to the address of the symbol's procedure linkage table entry. Additionally, R_SPARC_WPLT30 instructs the link-editor to create a procedure linkage table.

R_SPARC_COPY

Created by the link-editor for dynamic executables to preserve a read-only text segment. Its offset member refers to a location in a writable segment. The symbol table index specifies a symbol that should exist both in the current object file and in

a shared object. During execution, the runtime linker copies data associated with the shared object's symbol to the location specified by the offset. See "Copy Relocations" on page 122.

R_SPARC_GLOB_DAT

Resembles R_SPARC_32, except that it sets a global offset table entry to the address of the specified symbol. The special relocation type enables you to determine the correspondence between symbols and global offset table entries.

R_SPARC_JMP_SLOT

Created by the link-editor for dynamic objects to provide lazy binding. Its offset member gives the location of a procedure linkage table entry. The runtime linker modifies the procedure linkage table entry to transfer control to the designated symbol address.

R_SPARC_RELATIVE

Created by the link-editor for dynamic objects. Its offset member gives the location within a shared object that contains a value representing a relative address. The runtime linker computes the corresponding virtual address by adding the virtual address at which the shared object is loaded to the relative address. Relocation entries for this type must specify 0 for the symbol table index.

R_SPARC_UA32

Resembles R_SPARC_32, except that it refers to an unaligned word. The word to be relocated must be treated as four separate bytes with arbitrary alignment, not as a word aligned according to the architecture requirements.

R_SPARC_LM22

Resembles R_SPARC_HI22, except that it truncates rather than validates.

R_SPARC_PC_LM22

Resembles R_SPARC_PC22, except that it truncates rather than validates.

R_SPARC_HIX22

Used with R_SPARC_LOX10 for executables that are confined to the uppermost 4 gigabytes of the 64-bit address space. Similar to R_SPARC_HI22, but supplies ones complement of linked value.

R_SPARC_LOX10

Used with R_SPARC_HIX22. Similar to R_SPARC_LO10, but always sets bits 10 through 12 of the linked value.

R_SPARC_L44

Used with the R_SPARC_H44 and R_SPARC_M44 relocation types to generate a 44-bit absolute addressing model.

R_SPARC_REGISTER

Used to initialize a register symbol. Its offset member contains the register number to be initialized. There must be a corresponding register symbol for this register of type SHN_ABS.

R_SPARC_GOTDATA_OP_HIX22, R_SPARC_GOTDATA_OP_LOX22, and R_SPARC_GOTDATA_OP

These relocations provide for code transformations.

64-bit SPARC: Relocation Types

The following notation, used in relocation calculation, is unique to 64-bit SPARC.

- The secondary addend used to compute the value of the relocation field. This addend is extracted from the `r_info` field by applying the `ELF64_R_TYPE_DATA` macro.

The relocations listed in the following table extend, or alter, those define for 32-bit SPARC. See “SPARC: Relocation Types” on page 215.

TABLE 7-19 64-bit SPARC: ELF Relocation Types

Name	Value	Field	Calculation
R_SPARC_HI22	9	V-imm22	$(S + A) \gg 10$
R_SPARC_GLOB_DAT	20	V-xword64	$S + A$
R_SPARC_RELATIVE	22	V-xword64	$B + A$
R_SPARC_64	32	V-xword64	$S + A$
R_SPARC_OLO10	33	V-simm13	$((S + A) \& 0x3ff) + O$
R_SPARC_DISP64	46	V-xword64	$S + A - P$
R_SPARC_PLT64	47	V-xword64	$L + A$
R_SPARC_REGISTER	53	V-xword64	$S + A$
R_SPARC_UA64	54	V-xword64	$S + A$

The following relocation type has semantics beyond simple calculation:

R_SPARC_OLO10

Resembles R_SPARC_LO10, except that an extra offset is added to make full use of the 13-bit signed immediate field.

x86: Relocation Types

The relocations listed in the following table are defined for 32-bit x86.

TABLE 7-20 x86: ELF Relocation Types

Name	Value	Field	Calculation
R_386_NONE	0	None	None
R_386_32	1	word32	S + A
R_386_PC32	2	word32	S + A - P
R_386_GOT32	3	word32	G + A
R_386_PLT32	4	word32	L + A - P
R_386_COPY	5	None	None
R_386_GLOB_DAT	6	word32	S
R_386_JMP_SLOT	7	word32	S
R_386_RELATIVE	8	word32	B + A
R_386_GOTOFF	9	word32	S + A - GOT
R_386_GOTPC	10	word32	GOT + A - P
R_386_32PLT	11	word32	L + A

Note – Additional relocations are available for thread-local storage references. These relocations are covered in Chapter 8.

Some relocation types have semantics beyond simple calculation:

R_386_GOT32

Computes the distance from the base of the global offset table to the symbol's global offset table entry. It also instructs the link-editor to create a global offset table.

R_386_PLT32

Computes the address of the symbol's procedure linkage table entry and instructs the link-editor to create a procedure linkage table.

R_386_COPY

Created by the link-editor for dynamic executables to preserve a read-only text segment. Its offset member refers to a location in a writable segment. The symbol table index specifies a symbol that should exist both in the current object file and in a shared object. During execution, the runtime linker copies data associated with the shared object's symbol to the location specified by the offset. See "Copy Relocations" on page 122.

R_386_GLOB_DAT

Used to set a global offset table entry to the address of the specified symbol. The special relocation type enable you to determine the correspondence between symbols and global offset table entries.

R_386_JMP_SLOT

Created by the link-editor for dynamic objects to provide lazy binding. Its offset member gives the location of a procedure linkage table entry. The runtime linker modifies the procedure linkage table entry to transfer control to the designated symbol address.

R_386_RELATIVE

Created by the link-editor for dynamic objects. Its offset member gives the location within a shared object that contains a value representing a relative address. The runtime linker computes the corresponding virtual address by adding the virtual address at which the shared object is loaded to the relative address. Relocation entries for this type must specify 0 for the symbol table index.

R_386_GOTOFF

Computes the difference between a symbol's value and the address of the global offset table. It also instructs the link-editor to create the global offset table.

R_386_GOTPC

Resembles R_386_PC32, except that it uses the address of the global offset table in its calculation. The symbol referenced in this relocation normally is `_GLOBAL_OFFSET_TABLE_`, which also instructs the link-editor to create the global offset table.

String Table Section

String table sections hold null-terminated character sequences, commonly called strings. The object file uses these strings to represent symbol and section names. You reference a string as an index into the string table section.

The first byte, which is index zero, holds a null character. Likewise, a string table's last byte holds a null character, ensuring null termination for all strings. A string whose index is zero specifies either no name or a null name, depending on the context.

An empty string table section is permitted. The section header's `sh_size` member contains zero. Nonzero indexes are invalid for an empty string table.

A section header's `sh_name` member holds an index into the section header string table section, as designated by the `e_shstrndx` member of the ELF header. The following figure shows a string table with 25 bytes and the strings associated with various indexes.

Index	+0	+1	+2	+3	+4	+5	+6	+7	+8	+9
0	\0	n	a	m	e	.	\0	V	a	r
10	i	a	b	l	e	\0	a	b	l	e
20	\0	\0	x	x	\0					

FIGURE 7-7 ELF String Table

The table below shows the strings of the string table shown in the preceding figure.

TABLE 7-21 ELF String Table Indexes

Index	String
0	None
1	name
7	Variable
11	able
16	able
24	<i>null string</i>

As the example shows, a string table index can refer to any byte in the section. A string can appear more than once. References to substrings can exist. A single string can be referenced multiple times. Unreferenced strings also are allowed.

Symbol Table Section

An object file's symbol table holds information needed to locate and relocate a program's symbolic definitions and references. A symbol table index is a subscript into this array. Index 0 both designates the first entry in the table and serves as the undefined symbol index. See Table 7-25.

A symbol table entry has the following format, defined in `sys/elf.h`:

```
typedef struct {
    Elf32_Word      st_name;
    Elf32_Addr      st_value;
    Elf32_Word      st_size;
    unsigned char   st_info;
    unsigned char   st_other;
    Elf32_Half      st_shndx;
} Elf32_Sym;

typedef struct {
    Elf64_Word      st_name;
```

```

        unsigned char    st_info;
        unsigned char    st_other;
        Elf64_Half       st_shndx;
        Elf64_Addr       st_value;
        Elf64_Xword      st_size;
} Elf64_Sym;

```

The elements of this structure are:

st_name

An index into the object file's symbol string table, which holds the character representations of the symbol names. If the value is nonzero, it represents a string table index that gives the symbol name. Otherwise, the symbol table entry has no name.

st_value

The value of the associated symbol. Depending on the context, this can be an absolute value, an address, and so forth. See "Symbol Values" on page 228.

st_size

Many symbols have associated sizes. For example, a data object's size is the number of bytes contained in the object. This member holds 0 if the symbol has no size or an unknown size.

st_info

The symbol's type and binding attributes. A list of the values and meanings appears in Table 7-22. The following code shows how to manipulate the values, defined in `sys/elf.h`:

```

#define ELF32_ST_BIND(info)      ((info) >> 4)
#define ELF32_ST_TYPE(info)     ((info) & 0xf)
#define ELF32_ST_INFO(bind, type)  (((bind)<<4) + ((type) &0xf))

#define ELF64_ST_BIND(info)      ((info) >> 4)
#define ELF64_ST_TYPE(info)     ((info) & 0xf)
#define ELF64_ST_INFO(bind, type)  (((bind)<<4) + ((type) &0xf))

```

st_other

A symbol's visibility. A list of the values and meanings appears in Table 7-24. The following code shows how to manipulate the values for both 32-bit and 64-bit objects. Other bits contain 0 and have no defined meaning.

```

#define ELF32_ST_VISIBILITY(o)    ((o) &0x3)
#define ELF64_ST_VISIBILITY(o)    ((o) &0x3)

```

st_shndx

Every symbol table entry is defined in relation to some section. This member holds the relevant section header table index. Some section indexes indicate special meanings. See Table 7-11.

If this member contains `SHN_XINDEX`, then the actual section header index is too large to fit in this field. The actual value is contained in the associated section of type `SHT_SYMTAB_SHNDX`.

A symbol's binding, determined from its `st_info` field, determines the linkage visibility and behavior.

TABLE 7-22 ELF Symbol Binding, `ELF32_ST_BIND` and `ELF64_ST_BIND`

Name	Value
<code>STB_LOCAL</code>	0
<code>STB_GLOBAL</code>	1
<code>STB_WEAK</code>	2
<code>STB_LOOS</code>	10
<code>STB_HIOS</code>	12
<code>STB_LOPROC</code>	13
<code>STB_HIPROC</code>	15

`STB_LOCAL`

Local symbol. These symbols are not visible outside the object file containing their definition. Local symbols of the same name can exist in multiple files without interfering with each other.

`STB_GLOBAL`

Global symbols. These symbols are visible to all object files being combined. One file's definition of a global symbol satisfies another file's undefined reference to the same global symbol.

`STB_WEAK`

Weak symbols. These symbols resemble global symbols, but their definitions have lower precedence.

`STB_LOOS` - `STB_HIOS`

Values in this inclusive range are reserved for operating system-specific semantics.

`STB_LOPROC` - `STB_HIPROC`

Values in this inclusive range are reserved for processor-specific semantics.

Global symbols and weak symbols differ in two major ways:

- When the link-editor combines several relocatable object files, it does not allow multiple definitions of `STB_GLOBAL` symbols with the same name. On the other hand, if a defined global symbol exists, the appearance of a weak symbol with the same name does not cause an error. The link-editor honors the global definition and ignores the weak ones.

Similarly, if a common symbol exists, the appearance of a weak symbol with the same name does not cause an error. The link-editor uses the common definition and ignores the weak one. A common symbol has the `st_shndx` field holding `SHN_COMMON`. See "Symbol Resolution" on page 38.

- When the link-editor searches archive libraries it extracts archive members that contain definitions of undefined or tentative global symbols. The member's definition can be either a global or a weak symbol.

The link-editor, by default, does not extract archive members to resolve undefined weak symbols. Unresolved weak symbols have a zero value. The use of `-z weakextract` overrides this default behavior. It enables weak references to cause the extraction of archive members.

Note – Weak symbols are intended primarily for use in system software. Their use in application programs is discouraged.

In each symbol table, all symbols with `STB_LOCAL` binding precede the weak and global symbols. As “Sections” on page 187 describes, a symbol table section's `sh_info` section header member holds the symbol table index for the first non-local symbol.

A symbol's type, determined from its `st_info` field, provides a general classification for the associated entity.

TABLE 7-23 ELF Symbol Types, `ELF32_ST_TYPE` and `ELF64_ST_TYPE`

Name	Value
<code>STT_NOTYPE</code>	0
<code>STT_OBJECT</code>	1
<code>STT_FUNC</code>	2
<code>STT_SECTION</code>	3
<code>STT_FILE</code>	4
<code>STT_COMMON</code>	5
<code>STT_TLS</code>	6
<code>STT_LOOS</code>	10
<code>STT_HIOS</code>	12
<code>STT_LOPROC</code>	13
<code>STT_SPARC_REGISTER</code>	13
<code>STT_HIPROC</code>	15

`STT_NOTYPE`
The symbol type is not specified.

STT_OBJECT

This symbol is associated with a data object, such as a variable, an array, and so forth.

STT_FUNC

This symbol is associated with a function or other executable code.

STT_SECTION

This symbol is associated with a section. Symbol table entries of this type exist primarily for relocation and normally have `STB_LOCAL` binding.

STT_FILE

Conventionally, the symbol's name gives the name of the source file associated with the object file. A file symbol has `STB_LOCAL` binding and its section index is `SHN_ABS`. This symbol, if present, precedes the other `STB_LOCAL` symbols for the file. Symbol index 1 of the `SHT_SYMTAB` is an `STT_FILE` symbol representing the file itself. Conventionally, this symbol is followed by the files `STT_SECTION` symbols, and any global symbols that have been reduced to locals.

STT_COMMON

This symbol labels an uninitialized common block. It is treated exactly the same as `STT_OBJECT`.

STT_TLS

The symbol specifies a thread-local storage entity. When defined, it gives the assigned offset for the symbol, not the actual address. Symbols of type `STT_TLS` can be referenced by only special thread-local storage relocations and thread-local storage relocations can only reference symbols with type `STT_TLS`. See Chapter 8 for details.

STT_LOOS - STT_HIOS

Values in this inclusive range are reserved for operating system-specific semantics.

STT_LOPROC - STT_HIPROC

Values in this inclusive range are reserved for processor-specific semantics.

A symbol's visibility, determined from its `st_other` field, may be specified in a relocatable object. This visibility defines how that symbol may be accessed once the symbol has become part of an executable or shared object.

TABLE 7-24 ELF Symbol Visibility

Name	Value
<code>STV_DEFAULT</code>	0
<code>STV_INTERNAL</code>	1
<code>STV_HIDDEN</code>	2
<code>STV_PROTECTED</code>	3

STV_DEFAULT

The visibility of symbols with the `STV_DEFAULT` attribute is as specified by the symbol's binding type. That is, global and weak symbols are visible outside of their defining component, the executable file or shared object. Local symbols are hidden. Global and weak symbols can also be preempted, that is, they may be interposed by definitions of the same name in another component.

STV_PROTECTED

A symbol defined in the current component is protected if it is visible in other components but cannot be preempted. Any reference to such a symbol from within the defining component must be resolved to the definition in that component, even if there is a definition in another component that would interpose by the default rules. A symbol with `STB_LOCAL` binding will not have `STV_PROTECTED` visibility.

STV_HIDDEN

A symbol defined in the current component is hidden if its name is not visible to other components. Such a symbol is necessarily protected. This attribute is used to control the external interface of a component. An object named by such a symbol may still be referenced from another component if its address is passed outside.

A hidden symbol contained in a relocatable object is either removed or converted to `STB_LOCAL` binding by the link-editor when the relocatable object is included in an executable file or shared object.

STV_INTERNAL

This visibility attribute is currently reserved.

None of the visibility attributes affects the resolution of symbols within an executable or shared object during link-editing. Such resolution is controlled by the binding type. Once the link-editor has chosen its resolution, these attributes impose two requirements. Both requirements are based on the fact that references in the code being linked may have been optimized to take advantage of the attributes.

- First, all of the non-default visibility attributes, when applied to a symbol reference, imply that a definition to satisfy that reference must be provided within the current executable or shared object. If this type of symbol reference has no definition within the component being linked, then the reference must have `STB_WEAK` binding and is resolved to zero.
- Second, if any reference to or definition of a name is a symbol with a non-default visibility attribute, the visibility attribute must be propagated to the resolving symbol in the linked object. If different visibility attributes are specified for distinct references to or definitions of a symbol, the most constraining visibility attribute must be propagated to the resolving symbol in the linked object. The attributes, ordered from least to most constraining, are `STV_PROTECTED`, `STV_HIDDEN` and `STV_INTERNAL`.

If a symbol's value refers to a specific location within a section, its section index member, `st_shndx`, holds an index into the section header table. As the section moves during relocation, the symbol's value changes as well. References to the symbol continue to point to the same location in the program. Some special section index values give other semantics:

SHN_ABS

This symbol has an absolute value that does not change because of relocation.

SHN_COMMON

This symbol labels a common block that has not yet been allocated. The symbol's value gives alignment constraints, similar to a section's `sh_addralign` member. The link-editor allocates the storage for the symbol at an address that is a multiple of `st_value`. The symbol's size tells how many bytes are required.

SHN_UNDEF

This section table index means the symbol is undefined. When the link-editor combines this object file with another that defines the indicated symbol, this file's references to the symbol will be bound to the actual definition.

As mentioned above, the symbol table entry for index 0 (`STN_UNDEF`) is reserved. This entry holds the values listed in the following table.

TABLE 7-25 ELF Symbol Table Entry: Index 0

Name	Value	Note
<code>st_name</code>	0	No name
<code>st_value</code>	0	Zero value
<code>st_size</code>	0	No size
<code>st_info</code>	0	No type, local binding
<code>st_other</code>	0	
<code>st_shndx</code>	<code>SHN_UNDEF</code>	No section

Symbol Values

Symbol table entries for different object file types have slightly different interpretations for the `st_value` member.

- In relocatable files, `st_value` holds alignment constraints for a symbol whose section index is `SHN_COMMON`.
- In relocatable files, `st_value` holds a section offset for a defined symbol. `st_value` is an offset from the beginning of the section that `st_shndx` identifies.
- In executable and shared object files, `st_value` holds a virtual address. To make these files' symbols more useful for the runtime linker, the section offset (file interpretation) gives way to a virtual address (memory interpretation) for which

the section number is irrelevant.

Although the symbol table values have similar meanings for different object files, the data allow efficient access by the appropriate programs.

Register Symbols

The SPARC architecture supports register symbols, which are symbols that initialize a global register. A symbol table entry for a register symbol contains the entries listed in the following table.

TABLE 7-26 SPARC: ELF Symbol Table Entry: Register Symbol

Field	Meaning
st_name	Index into string table of the name of the symbol, or 0 for a scratch register.
st_value	Register number. See the ABI manual for integer register assignments.
st_size	Unused (0).
st_info	Bind is typically STB_GLOBAL, type must be STT_SPARC_REGISTER.
st_other	Unused (0).
st_shndx	SHN_ABS if this object initializes this register symbol; SHN_UNDEF otherwise.

The register values defined for SPARC are listed in the following table.

TABLE 7-27 SPARC: ELF Register Numbers

Name	Value	Meaning
STO_SPARC_REGISTER_G2	0x2	%g2
STO_SPARC_REGISTER_G3	0x3	%g3

Absence of an entry for a particular global register means that the particular global register is not used at all by the object.

Syminfo Table Section

The syminfo section contains multiple entries of the type Elf32_Syminfo or Elf64_Syminfo. There is one entry in the .SUNW_syminfo section for every entry in the associated symbol table (sh_link).

If this section is present in an object, additional symbol information is to be found by taking the symbol index from the associated symbol table and using that to find the corresponding `Elf32_Syminfo` or `Elf64_Syminfo` entry in this section. The associated symbol table and the `Syminfo` table will always have the same number of entries.

Index 0 is used to store the current version of the `Syminfo` table, which is `SYMINFO_CURRENT`. Since symbol table entry 0 is always reserved for the `UNDEF` symbol table entry, this does not pose any conflicts.

An `Syminfo` entry has the following format, defined in `sys/link.h`:

```
typedef struct {
    Elf32_Half    si_boundto;
    Elf32_Half    si_flags;
} Elf32_Syminfo;

typedef struct {
    Elf64_Half    si_boundto;
    Elf64_Half    si_flags;
} Elf64_Syminfo;
```

The elements of this structure are:

`si_boundto`

This index is to an entry in the `.dynamic` section, identified by the `sh_info` field, that augments the `Syminfo` flags. For example, a `DT_NEEDED` entry identifies a dynamic object associated with the `Syminfo` entry. The entries that follow are reserved values for `si_boundto`.

TABLE 7-28 ELF `si_boundto` Reserved Values

Name	Value	Meaning
<code>SYMINFO_BT_SELF</code>	<code>0xffff</code>	Symbol bound to self.
<code>SYMINFO_BT_PARENT</code>	<code>0xffffe</code>	Symbol bound to parent. The parent is the first object to cause this dynamic object to be loaded.
<code>SYMINFO_BT_NONE</code>	<code>0xfffd</code>	Symbol has no special symbol binding.

`si_flags`

This bit-field can have flags set, as shown in the following table.

TABLE 7-29 ELF `Syminfo` Flags

Name	Value	Meaning
<code>SYMINFO_FLG_DIRECT</code>	<code>0x01</code>	Symbol reference has direct association to object containing definition.

TABLE 7-29 ELF Syminfo Flags (Continued)

Name	Value	Meaning
SYMINFO_FLG_COPY	0x04	Symbol definition is the result of a copy-relocation.
SYMINFO_FLG_LAZYLOAD	0x08	Symbol reference is to an object that should be lazily loaded.
SYMINFO_FLG_DIRECTBIND	0x10	Symbol reference should be bound directly to the definition.
SYMINFO_FLG_NOEXTDIRECT	0x20	Do not allow external reference to directly bind to this symbol definition.

Versioning Sections

Objects created by the link-editor can contain two types of versioning information:

- *Version definitions* provide associations of global symbols and are implemented using sections of type SHT_SUNW_verdef and SHT_SUNW_versym.
- *Version dependencies* indicate the version definition requirements from other object dependencies and are implemented using sections of type SHT_SUNW_verneed.

The structures that form these sections are defined in `sys/link.h`. Sections that contain versioning information are named `.SUNW_version`.

Version Definition Section

This section is defined by the type SHT_SUNW_verdef. If this section exists, a SHT_SUNW_versym section must also exist. Using these two structures, an association of symbols-to-version definitions is maintained within the file. See “Creating a Version Definition” on page 131. Elements of this section have the following structure:

```
typedef struct {
    Elf32_Half    vd_version;
    Elf32_Half    vd_flags;
    Elf32_Half    vd_ndx;
    Elf32_Half    vd_cnt;
    Elf32_Word    vd_hash;
    Elf32_Word    vd_aux;
    Elf32_Word    vd_next;
} Elf32_Verdef;

typedef struct {
    Elf32_Word    vda_name;
    Elf32_Word    vda_next;
} Elf32_Verdaux;
```

```

typedef struct {
    Elf64_Half    vd_version;
    Elf64_Half    vd_flags;
    Elf64_Half    vd_ndx;
    Elf64_Half    vd_cnt;
    Elf64_Word    vd_hash;
    Elf64_Word    vd_aux;
    Elf64_Word    vd_next;
} Elf64_Verdef;

typedef struct {
    Elf64_Word    vda_name;
    Elf64_Word    vda_next;
} Elf64_Verdaux;

```

The elements of this structure are:

vd_version

This member identifies the version of the structure itself, as listed in the following table.

TABLE 7-30 ELF Version Definition Structure Versions

Name	Value	Meaning
VER_DEF_NONE	0	Invalid version.
VER_DEF_CURRENT	>=1	Current version.

The value 1 signifies the original section format. Extensions will create new versions with higher numbers. The value of VER_DEF_CURRENT changes as necessary to reflect the current version number.

vd_flags

This member holds version definition-specific information, as listed in the following table.

TABLE 7-31 ELF Version Definition Section Flags

Name	Value	Meaning
VER_FLG_BASE	0x1	Version definition of the file itself.
VER_FLG_WEAK	0x2	Weak version identifier.

The base version definition is always present when version definitions, or symbol auto-reduction, have been applied to the file. The base version provides a default version for the files reserved symbols. A weak version definition has no symbols associated with it. See “Creating a Weak Version Definition” on page 134.

vd_ndx

The version index. Each version definition has a unique index that is used to associate SHT_SUNW_versym entries to the appropriate version definition.

`vd_cnt`
The number of elements in the `Elf32_Verdaux` array.

`vd_hash`
The hash value of the version definition name. This value is generated using the same hashing function described in “Hash Table Section” on page 205.

`vd_aux`
The byte offset from the start of this `Elf32_Verdef` entry to the `Elf32_Verdaux` array of version definition names. The first element of the array must exist. It points to the version definition string this structure defines. Additional elements can be present. The number of elements is indicated by the `vd_cnt` value. These elements represent the dependencies of this version definition. Each of these dependencies will have its own version definition structure.

`vd_next`
The byte offset from the start of this `Elf32_Verdef` structure to the next `Elf32_Verdef` entry.

`vda_name`
The string table offset to a null-terminated string, giving the name of the version definition.

`vda_next`
The byte offset from the start of this `Elf32_Verdaux` entry to the next `Elf32_Verdaux` entry.

Version Symbol Section

The version symbol section is defined by the type `SHT_SUNW_versym`, and consists of an array of elements having the following structure:

```
typedef Elf32_Half    Elf32_Versym;
typedef Elf64_Half    Elf64_Versym;
```

The number of elements of the array must equal the number of symbol table entries contained in the associated symbol table. This number is determined by the section’s `sh_link` value. Each element of the array contains a single index that can have the values shown in the following table.

TABLE 7-32 ELF Version Dependency Indexes

Name	Value	Meaning
<code>VER_NDX_LOCAL</code>	0	Symbol has local scope.
<code>VER_NDX_GLOBAL</code>	1	Symbol has global scope (assigned to base version definition).

TABLE 7-32 ELF Version Dependency Indexes (Continued)

Name	Value	Meaning
	>1	Symbol has global scope (assigned to user-defined version definition).

Any index values greater than `VER_NDX_GLOBAL` must correspond to the `vd_ndx` value of an entry in the `SHT_SUNW_verdef` section. If no index values greater than `VER_NDX_GLOBAL` exist, then no `SHT_SUNW_verdef` section need be present.

