

ARM[®] Compiler toolchain v5.02 for μVision

Using the Linker



ARM Compiler toolchain v5.02 for μ Vision

Using the Linker

Copyright © 2007-2008, 2011-2012 ARM. All rights reserved.

Release Information

The following changes have been made to this book.

Change History

Date	Issue	Confidentiality	Change
May 2007	A	Non-Confidential	Release for RVCT v3.1 for μ Vision
December 2008	B	Non-Confidential	Release for RVCT v4.0 for μ Vision
June 2011	C	Non-Confidential	Release for ARM Compiler toolchain v4.1 for μ Vision
July 2012	D	Non-Confidential	Release for ARM Compiler toolchain v5.02 for μ Vision

Proprietary Notice

Words and logos marked with [®] or [™] are registered trademarks or trademarks of ARM in the EU and other countries, except as otherwise stated below in this proprietary notice. Other brands and names mentioned herein may be the trademarks of their respective owners.

Neither the whole nor any part of the information contained in, or the product described in, this document may be adapted or reproduced in any material form except with the prior written permission of the copyright holder.

The product