Version Dependency Section

The version dependency section is defined by the type `SHT_SUNW_verneed`. This section complements the dynamic dependency requirements of the file by indicating the version definitions required from these dependencies. A recording is made in this section only if a dependency contains version definitions. Elements of this section have the following structure:

```
typedef struct {
    Elf32_Half    vn_version;
    Elf32_Half    vn_cnt;
    Elf32_Word    vn_file;
    Elf32_Word    vn_aux;
    Elf32_Word    vn_next;
} Elf32_Verneed;

typedef struct {
    Elf32_Word    vna_hash;
    Elf32_Half    vna_flags;
    Elf32_Half    vna_other;
    Elf32_Word    vna_name;
    Elf32_Word    vna_next;
} Elf32_Vernaux;

typedef struct {
    Elf64_Half    vn_version;
    Elf64_Half    vn_cnt;
    Elf64_Word    vn_file;
    Elf64_Word    vn_aux;
    Elf64_Word    vn_next;
} Elf64_Verneed;

typedef struct {
    Elf64_Word    vna_hash;
    Elf64_Half    vna_flags;
    Elf64_Half    vna_other;
    Elf64_Word    vna_name;
    Elf64_Word    vna_next;
} Elf64_Vernaux;
```

The elements of this structure are:

`vn_version`

This member identifies the version of the structure itself, as listed in the following table.

TABLE 7-33 ELF Version Dependency Structure Versions

Name	Value	Meaning
<code>VER_NEED_NONE</code>	0	Invalid version.
<code>VER_NEED_CURRENT</code>	≥ 1	Current version.

The value 1 signifies the original section format. Extensions will create new versions with higher numbers. The value of `VER_NEED_CURRENT` changes as necessary to reflect the current version number.

`vn_cnt`

The number of elements in the `Elf32_VerNaux` array.

`vn_file`

The string table offset to a null-terminated string, that provides the file name having a version dependency. This name matches one of the `.dynamic` dependencies found in the file. See “Dynamic Section” on page 248.

`vn_aux`

The byte offset, from the start of this `Elf32_Verneed` entry, to the `Elf32_VerNaux` array of version definitions required from the associated file dependency. There must exist at least one version dependency. Additional version dependencies can be present, the number being indicated by the `vn_cnt` value.

`vn_next`

The byte offset, from the start of this `Elf32_Verneed` entry, to the next `Elf32_Verneed` entry.

`vna_hash`

The hash value of the version dependency name. This value is generated using the same hashing function described in “Hash Table Section” on page 205.

`vna_flags`

Version dependency specific information, as listed in the following table.

TABLE 7-34 ELF Version Dependency Structure Flags

Name	Value	Meaning
<code>VER_FLG_WEAK</code>	0x2	Weak version identifier.

A weak version dependency indicates an original binding to a weak version definition.

`vna_other`

Presently unused.

vna_name

The string table offset to a null-terminated string, giving the name of the version dependency.

vna_next

The byte offset from the start of this `Elf32_Vernaux` entry to the next `Elf32_Vernaux` entry.

Dynamic Linking

This section describes the object file information and system actions that create running programs. Most information here applies to all systems. Information specific to one processor resides in sections marked accordingly.

Executable and shared object files statically represent application programs. To execute such programs, the system uses the files to create dynamic program representations, or process images. A process image has segments that contain its text, data, stack, and so on. The major subsections of this section are:

- “Program Header” on page 236 describes object file structures that are directly involved in program execution. The primary data structure, a program header table, locates segment images in the file and contains other information needed to create the memory image of the program.
- “Program Loading (Processor-Specific)” on page 241 describes the information used to load a program into memory.
- “Runtime Linker” on page 247 describes the information used to specify and resolve symbolic references among the object files of the process image.

Program Header

An executable or shared object file’s program header table is an array of structures, each describing a segment or other information that the system needs to prepare the program for execution. An object file segment contains one or more sections, as described in “Segment Contents” on page 241.

Program headers are meaningful only for executable and shared object files. A file specifies its own program header size with the ELF header’s `e_phentsize` and `e_phnum` members..

A program header has the following structure, defined in `sys/elf.h`:

```
typedef struct {
    Elf32_Word    p_type;
    Elf32_Off     p_offset;
```

```

        Elf32_Addr    p_vaddr;
        Elf32_Addr    p_paddr;
        Elf32_Word    p_filesz;
        Elf32_Word    p_memsz;
        Elf32_Word    p_flags;
        Elf32_Word    p_align;
    } Elf32_Phdr;

typedef struct {
    Elf64_Word    p_type;
    Elf64_Word    p_flags;
    Elf64_Off     p_offset;
    Elf64_Addr    p_vaddr;
    Elf64_Addr    p_paddr;
    Elf64_Xword   p_filesz;
    Elf64_Xword   p_memsz;
    Elf64_Xword   p_align;
} Elf64_Phdr;

```

The elements of this structure are:

p_type

The kind of segment this array element describes or how to interpret the array element's information. Type values and their meanings are specified in Table 7-35.

p_offset

The offset from the beginning of the file at which the first byte of the segment resides.

p_vaddr

The virtual address at which the first byte of the segment resides in memory.

p_paddr

The segment's physical address for systems in which physical addressing is relevant. Because the system ignores physical addressing for application programs, this member has unspecified contents for executable files and shared objects.

p_filesz

The number of bytes in the file image of the segment, which can be zero.

p_memsz

The number of bytes in the memory image of the segment, which can be zero.

p_flags

Flags relevant to the segment. Type values and their meanings are specified in Table 7-36.

p_align

Loadable process segments must have congruent values for `p_vaddr` and `p_offset`, modulo the page size. This member gives the value to which the segments are aligned in memory and in the file. Values 0 and 1 mean no alignment is required. Otherwise, `p_align` should be a positive, integral power of 2, and `p_vaddr` should equal `p_offset`, modulo `p_align`. See "Program Loading (Processor-Specific)" on page 241.

Some entries describe process segments. Other entries give supplementary information and do not contribute to the process image. Segment entries can appear in any order, except as explicitly noted. Defined type values are listed in the following table.

TABLE 7-35 ELF Segment Types

Name	Value
PT_NULL	0
PT_LOAD	1
PT_DYNAMIC	2
PT_INTERP	3
PT_NOTE	4
PT_SHLIB	5
PT_PHDR	6
PT_TLS	7
PT_LOSUNW	0x6fffffff
PT_SUNWBSS	0x6fffffff
PT_SUNWSTACK	0x6fffffff
PT_HISUNW	0x6fffffff
PT_LOPROC	0x70000000
PT_HIPROC	0x7fffffff

PT_NULL
 Unused; other members' values are undefined. This type enables the program header table to contain ignored entries.

PT_LOAD
 Specifies a loadable segment, described by `p_filesz` and `p_memsz`. The bytes from the file are mapped to the beginning of the memory segment. If the segment's memory size (`p_memsz`) is larger than the file size (`p_filesz`), the extra bytes are defined to hold the value 0 and to follow the segment's initialized area. The file size can not be larger than the memory size. Loadable segment entries in the program header table appear in ascending order, sorted on the `p_vaddr` member.

PT_DYNAMIC
 Specifies dynamic linking information. See "Dynamic Section" on page 248.

PT_INTERP
 Specifies the location and size of a null-terminated path name to invoke as an interpreter. This segment type is mandatory for dynamic executable files and can

occur in shared objects. It cannot occur more than once in a file. This type, if present, it must precede any loadable segment entry. See “Program Interpreter” on page 247 for details.

PT_NOTE

Specifies the location and size of auxiliary information. See “Note Section” on page 208 for details.

PT_SHLIB

Reserved but has unspecified semantics.

PT_PHDR

Specifies the location and size of the program header table itself, both in the file and in the memory image of the program. This segment type cannot occur more than once in a file. Moreover, it can occur only if the program header table is part of the memory image of the program. This type, if present, must precede any loadable segment entry. See “Program Interpreter” on page 247 for details.

PT_TLS

Specifies a thread-local storage template. See “Thread-Local Storage Section” on page 272 for details.

PT_LOSUNW - PT_HISUNW

Values in this inclusive range are reserved for Sun-specific semantics.

PT_SUNWBSS

The same attributes as a **PT_LOAD** element and used to describe a `.SUNW_bss` section.

PT_SUNWSTACK

Describes a process stack. Presently only one such element may exist, and only access permissions, as defined in the `p_flags` field, are meaningful.

PT_LOPROC - PT_HIPROC

Values in this inclusive range are reserved for processor-specific semantics.

Note – Unless specifically required elsewhere, all program header segment types are optional. A file’s program header table can contain only those elements relevant to its contents.

Base Address

Executable and shared object files have a base address, which is the lowest virtual address associated with the memory image of the program’s object file. One use of the base address is to relocate the memory image of the program during dynamic linking.

An executable or shared object file's base address is calculated during execution from three values: the memory load address, the maximum page size, and the lowest virtual address of a program's loadable segment. The virtual addresses in the program headers might not represent the actual virtual addresses of the program's memory image. See "Program Loading (Processor-Specific)" on page 241.

To compute the base address, you determine the memory address associated with the lowest `p_vaddr` value for a `PT_LOAD` segment. You then obtain the base address by truncating the memory address to the nearest multiple of the maximum page size. Depending on the kind of file being loaded into memory, the memory address might not match the `p_vaddr` values.

Segment Permissions

A program to be loaded by the system must have at least one loadable segment, although this is not required by the file format. When the system creates loadable segment memory images, it gives access permissions, as specified in the `p_flags` member. All bits included in the `PF_MASKPROC` mask are reserved for processor-specific semantics.

TABLE 7-36 ELF Segment Flags

Name	Value	Meaning
<code>PF_X</code>	<code>0x1</code>	Execute
<code>PF_W</code>	<code>0x2</code>	Write
<code>PF_R</code>	<code>0x4</code>	Read
<code>PF_MASKPROC</code>	<code>0xf0000000</code>	Unspecified

If a permission bit is 0, that bit's type of access is denied. Actual memory permissions depend on the memory management unit, which can vary from one system to another. Although all flag combinations are valid, the system can grant more access than requested. In no case, however, will a segment have write permission unless it is specified explicitly. The following table lists both the exact flag interpretation and the allowable flag interpretation.

TABLE 7-37 ELF Segment Permissions

Flags	Value	Exact	Allowable
None	0	All access denied	All access denied
<code>PF_X</code>	1	Execute only	Read, execute
<code>PF_W</code>	2	Write only	Read, write, execute

TABLE 7-37 ELF Segment Permissions *(Continued)*

Flags	Value	Exact	Allowable
PF_W + PF_X	3	Write, execute	Read, write, execute
PF_R	4	Read only	Read, execute
PF_R + PF_X	5	Read, execute	Read, execute
PF_R + PF_W	6	Read, write	Read, write, execute
PF_R + PF_W + PF_X	7	Read, write, execute	Read, write, execute

For example, typical text segments have read and execute, but not write permissions. Data segments normally have read, write, and execute permissions.

Segment Contents

An object file segment consists of one or more sections, though this fact is transparent to the program header. Whether the file segment holds one or many sections also is immaterial to program loading. Nonetheless, various data must be present for program execution, dynamic linking, and so on. The diagrams below illustrate segment contents in general terms. The order and membership of sections within a segment can vary.

Text segments contain read-only instructions and data. Data segments contain writable data and instructions. See Table 7-16 for a list of all special sections.

A `PT_DYNAMIC` program header element points at the `.dynamic` section. The `.got` and `.plt` sections also hold information related to position-independent code and dynamic linking.

The `.plt` can reside in a text or a data segment, depending on the processor. See “Global Offset Table (Processor-Specific)” on page 260 and “Procedure Linkage Table (Processor-Specific)” on page 261 for details.

The `.bss` section has the type `SHT_NOBITS`. Although it occupies no space in the file, it contributes to the segment’s memory image. Normally, these uninitialized data reside at the end of the segment, thereby making `p_memsz` larger than `p_filesz` in the associated program header element.

Program Loading (Processor-Specific)

As the system creates or augments a process image, it logically copies a file’s segment to a virtual memory segment. When, and if, the system physically reads the file depends on the program’s execution behavior, system load, and so forth.

A process does not require a physical page unless it references the logical page during execution, and processes commonly leave many pages unreferenced. Therefore, delaying physical reads frequently obviates them, improving system performance. To obtain this efficiency in practice, executable and shared object files must have segment images whose file offsets and virtual addresses are congruent, modulo the page size.

Virtual addresses and file offsets for 32-bit segments are congruent modulo 64K (0x10000). Virtual addresses and file offsets for 64-bit segments are congruent modulo 1 megabyte (0x100000). By aligning segments to the maximum page size, the files are suitable for paging regardless of physical page size.

By default, 64-bit SPARC programs are linked with a starting address of 0x100000000. The whole program is above 4 gigabytes, including its text, data, heap, stack, and shared object dependencies. This helps ensure that 64-bit programs are correct because the program will fault in the least significant 4 gigabytes of its address space if it truncates any of its pointers. While 64-bit programs are linked above 4 gigabytes, you can still link them below 4 gigabytes by using a `mapfile` and the `-M` option to the compiler or link-editor. See `/usr/lib/ld/sparcv9/map.below4G`.

The following figure presents the SPARC version of the executable file.

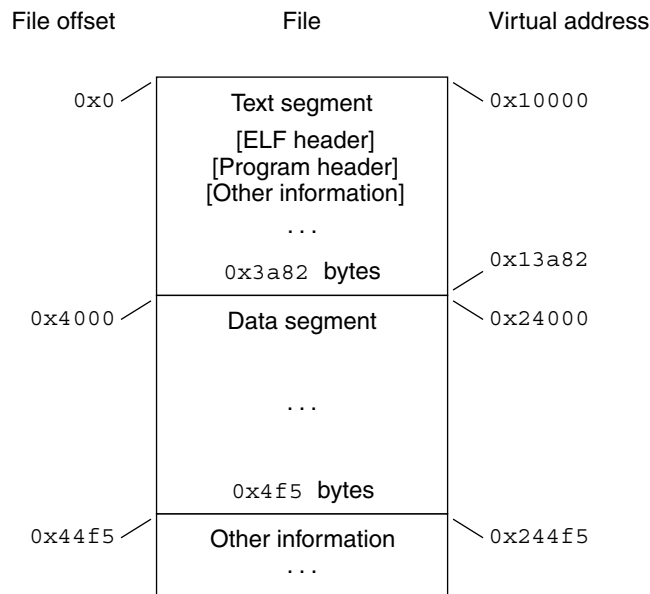


FIGURE 7-8 SPARC: Executable File (64K alignment)

The following table defines the loadable segment elements for the previous figure.

TABLE 7-38 SPARC: ELF Program Header Segments (64K alignment)

Member	Text	Data
p_type	PT_LOAD	PT_LOAD
p_offset	0x0	0x4000
p_vaddr	0x10000	0x24000
p_paddr	Unspecified	Unspecified
p_filesz	0x3a82	0x4f5
p_memsz	0x3a82	0x10a4
p_flags	PF_R + PF_X	PF_R + PF_W + PF_X
p_align	0x10000	0x10000

The following figure presents the x86 version of the executable file.

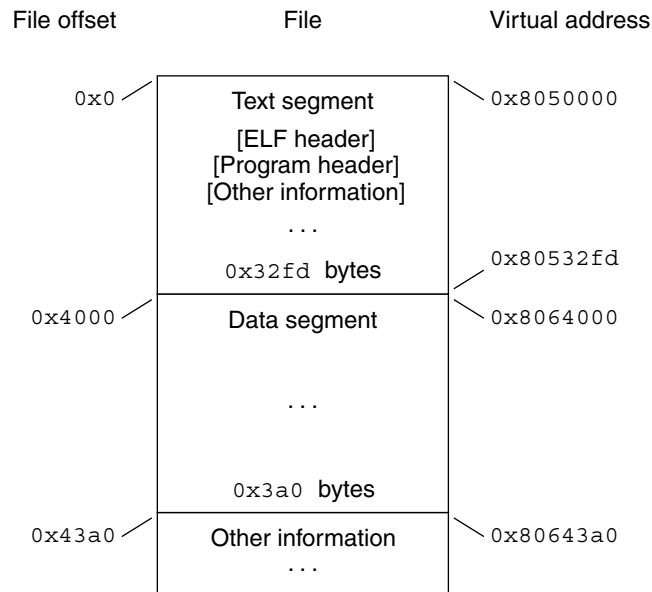


FIGURE 7-9 x86: Executable File (64K alignment)

The following table defines the loadable segment elements for the previous figure.

TABLE 7-39 x86: ELF Program Header Segments (64K alignment)

Member	Text	Data
p_type	PT_LOAD	PT_LOAD
p_offset	0x0	0x4000
p_vaddr	0x8050000	0x8064000
p_paddr	Unspecified	Unspecified
p_filesz	0x32fd	0x3a0
p_memsz	0x32fd	0xdc4
p_flags	PF_R + PF_X	PF_R + PF_W + PF_X
p_align	0x10000	0x10000

The example's file offsets and virtual addresses are congruent modulo the maximum page size for both text and data. Up to four file pages hold impure text or data depending on page size and file system block size.

- The first text page contains the ELF header, the program header table, and other information.
- The last text page holds a copy of the beginning of data.
- The first data page has a copy of the end of text.
- The last data page can contain file information not relevant to the running process. Logically, the system enforces the memory permissions as if each segment were complete and separate. The segments addresses are adjusted to ensure that each logical page in the address space has a single set of permissions. In the examples above, the region of the file holding the end of text and the beginning of data is mapped twice: at one virtual address for text and at a different virtual address for data.

Note – The examples above reflect typical Solaris system binaries that have their text segments rounded.

The end of the data segment requires special handling for uninitialized data, which the system defines to begin with zero values. If a file's last data page includes information not in the logical memory page, the extraneous data must be set to zero, not the unknown contents of the executable file.

Impurities in the other three pages are not logically part of the process image. Whether the system expunges these impurities is unspecified. The memory image for this program is shown in the following figures, assuming 4 Kbyte (0x1000) pages. For simplicity, these figures illustrate only one page size.

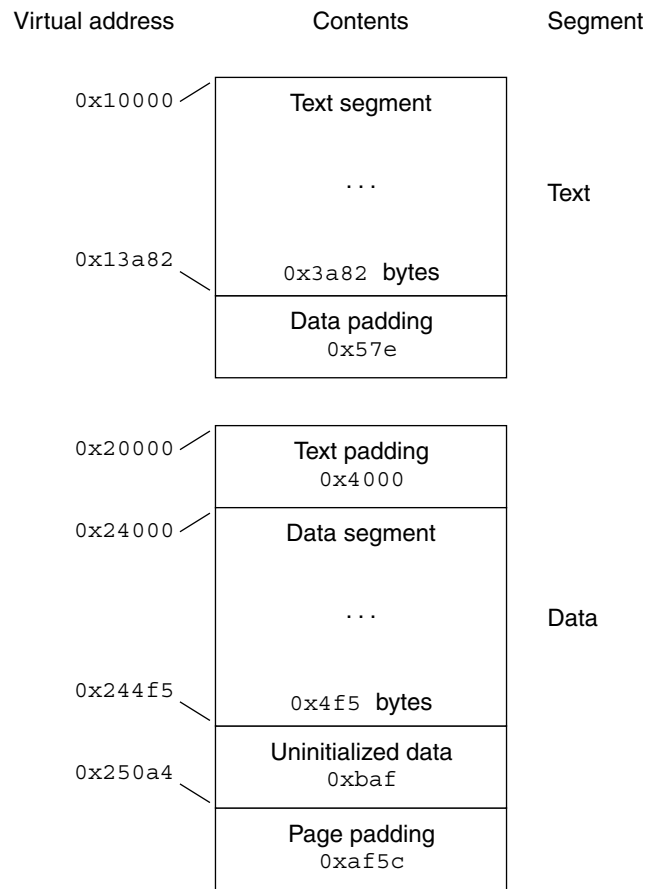


FIGURE 7-10 SPARC: Process Image Segments

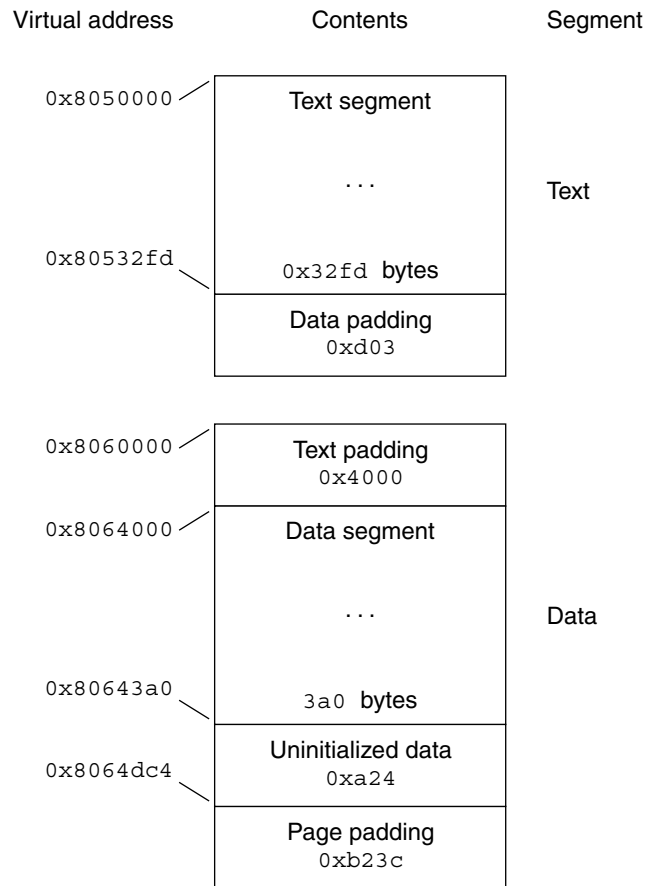


FIGURE 7-11 x86: Process Image Segments

One aspect of segment loading differs between executable files and shared objects. Executable file segments typically contain absolute code. For the process to execute correctly, the segments must reside at the virtual addresses used to create the executable file. The system uses the `p_vaddr` values unchanged as virtual addresses.

On the other hand, shared object segments typically contain position-independent code. This code enables a segment's virtual address change from one process to another, without invalidating execution behavior.

Though the system chooses virtual addresses for individual processes, it maintains the relative positions of the segments. Because position-independent code uses relative addressing between segments, the difference between virtual addresses in memory must match the difference between virtual addresses in the file.

The following tables show possible shared object virtual address assignments for several processes, illustrating constant relative positioning. The tables also include the base address computations.

TABLE 7-40 SPARC: ELF Example Shared Object Segment Addresses

Source	Text	Data	Base Address
File	0x0	0x4000	0x0
Process 1	0xc0000000	0xc0024000	0xc0000000
Process 2	0xc0010000	0xc0034000	0xc0010000
Process 3	0xd0020000	0xd0024000	0xd0020000
Process 4	0xd0030000	0xd0034000	0xd0030000

TABLE 7-41 x86: ELF Example Shared Object Segment Addresses

Source	Text	Data	Base Address
File	0x0	0x4000	0x0
Process 1	0x8000000	0x8004000	0x80000000
Process 2	0x80081000	0x80085000	0x80081000
Process 3	0x900c0000	0x900c4000	0x900c0000
Process 4	0x900c6000	0x900ca000	0x900c6000

Program Interpreter

A dynamic executable or shared object that initiates dynamic linking can have one `PT_INTERP` program header element. During `exec(2)`, the system retrieves a path name from the `PT_INTERP` segment and creates the initial process image from the interpreter file's segments. The interpreter is responsible for receiving control from the system and providing an environment for the application program.

In the Solaris operating environment the interpreter is known as the runtime linker, `ld.so.1(1)`.

Runtime Linker

When creating a dynamic object that initiates dynamic linking, the link-editor adds a program header element of type `PT_INTERP` to an executable file. This element instructing the system to invoke the runtime linker as the program interpreter. `exec(2)` and the runtime linker cooperate to create the process image for the program.

The link-editor constructs various data for executable and shared object files that assist the runtime linker. These data reside in loadable segments, making them available during execution. These segments include:

- A `.dynamic` section with type `SHT_DYNAMIC` that holds various data. The structure residing at the beginning of the section holds the addresses of other dynamic linking information.
- The `.got` and `.plt` sections with type `SHT_PROGBITS` that hold two separate tables: the global offset table and the procedure linkage table. Sections below explain how the runtime linker uses and changes the tables to create memory images for object files.
- The `.hash` section with type `SHT_HASH` that holds a symbol hash table.

Shared objects can occupy virtual memory addresses that are different from the addresses recorded in the file's program header table. The runtime linker relocates the memory image, updating absolute addresses before the application gains control.

Dynamic Section

If an object file participates in dynamic linking, its program header table will have an element of type `PT_DYNAMIC`. This segment contains the `.dynamic` section. A special symbol, `_DYNAMIC`, labels the section, which contains an array of the following structures, defined in `sys/link.h`:

```
typedef struct {
    Elf32_Sword d_tag;
    union {
        Elf32_Word    d_val;
        Elf32_Addr    d_ptr;
        Elf32_Off     d_off;
    } d_un;
} Elf32_Dyn;

typedef struct {
    Elf64_Xword d_tag;
    union {
        Elf64_Xword    d_val;
        Elf64_Addr     d_ptr;
    } d_un;
} Elf64_Dyn;
```

For each object with this type, `d_tag` controls the interpretation of `d_un`.

`d_val`

These objects represent integer values with various interpretations.

`d_ptr`

These objects represent program virtual addresses. A file's virtual addresses might not match the memory virtual addresses during execution. When interpreting addresses contained in the dynamic structure, the runtime linker computes actual

addresses, based on the original file value and the memory base address. For consistency, files do not contain relocation entries to *correct* addresses in the dynamic structure.

To make interpreting the contents of dynamic section entries simpler for tools, the value of each tag, except for those in two special compatibility ranges, will determine the interpretation of the `d_un` union. A tag whose value is an even number indicates a dynamic section entry that uses `d_ptr`. A tag whose value is an odd number indicates a dynamic section entry that uses `d_val` or that uses neither `d_ptr` nor `d_val`. Tags whose values are less than the special value `DT_ENCODING` and tags whose values fall between `DT_HIOS` and `DT_LOPROC` do not follow these rules.

The following table summarizes the tag requirements for executable and shared object files. If a tag is marked *mandatory*, then the dynamic linking array must have an entry of that type. Likewise, *optional* means an entry for the tag can appear but is not required.

TABLE 7-42 ELF Dynamic Array Tags

Name	Value	d_un	Executable	Shared Object
DT_NULL	0	Ignored	Mandatory	Mandatory
DT_NEEDED	1	d_val	Optional	Optional
DT_PLTRELSZ	2	d_val	Optional	Optional
DT_PLTGOT	3	d_ptr	Optional	Optional
DT_HASH	4	d_ptr	Mandatory	Mandatory
DT_STRTAB	5	d_ptr	Mandatory	Mandatory
DT_SYMTAB	6	d_ptr	Mandatory	Mandatory
DT_RELA	7	d_ptr	Mandatory	Optional
DT_RELASZ	8	d_val	Mandatory	Optional
DT_RELAENT	9	d_val	Mandatory	Optional
DT_STRSZ	10	d_val	Mandatory	Mandatory
DT_SYMENT	11	d_val	Mandatory	Mandatory
DT_INIT	12	d_ptr	Optional	Optional
DT_FINI	13	d_ptr	Optional	Optional
DT_SONAME	14	d_val	Ignored	Optional
DT_RPATH	15	d_val	Optional	Optional
DT_SYMBOLIC	16	Ignored	Ignored	Optional

TABLE 7-42 ELF Dynamic Array Tags (Continued)

Name	Value	d_un	Executable	Shared Object
DT_REL	17	d_ptr	Mandatory	Optional
DT_RELSZ	18	d_val	Mandatory	Optional
DT_RELENT	19	d_val	Mandatory	Optional
DT_PLTREL	20	d_val	Optional	Optional
DT_DEBUG	21	d_ptr	Optional	Ignored
DT_TEXTREL	22	Ignored	Optional	Optional
DT_JMPREL	23	d_ptr	Optional	Optional
DT_BIND_NOW	24	Ignored	Optional	Optional
DT_INIT_ARRAY	25	d_ptr	Optional	Optional
DT_FINI_ARRAY	26	d_ptr	Optional	Optional
DT_INIT_ARRAYSZ	27	d_val	Optional	Optional
DT_FINI_ARRAYSZ	28	d_val	Optional	Optional
DT_RUNPATH	29	d_val	Optional	Optional
DT_FLAGS	30	d_val	Optional	Optional
DT_ENCODING	32	Unspecified	Unspecified	Unspecified
DT_PREINIT_ARRAY	32	d_ptr	Optional	Ignored
DT_PREINIT_ARRAYSZ	33	d_val	Optional	Ignored
DT_LOOS	0x6000000d	Unspecified	Unspecified	Unspecified
DT_SUNW_RTLDINF	0x6000000e	d_ptr	Optional	Optional
DT_HIOS	0x6ffff000	Unspecified	Unspecified	Unspecified
DT_VALRNGLO	0x6ffffd00	Unspecified	Unspecified	Unspecified
DT_CHECKSUM	0x6ffffdf8	d_val	Optional	Optional
DT_PLTPADSZ	0x6ffffdf9	d_val	Optional	Optional
DT_MOVEENT	0x6ffffdfa	d_val	Optional	Optional
DT_MOVESZ	0x6ffffdfb	d_val	Optional	Optional
DT_FEATURE_1	0x6ffffdfc	d_val	Optional	Optional
DT_POSFLAG_1	0x6ffffdfd	d_val	Optional	Optional
DT_SYMINSZ	0x6ffffdfe	d_val	Optional	Optional

TABLE 7-42 ELF Dynamic Array Tags (Continued)

Name	Value	d_un	Executable	Shared Object
DT_SYMINENT	0x6ffffdff	d_val	Optional	Optional
DT_VALRNGHI	0x6ffffdff	Unspecified	Unspecified	Unspecified
DT_ADDRNGLO	0x6ffffe00	Unspecified	Unspecified	Unspecified
DT_CONFIG	0x6ffffefa	d_ptr	Optional	Optional
DT_DEPAUDIT	0x6ffffefb	d_ptr	Optional	Optional
DT_AUDIT	0x6ffffefc	d_ptr	Optional	Optional
DT_PLTPAD	0x6ffffefd	d_ptr	Optional	Optional
DT_MOVETAB	0x6ffffefe	d_ptr	Optional	Optional
DT_SYMINFO	0x6ffffeff	d_ptr	Optional	Optional
DT_ADDRNGHI	0x6ffffeff	Unspecified	Unspecified	Unspecified
DT_RELACOUNT	0x6ffffff9	d_val	Optional	Optional
DT_RELCOUNT	0x6ffffffa	d_val	Optional	Optional
DT_FLAGS_1	0x6ffffffb	d_val	Optional	Optional
DT_VERDEF	0x6ffffffc	d_ptr	Optional	Optional
DT_VERDEFNUM	0x6ffffffd	d_val	Optional	Optional
DT_VERNEED	0x6ffffffe	d_ptr	Optional	Optional
DT_VERNEEDNUM	0x6fffffff	d_val	Optional	Optional
DT_LOPROC	0x70000000	Unspecified	Unspecified	Unspecified
DT_SPARC_REGISTER	0x70000001	d_val	Optional	Optional
DT_AUXILIARY	0x7fffffff	d_val	Unspecified	Optional
DT_USED	0x7fffffff	d_val	Optional	Optional
DT_FILTER	0x7fffffff	d_val	Unspecified	Optional
DT_HIPROC	0x7fffffff	Unspecified	Unspecified	Unspecified

DT_NULL

Marks the end of the `__DYNAMIC` array.

DT_NEEDED

The `DT_STRTAB` string table offset of a null-terminated string, giving the name of a needed dependency. The dynamic array can contain multiple entries of this type. The relative order of these entries is significant, though their relation to entries of other types is not. See “Shared Object Dependencies” on page 64.

DT_PLTRELSZ

The total size, in bytes, of the relocation entries associated with the procedure linkage table. See “Procedure Linkage Table (Processor-Specific)” on page 261.

DT_PLTGOT

An address associated with the procedure linkage table or the global offset table. See “Procedure Linkage Table (Processor-Specific)” on page 261 and “Global Offset Table (Processor-Specific)” on page 260.

DT_HASH

The address of the symbol hash table. This table refers to the symbol table indicated by the DT_SYMTAB element. See “Hash Table Section” on page 205.

DT_STRTAB

The address of the string table. Symbol names, dependency names, and other strings required by the runtime linker reside in this table. See “String Table Section” on page 221.

DT_SYMTAB

The address of the symbol table. See “Symbol Table Section” on page 222.

DT_RELA

The address of a relocation table. See “Relocation Sections” on page 210.

An object file can have multiple relocation sections. When creating the relocation table for an executable or shared object file, the link-editor concatenates those sections to form a single table. Although the sections may remain independent in the object file, the runtime linker sees a single table. When the runtime linker creates the process image for an executable file or adds a shared object to the process image, it reads the relocation table and performs the associated actions.

This element requires the DT_RELASZ and DT_RELAENT elements also be present. When relocation is mandatory for a file, either DT_RELA or DT_REL can occur.

DT_RELASZ

The total size, in bytes, of the DT_RELA relocation table.

DT_RELAENT

The size, in bytes, of the DT_RELA relocation entry.

DT_STRSZ

The total size, in bytes, of the DT_STRTAB string table.

DT_SYMENT

The size, in bytes, of the DT_SYMTAB symbol entry.

DT_INIT

The address of an initialization function. See “Initialization and Termination Sections” on page 36.

DT_FINI

The address of a termination function. See “Initialization and Termination Sections” on page 36.

- DT_SONAME**
The DT_STRTAB string table offset of a null-terminated string, identifying the name of the shared object. See “Recording a Shared Object Name” on page 103.
- DT_RPATH**
The DT_STRTAB string table offset of a null-terminated library search path string. This element’s use has been superseded by DT_RUNPATH. See “Directories Searched by the Runtime Linker” on page 64.
- DT_SYMBOLIC**
Indicates the object contains symbolic bindings that were applied during its link-edit. This elements use has been superseded by the DF_SYMBOLIC flag. See “Using -B symbolic” on page 125.
- DT_REL**
Similar to DT_RELA, except its table has implicit addends. This element requires that the DT_RELSZ and DT_RELENT elements also be present.
- DT_RELSZ**
The total size, in bytes, of the DT_REL relocation table.
- DT_RELENT**
The size, in bytes, of the DT_REL relocation entry.
- DT_PLTREL**
Indicates the type of relocation entry to which the procedure linkage table refers, either DT_REL or DT_RELA. All relocations in a procedure linkage table must use the same relocation. See “Procedure Linkage Table (Processor-Specific)” on page 261. This element requires a DT_JMPREL element also be present.
- DT_DEBUG**
Used for debugging.
- DT_TEXTREL**
Indicates that one or more relocation entries might request modifications to a non-writable segment, and the runtime linker can prepare accordingly. This element’s use has been superseded by the DF_TEXTREL flag. See “Position-Independent Code” on page 115.
- DT_JMPREL**
The address of relocation entries associated solely with the procedure linkage table. See “Procedure Linkage Table (Processor-Specific)” on page 261. Separating these relocation entries enables the runtime linker to ignore them when the object is loaded if lazy binding is enabled. This element requires the DT_PLTRELSZ and DT_PLTREL elements also be present.
- DT_POSFLAG_1**
Various state flags which are applied to the DT_ element immediately following. See Table 7-45.
- DT_BIND_NOW**
Indicates that all relocations for this object must be processed before returning control to the program. The presence of this entry takes precedence over a directive

to use lazy binding when specified through the environment or via `dlopen(3DL)`. This element's use has been superseded by the `DF_BIND_NOW` flag. See "When Relocations Are Performed" on page 71.

`DT_INIT_ARRAY`

The address of an array of pointers to initialization functions. This element requires that a `DT_INIT_ARRAYSZ` element also be present. See "Initialization and Termination Sections" on page 36.

`DT_FINI_ARRAY`

The address of an array of pointers to termination functions. This element requires that a `DT_FINI_ARRAYSZ` element also be present. See "Initialization and Termination Sections" on page 36.

`DT_INIT_ARRAYSZ`

The total size, in bytes, of the `DT_INIT_ARRAY` array.

`DT_FINI_ARRAYSZ`

The total size, in bytes, of the `DT_FINI_ARRAY` array.

`DT_RUNPATH`

The `DT_STRTAB` string table offset of a null-terminated library search path string. See "Directories Searched by the Runtime Linker" on page 64.

`DT_FLAGS`

Flag values specific to this object. See Table 7-43.

`DT_ENCODING`

Values greater than or equal to `DT_ENCODING` and less than or equal to `DT_HIOS` follow the rules for the interpretation of the `d_un` union.

`DT_PREINIT_ARRAY`

The address of an array of pointers to pre-initialization functions. This element requires that a `DT_PREINIT_ARRAYSZ` element also be present. This array is processed only in an executable file. It is ignored if contained in a shared object. See "Initialization and Termination Sections" on page 36.

`DT_PREINIT_ARRAYSZ`

The total size, in bytes, of the `DT_PREINIT_ARRAY` array.

`DT_LOOS - DT_HIOS`

Values in this inclusive range are reserved for operating system-specific semantics. All such values follow the rules for the interpretation of the `d_un` union.

`DT_SUNW_RTLDINF`

Reserved for internal use by the runtime-linker.

`DT_SYMINFO`

The address of the symbol information table. This element requires that the `DT_SYMINENT` and `DT_SYMINSZ` elements also be present. See "Syminfo Table Section" on page 229.

`DT_SYMINENT`

The size, in bytes, of the `DT_SYMINFO` information entry.

DT_SYMINSZ
The total size, in bytes, of the DT_SYMINFO table.

DT_VERDEF
The address of the version definition table. Elements within this table contain indexes into the string table DT_STRTAB. This element requires that the DT_VERDEFNUM element also be present. See “Version Definition Section” on page 231.

DT_VERDEFNUM
The number of entries in the DT_VERDEF table.

DT_VERNEED
The address of the version dependency table. Elements within this table contain indexes into the string table DT_STRTAB. This element requires that the DT_VERNEEDNUM element also be present. See “Version Dependency Section” on page 234.

DT_VERNEEDNUM
The number of entries in the DT_VERNEEDNUM table.

DT_RELACOUNT
Indicates that all Elf32_Rela (or Elf64_Rela) RELATIVE relocations have been concatenated together, and specifies the RELATIVE relocation count. See “Combined Relocation Sections” on page 121.

DT_RELCOUNT
Indicates that all Elf32_Rel RELATIVE relocations have been concatenated together, and specifies the RELATIVE relocation count. See “Combined Relocation Sections” on page 121.

DT_AUXILIARY
The DT_STRTAB string table offset of a null-terminated string that names one or more auxiliary filtees. See “Generating Auxiliary Filters” on page 110.

DT_FILTER
The DT_STRTAB string table offset of a null-terminated string that names one or more standard filtees. See “Generating Standard Filters” on page 108.

DT_CHECKSUM
A simple checksum of selected sections of the object. See `gelf_checksum(3ELF)`.

DT_MOVEENT
The size, in bytes, of the DT_MOVE TAB move entries.

DT_MOVESZ
The total size, in bytes, of the DT_MOVE TAB table.

DT_MOVE TAB
The address of a move table. This element requires that the DT_MOVE ENT and DT_MOVE SZ elements also be present. See “Move Section” on page 206.

DT_CONFIG

The DT_STRTAB string table offset of a null-terminated string defining a configuration file. The configuration file is only meaningful in an executable, and is typically unique to this object. See “Configuring the Default Search Paths” on page 67.

DT_DEPAUDIT

The DT_STRTAB string table offset of a null-terminated string defining one or more audit libraries. See “Runtime Linker Auditing Interface” on page 155.

DT_AUDIT

The DT_STRTAB string table offset of a null-terminated string defining one or more audit libraries. See “Runtime Linker Auditing Interface” on page 155.

DT_FLAGS_1

Flag values specific to this object. See Table 7-44.

DT_FEATURE_1

Feature values specific to this object. See “Feature Checking” on page 92.

DT_VALRNGLO - DT_VALRNGHI

Values in this inclusive range use the `d_un.d_val` field of the dynamic structure.

DT_ADDRNGLO - DT_ADDRNGHI

Values in this inclusive range use the `d_un.d_ptr` field of the dynamic structure. If any adjustment is made to the ELF object after it has been built, these entries must be updated accordingly.

DT_SPARC_REGISTER

The index of an `STT_SPARC_REGISTER` symbol within the `DT_SYMTAB` symbol table. There is one entry for every `STT_SPARC_REGISTER` symbol in the symbol table. See “Register Symbols” on page 229.

DT_LOPROC - DT_HIPROC

Values in this inclusive range are reserved for processor-specific semantics.

Except for the `DT_NULL` element at the end of the dynamic array and the relative order of `DT_NEEDED` and `DT_POSFLAG_1` elements, entries can appear in any order. Tag values not appearing in the table are reserved.

TABLE 7-43 ELF Dynamic Flags, `DT_FLAGS`

Name	Value	Meaning
<code>DF_ORIGIN</code>	0x1	\$ORIGIN processing required
<code>DF_SYMBOLIC</code>	0x2	Symbolic symbol resolution required
<code>DF_TEXTREL</code>	0x4	Text relocations exist
<code>DF_BIND_NOW</code>	0x8	Non-lazy binding required
<code>DF_STATIC_TLS</code>	0x10	Object uses static thread-local storage scheme

- DF_ORIGIN**
Indicates that the object requires \$ORIGIN processing. See “Locating Associated Dependencies” on page 323.
- DF_SYMBOLIC**
Indicates that the object contains symbolic bindings that were applied during its link-edit. See “Using -B symbolic” on page 125.
- DF_TEXTREL**
Indicates that one or more relocation entries might request modifications to a non-writable segment, and the runtime linker can prepare accordingly. See “Position-Independent Code” on page 115.
- DF_BIND_NOW**
Indicates that all relocations for this object must be processed before returning control to the program. The presence of this entry takes precedence over a directive to use lazy binding when specified through the environment or via dlopen(3DL). See “When Relocations Are Performed” on page 71.
- DF_STATIC_TLS**
Indicates that the object contains code using a static thread-local storage scheme. Static thread-local storage can not be used in objects that are dynamically loaded, either using dlopen(3DL), or using lazy loading. Because of this restriction, the link-editor does not support the creation of a shared object that requires static thread-local storage.

TABLE 7-44 ELF Dynamic Flags, DT_FLAGS_1

Name	Value	Meaning
DF_1_NOW	0x1	Perform complete relocation processing.
DF_1_GLOBAL	0x2	Unused.
DF_1_GROUP	0x4	Indicate object is a member of a group.
DF_1_NODELETE	0x8	Object cannot be deleted from a process.
DF_1_LOADFLTR	0x10	Ensure immediate loading of filtees.
DF_1_INITFIRST	0x20	Objects’ initialization occurs first.
DF_1_NOOPEN	0x40	Object can not be used with dlopen(3DL).
DF_1_ORIGIN	0x80	\$ORIGIN processing required.
DF_1_DIRECT	0x100	Direct bindings enabled.
DF_1_INTERPOSE	0x400	Object is an interposer.
DF_1_NODEFLIB	0x800	Ignore default library search path.
DF_1_NODUMP	0x1000	Object cannot be dumped with dldump(3DL).

TABLE 7-44 ELF Dynamic Flags, DT_FLAGS_1 (Continued)

Name	Value	Meaning
DF_1_CONFALT	0x2000	Object is a configuration alternative.
DF_1_ENDFILTEE	0x4000	Filtee terminates filter's search.
DF_1_DISPRELDNE	0x8000	Displacement relocation done.
DF_1_DISPRELPND	0x10000	Displacement relocation pending.
DF_1_NODIRECT	0x20000	Object contains non-direct bindings.

DF_1_NOW

Indicates that all relocations for this object must be processed before returning control to the program. The presence of this flag takes precedence over a directive to use lazy binding when specified through the environment or via `dlopen(3DL)`. See "When Relocations Are Performed" on page 71.

DF_1_GROUP

Indicates that the object is a member of a group. This flag is recorded in the object using the link-editor's `-B group` option. See "Object Hierarchies" on page 88.

DF_1_NODELETE

Indicates that the object cannot be deleted from a process. If the object is loaded in a process, either directly or as a dependency, with `dlopen(3DL)`, it cannot be unloaded with `dldclose(3DL)`. This flag is recorded in the object using the link-editor's `-z nodelete` option.

DF_1_LOADFLTR

Meaningful only for filters. Indicates that all associated filtees be processed immediately. This flag is recorded in the object using the link-editor's `-z loadfltr` option. See "Filtee Processing" on page 111.

DF_1_INITFIRST

Indicates that this object's initialization section be run before any other objects loaded with it. This flag is intended for specialized system libraries only, and is recorded in the object using the link-editor's `-z initfirst` option.

DF_1_NOOPEN

Indicates that the object cannot be added to a running process with `dlopen(3DL)`. This flag is recorded in the object using the link-editor's `-z nodlopen` option.

DF_1_ORIGIN

Indicates that the object requires `$ORIGIN` processing. See "Locating Associated Dependencies" on page 323.

DF_1_DIRECT

Indicates that the object should use direct binding information. See "Direct Binding" on page 69.

DF_1_INTERPOSE

Indicates that the object's symbol table is to interpose before all symbols except the primary load object, which is typically the executable. This flag is recorded with the link-editor's `-z interpose` option. See "Direct Binding" on page 69.

DF_1_NODEFLIB

Indicates that the search for dependencies of this object ignores any default library search paths. This flag is recorded in the object using the link-editor's `-z nodefaultlib` option. See "Directories Searched by the Runtime Linker" on page 35.

DF_1_NODUMP

Indicates that this object is not dumped by `dldump(3DL)`. Candidates for this option include objects with no relocations that might get included when generating alternative objects using `crle(1)`. This flag is recorded in the object using the link-editor's `-z nodump` option.

DF_1_CONFALT

Identifies this object as a configuration alternative object generated by `crle(1)`. This flag triggers the runtime linker to search for a configuration file `$(ORIGIN)/ld.config.app-name`.

DF_1_ENDFILTEE

Meaningful only for filterees. Terminates a filter's search for any further filterees. This flag is recorded in the object using the link-editor's `-z endfiltee` option. See "Reducing Filtee Searches" on page 322.

DF_1_DISPRELDNE

Indicates that this object has displacement relocations applied. The displacement relocation records no longer exist within the object as they were discarded once the relocation was applied. See "Displacement Relocations" on page 58.

DF_1_DISPRELPND

Indicates that this object has displacement relocations pending. The displacement relocations exist within the object so they can be completed at runtime. See "Displacement Relocations" on page 58.

DF_1_NODIRECT

Indicates that this object contains symbols that can not be directly bound to. See "Defining Additional Symbols" on page 46.

TABLE 7-45 ELF Dynamic Position Flags, `DT_POSFLAG_1`

Name	Value	Meaning
<code>DF_P1_LAZYLOAD</code>	<code>0x1</code>	Identify lazy loaded dependency.
<code>DF_P1_GROUPEPERM</code>	<code>0x2</code>	Identify group dependency.

DF_P1_LAZYLOAD

Identifies the following DT_NEEDED entry as an object to be lazy loaded. This flag is recorded in the object using the link-editor's `-z lazyload` option. See "Lazy Loading of Dynamic Dependencies" on page 74.

DF_P1_GROUPPERM

Identifies the following DT_NEEDED entry as an object to be loaded as a group. This flag is recorded in the object using the link-editor's `-z groupperm` option. See "Isolating a Group" on page 88.

TABLE 7-46 ELF Dynamic Feature Flags, DT_FEATURE_1

Name	Value	Meaning
DTF_1_PARINIT	0x1	Partial initialization is required.
DTF_1_CONFEXP	0x2	A Configuration file is expected.

DTF_1_PARINIT

Indicates that the object requires partial initialization. See "Move Section" on page 206.

DTF_1_CONFEXP

Identifies this object as a configuration alternative object generated by `crle(1)`. This flag triggers the runtime linker to search for a configuration file `$_ORIGIN/ld.config.app-name`. This flag has the same affect as `DF_1_CONFALT`.

Global Offset Table (Processor-Specific)

Position-independent code cannot, in general, contain absolute virtual addresses. Global offset tables hold absolute addresses in private data. Addresses are therefore available without compromising the position-independence and shareability of a program's text. A program references its global offset table using position-independent addressing and extracts absolute values. This technique redirects position-independent references to absolute locations.

Initially, the global offset table holds information as required by its relocation entries. After the system creates memory segments for a loadable object file, the runtime linker processes the relocation entries, some of which will be type `R_SPARC_GLOB_DAT` (for SPARC), or `R_386_GLOB_DAT` (for x86), referring to the global offset table.

The runtime linker determines the associated symbol values, calculates their absolute addresses, and sets the appropriate memory table entries to the proper values. Although the absolute addresses are unknown when the link-editor creates an object file, the runtime linker knows the addresses of all memory segments and can thus calculate the absolute addresses of the symbols contained therein.

If a program requires direct access to the absolute address of a symbol, that symbol will have a global offset table entry. Because the executable file and shared objects have separate global offset tables, a symbol's address can appear in several tables. The runtime linker processes all the global offset table relocations before giving control to any code in the process image. This processing ensures that absolute addresses are available during execution.

The table's entry zero is reserved to hold the address of the dynamic structure, referenced with the symbol `_DYNAMIC`. This symbol enables a program, such as the runtime linker, to find its own dynamic structure without having yet processed its relocation entries. This method is especially important for the runtime linker, because it must initialize itself without relying on other programs to relocate its memory image.

The system can choose different memory segment addresses for the same shared object in different programs. It can even choose different library addresses for different executions of the same program. Nonetheless, memory segments do not change addresses once the process image is established. As long as a process exists, its memory segments reside at fixed virtual addresses.

A global offset table's format and interpretation are processor-specific. For SPARC and x86 processors, the symbol `_GLOBAL_OFFSET_TABLE_` can be used to access the table. This symbol can reside in the middle of the `.got` section, allowing both negative and nonnegative subscripts into the array of addresses. The symbol type is an array of `Elf32_Addr` for 32-bit code, and an array of `Elf64_Addr` for 64-bit code:

```
extern Elf32_Addr _GLOBAL_OFFSET_TABLE_[];
extern Elf64_Addr _GLOBAL_OFFSET_TABLE_[];
```

Procedure Linkage Table (Processor-Specific)

The global offset table converts position-independent address calculations to absolute locations. Similarly the procedure linkage table converts position-independent function calls to absolute locations. The link-editor cannot resolve execution transfers such as function calls from one executable or shared object to another. So, the link-editor arranges to have the program transfer control to entries in the procedure linkage table. The runtime linker thus redirects the entries without compromising the position-independence and shareability of the program's text. Executable files and shared object files have separate procedure linkage tables.

SPARC: 32-bit Procedure Linkage Table

For 32-bit SPARC dynamic objects, the procedure linkage table resides in private data. The runtime linker determines the absolute addresses of the destinations and modifies the procedure linkage table's memory image accordingly.

The first four procedure linkage table entries are reserved. The original contents of these entries are unspecified, despite the example shown in Table 7-47. Each entry in the table occupies 3 words (12 bytes), and the last table entry is followed by a `nop` instruction.

A relocation table is associated with the procedure linkage table. The `DT_JMP_REL` entry in the `_DYNAMIC` array gives the location of the first relocation entry. The relocation table has one entry, in the same sequence, for each non-reserved procedure linkage table entry. The relocation type of each of these entries is `R_SPARC_JMP_SLOT`. The relocation offset specifies the address of the first byte of the associated procedure linkage table entry. The symbol table index refers to the appropriate symbol.

To illustrate procedure linkage tables, Table 7-47 shows four entries: two of the four initial reserved entries, the third is a call to `name101`, and the fourth entry is a call to `name102`. The example assumes that the entry for `name102` is the table's last entry and shows the following `nop` instruction. The left column shows the instructions from the object file before dynamic linking. The right column demonstrates a possible way the runtime linker might fix the procedure linkage table entries.

TABLE 7-47 SPARC: Procedure Linkage Table Example

<i>Object File</i>	<i>Memory Segment</i>
<code>.PLT0:</code>	<code>.PLT0:</code>
<code>unimp</code>	<code>save %sp, -64, %sp</code>
<code>unimp</code>	<code>call runtime_linker</code>
<code>unimp</code>	<code>nop</code>
<code>.PLT1:</code>	<code>.PLT1:</code>
<code>unimp</code>	<code>.word identification</code>
<code>unimp</code>	<code>unimp</code>
<code>unimp</code>	<code>unimp</code>
<code>.PLT101:</code>	<code>.PLT101:</code>
<code>sethi (.-.PLT0), %g1</code>	<code>nop</code>
<code>ba,a .PLT0</code>	<code>ba,a name101</code>
<code>nop</code>	<code>nop</code>
<code>.PLT102:</code>	<code>.PLT102:</code>
<code>sethi (.-.PLT0), %g1</code>	<code>sethi (.-.PLT0), %g1</code>
<code>ba,a .PLT0</code>	<code>sethi %hi(name102), %g1</code>
<code>nop</code>	<code>jmp1 %g1+%lo(name102), %g0</code>
<code>nop</code>	<code>nop</code>

Following the steps below, the runtime linker and program jointly resolve the symbolic references through the procedure linkage table. Again, the steps described below are for explanation only. The precise execution-time behavior of the runtime linker is not specified.

1. When first creating the memory image of the program, the runtime linker changes the initial procedure linkage table entries, making them transfer control to one of the runtime linker's own routines. The runtime linker also stores a word of identification information in the second entry. When the runtime linker receives control, it can examine this word to find which object called it.
2. All other procedure linkage table entries initially transfer to the first entry, letting the runtime linker to gain control at the first execution of each table entry. For example, the program calls `name101`, which transfers control to the label `.PLT101`.
3. The `sethi` instruction computes the distance between the current and the initial procedure linkage table entries, `.PLT101` and `.PLT0`, respectively. This value occupies the most significant 22 bits of the `%g1` register.
4. Next, the `ba, a` instruction jumps to `.PLT0`, establishing a stack frame and calls the runtime linker.
5. With the identification value, the runtime linker gets its data structures for the object, including the relocation table.
6. By shifting the `%g1` value and dividing by the size of the procedure linkage table entries, the runtime linker calculates the index of the relocation entry for `name101`. Relocation entry `101` has type `R_SPARC_JMP_SLOT`, its offset specifies the address of `.PLT101`, and its symbol table index refers to `name101`. Thus, the runtime linker gets the symbol's real value, unwinds the stack, modifies the procedure linkage table entry, and transfers control to the desired destination.

The runtime linker does not have to create the instruction sequences under the memory segment column. If it does, some points deserve more explanation.

- To make the code re-entrant, the procedure linkage table's instructions are changed in a particular sequence. If the runtime linker is fixing a function's procedure linkage table entry and a signal arrives, the signal handling code must be able to call the original function with predictable and correct results.
- The runtime linker changes three words to convert an entry. The runtime linker can update only a single word atomically with regard to instruction execution. Therefore, re-entrancy is achieved by updating each word in reverse order. If a re-entrant function call occurs just prior to the last patch, the runtime linker gains control a second time. Although both invocations of the runtime linker modify the same procedure linkage table entry, their changes do not interfere with each other.
- The first `sethi` instruction of a procedure linkage table entry can fill the delay slot of the previous entry's `jmp1` instruction. Although the `sethi` changes the value of the `%g1` register, the previous contents can be safely discarded.
- After conversion, the last procedure linkage table entry, `.PLT102`, needs a delay instruction for its `jmp1`. The required, trailing `nop` fills this delay slot.

Note – The different instruction sequences shown for `.PLT101`, and `.PLT102` demonstrate how the update may be optimized for the associated destination.

The `LD_BIND_NOW` environment variable changes dynamic linking behavior. If its value is non-null, the runtime linker processes `R_SPARC_JMP_SLOT` relocation entries (procedure linkage table entries) before transferring control to the program.

SPARC: 64-bit Procedure Linkage Table

For 64-bit SPARC dynamic objects, the procedure linkage table resides in private data. The runtime linker determines the absolute addresses of the destinations and modifies the procedure linkage table's memory image accordingly.

The first four procedure linkage table entries are reserved. The original contents of these entries are unspecified, despite the example shown in Table 7-48. Each of the first 32,768 entries in the table occupies 8 words (32 bytes), and must be aligned on a 32-byte boundary. The table as a whole must be aligned on a 256-byte boundary. If more than 32,768 entries are required, the remaining entries consist of 6 words (24 bytes) and 1 pointer (8 bytes). The instructions are collected together in blocks of 160 entries followed by 160 pointers. The last group of entries and pointers may contain less than 160 items. No padding is required.

Note – The numbers 32,768 and 160 are based on the limits of branch and load displacements respectively with the second rounded down to make the divisions between code and data fall on 256-byte boundaries so as to improve cache performance.

A relocation table is associated with the procedure linkage table. The `DT_JMP_REL` entry in the `_DYNAMIC` array gives the location of the first relocation entry. The relocation table has one entry, in the same sequence, for each non-reserved procedure linkage table entry. The relocation type of each of these entries is `R_SPARC_JMP_SLOT`. For the first 32,767 slots, the relocation offset specifies the address of the first byte of the associated procedure linkage table entry, the addend field is zero. The symbol table index refers to the appropriate symbol. For slots 32,768 and beyond, the relocation offset specifies the address of the first byte of the associated pointer. The addend field is the unrelocated value - (`.PLTN` + 4). The symbol table index refers to the appropriate symbol.

To illustrate procedure linkage tables, Table 7-48 shows several entries. The first three show initial reserved entries. The following three show examples of the initial 32,768 entries together with possible resolved forms that might apply if the target address was +/- 2 Gbytes of the entry, within the lower 4 Gbytes of the address space, or anywhere respectively. The final two show examples of later entries, which consist of instruction and pointer pairs. The left column shows the instructions from the object file before dynamic linking. The right column demonstrates a possible way the runtime linker might fix the procedure linkage table entries.

TABLE 7-48 64-bit SPARC: Procedure Linkage Table Example

<i>Object File</i>	<i>Memory Segment</i>
.PLT0: unimp unimp unimp unimp unimp unimp unimp unimp	.PLT0: save %sp, -176, %sp sethi %hh(runtime_linker_0), %l0 sethi %lm(runtime_linker_0), %l1 or %l0, %hm(runtime_linker_0), %l0 sllx %l0, 32, %l0 or %l0, %l1, %l0 jmp1 %l0+%l0(runtime_linker_0), %o1 mov %g1, %o0
.PLT1: unimp unimp unimp unimp unimp unimp unimp unimp	.PLT1: save %sp, -176, %sp sethi %hh(runtime_linker_1), %l0 sethi %lm(runtime_linker_1), %l1 or %l0, %hm(runtime_linker_1), %l0 sllx %l0, 32, %l0 or %l0, %l1, %l0 jmp1 %l0+%l0(runtime_linker_0), %o1 mov %g1, %o0
.PLT2: unimp	.PLT2: .xword identification
.PLT101: sethi (.-.PLT0), %g1 ba,a %xcc, .PLT1 nop nop nop; nop nop; nop	.PLT101: nop mov %o7, %g1 call name101 mov %g1, %o7 nop; nop; nop
.PLT102: sethi (.-.PLT0), %g1 ba,a %xcc, .PLT1 nop nop nop; nop nop; nop	.PLT102: nop sethi %hi(name102), %g1 jmp1 %g1+%l0(name102), %g0 nop nop; nop; nop
.PLT103: sethi (.-.PLT0), %g1 ba,a %xcc, .PLT1 nop nop nop nop nop nop	.PLT103: nop sethi %hh(name103), %g1 sethi %lm(name103), %g5 or %hm(name103), %g1 sllx %g1, 32, %g1 or %g1, %g5, %g5 jmp1 %g5+%l0(name103), %g0 nop

TABLE 7-48 64-bit SPARC: Procedure Linkage Table Example (Continued)

Object File	Memory Segment
.PLT32768:	.PLT32768:
mov %o7, %g5	<unchanged>
call .+8	<unchanged>
nop	<unchanged>
ldx [%o7+.PLTP32768 - (.PLT32768+4)], %g1	<unchanged>
jmp1 %o7+%g1, %g1	<unchanged>
mov %g5, %o7	<unchanged>
...	...
.PLT32927:	.PLT32927:
mov %o7, %g5	<unchanged>
call .+8	<unchanged>
nop	<unchanged>
ldx [%o7+.PLTP32927 - (.PLT32927+4)], %g1	<unchanged>
jmp1 %o7+%g1, %g1	<unchanged>
mov %g5, %o7	<unchanged>
.PLTP32768	.PLTP32768
.xword .PLT0 - (.PLT32768+4)	.xword name32768 - (.PLT32768+4)
...	...
.PLTP32927	.PLTP32927
.xword .PLT0 - (.PLT32927+4)	.xword name32927 - (.PLT32927+4)

Following the steps below, the runtime linker and program jointly resolve the symbolic references through the procedure linkage table. Again, the steps described below are for explanation only. The precise execution-time behavior of the runtime linker is not specified.

1. When first creating the memory image of the program, the runtime linker changes the initial procedure linkage table entries, making them transfer control to one of the runtime linker's own routines. The runtime linker also stores an extended word of identification information in the third entry. When the runtime linker receives control, it can examine this extended word to find which object called it.
2. All other procedure linkage table entries initially transfer to the first or second entry. Those entries establish a stack frame and call the runtime linker.
3. With the identification value, the runtime linker gets its data structures for the object, including the relocation table.
4. The runtime linker computes the index of the relocation entry for the table slot.

5. With the index information, the runtime linker gets the symbol's real value, unwinds the stack, modifies the procedure linkage table entry, and transfers control to the desired destination.

The runtime linker does not have to create the instruction sequences under the memory segment column, it might. If it does, some points deserve more explanation.

- To make the code re-entrant, the procedure linkage table's instructions are changed in a particular sequence. If the runtime linker is fixing a function's procedure linkage table entry and a signal arrives, the signal handling code must be able to call the original function with predictable and correct results.
- The runtime linker may change up to eight words to convert an entry. The runtime linker can update only a single word atomically with regard to instruction execution. Therefore, re-entrancy is achieved by first overwriting the `nop` instructions with their replacement instructions, and then patching the `ba`, `a`, and the `sethi` if using a 64-bit store. If a re-entrant function call occurs just prior to the last patch, the runtime linker gains control a second time. Although both invocations of the runtime linker modify the same procedure linkage table entry, their changes do not interfere with each other.
- If the initial `sethi` instruction is changed, it can only be replaced by a `nop`.

Changing the pointer as done for the second form of entry is done using a single atomic 64-bit store.

Note – The different instruction sequences shown for `.PLT101`, `.PLT102`, and `.PLT103` demonstrate how the update may be optimized for the associated destination.

The `LD_BIND_NOW` environment variable changes dynamic linking behavior. If its value is non-null, the runtime linker processes `R_SPARC_JMP_SLOT` relocation entries (procedure linkage table entries) before transferring control to the program.

x86: 32-bit Procedure Linkage Table

For 32-bit x86 dynamic objects, the procedure linkage table resides in shared text but uses addresses in the private global offset table. The runtime linker determines the absolute addresses of the destinations and modifies the global offset table's memory image accordingly. The runtime linker thus redirects the entries without compromising the position-independence and shareability of the program's text. Executable files and shared object files have separate procedure linkage tables.

TABLE 7-49 x86: Absolute Procedure Linkage Table Example

```
.PLT0 :
    pushl    got_plus_4
    jmp     *got_plus_8
    nop;    nop
    nop;    nop
.PLT1 :
    jmp     *name1_in_GOT
    pushl   $offset
    jmp     .PLT0@PC
.PLT2 :
    jmp     *name2_in_GOT
    pushl   $offset
    jmp     .PLT0@PC
```

TABLE 7-50 x86: Position-Independent Procedure Linkage Table Example

```
.PLT0 :
    pushl    4(%ebx)
    jmp     *8(%ebx)
    nop;    nop
    nop;    nop
.PLT1 :
    jmp     *name1@GOT(%ebx)
    pushl   $offset
    jmp     .PLT0@PC
.PLT2 :
    jmp     *name2@GOT(%ebx)
    pushl   $offset
    jmp     .PLT0@PC
```

Note – As the preceding examples show, the procedure linkage table instructions use different operand addressing modes for absolute code and for position-independent code. Nonetheless, their interfaces to the runtime linker are the same.

Following the steps below, the runtime linker and program cooperate to resolve the symbolic references through the procedure linkage table and the global offset table.

1. When first creating the memory image of the program, the runtime linker sets the second and third entries in the global offset table to special values. The steps below explain these values.
2. If the procedure linkage table is position-independent, the address of the global offset table must be in `%ebx`. Each shared object file in the process image has its own procedure linkage table, and control transfers to a procedure linkage table entry only from within the same object file. So, the calling function must set the global offset table base register before it calls the procedure linkage table entry.

3. For example, the program calls `name1`, which transfers control to the label `.PLT1`.
4. The first instruction jumps to the address in the global offset table entry for `name1`. Initially, the global offset table holds the address of the following `pushl` instruction, not the real address of `name1`.
5. The program pushes a relocation offset (`offset`) on the stack. The relocation offset is a 32-bit, nonnegative byte offset into the relocation table. The designated relocation entry has the type `R_386_JMP_SLOT`, and its offset specifies the global offset table entry used in the previous `jmp` instruction. The relocation entry also contains a symbol table index, which the runtime linker uses to get the referenced symbol, `name1`.
6. After pushing the relocation offset, the program jumps to `.PLT0`, the first entry in the procedure linkage table. The `pushl` instruction pushes the value of the second global offset table entry (`got_plus_4` or `4(%ebx)`) on the stack, giving the runtime linker one word of identifying information. The program then jumps to the address in the third global offset table entry (`got_plus_8` or `8(%ebx)`), to jump to the runtime linker.
7. The runtime linker unwinds the stack, checks the designated relocation entry, gets the symbol's value, stores the actual address of `name1` in its global offset entry table, and jumps to the destination.
8. Subsequent executions of the procedure linkage table entry transfer directly to `name1`, without calling the runtime linker again. The `jmp` instruction at `.PLT1` jumps to `name1` instead of falling through to the `pushl` instruction.

The `LD_BIND_NOW` environment variable changes dynamic linking behavior. If its value is non-null, the runtime linker processes `R_386_JMP_SLOT` relocation entries (procedure linkage table entries) before transferring control to the program.

Thread-Local Storage

The compilation environment supports the declaration of thread-local data. This data is sometime referred to as thread-specific, or thread-private data, but more typically by the acronym TLS. By declaring variables to be thread-local, the compiler automatically arranges for these variables to be allocated on a per-thread basis.

The built-in support for this feature serves three purposes:

- It provides a foundation upon which the POSIX interfaces for allocating thread specific data are built.
- It offers a more convenient and more efficient mechanism for direct use by applications and libraries.
- It allows compilers to allocate TLS as necessary when performing loop-parallelizing optimizations.

C/C++ Programming Interface

Variables are declared thread-local using the `__thread` keyword, as in the following examples:

```
__thread int i;  
__thread char *p;  
__thread struct state s;
```

During loop optimizations, the compiler may choose to create thread-local temporaries as needed.

Applicability

The `__thread` keyword may be applied to any global, file-scoped static, or function-scoped static variable. It has no effect on automatic variables, which are always thread-local.

Initialization

In C++, a thread-local variable may not be initialized if the initialization requires a static constructor. Otherwise, a thread-local variable may be initialized to any value that would be legal for an ordinary static variable.

No variable, thread-local or otherwise, may be statically initialized to the address of a thread-local variable.

Binding

Thread-local variables may be declared and referenced externally, and they are subject to the same interposition rules as normal symbols.

Dynamic loading restrictions

A shared library can be dynamically loaded during process startup, or after process startup via lazy loading, filters, or `dlopen(3DL)`. A shared library containing a reference to a thread-local variable, may be loaded post-startup if every translation unit containing the reference is compiled with a dynamic TLS model.

Static TLS models generates faster code. However, code compiled to use this model cannot reference thread-local variables in post-startup dynamically loaded libraries. A dynamic TLS model is able to reference all TLS. These models are described in “Thread-Local Storage Access Models” on page 277.

Address-of operator

The address-of operator, `&`, can be applied to a thread-local variable. This operator is evaluated at runtime, and returns the address of the variable within the current thread. The address obtained by this operator may be used freely by any thread in the process as long as the thread that evaluated the address remains in existence. When a thread terminates, any pointers to thread-local variables in that thread become invalid.

When `dlsym(3DL)` is used to obtain the address of a thread-local variable, the address returned is the address of the instance of that variable in the thread that called `dlsym()`.

Thread-Local Storage Section

Separate copies of thread-local data, allocated at compile-time, must be associated with individual threads of execution. To provide this data, TLS sections are used to specify the size and initial contents.

The compilation environment allocates TLS in sections identified with the `SHF_TLS` flag. These sections provide initialized and uninitialized TLS based on how the storage is declared:

- An initialized thread-local variable is allocated in a `.tdata`, or `.tdata1` section. This initialization may require relocation.

- An uninitialized thread-local variable is defined as a `COMMON` symbol. The resulting allocation is made in a `.tbss` section.

The uninitialized section is allocated immediately following any initialized sections, subject to padding for proper alignment. Together, the combined sections form a TLS template that is used to allocate TLS whenever a new thread is created.

The initialized portion of this template is called the TLS initialization image. All relocations generated as a result of initialized thread-local variables are applied to this template. These relocated values are then used when a new thread requires the initial values.

TLS symbols have the symbol type `STT_TLS`. These symbols are assigned offsets relative to the beginning of the TLS template. The actual virtual address associated with these symbols is irrelevant. The address refers only to the template, and not to the per-thread copy of each data item.

In dynamic executables and shared objects, the `st_value` field of a `STT_TLS` symbol contains the assigned offset for defined symbols, or zero for undefined symbols.

Several relocations are defined to support access to TLS. See “SPARC: Thread-Local Storage Relocation Types” on page 283 and “x86: Thread-Local Storage Relocation Types” on page 290. Symbols of type `STT_TLS` are only referenced by TLS relocations. TLS relocations only reference symbols of type `STT_TLS`.

In dynamic executables and shared objects, a `PT_TLS` program entry describes a TLS template, and has the following members:

TABLE 8-1 ELF `PT_TLS` Program Header Entry

Member	Value
<code>p_offset</code>	File offset of the TLS initialization image
<code>p_vaddr</code>	Virtual memory address of the TLS initialization image
<code>p_paddr</code>	Reserved
<code>p_filesz</code>	Size of the TLS initialization image
<code>p_memsz</code>	Total size of the TLS template
<code>p_flags</code>	<code>PF_R</code>
<code>p_align</code>	Alignment of the TLS template

Runtime Allocation of Thread-Local Storage

TLS is created at three occasions during the lifetime of a program:

- At program startup.
- When a new thread is created.
- When a thread references a TLS block for the first time after a shared library is loaded following program startup.

Thread-local data storage is layed out at runtime as illustrated in Figure 8-1.

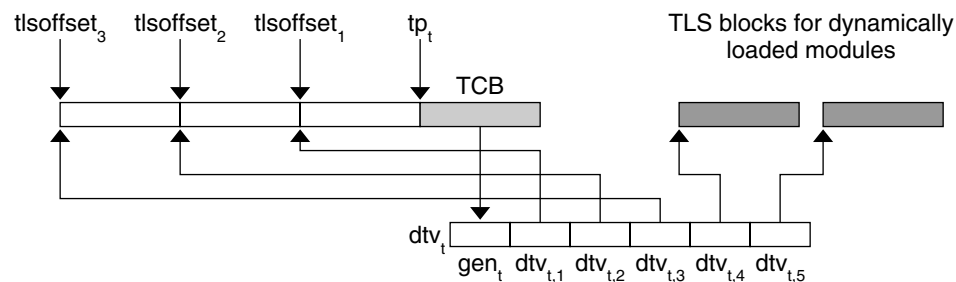


FIGURE 8-1 Runtime Storage Layout of Thread-Local Storage

Program Startup

At program startup, the runtime system creates TLS for the main thread.

First, the runtime linker logically combines the TLS templates for all loaded dynamic objects, including the dynamic executable, into a single static template. Each dynamic object's TLS template is assigned an offset within the combined template, $tloffset_m$ as follows:

- $tloffset_1 = \text{round}(tlssize_1, \text{align}_1)$
- $tloffset_{m+1} = \text{round}(tloffset_m + tlssize_{m+1}, \text{align}_{m+1})$

$tlssize_{m+1}$ and align_{m+1} are the size and alignment, respectively, for the allocation template for dynamic object m ($1 \leq m \leq M$, where M is the total number of loaded dynamic objects). The $\text{round}(\text{offset}, \text{align})$ function returns an offset rounded up to the next multiple of align . The TLS template is placed immediately preceding the thread pointer tp_t . Accesses to the TLS data are based off of subtractions from tp_t .

Next, the runtime linker computes the total startup TLS allocation size, $tlssize_s$, which is equal to $tlsoffset_m$.

The runtime linker then constructs a linked list of initialization records. Each record in this list describes the TLS initialization image for one loaded dynamic object, and contains the following fields:

- A pointer to the TLS initialization image.
- The size of the TLS initialization image.
- The $tlsoffset_m$ of the object.
- A flag indicating whether the object uses a static TLS model.

The thread library uses this information to allocate storage for the initial thread. This storage is initialized, and a dynamic TLS vector for the initial thread is created.

Thread Creation

For the initial thread, and for each new thread created, the thread library allocates a new TLS block for each loaded dynamic object. Blocks may be allocated separately, or as a single contiguous block.

Each thread t has an associated thread pointer tp_t , which points to the thread control block, TCB. The thread pointer, tp , always contains the value of tp_t for the current running thread.

The thread library then creates a vector of pointers, dtv_t , for the current thread t . The first element of each vector contains a generation number gen_t , which is used to determine when the vector needs to be extended. See “Deferred Allocation of Thread-Local Storage Blocks” on page 276.

Each remaining element in the vector, $dtv_{t,m}$, is a pointer to the block reserved for the TLS belonging to dynamic object m .

For dynamically loaded, post-startup objects, the thread library defers the allocation of TLS blocks. This allocation occurs when the first reference is made to a TLS variable within the loaded object. All references to TLS defined in a post-startup, dynamically loaded object, must use a dynamic TLS model. For blocks whose allocation has been deferred, the pointer $dtv_{t,m}$ is set to an implementation-defined special value.

Note – The runtime linker may group TLS templates for all startup objects such that they share a single element in the vector, $dtv_{t,1}$. This does not affect the offset calculations described above or the creation of the list of initialization records. For the following sections, however, the value of M , the total number of objects, start with the value of 1.

The thread library then copies the initialization images to the corresponding locations within the new block of storage.

Post-Startup Dynamic Loading

When a shared library containing TLS is loaded following process startup, the runtime linker extends the list of initialization records to include the initialization template of new library. The new object is given an index of $m = M + 1$, and the counter M is incremented by one. However, the allocation of new TLS blocks is deferred until they are actually referenced.

When a library containing TLS is unloaded, the TLS blocks used by that library are freed.

Deferred Allocation of Thread-Local Storage Blocks

In a dynamic TLS model, when a thread t needs to access a TLS block for object m , the code updates the dtv_t and performs the initial allocation of the TLS block. The thread library provides the following interface to provide for dynamic TLS allocation:

```
typedef struct {
    unsigned long ti_moduleid;
    unsigned long ti_tlsoffset;
} TLS_index;

extern void * __tls_get_addr(TLS_index * ti);      (SPARC)
extern void * ___tls_get_addr(TLS_index * ti);   (x86)
```

Note – The SPARC and x86 definitions of this function have the same function signature. However, the x86 version does not use the default x86 calling convention of passing arguments on the stack. Instead, the x86 version passes its argument via the `%eax` register which is more efficient. To denote that this alternate calling method is used, the x86 function name has three leading underscores in its name.

Both versions of `tls_get_addr()` check the per-thread generation counter, gen_t , to determine whether the vector needs to be updated. If the vector dtv_t is out of date, the routine updates the vector, possibly reallocating it to make room for more entries. The routine then checks to see if the TLS block corresponding to $dtv_{t,m}$ has been allocated. If it has not been allocated, the routine allocates and initializes the block, using the information in the list of initialization records provided by the runtime linker. The pointer $dtv_{t,m}$ is set to point to the allocated block. The routine returns a pointer to the given offset within the block.

Thread-Local Storage Access Models

Each TLS reference follows one of the following access models. These models are listed from the most general, but least optimized, to the fastest, but most restrictive.

General Dynamic (GD) - dynamic TLS

This model allows reference of all TLS variables, from either a shared object or a dynamic executable. This model also supports the deferred allocation of a TLS block when it is first referenced from a specific thread.

Local Dynamic (LD) - dynamic TLS of local symbols

This model is an optimization of the *GD* model. If the compiler determines that a variable is bound locally, or protected, within the dynamic object being built, it instructs the link-editor to statically bind the dynamic `tls_offset` and use this model. This provides a performance benefit over the *GD* model. Only one call to `tls_get_addr()` is required per function, to determine the address of `dtv0,m`. The dynamic TLS offset, bound at link-edit time, is added to the `dtv0,m` address for each reference.

Initial Executable (IE) - static TLS with assigned offsets

This model can only reference TLS variables which are available as part of the initial static TLS template. This template is composed of all TLS blocks available at process startup. In this model, the thread pointer-relative offset for a given variable *x* is stored in the global offset table entry for *x*. This model can not reference TLS variables from shared libraries loaded after initial process startup, such as via lazy loading, filters, or `dlopen(3DL)`. This model can not access TLS blocks which use deferred allocation.

Local Executable (LE) - static TLS

This model can only reference TLS variables which are part of the TLS block of the dynamic executable itself. The link-editor calculates the thread pointer-relative offsets statically, without the need for dynamic relocations, or the extra reference to the global offset table. This model can not be used to reference variables outside of the dynamic executable.

The link-editor can transition code from the more general access models to the more optimized models, if it is determined appropriate to do so. This transitioning is achievable through the use of unique TLS relocations. These relocations, not only request updates be performed, but identify which TLS access model is being used.

Knowing the TLS access model, and the type of object being created, allows the link-editor to perform translations. For example, if a relocatable object using the *GD* access model is being linked into a dynamic executable, the link-editor can transition the references using the *IE* or *LE* access models, as appropriate. The relocations required for the model are then performed.

The following diagram illustrates the different access models, and when one model can be transitioned from one to the other.




FIGURE 8-2 Thread-Local Storage Access Models and Transitions

SPARC: Thread-Local Variable Access

On SPARC, the following code sequence models are available for accessing thread-local variables.

SPARC: 32-bit and 64-bit General Dynamic (GD)

This code sequence is the most general, and can be included in both shared objects and dynamic executables. This code sequence can also reference an external TLS variable in either a shared object or dynamic executable.

TABLE 8-2 SPARC: 32-bit and 64-bit General Dynamic Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
# %l7 - initialized to GOT pointer		
0x00 sethi %hi(@dtlndx(x)), %o0	R_SPARC_TLS_GD_HI22	x
0x04 add %o0,%lo(@dtlndx(x)),%o0	R_SPARC_TLS_GD_LO22	x
0x08 add %l7, %o0, %o0	R_SPARC_TLS_GD_ADD	x
0x0c call x@TLSPLT	R_SPARC_TLS_GD_CALL	x
# %o0 - contains address of TLS variable		
Outstanding Relocations: 32-bit		Symbol
GOT[n]	R_SPARC_TLS_DTPMOD32	x
GOT[n + 1]	R_SPARC_TLS_DTPOFF32	x
Outstanding Relocations: 64-bit		Symbol
GOT[n]	R_SPARC_TLS_DTPMOD64	x
GOT[n + 1]	R_SPARC_TLS_DTPOFF64	x

The `sethi`, and `add` instructions generate `R_SPARC_TLS_GD_HI22` and `R_SPARC_TLS_GD_LO10` relocations respectively. These relocations instruct the link-editor to allocate space in the global offset table to hold a `TLS_index` structure for variable `x`. The link-editor processes this relocation by substituting the GOT-relative offset for the new GOT entry.

Since the load object index and TLS block index for `x` are not known until runtime, the link-editor places the `R_SPARC_TLS_DTPMOD32` and `R_SPARC_TLS_DTPOFF32` relocations against the GOT for processing by the runtime linker.

The `add` instruction causes the generation of the `R_SPARC_TLS_GD_ADD` relocation. This relocation is used only if the GD code sequence is changed to another sequence by the link-editor.

The `call` instruction generates the `R_SPARC_TLS_GD_CALL` relocation. This relocation instructs the link-editor to bind the call to the `__tls_get_addr()` function, and associates the `call` instruction with the GD code sequence.

Note – The `add` instruction must appear before the `call` instruction. It cannot be placed into the delay slot for the call. This is required as the code-transformations that can occur later require a known order.

The register used as the GOT-pointer for the `add` instruction tagged by the `R_SPARC_TLS_GD_ADD` relocation, must be the first register in the `add` instruction. This permits the link-editor to identify the GOT-pointer register during a code transformation.

SPARC: 32-bit and 64-bit Local Dynamic (LD)

This code sequence can be used in either a shared object or dynamic executable. This sequence is used when referencing a TLS variable bound within the same object as the reference. Because the dynamic `tlsoffset` can be bound at link-edit time, only one call to `__tls_get_addr()` is required per function call for all symbols referenced via the `LD` code sequence.

TABLE 8-3 SPARC: 32-bit and 64-bit Local Dynamic Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
# %17 - initialized to GOT pointer		
0x00 sethi %hi(@tmdx(x1)), %0	R_SPARC_TLS_LDM_HI22	x1
0x04 add %0,%lo(@tmdx(x1)),%0	R_SPARC_TLS_LDM_LO10	x1
0x08 add %17, %0, %0	R_SPARC_TLS_LDM_ADD	x1
0x0c call x@TLSPLT	R_SPARC_TLS_LDM_CALL	x1
# %0 - contains address of TLS block of current object		
0x10 sethi %hi(@dtpoff(x1)), %11	R_SPARC_TLS_LDO_HIX22	x1
0x14 xor %11, %lo(@dtpoff(x1)), %11	R_SPARC_TLS_LDO_LOX10	x1
0x18 add %0, %11, %11	R_SPARC_TLS_LDO_ADD	x1
# %11 - contains address of local TLS variable x1		
0x20 sethi %hi(@dtpoff(x2)), %12	R_SPARC_TLS_LDO_HIX22	x2
0x24 xor %12, %lo(@dtpoff(x2)), %12	R_SPARC_TLS_LDO_LOX10	x2
0x28 add %0, %12, %12	R_SPARC_TLS_LDO_ADD	x2
# %12 - contains address of local TLS variable x2		

	Outstanding Relocations: 32-bit	Symbol
GOT [n] GOT [n + 1]	R_SPARC_TLS_DTPMOD32 <none>	x1

	Outstanding Relocations: 64-bit	Symbol
GOT [n] GOT [n + 1]	R_SPARC_TLS_DTPMOD64 <none>	x1

The first `sethi` and `add` instructions generate `R_SPARC_TLS_LDM_HI22` and `R_SPARC_TLS_LDM_LO10` relocations respectively. These relocations instruct the link-editor to allocate space in the global offset table to hold a `TLS_index` structure for the current object. The link-editor processes this relocation by substituting the GOT-relative offset for the new GOT entry.

Since the load object index is not known until runtime, a `R_SPARC_TLS_DTPMOD32` relocation is created, and the `ti_tlsoffset` field of the `TLS_index` structure is zero filled.

The second `add` and the `call` instruction are tagged with the `R_SPARC_TLS_LDM_ADD` and `R_SPARC_TLS_LDM_CALL` relocations respectively.

The following `sethi` and `or` instructions generate the `R_SPARC_LDO_HIX22` and `R_SPARC_TLS_LDO_LOX10` relocations, respectively. The TLS offset for each local symbol is known at link-edit time, therefore these values are filled in directly. The `add` instruction is tagged with the `R_SPARC_TLS_LDO_ADD` relocation.

When a procedure references more than one local symbol, the compiler generates code to obtain the base address of the TLS block once. This base address is then used to calculate the address of each symbol without a separate library call.

Note – The register containing the TLS object address in the `add` instruction tagged by the `R_SPARC_TLS_LDO_ADD` must be the first register in the instruction sequence. This permits the link-editor to identify the register during a code transformation.

SPARC: 32-bit Initial Executable (IE)

This code sequence can only be used in a dynamic executable. It can reference a TLS variable defined in either the executable or any shared libraries loaded at process startup. This model can not reference TLS variables from shared libraries loaded after process startup.

TABLE 8-4 SPARC: 32-bit Initial Executable Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
# %l7 - initialized to GOT pointer		
0x00 sethi %hi(@tpoff(x)), %o0	R_SPARC_TLS_IE_HI22	x
0x04 or %o0,%lo(@tpoff(x)),%o0	R_SPARC_TLS_IE_LO10	x
0x08 ld [%l7 + %o0], %o0	R_SPARC_TLS_IE_LD	x
0x0c add %g7, %o0, %o0	R_SPARC_TLS_IE_ADD	x
# %o0 - contains address of TLS variable		
Outstanding Relocations		Symbol
GOT[n]	R_SPARC_TLS_TPOFF32	x

The `sethi` and `or` instructions generate `R_SPARC_TLS_IE_HI22` and `R_SPARC_TLS_IE_LO10` relocations, respectively. These relocations instruct the link-editor to create space in the global offset table to store the static TLS offset for symbol `x`. A `R_SPARC_TLS_TPOFF32` relocation is left outstanding against the `GOT` for the runtime linker to fill in with the negative static TLS offset for symbol `x`. The `ld` and the `add` instructions are tagged with the `R_SPARC_TLS_IE_LD` and `R_SPARC_TLS_IE_ADD` relocations respectively.

Note – The register used as the GOT-pointer for the `add` instruction tagged by the `R_SPARC_TLS_IE_ADD` relocation must be the first register in the instruction. This permits the link-editor to identify the GOT-pointer register during a code transformation.

SPARC: 64-bit Initial Executable (IE)

This sequence is identical to the SPARC 32-bit sequence, except that an `ldx` instruction is used to load the 64-bit address instead of an `ld` instruction.

TABLE 8-5 SPARC: 64-bit Initial Executable Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
# %l7 - initialized to GOT pointer		
0x00 sethi %hi(@tpoff(x)), %o0	R_SPARC_TLS_IE_HI22	x
0x04 or %o0,%lo(@tpoff(x)),%o0	R_SPARC_TLS_IE_LO10	x
0x08 ldx [%l7 + %o0], %o0	R_SPARC_TLS_IE_LD	x
0x0c add %g7, %o0, %o0	R_SPARC_TLS_IE_ADD	x
# %o0 - contains address of TLS variable		
Outstanding Relocations		Symbol
GOT[n]		R_SPARC_TLS_TPOFF64 x

SPARC: 32-bit and 64-bit Local Executable (LE)

This code sequence can only be used from within a dynamic executable to reference a TLS variable defined within the executable. If this is the case, the static `tlsoffset` is known at link-edit time and no runtime relocations are required.

TABLE 8-6 SPARC: 32-bit and 64-bit Local Executable Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 sethi %hix(@tpoff(x)), %o0	R_SPARC_TLS_LE_HIX22	x
0x04 xor %o0,%lo(@tpoff(x)),%o0	R_SPARC_TLS_LE_LOX10	x
0x08 add %g7, %o0, %o0	<none>	
# %o0 - contains address of TLS variable		

The `sethi` and `or` instructions generate `R_SPARC_TLS_LE_HIX22` and `R_SPARC_TLS_LE_LOX10` relocations respectively. The link-editor binds these relocations directly to the static TLS offset for the symbol defined in the executable. No relocation processing is required at runtime.

SPARC: Thread-Local Storage Relocation Types

The TLS relocations listed in the following table are defined for SPARC. Descriptions in the table use the following notation:

`@dtlndx(x)`

Allocates two contiguous entries in the global offset table to hold a `TLS_index` structure. This information is passed to `__tls_get_addr()`. The instruction referencing this entry is bound to the address of the first of the two GOT entries.

@tmndx(x)

Allocates two contiguous entries in the global offset table to hold a `TLS_index` structure. This information is passed to `__tls_get_addr()`. The `ti_tlsoffset` field of this structure is set to 0, and the `ti_moduleid` is filled in at runtime. The call to `__tls_get_addr()` returns the starting offset of the dynamic TLS block.

@dtpoff(x)

Calculates the `tlsoffset` relative to the TLS block.

@tpoff(x)

Calculates the negative `tlsoffset` relative to the static TLS block. This value is added to the thread-pointer to calculate the TLS address.

@dtpmod(x)

Calculates the object identifier of the object containing symbol `S`.

TABLE 8-7 SPARC: Thread-Local Storage Relocation Types

Name	Value	Field	Calculation
R_SPARC_TLS_GD_HI22	56	T-simm22	@dtlndx(S + A) >> 10
R_SPARC_TLS_GD_LO10	57	T-simm13	@dtlndx(S + A) & 0x3ff
R_SPARC_TLS_GD_ADD	58	None	See R_SPARC_R_SPARC_TLS_GD_ADD
R_SPARC_TLS_GD_CALL	59	V-disp30	See R_SPARC_R_SPARC_TLS_GD_CALL
R_SPARC_TLS_LDM_HI22	60	T-simm22	@tmndx(S + A) >> 10
R_SPARC_TLS_LDM_LO10	61	T-simm13	@tmndx(S + A) & 0x3ff
R_SPARC_TLS_LDM_ADD	62	None	See R_SPARC_R_SPARC_TLS_LDM_ADD
R_SPARC_TLS_LDM_CALL	63	V-disp30	See R_SPARC_R_SPARC_TLS_LDM_CALL
R_SPARC_TLS_LDO_HIX22	64	T-simm22	@dtpoff(S + A) >> 10
R_SPARC_TLS_LDO_LOX10	65	T-simm13	@dtpoff(S + A) & 0x3ff
R_SPARC_TLS_LDO_ADD	66	None	See R_SPARC_R_SPARC_TLS_LDO_ADD
R_SPARC_TLS_IE_HI22	67	T-simm22	@got(@tpoff(S + A)) >> 10
R_SPARC_TLS_IE_LO10	68	T-simm13	@got(@tpoff(S + A)) & 0x3ff
R_SPARC_TLS_IE_LD	69	None	See R_SPARC_R_SPARC_TLS_IE_LD
R_SPARC_TLS_IE_LDX	70	None	See R_SPARC_R_SPARC_TLS_IE_LDX
R_SPARC_TLS_IE_ADD	71	None	See R_SPARC_R_SPARC_TLS_IE_ADD
R_SPARC_TLS_LE_HIX22	72	T-imm22	(@tpoff(S + A) ^0xffffffffffffffff) >> 10
R_SPARC_TLS_LE_LOX10	73	T-simm13	(@tpoff(S + A) & 0x3ff) 0x1c00

TABLE 8-7 SPARC: Thread-Local Storage Relocation Types (Continued)

Name	Value	Field	Calculation
R_SPARC_TLS_DTPMOD32	74	V-word32	@dtpmod(S + A)
R_SPARC_TLS_DTPMOD64	75	V-word64	@dtpmod(S + A)
R_SPARC_TLS_DTPOFF32	76	V-word32	@dtpoff(S + A)
R_SPARC_TLS_DTPOFF64	77	V-word64	@dtpoff(S + A)
R_SPARC_TLS_TPOFF32	78	V-word32	@tpoff(S + A)
R_SPARC_TLS_TPOFF64	79	V-word64	@tpoff(S + A)

Some relocation types have semantics beyond simple calculations:

R_SPARC_TLS_GD_ADD

This relocation tags the add instruction of a GD code sequence. The register used for the GOT-pointer is the first register in the sequence. The instruction tagged by this relocation comes before the call instruction tagged by the R_SPARC_TLS_GD_CALL relocation. This is used to transition between TLS models at link-edit time.

R_SPARC_TLS_GD_CALL

This relocation is handled as if it were a R_SPARC_WPLT30 relocation referencing the __tls_get_addr() function. This relocation is part of a GD code sequence.

R_SPARC_LDM_ADD

This relocation tags the first add instruction of a LD code sequence. The register used for the GOT-pointer is the first register in the sequence. The instruction tagged by this relocation comes before the call instruction tagged by the R_SPARC_TLS_GD_CALL relocation. This is used to transition between TLS models at link-edit time.

R_SPARC_LDM_CALL

This relocation is handled as if it were a R_SPARC_WPLT30 relocation referencing the __tls_get_addr() function. This relocation is part of a LD code sequence.

R_SPARC_LDO_ADD

This relocation tags the final add instruction in a LD code sequence. The register which contains the object address computed in the initial part of the code sequence is the first register in this instruction. This permits the link-editor to identify this register for code transformations.

R_SPARC_TLS_IE_LD

This relocation tags the ld instruction in the 32-bit IE code sequence. This is used to transition between TLS models at link-edit time.

R_SPARC_TLS_IE_LDX

This relocation tags the ldx instruction in the 64-bit IE code sequence. This is used to transition between TLS models at link-edit time.

R_SPARC_TLS_IE_ADD

This relocation tags the add instruction in the IE code sequence. The register that is used for the GOT-pointer is the first register in the sequence.

x86: Thread-Local Variable Access

On x86, the following code sequence models are available for accessing TLS

x86: General Dynamic (GD)

This code sequence is the most general, and can be included both in a shared objects and dynamic executables. This code sequence can reference an external TLS variable in either a shared object or dynamic executable.

TABLE 8-8 x86: General Dynamic Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 leal x@tlsgd(,%ebx,1),%eax	R_386_TLS_GD	x
0x07 call x@tlsgdplt	R_386_TLS_GD_PLT	x
# %eax - contains address of TLS variable		

	Outstanding Relocations	Symbol
GOT[n]	R_386_TLS_DTPMOD32	x
GOT[n + 1]	R_386_TLS_DTPOFF32	

The `leal` instruction generates a `R_386_TLS_GD` relocation which instructs the link-editor to allocate space in the global offset table to hold a `TLS_index` structure for variable `x`. The link-editor processes this relocation by substituting the GOT-relative offset for the new GOT entry.

Since the load object index and TLS block index for `x` are not known until runtime, the link-editor places the `R_386_TLS_DTPMOD32` and `R_386_TLS_DTPOFF32` relocations against the GOT for processing by the runtime linker. The address of the generated GOT entry is loaded into register `%eax` for the call to `__tls_get_addr()`.

The `call` instruction causes the generation of the `R_386_TLS_GD_PLT` relocation. This instructs the link-editor to bind the call to the `__tls_get_addr()` function and associates the `call` instruction with the GD code sequence.

The `call` instruction must immediately follow the `leal` instruction. This is a required to permit the code transformations.

x86: Local Dynamic (LD)

This code sequence can be used in either a shared object or dynamic executable. This sequence is used when referencing a TLS variable bound to the same object as the reference. Because the dynamic `tlsoffset` can be bound at link-edit time, only one call to `__tls_get_addr()` is required per function call for all symbols which are referenced via the LD code sequence.

TABLE 8-9 x86: Local Dynamic Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 leal x1@tlsldm(%ebx), %eax	R_386_TLS_LDM	x1
0x06 call x@tlsldmplt	R_386_TLS_LDM_PLT	x1
# %eax - contains address of TLS block of current object		
0x10 leal x1@dtppoff(%eax), %edx	R_386_TLS_LDO_32	x1
# %edx - contains address of local TLS variable x1		
0x20 leal x2@dtppoff(%eax), %edx	R_386_TLS_LDO_32	x2
# %edx - contains address of local TLS variable x2		
	Outstanding Relocations	Symbol
GOT[n]	R_386_TLS_DTPMOD32	x
GOT[n + 1]	<none>	

The first `leal` instruction generates a `R_386_TLS_LDM` relocation that instructs the link-editor to allocate space in the global offset table to hold a `TLS_index` structure for the current object. The link-editor process this relocation by substituting the GOT-relative offset for the new linkage table entry.

Since the load object index is not known until runtime, a `R_386_TLS_DTPMOD32` relocation is created, and the `ti_tlsoffset` field of the structure is zero filled. The `call` instruction is tagged with the `R_386_TLS_LDM_PLT` relocation.

The TLS offset for each local symbol is known at link-edit time so the link-editor fills these values in directly.

When a procedure references more than one local symbol, the compiler generates code to obtain the base address of the TLS block once. This base address is then used to calculate the address of each symbol without a separate library call.

x86: Initial Executable (IE)

There are two code-sequences for the IE model. One sequence is for position independent code which uses a GOT-pointer. The other sequence is for position dependent code which does not use a GOT-pointer. Both of these code sequences can only be used in a dynamic executable. These code sequences can reference a TLS variable defined in either the executable or any of the shared libraries loaded at process startup. This model can not reference TLS variables from shared libraries loaded after process startup.

TABLE 8-10 x86: Initial Executable, Position Independent, Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 movl %gs:0, %eax 0x06 addl x@gotntpoff(%ebx), %eax # %eax - contains address of TLS variable	<none> R_386_TLS_GOTIE	x
	Outstanding Relocations	Symbol
GOT[n]	R_386_TLS_TPOFF	x

The `addl` instruction generates a `R_386_TLS_GOTIE` relocation that instructs the link-editor to create space in the global offset table to store the static TLS offset for symbol `x`. A `R_386_TLS_TPOFF` relocation is left outstanding against the GOT table for the runtime linker to fill in with the static TLS offset for symbol `x`.

TABLE 8-11 x86: Initial Executable, Position Dependent, Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 movl %gs:0, %eax 0x06 addl x@indntpoff, %eax # %eax - contains address of TLS variable	<none> R_386_TLS_IE	x
	Outstanding Relocations	Symbol
GOT[n]	R_386_TLS_TPOFF	x

The `addl` instruction generates a `R_386_TLS_IE` relocation, that instructs the link-editor to create space in the global offset table to store the static TLS offset for symbol `x`. The main difference between this sequence and the position independent

form, is that the instruction is bound directly to the GOT entry created, instead of via an offset off of the GOT-pointer register. A `R_386_TLS_TPOFF` relocation is left outstanding against the GOT for the runtime linker to fill in with the static TLS offset for symbol `x`.

The contents of variable `x`, rather than the address, can be loaded by embedding the offset directly into the memory reference as shown in the next two sequences.

TABLE 8-12 x86: Initial Executable, Position Independent, Dynamic Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
<pre>0x00 movl x@gotntpoff(%ebx), %eax 0x06 movl %gs:(%eax), %eax # %eax - contains address of TLS variable</pre>	<pre>R_386_TLS_GOTIE <none></pre>	<code>x</code>
Outstanding Relocations		Symbol
GOT[n]	<code>R_386_TLS_TPOFF</code>	<code>x</code>

TABLE 8-13 x86: Initial Executable, Position Independent, Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
<pre>0x00 movl x@indntpoff, %ecx 0x06 movl %gs:(%ecx), %eax # %eax - contains address of TLS variable</pre>	<pre>R_386_TLS_IE <none></pre>	<code>x</code>
Outstanding Relocations		Symbol
GOT[n]	<code>R_386_TLS_TPOFF</code>	<code>x</code>

In the last sequence, if the `%eax` register is used instead of the `%ecx` above, the first instruction may be either 5 or 6 bytes long.

x86: Local Executable (LE)

This code sequence can only be used from within a dynamic executable and referencing a TLS variable defined within the executable. If this is the case the static `tlsoffset` is known at link-edit time and no runtime relocations are required.

TABLE 8-14 x86: Local Executable Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 movl %gs:0, %eax	<none>	
0x06 leal x@ntpoff(%eax), %eax	R_386_TLS_LE	x
# %eax - contains address of TLS variable		

The `movl` instruction generates `aR_386_TLS_LE_32` relocation. The link-editor binds this relocation directly to the static TLS offset for the symbol defined in the executable. No processing is required at runtime.

The contents of variable `x`, rather than the address, can be accessed with the same relocation by using the following instruction sequence.

TABLE 8-15 x86: Local Executable Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 movl %gs:0, %eax	<none>	
0x06 movl x@ntpoff(%eax), %eax	R_386_TLS_LE	x
# %eax - contains address of TLS variable		

If instead of computing the address of the variable we want to load from it or store in it the following sequence can be used. Note that in this case we use the `x@ntpoff` expression not as an immediate value, but instead as an absolute address.

TABLE 8-16 x86: Local Executable Thread-Local Variable Access Codes

Code Sequence	Initial Relocations	Symbol
0x00 movl %gs:x@ntpoff, %eax	R_386_TLS_LE	x
# %eax - contains address of TLS variable		

x86: Thread-Local Storage Relocation Types

The TLS relocations listed in the following table are defined for x86. Descriptions in the table use the following notation:

@t1sgd(x)
 Allocates two contiguous entries in the GOT to hold a `TLS_index` structure. This structure is passed to `__tls_get_addr()`. The instruction referencing this entry will be bound to the first of the two GOT entries.

@t1sgdplt(x)
 This relocation is handled as if it were a `R_386_PLT32` relocation referencing the `__tls_get_addr()` function.

@t1sldm(x)
 Allocates two contiguous entries in the GOT to hold a `TLS_index` structure. This structure is passed to the `__tls_get_addr()`. The `ti_tlsoffset` field of the `TLS_index` is set to 0, and the `ti_moduleid` is filled in at runtime. The call to `__tls_get_addr()` returns the starting offset of the dynamic TLS block.

@gotntpoff(x)
 Allocates a entry in the GOT, and initializes it with the negative `tlsoffset` relative to the static TLS block. This is performed at runtime via the `R_386_TLS_TPOFF` relocation.

@indntpoff(x)
 This expression is similar to `@gotntpoff`, but used in position dependent code. `@gotntpoff` resolves to a GOT slot address relative to the start of the GOT in the `movl` or `addl` instructions. `@indntpoff` resolves to the absolute GOT slot address.

@ntpoff(x)
 Calculates the negative offset of the variable it is added to relative to the static TLS block.

@dtpoff(x)
 Calculates the `tlsoffset` relative to the TLS block. The value is used as an immediate value of an addend and is not associated with a specific register.

@dtpmod(x)
 Calculates the object identifier of the object containing symbol `S`.

TABLE 8-17 x86: Thread-Local Storage Relocation Types

Name	Value	Field	Calculation
<code>R_386_TLS_GD_PLT</code>	12	Word32	<code>@t1sgdplt</code>
<code>R_386_TLS_LDM_PLT</code>	13	Word32	<code>@t1sldmplt</code>
<code>R_386_TLS_TPOFF</code>	14	Word32	<code>@ntpoff(S)</code>
<code>R_386_TLS_IE</code>	15	Word32	<code>@indntpoff(S)</code>
<code>R_386_TLS_GOTIE</code>	16	Word32	<code>@gotntpoff(S)</code>
<code>R_386_TLS_LE</code>	17	Word32	<code>@ntpoff(S)</code>
<code>R_386_TLS_GD</code>	18	Word32	<code>@t1sgd(S)</code>

TABLE 8-17 x86: Thread-Local Storage Relocation Types *(Continued)*

Name	Value	Field	Calculation
R_386_TLS_LDM	19	Word32	@tlsldm(S)
R_386_TLS_LDO_32	32	Word32	@dtpoff(S)
R_386_TLS_DTPMOD32	35	Word32	@dtpmod(S)
R_386_TLS_DTPOFF32	36	Word32	@dtpoff(S)

Mapfile Option

The link-editor automatically and intelligently maps input sections from relocatable objects to segments in the output file being created. The `-M` option with an associated `mapfile` enables you to change the default mapping provided by the link-editor. In addition, new segments can be created, attributes modified, and symbol versioning information can be supplied with the `mapfile`.

Note – When using a `mapfile` option, you can easily create an output file that does not execute. The link-editor knows how to produce a correct output file without the use of the `mapfile` option.

Sample `mapfiles` provided on the system reside in the `/usr/lib/ld` directory.

Mapfile Structure and Syntax

You can enter four basic types of directives into a `mapfile`:

- Segment declarations.
- Mapping directives.
- Section-to-segment ordering.
- Size-symbol declarations.
- File control directives.

Each directive can span more than one line and can have any amount of white space, including new lines, as long as that white space is followed by a semicolon.

Typically, segment declarations are followed by mapping directives. You declare a segment and then define the criteria by which a section becomes part of that segment. If you enter a mapping directive or size-symbol declaration without first declaring the segment to which you are mapping, except for built-in segments, the segment is given default attributes. Such segment is an *implicitly* declared segment.

Size-symbol declarations and file control directives can appear anywhere in a `mapfile`.

The following sections describe each directive type. For all syntax discussions, the following notations apply:

- All entries in constant width, all colons, semicolons, equal signs, and at (@) signs are typed in literally.
- All entries in *italics* are substitutable.
- { ... }* means “zero or more.”
- { ... }+ means “one or more.”
- [...] means “optional.”
- `section_names` and `segment_names` follow the same rules as C identifiers, where a period (.) is treated as a letter. For example, `.bss` is a legal name.
- `section_names`, `segment_names`, `file_names`, and `symbol_names` are case sensitive. Everything else is not case sensitive.
- Spaces, or new-lines, can appear anywhere except before a number or in the middle of a name or value.
- Comments beginning with # and ending at a newline can appear anywhere that a space can appear.

Segment Declarations

A segment declaration creates a new segment in the output file, or changes the attribute values of an existing segment. An existing segment is one that you previously defined or one of the four built-in segments described immediately following.

A segment declaration has the following syntax:

```
segment_name = {segment_attribute_value}*;
```

For each `segment_name`, you can specify any number of `segment_attribute_values` in any order, each separated by a space. Only one attribute value is allowed for each segment attribute. The segment attributes and their valid values are as shown in the following table.

TABLE 9-1 Mapfile Segment Attributes

Attribute	Value
<code>segment_type</code>	LOAD NOTE STACK
<code>segment_flags</code>	? [E] [N] [O] [R] [W] [X]
<code>virtual_address</code>	<i>Vnumber</i>
<code>physical_address</code>	<i>Pnumber</i>
<code>length</code>	<i>Lnumber</i>
<code>rounding</code>	<i>Rnumber</i>
<code>alignment</code>	<i>Anumber</i>

There are four built-in segments with the following default attribute values:

- `text` – LOAD, ?RX, no `virtual_address`, `physical_address`, or `length` specified, `alignment` values set to defaults per CPU type.
- `data` – LOAD, ?RWX, no `virtual_address`, `physical_address`, or `length` specified, `alignment` values set to defaults per CPU type.
- `bss` – disabled, LOAD, ?RWX, no `virtual_address`, `physical_address`, or `length` specified, `alignment` values set to defaults per CPU type.
- `note` – NOTE.

By default, the `bss` segment is disabled. Any sections of type `SHT_NOBITS`, which are its sole input, are captured in the `data` segment. See Table 7-12 for a full description of `SHT_NOBITS` sections. The simplest `bss` declaration:

```
bss =;
```

is sufficient to enable the creation of a `bss` segment. Any `SHT_NOBITS` sections is captured by this segment, rather than captured in the `data` segment. In its simplest form, this segment is aligned using the same defaults as applied to any other segment. The declaration can also provide additional segment attributes that both enable the segment creation and assign the specified attributes.

The link-editor behaves as if these segments are declared before your `mapfile` is read in. See “Mapfile Option Defaults” on page 302.

Note the following when entering segment declarations:

- A number can be hexadecimal, decimal, or octal, following the same rules as in the C language.
- No space is allowed between the `V`, `P`, `L`, `R`, or `A` and the number.
- The `segment_type` value can be either `LOAD`, `NOTE` or `STACK`. If unspecified it defaults to `LOAD`.

- The `segment_flags` values are R for readable, W for writable, X for executable, and O for order. No spaces are allowed between the question mark (?) and the individual flags that make up the `segment_flags` value.
- The `segment_flags` value for a LOAD segment defaults to RWX.
- NOTE segments cannot be assigned any segment attribute value other than a `segment_type`.
- One `segment_type` of value STACK is permitted. Only the access requirements of the segment, selected from the `segment_flags`, can be specified.
- Implicitly declared segments default to `segment_type` value LOAD, `segment_flags` value RWX, a default `virtual_address`, `physical_address`, and `alignment` value, and have no `length` limit.

Note – The link-editor calculates the addresses and length of the current segment based on the previous segment’s attribute values.

- LOAD segments can have an explicitly specified `virtual_address` value or `physical_address` value, as well as a maximum segment length value.
- If a segment has a `segment_flags` value of ? with nothing following, the value defaults to not readable, not writable, and not executable.
- The `alignment` value is used in calculating the virtual address of the beginning of the segment. This alignment only affects the segment for which it is specified. Other segments still have the default alignment unless their alignments are also changed.
- If any of the `virtual_address`, `physical_address`, or `length` attribute values are not set, the link-editor calculates these values as it creates the output file.
- If an `alignment` value is not specified for a segment, it is set to the built-in default. This default differs from one CPU to another and might even differ between software revisions.
- If both a `virtual_address` and an `alignment` value are specified for a segment, the `virtual_address` value takes priority.
- If a `virtual_address` value is specified for a segment, the `alignment` field in the program header contains the default alignment value.
- If the `rounding` value is set for a segment, that segment’s virtual address will be rounded to the next address that conforms to the value given. This value only effects the segments that it is specified for. If no value is given, no rounding is performed.

Note – If a `virtual_address` value is specified, the segment is placed at that virtual address. For the system kernel, this method creates a correct result. For files that start through `exec(2)`, this method creates an incorrect output file because the segments do not have correct offsets relative to their page boundaries.

The `?E` flag allows the creation of an empty segment. This empty segment has no sections associated with it. This segment can only be specified for executables, and must be of type `LOAD` with a specified size and alignment. Multiple segment definitions of this type are permitted.

The `?N` flag enables you control whether the ELF header, and any program headers are included as part of the first loadable segment. By default, the ELF header and program headers are included with the first segment. The information in these headers is used within the mapped image, typically by the runtime linker. The use of the `?N` option causes the virtual address calculations for the image to start at the first section of the first segment.

The `?O` flag enables you control the order of sections in the output file. This flag is intended for use in conjunction with the `-xF` option to the compilers. When a file is compiled with the `-xF` option, each function in that file is placed in a separate section with the same attributes as the `.text` section. These sections are called `.text%function_name`.

For example, a file containing three functions, `main()`, `foo()` and `bar()`, when compiled with the `-xF` option, yields a relocatable object file with text for the three functions being placed in sections called `.text%main`, `.text%foo`, and `.text%bar`. Because the `-xF` option forces one function per section, the use of the `?O` flag to control the order of sections in effect controls the order of functions.

Consider the following user-defined `mapfile`:

```
text = LOAD ?RXO;
text: .text%foo;
text: .text%bar;
text: .text%main;
```

The first declaration associates the `?O` flag with the default text segment.

If the order of function definitions in the source file is `main`, `foo`, and `bar`, then the final executable contains functions in the order `foo`, `bar`, and `main`.

For static functions with the same name, the file names must also be used. The `?O` flag forces the ordering of sections as requested in the `mapfile`. For example, if a static function `bar()` exists in files `a.o` and `b.o`, and function `bar()` from file `a.o` is to be placed before function `bar()` from file `b.o`, then the `mapfile` entries should read:

```
text: .text%bar: a.o;
text: .text%bar: b.o;
```

Although the syntax allows for the entry:

```
text: .text%bar: a.o b.o;
```

this entry does not guarantee that function `bar ()` from file `a.o` is placed before function `bar ()` from file `b.o`. The second format is not recommended as the results are not reliable.

Mapping Directives

A mapping directive instructs the link-editor how to map input sections to output segments. Basically, you name the segment that you are mapping to and indicate what the attributes of a section must be in order to map into the named segment. The set of `section_attribute_values` that a section must have to map into a specific segment is called the *entrance criteria* for that segment. In order to be placed in a specified segment of the output file, a section must meet the entrance criteria for a segment exactly.

A mapping directive has the following syntax:

```
segment_name : {section_attribute_value}* [: {file_name}+];
```

For a `segment_name`, you specify any number of `section_attribute_values` in any order, each separated by a space. At most, one section attribute value is allowed for each section attribute. You can also specify that the section must come from a certain `.o` file through a `file_name` declaration. The section attributes and their valid values are shown in the following table.

TABLE 9-2 Section Attributes

Section Attribute	Value
<code>section_name</code>	Any valid section name
<code>section_type</code>	<code>\$PROGBITS</code> <code>\$SYMTAB</code> <code>\$STRTAB</code> <code>\$REL</code> <code>\$RELA</code> <code>\$NOTE</code> <code>\$NOBITS</code>
<code>section_flags</code>	? [!]A [!]W [!]X

Note the following points when entering mapping directives:

- You must choose at most one `section_type` from the `section_types` listed above. The `section_types` listed above are built-in types. For more information on `section_types`, see “Sections” on page 187.
- The `section_flags` values are A for allocatable, W for writable, or X for executable. If an individual flag is preceded by an exclamation mark (!), the link-editor checks that the flag is not set. No spaces are allowed between the question mark, exclamation marks, and the individual flags that make up the `section_flags` value.
- `file_name` can be any legal file name, of the form `*filename`, or of the form `archive_name(component_name)`, for example, `/usr/lib/libc.a` (`printf.o`). The link-editor does not check the syntax of file names.
- If a `file_name` is of the form `*filename`, the link-editor simulates a `basename(1)` on the file from the command line and uses it to match against the specified file name. In other words, the `filename` from the `mapfile` only needs to match the last part of the file name from the command line. See “Mapping Example” on page 300.
- If you use the `-l` option during a link-edit, and the library after the `-l` option is in the current directory, you must precede the library with `./`, or the entire path name, in the `mapfile` in order to create a match.
- More than one directive line can appear for a particular output segment. For example, the following set of directives is legal:

```
S1 : $PROGBITS;
S1 : $NOBITS;
```

Entering more than one mapping directive line for a segment is the only way to specify multiple values of a section attribute.

- A section can match more than one entrance criteria. In this case, the first segment encountered in the `mapfile` with that entrance criteria is used. For example, if a `mapfile` reads:

```
S1 : $PROGBITS;
S2 : $PROGBITS;
```

the `$PROGBITS` sections are mapped to segment `S1`.

Section-Within-Segment Ordering

By using the following notation you can specify the order that sections are placed within a segment:

```
segment_name | section_name1;
segment_name | section_name2;
segment_name | section_name3;
```

The sections that are named in the above form are placed before any unnamed sections, and in the order they are listed in the `mapfile`.

Size-Symbol Declarations

Size-symbol declarations enable you to define a new global-absolute symbol that represents the size, in bytes, of the specified segment. This symbol can be referenced in your object files. A size-symbol declaration has the following syntax:

```
segment_name @ symbol_name;
```

`symbol_name` can be any legal C identifier. The link-editor does not check the syntax of the `symbol_name`.

File Control Directives

File control directives enable you to specify which version definitions within shared objects are to be made available during a link-edit. The file control definition has the following syntax:

```
shared_object_name - version_name [ version_name ... ];
```

`version_name` is a version definition name contained within the specified `shared_object_name`.

Mapping Example

The following example is a user-defined `mapfile`. The numbers on the left are included in the example for tutorial purposes. Only the information to the right of the numbers actually appears in the `mapfile`.

EXAMPLE 9-1 User-Defined Mapfile

```
1. elephant : .data : peanuts.o *popcorn.o;
2. monkey : $PROGBITS ?AX;
3. monkey : .data;
4. monkey = LOAD V0x80000000 L0x4000;
5. donkey : .data;
6. donkey = ?RX A0x1000;
7. text = V0x80008000;
```

Four separate segments are manipulated in this example. The implicitly declared segment `elephant` (line 1) receives all of the `.data` sections from the files `peanuts.o` and `popcorn.o`. Notice that `*popcorn.o` matches any `popcorn.o` file that can be supplied to the link-edit. The file need not be in the current directory. On the other hand, if `/var/tmp/peanuts.o` was supplied to the link-edit, it does not match `peanuts.o` because it is not preceded by an `*`.

The implicitly declared segment `monkey` (line 2) receives all sections that are both `$PROGBITS` and `allocatable-executable (?AX)`, as well as all sections not already in the segment `elephant` with the name `.data` (line 3). The `.data` sections entering the `monkey` segment need not be `$PROGBITS` or `allocatable-executable` because the `section_type` and `section_flags` values are entered on a separate line from the `section_name` value.

An “and” relationship exists between attributes on the same line as illustrated by `$PROGBITS “and” ?AX` on line 2. An “or” relationship exists between attributes for the same segment that span more than one line, as illustrated by `$PROGBITS ?AX` on line 2 “or” `.data` on line 3.

The `monkey` segment is implicitly declared in line 2 with `segment_type` value `LOAD`, `segment_flags` value `RWX`, and no `virtual_address`, `physical_address`, `length` or `alignment` values specified (defaults are used). In line 4 the `segment_type` value of `monkey` is set to `LOAD`. Because the `segment_type` attribute value does not change, no warning is issued. The `virtual_address` value is set to `0x80000000` and the maximum length value to `0x4000`.

Line 5 implicitly declares the `donkey` segment. The entrance criteria are designed to route all `.data` sections to this segment. Actually, no sections fall into this segment because the entrance criteria for `monkey` in line 3 capture all of these sections. In line 6, the `segment_flags` value is set to `?RX` and the `alignment` value is set to `0x1000`. Because both of these attribute values changed, a warning is issued.

Line 7 sets the `virtual_address` value of the `text` segment to `0x80008000`.

The example of a user-defined `mapfile` is designed to cause warnings for illustration purposes. If you want to change the order of the directives to avoid warnings, use the following example:

```
1. elephant : .data : peanuts.o *popcorn.o;
4. monkey = LOAD V0x80000000 L0x4000;
2. monkey : $PROGBITS ?AX;
3. monkey : .data;
6. donkey = ?RX A0x1000;
5. donkey : .data;
7. text = V0x80008000;
```

The following `mapfile` example uses the segment-within-section ordering:

```
1. text = LOAD ?RXN V0xf0004000;
2. text | .text;
3. text | .rodata;
4. text : $PROGBITS ?A!W;
5. data = LOAD ?RWX R0x1000;
```

The `text` and `data` segments are manipulated in this example. Line 1 declares the `text` segment to have a `virtual_address` of `0xf0004000` and to *not* include the ELF header or any program headers as part of this segment’s address calculations.

Lines 2 and 3 turn on section-within-segment ordering and specify that the `.text` and `.rodata` sections are the first two sections in this segment. The result is that the `.text` section have a virtual address of `0xf0004000`, and the `.rodata` section immediately follows that address.

Any other `$PROGBITS` section that makes up the `text` segment follows the `.rodata` section. Line 5 declares the data segment and specifies that its virtual address must begin on a `0x1000` byte boundary. The first section that constitutes the data segment also resides on a `0x1000` byte boundary within the file image.

Mapfile Option Defaults

The link-editor defines four built-in segments (`text`, `data`, `bss` and `note`) with default `segment_attribute_values` and corresponding default mapping directives. Even though the link-editor does not use an actual `mapfile` to provide the defaults, the model of a default `mapfile` helps illustrate what happens when the link-editor encounters your `mapfile`.

The following example shows how a `mapfile` would appear for the link-editor defaults. The link-editor begins execution behaving as if the `mapfile` has already been read in. Then the link-editor reads your `mapfile` and either augments or makes changes to the defaults.

```
text = LOAD ?RX;
text : ?A!W;
data = LOAD ?RWX;
data : ?AW;
note = NOTE;
note : $NOTE;
```

As each segment declaration in your `mapfile` is read in, it is compared to the existing list of segment declarations as follows:

1. If the segment does not already exist in the `mapfile` but another with the same segment-type value exists, the segment is added before all of the existing segments of the same `segment_type`.
2. If none of the segments in the existing `mapfile` has the same `segment_type` value as the segment just read in, then the segment is added by `segment_type` value to maintain the following order:

```
INTERP
LOAD
DYNAMIC
NOTE
```

3. If the segment is of `segment_type` `LOAD` and you have defined a `virtual_address` value for this `LOADable` segment, the segment is placed before any `LOADable` segments without a defined `virtual_address` value or with a higher `virtual_address` value, but after any segments with a `virtual_address` value that is lower.

As each mapping directive in a `mapfile` is read in, the directive is added after any other mapping directives that you already specified for the same segment but before the default mapping directives for that segment.

Internal Map Structure

One of the most important data structures in the ELF-based link-editor is the map structure. A default map structure, corresponding to the model default `mapfile`, is used by the link-editor. Any user `mapfile` augments or overrides certain values in the default map structure.

A typical although somewhat simplified map structure is illustrated in Figure 9–1. The “Entrance Criteria” boxes correspond to the information in the default mapping directives. The “Segment Attribute Descriptors” boxes correspond to the information in the default segment declarations. The “Output Section Descriptors” boxes give the detailed attributes of the sections that fall under each segment. The sections themselves are shown in circles.




FIGURE 9-1 Simple Map Structure

The link-editor performs the following steps when mapping sections to segments:

1. When a section is read in, the link-editor checks the list of Entrance Criteria looking for a match. All specified criteria must be matched.

In Figure 9-1, a section that falls into the `text` segment must have a `section_type` value of `$PROGBITS` and have a `section_flags` value of `?A!W`. It need not have the name `.text` since no name is specified in the Entrance Criteria. The section can be either `X` or `!X` in the `section_flags` value because nothing was specified for the execute bit in the Entrance Criteria.

If no Entrance Criteria match is found, the section is placed at the end of the output file after all other segments. No program header entry is created for this information.

2. When the section falls into a segment, the link-editor checks the list of existing Output Section Descriptors in that segment as follows:

If the section attribute values match those of an existing Output Section Descriptor exactly, the section is placed at the end of the list of sections associated with that Output Section Descriptor.

For instance, a section with a `section_name` value of `.data1`, a `section_type` value of `$PROGBITS`, and a `section_flags` value of `?AWX` falls into the second Entrance Criteria box in Figure 9-1, placing it in the data segment. The section matches the second Output Section Descriptor box exactly (`.data1`, `$PROGBITS`, `?AWX`) and is added to the end of the list associated with that box. The `.data1` sections from `fido.o`, `rover.o`, and `sam.o` illustrate this point.

If no matching Output Section Descriptor is found but other Output Section Descriptors of the same `section_type` exist, a new Output Section Descriptor is created with the same attribute values as the section and that section is associated with the new Output Section Descriptor. The Output Section Descriptor and the section are placed after the last Output Section Descriptor of the same section type. The `.data2` section in Figure 9-1 was placed in this manner.

If no other Output Section Descriptors of the indicated section type exist, a new Output Section Descriptor is created and the section is placed in that section.

Note – If the input section has a user-defined section type value between `SHT_LOUSER` and `SHT_HIUSER`, it is treated as a `$PROGBITS` section. No method exists for naming this `section_type` value in the `mapfile`, but these sections can be redirected using the other attribute value specifications (`section_flags`, `section_name`) in the entrance criteria.

3. If a segment contains no sections after all of the command line object files and libraries are read in, no program header entry is produced for that segment.

Note – Input sections of type `$SYMTAB`, `$STRTAB`, `$REL`, and `$RELA` are used internally by the link-editor. Directives that refer to these section types can only map output sections produced by the link-editor to segments.

Link-Editor Quick Reference

The following sections provide a simple overview, or *cheat sheet*, of the most commonly used link-editor scenarios. See “Link-Editing” on page 20 for an introduction to the kinds of output modules generated by the link-editor.

The examples provided show the link-editor options as supplied to a compiler driver, this being the most common mechanism of invoking the link-editor. In these examples we use `cc(1)`. See “Using a Compiler Driver” on page 27.

The link-editor places no meaning on the name of any input file. Each file is opened and inspected to determine the type of processing it requires. See “Input File Processing” on page 28.

Shared objects that follow a naming convention of `libx.so`, and archive libraries that follow a naming convention of `libx.a`, can be input using the `-l` option. See “Library Naming Conventions” on page 31. This provides additional flexibility in allowing search paths to be specified using the `-L` option. See “Directories Searched by the Link-Editor” on page 33.

The link-editor basically operates in one of two modes, *static* or *dynamic*.

Static Mode

Static mode is selected when the `-dn` option is used, and enables you to create relocatable objects and static executables. Under this mode, only relocatable objects and archive libraries are acceptable forms of input. Use of the `-l` option results in a search for archive libraries.

Creating a Relocatable Object

- To create a relocatable object use the `-d n` and `-r` options:

```
$ cc -dn -r -o temp.o file1.o file2.o file3.o .....
```

Creating a Static Executable

The use of static executables is limited. Static executables usually contain platform-specific implementation details that restricts the ability of the executable to be run on an alternative platform. Many implementations of Solaris libraries depend on dynamic linking capabilities, such as `dlopen(3DL)` and `dlsym(3DL)`. See “Loading Additional Objects” on page 73. These capabilities are not available to static executables.

- To create a static executable use the `-d n` option *without* the `-r` option:

```
$ cc -dn -o prog file1.o file2.o file3.o .....
```

The `-a` option is available to indicate the creation of a static executable. The use of `-d n` *without* a `-r` implies `-a`.

Dynamic Mode

Dynamic mode is the default mode of operation for the link-editor. It can be enforced by specifying the `-d y` option, but is implied when not using the `-d n` option.

Under this mode, relocatable objects, shared objects and archive libraries are acceptable forms of input. Use of the `-l` option results in a directory search, where each directory is searched for a shared object. If no shared object is found, the same directory is then searched for an archive library. A search only for archive libraries can be enforced by using the `-B static` option. See “Linking With a Mix of Shared Objects and Archives” on page 32.

Creating a Shared Object

- To create a shared object use the `-G` option. `-d y` is optional as it is implied by default.
- Input relocatable objects should be built from position-independent code. For example, the C compiler generates position-independent code under the `-K pic` option. See “Position-Independent Code” on page 115. Use the `-z text` option to enforce this requirement.

- Avoid including unused relocatable objects. Or, use the `-z ignore` option, which instructs the link-editor to eliminate unreferenced ELF sections input as part of the link-edit. See “Remove Unused Material” on page 117.
- If the shared object is intended for external use, make sure it uses no application registers. Not using application registers provides the external user freedom to use these registers without fear of compromising the shared object’s implementation. For example, the SPARC C compiler does not use application registers under the `-xregs=no%appl` option.
- Establish the shared objects public interface by defining the global symbols that should be visible from the shared object, and reducing any other global symbols to local scope. This definition is provided by the `-M` option together with an associated `mapfile`. See Appendix B.
- Use a versioned name for the shared object to allow for future upgrades. See “Coordination of Versioned Filenames” on page 146.
- Self-contained shared objects offer maximum flexibility. They are produced when the object expresses all dependency needs. Use the `-z defs` to enforce this self containment. See “Generating a Shared Object Output File” on page 44.
- Avoid unneeded dependencies. Use `ldd` with the `-u` option to detect and remove unneeded dependencies. See “Shared Object Processing” on page 30. Or, use the `-z ignore` option, which instructs the link-editor to record dependencies only to objects that are referenced.
- If the shared object being generated has dependencies on other shared objects, indicate they should be lazily loaded using the `-z lazyload` option. See “Lazy Loading of Dynamic Dependencies” on page 74.
- If the shared object being generated has dependencies on other shared objects, and these dependencies do not reside in `/usr/lib` for 32-bit objects, or `/usr/lib/64` for 64-bit objects, record their path name in the output file using the `-R` option. See “Shared Objects With Dependencies” on page 105.
- Optimize relocation processing by combining relocation sections into a single `.SUNW_reloc` section. Use the `-z combrelloc` option.
- If interposing symbols are not used on this object or its dependencies, establish direct binding information with `-B direct`. See “Direct Binding” on page 69.

The following example combines the above points:

```
$ cc -c -o foo.o -K pic -xregs=no%appl foo.c
$ cc -M mapfile -G -o libfoo.so.1 -z text -z defs -B direct -z lazyload \
-z combrelloc -z ignore -R /home/lib foo.o -L. -lbar -lc
```

- If the shared object being generated is used as input to another link-edit, record within it the shared object’s runtime name using the `-h` option. See “Recording a Shared Object Name” on page 103.
- Make the shared object available to the compilation environment by creating a file system link to a non-versioned shared object name. See “Coordination of Versioned Filenames” on page 146.

The following example combines the above points:

```
$ cc -M mapfile -G -o libfoo.so.1 -z text -z defs -B direct -z lazyload \  
-z combrelloc -z ignore -R /home/lib -h libfoo.so.1 foo.o -L. -lbar -lc  
$ ln -s libfoo.so.1 libfoo.so
```

- Consider the performance implications of the shared object: Maximize shareability, as described in “Maximizing Shareability” on page 118; Minimize paging activity, as described in “Minimizing Paging Activity” on page 120; Reduce relocation overhead, especially by minimizing symbolic relocations, as described in “Reducing Symbol Scope” on page 52; Allow access to data via functional interfaces, as described in “Copy Relocations” on page 122.

Creating a Dynamic Executable

- To create a dynamic executable don’t use the `-G`, or `-d n` options.
- Indicate that the dependencies of the dynamic executable should be lazily loaded using the `-z lazyload` option. See “Lazy Loading of Dynamic Dependencies” on page 74.
- If the dependencies of the dynamic executable do not reside in `/usr/lib` for 32-bit objects, or `/usr/lib/64` for 64-bit objects, record their path name in the output file using the `-R` option. See “Directories Searched by the Runtime Linker” on page 35.
- Establish direct binding information using `-B direct`. See “Direct Binding” on page 69.

The following example combines the above points:

```
$ cc -o prog -R /home/lib -z ignore -z lazyload -B direct -L. \  
-lfoo file1.o file2.o file3.o .....
```

Versioning Quick Reference

ELF objects make available global symbols to which other objects can bind. Some of these global symbols can be identified as providing the object's *public interface*. Other symbols are part of the object's internal implementation and are not intended for external use. An object's interface can evolve from one software release to another. The ability to identify this evolution is desirable.

In addition, identifying the *internal implementation* changes of an object from one software release to another might be desirable.

Both interface and implementation identifications can be recorded within an object by establishing internal *version definitions*. See Chapter 5 for a more complete introduction to the concept of internal versioning.

Shared objects are prime candidates for internal versioning. This technique defines their evolution, provides for interface validation during runtime processing (see "Binding to a Version Definition" on page 136), and provides for the selective binding of applications (see "Specifying a Version Binding" on page 140). Shared objects are used as the examples throughout this appendix.

The following sections provide a simple overview, or *cheat sheet*, of the internal versioning mechanism provided by the link-editors as applied to shared objects. The examples recommend conventions and mechanisms for versioning shared objects, from their initial construction through several common update scenarios.

Naming Conventions

A shared object follows a naming convention that includes a *major* number file suffix. See "Naming Conventions" on page 102. Within this shared object, one or more *version definitions* can be created. Each version definition corresponds to one of the following categories:

- It defines an industry-standard interface (for example, the *System V Application Binary Interface*).
- It defines a vendor-specific public interface.
- It defines a vendor-specific private interface.
- It defines a vendor-specific change to the internal implementation of the object.

The following version definition naming conventions help indicate which of these categories the definition represents.

The first three of these categories indicate interface definitions. These definitions consist of an association of the global symbol names that make up the interface, with a version definition name. See “Creating a Version Definition” on page 131. Interface changes within a shared object are often referred to as minor revisions. Therefore, version definitions of this type are suffixed with a minor version number, which is based on the file names major version number suffix.

The last category indicates a change having occurred within the object. This definition consists of a version definition acting as a label and has no symbol name associated with it. This definition is referred to as being a weak version definition. See “Creating a Weak Version Definition” on page 134. Implementation changes within a shared object are often referred to as micro revisions. Therefore, version definitions of this type are suffixed with a micro version number based on the previous minor number to which the internal changes have been applied.

Any industry standard interface should use a version definition name that reflects the standard. Any vendor interfaces should use a version definition name unique to that vendor. The company’s stock symbol is often appropriate.

Private version definitions indicate symbols that have restricted or uncommitted use, and should have the word “private” clearly visible.

All version definitions result in the creation of associated version symbol names. The use of unique names and the minor/micro suffix convention reduces the chance of symbol collision within the object being built.

The following version definition examples show the possible use of these naming conventions:

```
SVABI . 1
    Defines the System V Application Binary Interface standards interface.

SUNW_1 . 1
    Defines a Solaris public interface.

SUNWprivate_1 . 1
    Defines a Solaris private interface.

SUNW_1 . 1 . 1
    Defines a Solaris internal implementation change.
```

Defining a Shared Object's Interface

When establishing a shared object's interface, you should first determine which global symbols provided by the shared object can be associated to one of the three interface version definition categories:

- Industry standard interface symbols conventionally are defined in publicly available header files and associated manual pages supplied by the vendor, and are also documented in recognized standards literature.
- Vendor public interface symbols conventionally are defined in publicly available header files and associated manual pages supplied by the vendor.
- Vendor private interface symbols can have little or no public definition.

By defining these interfaces, a vendor is indicating the commitment level of each interface of the shared object. Industry standard and vendor public interfaces remain stable from release to release. You are free to bind to these interfaces safe in the knowledge that your application will continue to function correctly from release to release.

Industry-standard interfaces might be available on systems provided by other vendors. You can achieve a higher level of binary compatibility by restricting your applications to use these interfaces.

Vendor public interfaces might not be available on systems provided by other vendors. However, these interfaces remain stable during the evolution of the system on which they are provided.

Vendor private interfaces are very unstable, and can change, or even be deleted, from release to release. These interfaces provide for uncommitted or experimental functionality, or are intended to provide access for vendor-specific applications only. If you want to achieve any level of binary compatibility, you should avoid using these interfaces.

Any global symbols that do not fall into one of the above categories should be reduced to local scope so that they are no longer visible for binding. See "Reducing Symbol Scope" on page 52.

Versioning a Shared Object

Having determined a shared object's available interfaces, the associated version definitions are created using a `mapfile` and the `link-editor's -M` option. See "Defining Additional Symbols" on page 46 for an introduction to this `mapfile` syntax.

The following example defines a vendor public interface in the shared object `libfoo.so.1`:

```
$ cat mapfile
SUNW_1.1 {                # Release X.
    global:
        foo2;
        foo1;
    local:
        *;
};
$ cc -G -o libfoo.so.1 -h libfoo.so.1 -z text -M mapfile foo.c
```

The global symbols `foo1` and `foo2` are assigned to the shared object's public interface `SUNW_1.1`. Any other global symbols supplied from the input files are reduced to local by the auto-reduction directive `"*"`. See "Reducing Symbol Scope" on page 52.

Note – Each version definition `mapfile` entry should be accompanied by a comment reflecting the release or date of the update. This information helps coordinate multiple updates of a shared object, possibly by different developers, into one version definition suitable for delivery of the shared object as part of a software release.

Versioning an Existing (Non-versioned) Shared Object

Versioning an existing, non-versioned shared object requires extra care. The shared object delivered in a previous software release has made available all its global symbols for others to bind with. Although you can determine the shared object's intended interfaces, others might have discovered and bound to other symbols. Therefore, the removal of any symbols might result in an application's failure on delivery of the new versioned shared object.

The internal versioning of an existing, non-versioned shared object can be achieved if the interfaces can be determined, and applied, without breaking any existing applications. The runtime linker's debugging capabilities can be useful to help verify the binding requirements of various applications. See "Debugging Library" on page 93. However, this determination of existing binding requirements assumes that all users of the shared object are known.

If the binding requirements of an existing, non-versioned shared object cannot be determined, then you should create a new shared object file using a new versioned name. See "Coordination of Versioned Filenames" on page 146. In addition to this new shared object, the original shared object must also be delivered so as to satisfy the dependencies of any existing applications.

If the implementation of the original shared object is to be frozen, then maintaining and delivering the shared object binary might be sufficient. If, however, the original shared object might require updating then an alternative source tree from which to generate the shared object can be more applicable. Updating might be necessary through patches, or because its implementation must evolve to remain compatible with new platforms.

Updating a Versioned Shared Object

The only changes that can be made to a shared object that can be absorbed by internal versioning are compatible changes. See “Interface Compatibility” on page 130. Any incompatible changes require producing a new shared object with a new external versioned name. See “Coordination of Versioned Filenames” on page 146.

Compatible updates that can be accommodated by internal versioning fall into three basic categories:

- Adding new symbols
- Creating new interfaces from existing symbols
- Internal implementation changes

The first two categories are achieved by associating an interface version definition with the appropriate symbols. The latter is achieved by creating a weak version definition that has no associated symbols.

Adding New Symbols

Any compatible new release of a shared object that contains new global symbols should assign these symbols to a new version definition. This new version definition should inherit the previous version definition.

The following `mapfile` example assigns the new symbol `foo3` to the new interface version definition `SUNW_1.2`. This new interface inherits the original interface `SUNW_1.1`.

```
$ cat mapfile
SUNW_1.2 {
    global:
        foo3;
} SUNW_1.1;

SUNW_1.1 {
    global:
        foo2;
```

```

        foo1;
    local:
        *;
};

```

The inheritance of version definitions reduces the amount of version information that must be recorded in any user of the shared object.

Internal Implementation Changes

Any compatible new release of the shared object that consists of an update to the implementation of the object, for example, a bug fix or performance improvement, should be accompanied by a *weak* version definition. This new version definition should inherit the latest version definition present at the time the update occurred.

The following `mapfile` example generates a weak version definition `SUNW_1.1.1`. This new interface indicates that the internal changes were made to the implementation offered by the previous interface `SUNW_1.1`.

```

$ cat mapfile
SUNW_1.1.1 { } SUNW_1.1;      # Release X+1.

SUNW_1.1 {                   # Release X.
    global:
        foo2;
        foo1;
    local:
        *;
};

```

New Symbols and Internal Implementation Changes

If both internal changes and the addition of a new interface have occurred during the same release, both a weak version and an interface version definition should be created. The following example shows the addition of a version definition `SUNW_1.2` and an interface change `SUNW_1.1.1`, which are added during the same release cycle. Both interfaces inherit the original interface `SUNW_1.1`.

```

$ cat mapfile
SUNW_1.2 {                   # Release X+1.
    global:
        foo3;
} SUNW_1.1;

SUNW_1.1.1 { } SUNW_1.1;    # Release X+1.

SUNW_1.1 {                   # Release X.

```

```

        global:
            foo2;
            foo1;
        local:
            *;
};

```

Note – The comments for the `SUNW_1.1` and `SUNW_1.1.1` version definitions indicate that they have both been applied to the same release.

Migrating Symbols to a Standard Interface

Occasionally, symbols offered by a vendor's interface become absorbed into a new industry standard. When creating a new standard interface, make sure to maintain the original interface definitions provided by the shared object. Create intermediate version definitions on which the new standard, and original interface definitions, can be built.

The following `mapfile` example shows the addition of a new industry standard interface `STAND.1`. This interface contains the new symbol `foo4` and the existing symbols `foo3` and `foo1`, which were originally offered through the interfaces `SUNW_1.2` and `SUNW_1.1` respectively.

```

$ cat mapfile
STAND.1 {                                # Release X+2.
    global:
        foo4;
} STAND.0.1 STAND.0.2;

SUNW_1.2 {                                # Release X+1.
    global:
        SUNW_1.2;
} STAND.0.1 SUNW_1.1;

SUNW_1.1.1 { } SUNW_1.1;                 # Release X+1.

SUNW_1.1 {                                # Release X.
    global:
        foo2;
    local:
        *;
} STAND.0.2;

STAND.0.1 {                                # Subversion - providing for
    global:                                # SUNW_1.2 and STAND.1 interfaces.
        foo3;
};

STAND.0.2 {                                # Subversion - providing for
    # SUNW_1.1 and STAND.1 interfaces.

```

```

        global:
            foo1;
};

```

The symbols `foo3` and `foo1` are pulled into their own intermediate interface definitions, which are used to create the original and new interface definitions.

The new definition of the `SUNW_1.2` interface has referenced its own version definition symbol. Without this reference, the `SUNW_1.2` interface would have contained no immediate symbol references and hence would be categorized as a weak version definition.

When migrating symbol definitions to a standards interface, any original interface definitions must continue to represent the same symbol list. This requirement can be validated using `pvs(1)`. The following example shows the symbol list of the `SUNW_1.2` interface as it existed in the software release `X+1`.

```

$ pvs -ds -N SUNW_1.2 libfoo.so.1
SUNW_1.2:
    foo3;
SUNW_1.1:
    foo2;
    foo1;

```

Although the introduction of the new standards interface in software release `X+2` has changed the interface version definitions available, the list of symbols provided by each of the original interfaces remains constant. The following example shows that interface `SUNW_1.2` still provides symbols `foo1`, `foo2` and `foo3`.

```

$ pvs -ds -N SUNW_1.2 libfoo.so.1
SUNW_1.2:
STAND.0.1:
    foo3;
SUNW_1.1:
    foo2;
STAND.0.2:
    foo1;

```

An application might only reference one of the new subversions. In this case, any attempt to run the application on a previous release results in a runtime versioning error. See “Binding to a Version Definition” on page 136.

An application’s version binding can be promoted by directly referencing an existing version name. See “Binding to Additional Version Definitions” on page 142. For example, if an application only references the symbol `foo1` from the shared object `libfoo.so.1`, then its version reference is to `STAND.0.2`. To enable this application to be run on previous releases, the version binding can be promoted to `SUNW_1.1` using a version control `mapfile` directive.

```

$ cat prog.c
extern void foo1();

main()

```

```

{
    fool();
}
$ cc -o prog prog.c -L. -R. -lfoo
$ pvs -r prog
    libfoo.so.1 (STAND.0.2);

$ cat mapfile
libfoo.so - SUNW_1.1 $ADDVERS=SUNW_1.1;
$ cc -M mapfile -o prog prog.c -L. -R. -lfoo
$ pvs -r prog
    libfoo.so.1 (SUNW_1.1);

```

In practice, you rarely have to promote a version binding in this manner. The introduction of new standards binary interfaces is rare, and most applications reference many symbols from an interface family.

Establishing Dependencies with Dynamic String Tokens

A dynamic object can establish dependencies explicitly or through filters. Each of these mechanisms can be augmented with a *runpath*, which directs the runtime linker to search for and load the required dependency. String names used to record filters, dependencies and runpath information can be augmented with the reserved dynamic string tokens:

- `$ISALIST`
- `$OSNAME`, `$OSREL` and `$PLATFORM`
- `$ORIGIN`

The following sections provide examples of how each of these tokens may be employed.

Instruction Set Specific Shared Objects

The dynamic token `$ISALIST` is expanded at runtime to reflect the native instruction sets executable on this platform, as displayed by the utility `isalist(1)`. This token is available for filters or runpath definitions.

Any string name that incorporates the `$ISALIST` token is effectively duplicated into multiple strings. Each string is assigned one of the available instruction sets.

The following example shows how the auxiliary filter `libfoo.so.1` can be designed to access an instruction set specific filtee `libbar.so.1`.

```
$ LD_OPTIONS='-f /opt/ISV/lib/$ISALIST/libbar.so.1' \  
cc -o libfoo.so.1 -G -K pic -h libfoo.so.1 -R. foo.c  
$ dump -Lv libfoo.so.1 | egrep "SONAME|AUXILIARY"  
[1] SONAME libfoo.so.1  
[2] AUXILIARY /opt/ISV/lib/$ISALIST/libbar.so.1
```

Or alternatively the runpath can be used.

```

$ LD_OPTIONS='-f libbar.so.1' \
cc -o libfoo.so.1 -G -K pic -h libfoo.so.1 -R'/opt/ISV/lib/$ISALIST' foo.c
$ dump -Lv libfoo.so.1 | egrep "RUNPATH|AUXILIARY"
[1]   RUNPATH   /opt/ISV/lib/$ISALIST
[2]   AUXILIARY libbar.so.1

```

In either case the runtime linker uses the platform available instruction list to construct multiple search paths. For example, the following application is dependent on `libfoo.so.1` and executed on a SUNW, Ultra-2:

```

$ ldd -ls prog
.....
find object=libbar.so.1; required by ./libfoo.so.1
search path=/opt/ISV/lib/$ISALIST (RPATH from file ./libfoo.so.1)
trying path=/opt/ISV/lib/sparcv9+vis/libbar.so.1
trying path=/opt/ISV/lib/sparcv9/libbar.so.1
trying path=/opt/ISV/lib/sparcv8plus+vis/libbar.so.1
trying path=/opt/ISV/lib/sparcv8plus/libbar.so.1
trying path=/opt/ISV/lib/sparcv8/libbar.so.1
trying path=/opt/ISV/lib/sparcv8-fsmuld/libbar.so.1
trying path=/opt/ISV/lib/sparcv7/libbar.so.1
trying path=/opt/ISV/lib/sparc/libbar.so.1

```

Or an application with similar dependencies is executed on an MMX configured Pentium Pro:

```

$ ldd -ls prog
.....
find object=libbar.so.1; required by ./libfoo.so.1
search path=/opt/ISV/lib/$ISALIST (RPATH from file ./libfoo.so.1)
trying path=/opt/ISV/lib/pentium_pro+mmx/libbar.so.1
trying path=/opt/ISV/lib/pentium_pro/libbar.so.1
trying path=/opt/ISV/lib/pentium+mmx/libbar.so.1
trying path=/opt/ISV/lib/pentium/libbar.so.1
trying path=/opt/ISV/lib/i486/libbar.so.1
trying path=/opt/ISV/lib/i386/libbar.so.1
trying path=/opt/ISV/lib/i86/libbar.so.1

```

Reducing Filtee Searches

The use of `$ISALIST` within a filter enables one or more filtees to provide implementations of interfaces defined within the filter.

Any interface defined in a filter can result in an exhaustive search of all potential filtees in an attempt to locate the required interface. If filtees are being employed to provide performance critical functions, this exhaustive filtee searching can be counterproductive.

A filtee can be built with the link-editor's `-z endfiltee` option to indicate that it is the last of the available filtees. This option terminates any further filtee searching for that filter. From the previous SPARC example, if the `sparcv9` filtee existed, and was tagged with `-z endfiltee`, the filtee searches would be:

```

$ ldd -ls prog
.....
find object=libbar.so.1; required by ./libfoo.so.1
  search path=/opt/ISV/lib/$ISALIST (RPATH from file ./libfoo.so.1)
    trying path=/opt/ISV/lib/sparcv9+vis/libbar.so.1
    trying path=/opt/ISV/lib/sparcv9/libbar.so.1

```

System Specific Shared Objects

The dynamic tokens `$OSNAME`, `$OSREL` and `$PLATFORM` are expanded at runtime to provide system specific information. These tokens are available for filters, `runpath`, or dependency definitions.

`$OSNAME` expands to reflect the name of the operating system, as displayed by the utility `uname(1)` with the `-s` option. `$OSREL` expands to reflect the operating system release level, as displayed by `uname -r`. `$PLATFORM` expands to reflect the underlying hardware implementation, as displayed by `uname -i`.

The following example shows how the auxiliary filter `libfoo.so.1` can be designed to access a platform specific filtee `libbar.so.1`.

```

$ LD_OPTIONS='-f /usr/platform/$PLATFORM/lib/libbar.so.1' \
cc -o libfoo.so.1 -G -K pic -h libfoo.so.1 -R. foo.c
$ dump -Lv libfoo.so.1 | egrep "SONAME|AUXILIARY"
[1]  SONAME    libfoo.so.1
[2]  AUXILIARY /usr/platform/$PLATFORM/lib/libbar.so.1

```

This mechanism is used in the Solaris operating environment to provide platform specific extensions to the shared object `/usr/lib/libc.so.1`.

Locating Associated Dependencies

Typically, an unbundled product is designed to be installed in a standalone, unique location. This product is composed of binaries, shared object dependencies, and associated configuration files. For example, the unbundled product ABC might have the layout shown in the following figure.




FIGURE C-1 Unbundled Dependencies

Assume that the product is designed for installation under `/opt`. Normally, you would augment the `PATH` with `/opt/ABC/bin` to locate the product's binaries. Each binary locates their dependencies using a hard-coded runpath within the binary. For the application `abc`, this runpath would be:

```
% cc -o abc abc.c -R/opt/ABC/lib -L/opt/ABC/lib -lA
% dump -Lv abc
[1]  NEEDED  libA.so.1
[2]  RUNPATH /opt/ABC/lib
```

and similarly for the dependency `libA.so.1` this would be:

```
% cc -o libA.so.1 -G -Kpic A.c -R/opt/ABC/lib -L/opt/ABC/lib -lB
% dump -Lv libA.so.1
[1]  NEEDED  libB.so.1
[2]  RUNPATH /opt/ABC/lib
```

This dependency representation works until the product is installed in some directory other than the recommended default.

The dynamic token `$ORIGIN` expands to the directory in which an object originated. This token is available for filters, runpath, or dependency definitions. Use this technology to redefine the unbundled application to locate its dependencies in terms of `$ORIGIN`:

```
% cc -o abc abc.c '-R$ORIGIN/../lib' -L/opt/ABC/lib -lA
% dump -Lv abc
[1]  NEEDED  libA.so.1
[2]  RUNPATH $ORIGIN/../lib
```

and the dependency `libA.so.1` can also be defined in terms of `$ORIGIN`:

```
% cc -o libA.so.1 -G -Kpic A.c '-R$ORIGIN' -L/opt/ABC/lib -lB
% dump -Lv libA.so.1
[1]  NEEDED  libB.so.1
[2]  RUNPATH $ORIGIN
```

If this product is now installed under `/usr/local/ABC` and the user's `PATH` is augmented with `/usr/local/ABC/bin`, invocation of the application `abc` result in a pathname lookup for its dependencies as follows:

```
% ldd -s abc
.....
```

```

find object=libA.so.1; required by abc
search path=$ORIGIN/../lib (RPATH from file abc)
trying path=/usr/local/ABC/lib/libA.so.1
libA.so.1 => /usr/local/ABC/lib/libA.so.1

find object=libB.so.1; required by /usr/local/ABC/lib/libA.so.1
search path=$ORIGIN (RPATH from file /usr/local/ABC/lib/libA.so.1)
trying path=/usr/local/ABC/lib/libB.so.1
libB.so.1 => /usr/local/ABC/lib/libB.so.1

```

Dependencies Between Unbundled Products

Another issue related to dependency location is how to establish a model whereby unbundled products express dependencies between themselves.

For example, the unbundled product XYZ might have dependencies on the product ABC. This dependency can be established by a host package installation script. This script generates a symbolic link to the installation point of the ABC product, as shown in the following figure.

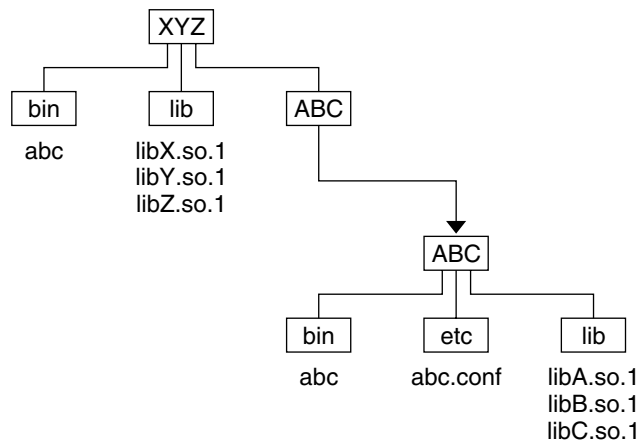


FIGURE C-2 Unbundled Co-Dependencies

The binaries and shared objects of the XYZ product can represent their dependencies on the ABC product using the symbolic link. This link is now a stable reference point. For the application xyz, this runpath would be:

```

% cc -o xyz xyz.c '-R$ORIGIN/../lib:$ORIGIN/../ABC/lib' \
-L/opt/ABC/lib -lX -lA
% dump -Lv xyz
[1]   NEEDED   libX.so.1
[2]   NEEDED   libA.so.1
[3]   RUNPATH  $ORIGIN/../lib:$ORIGIN/../ABC/lib

```

and similarly for the dependency `libX.so.1` this runpath would be:

```
% cc -o libX.so.1 -G -Kpic X.c '-R$ORIGIN:$ORIGIN/../../ABC/lib' \
-L/opt/ABC/lib -lY -lC
% dump -Lv libX.so.1
[1]   NEEDED  libY.so.1
[2]   NEEDED  libC.so.1
[3]   RUNPATH $ORIGIN:$ORIGIN/../../ABC/lib
```

If this product is now installed under `/usr/local/XYZ`, its post-install script would be required to establish a symbolic link of:

```
% ln -s ../../ABC /usr/local/XYZ/ABC
```

If the user's `PATH` is augmented with `/usr/local/XYZ/bin`, then invocation of the application `xyz` result in a pathname lookup for its dependencies as follows:

```
% ldd -s xyz
.....
find object=libX.so.1; required by xyz
  search path=$ORIGIN/../../lib:$ORIGIN/../../ABC/lib (RPATH from file xyz)
  trying path=/usr/local/XYZ/lib/libX.so.1
    libX.so.1 => /usr/local/XYZ/lib/libX.so.1

find object=libA.so.1; required by xyz
  search path=$ORIGIN/../../lib:$ORIGIN/../../ABC/lib (RPATH from file xyz)
  trying path=/usr/local/XYZ/lib/libA.so.1
  trying path=/usr/local/ABC/lib/libA.so.1
    libA.so.1 => /usr/local/ABC/lib/libA.so.1

find object=libY.so.1; required by /usr/local/XYZ/lib/libX.so.1
  search path=$ORIGIN:$ORIGIN/../../ABC/lib \
    (RPATH from file /usr/local/XYZ/lib/libX.so.1)
  trying path=/usr/local/XYZ/lib/libY.so.1
    libY.so.1 => /usr/local/XYZ/lib/libY.so.1

find object=libC.so.1; required by /usr/local/XYZ/lib/libX.so.1
  search path=$ORIGIN:$ORIGIN/../../ABC/lib \
    (RPATH from file /usr/local/XYZ/lib/libX.so.1)
  trying path=/usr/local/XYZ/lib/libC.so.1
  trying path=/usr/local/ABC/lib/libC.so.1
    libC.so.1 => /usr/local/ABC/lib/libC.so.1

find object=libB.so.1; required by /usr/local/ABC/lib/libA.so.1
  search path=$ORIGIN (RPATH from file /usr/local/ABC/lib/libA.so.1)
  trying path=/usr/local/ABC/lib/libB.so.1
    libB.so.1 => /usr/local/ABC/lib/libB.so.1
```

Security

In a secure process, the expansion of the `$ORIGIN` string is allowed only if it expands to a trusted directory. The occurrence of other relative path names, poses a security risk.

A path like `$ORIGIN/./lib` apparently points to a fixed location, fixed by the location of the executable. However, the location is not actually fixed. A writable directory in the same file system could exploit a secure program that uses `$ORIGIN`.

The following example shows this possible security breach if `$ORIGIN` was arbitrarily expanded within a secure process.

```
% cd /worldwritable/dir/in/same/fs
% mkdir bin lib
% ln $ORIGIN/bin/program bin/program
% cp ~/crooked-libc.so.1 lib/libc.so.1
% bin/program
..... using crooked-libc.so.1
```

You can use the utility `crle(1)` to specify trusted directories that enable secure applications to use `$ORIGIN`. Administrators who use this technique should ensure that the target directories are suitably protected from malicious intrusion.

New Linker and Libraries Features and Updates

This appendix provides an overview of new features and updates that have been added to the Solaris operating environment and indicates the release to which they were added:

Solaris 9 4/04 Release

- The new section types `SHT_SUNW_ANNOTATE`, `SHT_SUNW_DEBUGSTR`, `SHT_SUNW_DEBUG`, and `SHT_SPARC_GOTDATA` are supported. See Table 7–12.
- Greater control of direct bindings is provided with the link-editor options `-z direct` and `-z nodirect`, together with the `DIRECT` and `NODIRECT` `mapfile` directives. See “Defining Additional Symbols” on page 46, and “Direct Binding” on page 69.

Solaris 9 12/03 Release

- Performance improvements within `ld(1)` can significantly reduce the link-edit time of very large applications.

Solaris 9 8/03 Release

- `dlsym(3DL)` symbol processing can be reduced using a `dlopen(3DL)` handle that is created with the `RTLD_FIRST` flag. See “Obtaining New Symbols” on page 89.
- The signal used by the runtime linker to terminate an erroneous process can be managed using the `dldinfo(3DL)` flags `RTLD_DI_GETSIGNAL`, and `RTLD_DI_SETSIGNAL`.

Solaris 9 12/02 Release

- The link-editor provides string table compression, that can result in reduced `.dynstr` and `.strtab` sections. This default processing can be disabled using the link-editor's `-z nocompstrtab` option. See “String Table Compression” on page 56.
- The `-z ignore` option has been extended to eliminate unreferenced sections during a link-edit. See “Remove Unused Material” on page 117.
- Unreferenced dependencies can be determined using `ldd(1)`. See the `-U` option.
- The link-editors support extended ELF sections. See “ELF Header” on page 180, Table 7–12, “Sections” on page 187, Table 7–16 and “Symbol Table Section” on page 222.
- Greater flexibility in defining a symbols visibility is provided with the `protected` `mapfile` directive. See “Defining Additional Symbols” on page 46.

Solaris 9 Release

- Thread-Local Storage (TLS) support is provided. See Chapter 8.
- The `-z rescanner` option provides greater flexibility in specifying archive libraries to a link-edit. See “Position of an Archive on the Command Line” on page 32.
- The `-z ld32` and `-z ld64` options provide greater flexibility in using the link-editor support interfaces. See “32–Bit and 64–Bit Environments” on page 150.
- Additional link-editor support interfaces `ld_input_done()`, `ld_input_section()`, `ld_input_section64()` and `ld_version()` have been added. See “Support Interface Functions” on page 151.

- Environment variables interpreted by the runtime linker can now be established for multiple processes by specifying these variables within a configuration file. See the `-e` and `-E` options of `crle(1)`.
- Support for more than 32,768 procedure linkage table entries within 64-bit SPARC objects has been added. See “SPARC: 64-bit Procedure Linkage Table” on page 264.
- An `mdb(1)` debugger module enables you to inspect runtime linker data structures as part of process debugging. See “Debugger Module” on page 96.
- The `bss` segment declaration directive makes the creation of a `bss` segment easier. See “Segment Declarations” on page 294.

Solaris 8 07/01 Release

- Unused dependencies can be determined using `ldd(1)`. See the `-u` option.
- Various ELF ABI extensions have been added. See “Initialization and Termination Sections” on page 36, “Initialization and Termination Routines” on page 76, Table 7-4, Table 7-7, Table 7-14, Table 7-15, “Group Section” on page 204, Table 7-16, Table 7-24, Table 7-42, Table 7-43, and “Program Loading (Processor-Specific)” on page 241.
- Greater flexibility in the use of link-editor environment variables has been provided with the addition of `_32` and `_64` variants. See “Environment Variables” on page 23.

Solaris 8 01/01 Release

- The symbolic information available from `dladdr(3DL)` has been enhanced with the introduction of `dladdr1()`.
- The `$ORIGIN` of a dynamic object can be obtained from `dlinfo(3DL)`.
- The maintenance of runtime configuration files created with `crle(1)` has been simplified. Inspection of a configuration file displays the command-line options used to create the file. An update capability is provided with the `-u` option.
- The runtime linker and its debugger interface have been extended to detect procedure linkage table entry resolution. This update is identified by a new version number. See `rd_init()` under “Agent Manipulation Interfaces” on page 167. This update extends the `rd_plt_info_t` structure. See `rd_plt_resolution()` under “Procedure Linkage Table Skipping” on page 172.
- An application’s stack can be defined non-executable using the new `mapfile` segment descriptor `STACK`. See “Segment Declarations” on page 294.

Solaris 8 10/00 Release

- The environment variable `LD_BREADTH` is ignored by the runtime linker. See “Initialization and Termination Routines” on page 76.
- The runtime linker and its debugger interface have been extended for better runtime and core file analysis. This update is identified by a new version number. See `rd_init()` under “Agent Manipulation Interfaces” on page 167. This update extends the `rd_loadobj_t` structure. See “Scanning Loadable Objects” on page 168.
- You can now validate displacement relocated data in regard to its use, or possible use, with copy relocations. See “Displacement Relocations” on page 58.
- 64-bit filters can be built solely from a `mapfile` using the link-editor’s `-64` option. See “Generating Standard Filters” on page 108.
- The search paths used to locate the dependencies of dynamic objects can be inspected using `dlinfo(3DL)`.
- `dlsym(3DL)` and `dlinfo(3DL)` lookup semantics have been expanded with a new handle `RTLD_SELF`.
- The runtime symbol lookup mechanism used to relocate dynamic objects can be significantly reduced by establishing direct binding information within each dynamic object. See “Direct Binding” on page 69.

Solaris 8 Release

- The secure directory from which files can be preloaded is now `/usr/lib/secure` for 32-bit objects and `/usr/lib/secure/64` for 64-bit objects. See “Security” on page 80.
- Greater flexibility in modifying the runtime linker’s search paths can be achieved with the link-editor’s `-z nodefaultlib` option, and runtime configuration files created by the new utility `crle(1)`. See “Directories Searched by the Runtime Linker” on page 35 and “Configuring the Default Search Paths” on page 67.
- The new `EXTERN mapfile` directive enables you to use `-z defs` with externally defined symbols. See “Defining Additional Symbols” on page 46.
- The new `$ISALIST`, `$OSNAME`, and `$OSREL` dynamic string tokens provide greater flexibility in establishing instruction set specific, and system specific dependencies. See “Dynamic String Tokens” on page 67.
- The link-editor options `-p` and `-P` provide additional means of invoking runtime link auditing libraries. See “Recording Local Auditors” on page 157. The runtime link auditing interfaces `la_activity()` and `la_objsearch()` have been added.

See “Audit Interface Functions” on page 158.

- A new dynamic section tag, `DT_CHECKSUM`, enables you to coordinate ELF files with core images. See Table 7-42.

Solaris 7 Release

- The 64-bit ELF object format is now supported. See “File Format” on page 177 for details. Link-editor extensions and differences for 64-bit processing include the use of `/usr/lib/64` (see “Directories Searched by the Link-Editor” on page 33, “Directories Searched by the Runtime Linker” on page 35, and “Naming Conventions” on page 102), the environment variable `LD_LIBRARY_PATH_64` (see “Using an Environment Variable” on page 34, and “Directories Searched by the Runtime Linker” on page 64), and the runtime linker `/usr/lib/64/ld.so.1` (see Chapter 3).
- You can build shared objects with optimized relocation sections using the link-editor’s `-z combrelloc` option. See “Combined Relocation Sections” on page 121.
- The new `$ORIGIN` dynamic string token provides greater flexibility in establishing dependencies within unbundled software. See “Dynamic String Tokens” on page 67.
- The loading of a shared object can now be deferred until the object is actually referenced by the running program. See “Lazy Loading of Dynamic Dependencies” on page 114.
- The new `SHT_SUNW_COMDAT` section type enables the elimination of multiply-defined symbols. See “COMDAT Section” on page 203.
- The new `SHT_SUNW_move` section type enables partially initialized symbols. See “Move Section” on page 206.
- The runtime link auditing interfaces `la_symbind64()`, `la_sparcv9_pltenter()`, and `la_pltexit64()`, together with a new link-auditing flag `LA_SYMB_ALTVALUE`, have been added. See “Audit Interface Functions” on page 158.

Solaris 2.6 Release

- Weak symbol references can trigger archive member extraction by using the link-editor’s `-z weakextract` option. Extracting all archive members can be achieved using the `-z allextract` option. See “Archive Processing” on page 29.

- Shared objects specified as part of a link-edit that are not referenced by the object being built can be ignored, and hence their dependency recording suppressed, using the link-editor's `-z ignore` option. See "Shared Object Processing" on page 30.
- The link-editor generates the reserved symbols `_START_` and `_END_` to provide a means of establishing an object's address range. See "Generating the Output File" on page 57.
- Changes have been made to the runtime ordering of initialization and finalization code to better accommodate dependency requirements. See "Initialization and Termination Routines" on page 76.
- Symbol resolution semantics have been expanded for `dlopen(3DL)`. See "Symbol Lookup" on page 84, `RTLD_GROUP` in "Isolating a Group" on page 88, and `RTLD_PARENT` in "Object Hierarchies" on page 88.
- Symbol lookup semantics have been expanded with a new `dlsym(3DL)` handle `RTLD_DEFAULT`. See "Default Symbol Lookup Model" on page 85.
- Extensions have been made to filter processing that allow more than one filter to be defined, and provide for forcibly loading filters. See "Shared Objects as Filters" on page 107.
- You can record additional version dependencies using the `mapfile` file control directive `$ADDVERS`. See "Binding to Additional Version Definitions" on page 142.
- A runtime linker audit interface provides support for monitoring and modifying a dynamically linked application from within the process. See "Runtime Linker Auditing Interface" on page 155.
- A runtime linker debugger interface provides support for monitoring and modifying a dynamically linked application from an external process. See "Runtime Linker Debugger Interface" on page 164.
- Additional section information is supported. See Table 7-11 for `SHN_BEFORE` and `SHN_AFTER`. See Table 7-14 for `SHF_ORDERED` and `SHF_EXCLUDE`.
- A new dynamic section tag, `DT_1_FLAGS`, is supported. See Table 7-44 for the various flag values.
- A package of demonstration ELF programs is provided. See Chapter 7.
- The link-editors now support internationalized messages. All system errors are reported using `strerror(3C)`.
- The new `eliminate mapfile` directive, or the `-B eliminate` option, enable you to eliminate local symbol table entries. See "Symbol Elimination" on page 55.

Index

Numbers and Symbols

\$ADDVERS, *See* versioning
\$ISALIST, *See* search paths
\$ORIGIN, *See* search paths
\$OSNAME, *See* search paths
\$OSREL, *See* search paths
\$PLATFORM, *See* search paths
__thread, 271
32-bit/64-bit, 33, 35, 63, 64, 65, 67, 80, 82, 102,
105, 116, 150, 157, 164, 179, 180, 185, 187, 210,
230, 309
introduction, 23

A

ABI, *See* Application Binary Interface
Application Binary Interface, 22, 109, 129
ar(1), 29
archives, 31
inclusion of shared objects in, 104
link-editor processing, 29
multiple passes through, 29
naming conventions, 31
as(1), 20
atexit(3C), 77
auxiliary filters, 107, 110

B

base address, 239, 240
binding, 19

binding (Continued)

dependency ordering, 106
direct, 69, 70, 121
lazy, 71, 83, 95
to shared object dependencies, 103, 136
to version definitions, 136
to weak version definitions, 143

C

CC(1), 27
cc(1), 20, 27
COMDAT, 153, 204
COMMON, 38, 48, 51, 188
compilation environment, 21, 31, 101
See also link-editing and link-editor
compiler driver, 27
compiler options
-K PIC, 116
-K pic, 115, 308
-xF, 117, 204, 297
-xpg, 126
-xregs=no%appl, 309
crle(1), 67, 80, 125, 159, 259, 260, 331, 332
crle(1) options
-E, 331
-e, 125, 331
-l, 67
-s, 80
-u, 331

D

- data representation, 179
- debugging aids
 - link-editing, 60
 - runtime linking, 93
- demonstrations
 - prefcnt, 163
 - sotruss, 163
 - symbindrep, 163
 - whocalls, 163
- dependency
 - groups, 83, 84
- dependency ordering, 106
- direct binding, 69, 70, 121
- dladdr(3DL), 331
- dladdr1(3DL), 331
- dlclose(3DL), 77, 81
- dldump(3DL), 37
- dlerror(3DL), 81
- dlfcn.h, 81
- dlinfo(3DL), 330, 331, 332
- dlopen(3DL), 156
 - See also* dlopen(3DL)
- dlopen(3DL), 64, 81, 82, 88, 111, 139, 149, 150
 - effects of ordering, 87
 - group, 84
- dlopen(3DL)
 - group, 83
- dlopen(3DL)
 - modes
 - RTLD_GLOBAL, 83, 88
 - RTLD_GROUP, 88
 - RTLD_LAZY, 83
 - RTLD_NOLOAD, 156
 - RTLD_NOW, 72, 79, 83
 - RTLD_PARENT, 88, 89
 - of a dynamic executable, 88
- dlopen(3DL)
 - of a dynamic executable, 83
- dlopen(3DL)
 - shared object naming conventions, 102
- dlsym(3DL), 64, 81, 89, 92, 140, 150
 - special handle
 - RTLD_DEFAULT, 45, 89
 - RTLD_NEXT, 89
 - RTLD_SELF, 332
- dump(1), 23, 65, 68, 113, 115
- dynamic executables, 20, 21

dynamic information tags

- NEEDED, 65, 103
- RUNPATH, 65
- SONAME, 103
- SYMBOLIC, 125
- TEXTREL, 115

dynamic linking, 22

- implementation, 210, 244

E

ELF, 19, 25, 102, 112, 149

- See also* object files

elf(3E), 23, 149

environment variables, 23

- LD_AUDIT, 80, 157
- LD_BIND_NOT, 95
- LD_BIND_NOW, 72, 79, 95, 264, 267, 269
- LD_BREADTH, 78
- LD_CONFIG, 80
- LD_DEBUG, 93
- LD_DEBUG_OUTPUT, 94
- LD_LIBRARY_PATH, 34, 66, 80, 82, 106, 157
- LD_LOADFLTR, 111
- LD_NOAUDIT, 158
- LD_NOAUXFLTR, 111
- LD_NODIRECT, 70
- LD_NOLAZYLOAD, 76
- LD_NOVERSION, 139, 142
- LD_OPTIONS, 27, 61
- LD_PRELOAD, 70, 73, 80
- LD_PROFILE, 125
- LD_PROFILE_OUTPUT, 126
- LD_RUN_PATH, 36
- LD_SIGNAL, 80
- SGS_SUPPORT, 150

error messages

link-editor

- illegal argument to option, 28
- illegal option, 28
- incompatible options, 28
- multiple instances of an option, 28
- multiply-defined symbols, 42
- relocations against non-writable sections, 115
- shared object name conflicts, 105
- soname conflicts, 105

- error messages, link-editor (Continued)
 - symbol not assigned to version, 54
 - symbol warnings, 41
 - undefined symbols, 43
 - undefined symbols from an implicit reference, 44
 - version unavailable, 141
- runtime linker
 - copy relocation size differences, 59, 124
 - relocation errors, 72, 139
 - unable to find shared object, 66, 83
 - unable to find version definition, 138
 - unable to locate symbol, 90
- exec(2), 25, 63, 178
- executable and linking format, *See* ELF

F

- filtee, 107
- filters, 107
 - auxiliary, 107, 110
 - instruction set specific, 321
 - reducing filtee searches, 322
 - standard, 107, 108
 - system specific, 323

G

- generating a shared object, 44
- generating an executable, 43
- generating the output file image, 57
- global offset table, 58, 68, 115, 201, 214, 248, 252, 260
 - SPARC, 217, 218
 - x86, 220, 221, 267
- global symbols, 38, 129, 224, 227
- .got, *See* global offset table

I

- initialization and termination, 27, 36, 76
- input file processing, 28
- interface
 - private, 129
 - public, 129, 311

- interposition, 40, 52, 69, 74, 91, 130
- interpreter, *See* runtime linker

L

- lazy binding, 71, 83, 95, 155
- ld(1), 19
- LD_AUDIT, 80, 157
- LD_BIND_NOT, 95
- LD_BIND_NOW, 72, 79, 95, 264, 267, 269
- LD_BREADTH, 78
- LD_CONFIG, 80
- LD_DEBUG, 93
- LD_DEBUG_OUTPUT, 94
- LD_LIBRARY_PATH, 66, 80, 82, 106, 157
- LD_LOADFLTR, 111
- LD_NOAUDIT, 158
- LD_NOAUXFLTR, 111
- LD_NODIRECT, 70
- LD_NOLAZYLOAD, 76
- LD_NOVERSION, 139, 142
- LD_OPTIONS, 27, 61
- LD_PRELOAD, 70, 73, 80
- LD_PROFILE, 125
- LD_PROFILE_OUTPUT, 126
- LD_RUN_PATH, 36
- LD_SIGNAL, 80
- ld.so.1(1), *See* runtime linker
- ldd(1), 23, 65, 66, 69, 73, 111, 138, 139
- ldd(1) options
 - d, 59, 73, 124
 - i, 78
 - r, 59, 73, 124
 - u, 30
 - v, 138
- libdl.so.1, 81
- libelf.so.1, 151, 177
- libldstab.so.1, 150
- libraries
 - archives, 31
 - naming conventions, 31
 - shared, 210, 244
- link-editing, 20, 222, 244
 - adding additional libraries, 31
 - archive processing, 29
 - binding to a version definition, 136, 140
 - dynamic, 210, 244

link-editing (Continued)

- input file processing, 28
- library input processing, 29
- library linking options, 29
- mixing shared objects and archives, 32
- position of files on command line, 32
- search paths, 33
- shared object processing, 30

link-editor, 19, 25

- debugging aids, 60
- error messages
 - See error messages
- external bindings, 56
- invoking directly, 26
- invoking using compiler driver, 27
- overview, 25
- sections, 25
- segments, 25
- specifying options, 27

link-editor options

- 64, 23, 109
- a, 308
- B direct, 309, 310
- B dynamic, 32
- B eliminate, 56
- B group, 84, 88, 258
- B local, 54
- B reduce, 49, 55
- B static, 32, 308
- D, 60
- d n, 307, 310
- d y, 308
- e, 58
- F, 107
- f, 107
- G, 101, 308, 310
- h, 65, 103, 147, 309
- i, 34
- L, 33, 307
- l, 29, 31, 102, 146, 307
- M, 26, 46, 47, 130, 131, 140, 293, 309, 313
- m, 31, 40
- P, 157
- p, 157
- R, 35, 105, 309, 310
- r, 27, 308
- S, 150
- s, 55, 57

link-editor options (Continued)

- t, 41, 42
- u, 46, 47
- Y, 34
- z alleextract, 29
- z combrelloc, 309
- z defaultextract, 29
- z defs, 44, 48, 156, 309
- z direct, 329
- z endfiltee, 259
- z finiarray, 36
- z groupper, 260
- z ignore, 30, 117, 309
- z initarray, 36
- z initfirst, 258
- z interpose, 70, 259
- z lazyload, 75, 260, 309, 310
- z ld32, 150
- z ld64, 150
- z loadfltr, 111, 258
- z muldefs, 42
- z nocompstrtab, 56, 330
- z nodefaultlib, 35, 259
- z nodefs, 43, 72
- z nodelete, 258
- z nodirect, 329
- z nodlopen, 258
- z nodump, 259
- z nolazyload, 75
- z nopartial, 208
- z noversion, 54, 132, 138
- z now, 72, 79, 83
- z record, 118
- z rescan, 33
- z text, 115, 308
- z verbose, 59
- z weakextract, 29, 225

link-editor output

- dynamic executables, 20
- relocatable objects, 20
- shared objects, 20
- static executables, 20

link-editor support interface (ld-support), 149

- ld_atexit(), 153
- ld_atexit64(), 153
- ld_file(), 151
- ld_file64(), 151
- ld_input_done(), 153

link-editor support interface (ld-support)
(Continued)

- `ld_input_section()`, 152
- `ld_input_section64()`, 152
- `ld_section()`, 152
- `ld_section64()`, 152
- `ld_start()`, 151
- `ld_start64()`, 151
- `ld_version()`, 151

linker-editor, direct binding, 70

local symbols, 38, 224, 227

`lorder(1)`, 30, 61

M

mapfiles, 293

- defaults, 302

- example, 300

- map structure, 303

- mapping directives, 298

- segment declarations, 294

- size-symbol declarations, 300

- structure, 293

- syntax, 293

`mdb(1)`, 331

`mmap(2)`, 25, 57, 63, 112

multiply-defined data, 119, 203

multiply-defined symbols, 30, 40, 203

N

Namespace, 156

naming conventions

- archives, 31

- libraries, 31

- shared objects, 31, 102

NEEDED, 65, 103

`nm(1)`, 23, 113

O

object files, 19

- base address, 239, 240

- data representation, 179

object files (Continued)

- global offset table

 - See* global offset table

- note section, 208, 210

- preloading at runtime, 73

- procedure linkage table

 - See* procedure linkage table

- program header, 236, 239

- program interpreter, 247

- program loading, 241

- relocation, 210, 260

- section alignment, 190

- section attributes, 195, 203

- section group flags, 204

- section header, 187, 203

- section names, 203

- section types, 190, 203

- segment contents, 241

- segment permissions, 240, 241

- segment types, 237, 240

- string table, 221, 222

- symbol table, 222, 229

P

packages

- SUNWosdem, 163, 166, 177

- SUNWtoo, 163

paging, 241, 244

performance

- allocating buffers dynamically, 119

- collapsing multiple definitions, 119

- improving locality of references, 120, 125

- maximizing shareability, 118

- minimizing data segment, 118

- position-independent code

 - See* position-dependent code

- relocations, 120, 125

- the underlying system, 114

- using automatic variables, 119

PIC, *See* position-independent code

`.plt`, *See* procedure linkage table

position-independent code, 115, 253, 260

preloading objects, *See* LD_PRELOAD

procedure linkage table, 58, 71, 115, 201, 214,

- 248, 252, 253, 261

- 64-bit SPARC, 264

procedure linkage table (Continued)

- SPARC, 217, 218, 261, 264
- x86, 220, 221, 267
- profil(2), 126
- program interpreter, 247
 - See also* runtime linker
- pvs(1), 23, 132, 134, 136, 137

R

- relocatable objects, 20
- relocation, 67, 120, 124, 210
 - copy, 58, 122
 - displacement, 58
 - immediate, 71
 - lazy, 71
 - non-symbolic, 68, 120
 - runtime linker
 - symbol lookup, 68, 71, 83, 95
 - symbolic, 68, 120
- RTLD_DEFAULT, 45
 - See also* dependency ordering
- RTLD_GLOBAL, 83, 88
- RTLD_GROUP, 88
- RTLD_LAZY, 83
- RTLD_NEXT, 89
- RTLD_NOLOAD, 156
- RTLD_NOW, 72, 79, 83
- RTLD_PARENT, 88, 89
- runpath, 35, 65, 82, 105
- RUNPATH, *See* runpath
- runpath, security, 80
- runtime environment, 21, 31, 101
- runtime linker, 21, 63, 247
 - direct binding, 69, 70, 121
 - initialization and termination routines, 76
 - lazy binding, 71, 83, 95
 - link-maps, 156
 - loading additional objects, 73
 - namespace, 156
 - programming interface
 - See also* dlclose(3DL), dldump(3DL),
dlderror(3DL), dlmopen(3DL), and
dlopen(3DL)
 - relocation processing, 67
 - search paths, 35, 64
 - security, 80

runtime linker (Continued)

- shared object processing, 64
- version definition verification, 138
- runtime linker support interfaces
 - (rtld-audit), 149, 155
 - la_activity(), 159
 - la_i86_pltenter(), 161
 - la_objclose(), 162
 - la_objopen(), 159
 - la_objseach(), 159
 - la_pltexit(), 162
 - la_preinit(), 160
 - la_sparcv8_pltenter(), 161
 - la_sparcv9_pltenter(), 161
 - la_symbind32(), 160
 - la_symbind64(), 160
 - la_version(), 158
- runtime linker support interfaces
 - (rtld-debugger), 149, 164
 - ps_global_sym(), 175
 - ps_pglobal_sym(), 176
 - ps_plog(), 175
 - ps_pread(), 175
 - ps_pwrite(), 175
 - rd_delete(), 168
 - rd_errstr(), 168
 - rd_event_addr(), 171
 - rd_event_enable(), 171
 - rd_event_getmsg(), 172
 - rd_init(), 167
 - rd_loadobj_iter(), 170
 - rd_log(), 168
 - rd_new(), 167
 - rd_objpad_enable(), 174
 - rd_plt_resolution(), 173
 - rd_reset(), 167
- runtime linking, 21

S

- SCD, *See* Application Binary Interface
- search paths
 - link-editing, 33
 - runtime linker, 35, 64
 - \$ISALIST token, 321
 - \$ORIGIN token, 323
 - \$OSNAME token, 323

search paths, runtime linker (Continued)

 \$OSREL token, 323

 \$PLATFORM token, 323

section flags, 195

 SHF_ALLOC, 196, 202

 SHF_EXCLUDE, 153, 198

 SHF_EXECINSTR, 196

 SHF_GROUP, 197, 204

 SHF_INFO_LINK, 196

 SHF_LINK_ORDER, 188, 197

 SHF_MASKOS, 197

 SHF_MASKPROC, 198

 SHF_MERGE, 196

 SHF_ORDERED, 197

 SHF_OS_NONCONFORMING, 197

 SHF_STRINGS, 196

 SHF_WRITE, 196

 SHT_TLS, 197

section names

 .bss, 25, 122

 .data, 25, 118

 .dynamic, 57, 63, 125

 .dynstr, 57

 .dynsym, 57

 .fini, 36, 76

 .fini_array, 36, 76

 .got, 58, 68

 .init, 36, 76

 .init_array, 36, 76

 .interp, 63

 .picdata, 119

 .plt, 58, 71, 125

 .preinit_array, 36, 76

 .rela.text, 25

 .rodata, 118

 .strtab, 25, 57

 .SUNW_reloc, 122, 309

 .SUNW_version, 231

 .symtab, 25, 55, 57

 .tbss, 273

 .tdata, 273

 .tdata1, 273

 .text, 25

section numbers, 187, 273

 SHN_ABS, 188, 218, 226, 228

 SHN_AFTER, 188, 197

 SHN_BEFORE, 188, 197

 SHN_COMMON, 188, 224, 228

section numbers (Continued)

 SHN_HIOS, 188

 SHN_HIPROC, 188

 SHN_HIRESERVE, 189

 SHN_LOOS, 188

 SHN_LOPROC, 188

 SHN_LORESERVE, 188

 SHN_UNDEF, 183, 188, 195, 205, 228

 SHN_XINDEX, 189

section types, 192

 SHF_TLS, 272

 SHT_DYNAMIC, 192, 248

 SHT_DYNSTR, 192

 SHT_DYNSYM, 192

 SHT_FINI_ARRAY, 193

 SHT_GROUP, 193, 197, 204

 SHT_HASH, 192, 205, 248

 SHT_HIOS, 194

 SHT_HIPROC, 194

 SHT_HISUNW, 194

 SHT_HIUSER, 195, 305

 SHT_INIT_ARRAY, 193

 SHT_LOOS, 194

 SHT_LOPROC, 194

 SHT_LOSUNW, 194

 SHT_LOUSER, 194, 305

 SHT_NOBITS, 193, 200, 203, 241

 sh_offset, 190

 sh_size, 190

 SHT_NOTE, 193, 208

 SHT_NULL, 192

 SHT_PREINIT_ARRAY, 193

 SHT_PROGBITS, 192, 248

 SHT_REL, 193

 SHT_RELA, 192

 SHT_SHLIB, 193

 SHT_SPARC_GOTDATA, 194

 SHT_STRTAB, 192

 SHT_SUNW_ANNOTATE, 194

 SHT_SUNW_COMDAT, 153, 194, 204

 SHT_SUNW_DEBUG, 194

 SHT_SUNW_DEBUGSTR, 194

 SHT_SUNW_move, 194, 206

 SHT_SUNW_syminfo, 194

 SHT_SUNW_verdef, 194, 231, 234

 SHT_SUNW_verneed, 194, 231, 234

 SHT_SUNW_versym, 194, 231, 232, 233

 SHT_SYMTAB, 192, 226

- section types (Continued)
 - SHT_SYMTAB_SHNDX, 193
- sections, 25, 112
 - See also* section flags, section names, section numbers and section types
- security, 80, 326
- segments, 25, 112
 - data, 113, 114
 - text, 113, 114
- SGS_SUPPORT, 150
- shared libraries, *See* shared objects
- shared objects, 19, 20, 21, 64, 101
 - as filters, 107
 - dependency ordering, 106
 - explicit definition, 44
 - implementation, 210, 244
 - implicit definition, 44
 - link-editor processing, 30
 - naming conventions, 31, 102
 - recording a runtime name, 103
 - with dependencies, 105
- size(1), 112
- Solaris ABI, *See* Application Binary Interface
- Solaris Application Binary Interface, *See* Application Binary Interface
- SONAME, 103
- SPARC Compliance Definition, *See* Application Binary Interface
- standard filters, 107, 108
- static executables, 20
- strings(1), 119
- strip(1), 55, 57
- SUNWosdem, 163, 166, 177
- SUNWtoo, 163
- support interfaces
 - link-editor (ld-support), 149
 - runtime linker (rtld-audit), 149, 155
 - runtime linker (rtld-debugger), 149, 164
- symbol reserved names, 57
 - _DYNAMIC, 57
 - _edata, 57
 - _end, 57
 - _END_, 57
 - _etext, 57
 - _fini, 36
 - _GLOBAL_OFFSET_TABLE_, 58, 116, 261
 - _init, 36
 - main, 58
- symbol reserved names (Continued)
 - _PROCEDURE_LINKAGE_TABLE_, 58
 - _start, 58
 - _START_, 58
- symbol resolution, 38, 57
 - complex, 41
 - fatal, 42
 - interposition, 69
 - multiple definitions, 30
 - search scope
 - group, 84
 - simple, 39
 - symbol visibility, 224, 226
 - global, 84
 - local, 84
- SYMBOLIC, 125
- symbols
 - absolute, 48, 188
 - archive extraction, 29
 - auto-elimination, 55
 - auto-reduction, 48, 132, 314
 - COMMON, 38, 48, 51, 188
 - defined, 38
 - definition, 29, 42
 - elimination, 55
 - existence test, 45
 - global, 38, 39, 129, 224, 227
 - local, 38, 224, 227
 - multiply-defined, 30, 40, 203
 - ordered, 188
 - private interface, 129
 - public interface, 129
 - reference, 29, 42
 - registers, 229
 - runtime lookup, 84, 92
 - deferred, 71, 83, 95
 - scope, 84, 88
 - symbol visibility, 84
 - tentative, 29, 38, 46, 48, 51, 188
 - ordering in the output file, 46
 - realignment, 51
 - type, 225
 - undefined, 29, 38, 42, 43, 44, 188
 - visibility, 224, 226
 - weak, 29, 39, 45, 224, 227
- System V Application Binary Interface, 312

System V Application Binary Interface
(Continued)
See Application Binary Interface

W
weak symbols, 39, 224, 227
undefined, 29, 45

T

tentative symbols, 29, 38, 48, 51
TEXTREL, 115
Thread Local Storage, 330
thread-local storage, 271
 section definition, 272
TLS, *See* thread-local storage
tsort(1), 30, 61

U

undefined symbols, 42
/usr/ccs/bin/ld, *See* link-editor
/usr/ccs/lib, 33
/usr/lib, 33, 35, 64, 65, 82
/usr/lib/64, 33, 35, 64, 65, 82
/usr/lib/64/ld.so.1, 63, 164
/usr/lib/ld.so.1, 63, 164
/usr/lib/secure, 80, 157
/usr/lib/secure/64, 80, 157

V

versioning, 129
 base version definition, 132
 binding to a definition, 136, 140
 \$ADDVERS, 140
 defining a public interface, 54, 131
 definitions, 130, 131, 136
 file control directive, 140
 file name, 131, 315
 generating definitions within an image, 47,
 54, 131
 normalization, 138
 overview, 129
 runtime verification, 138, 139
virtual addressing, 241

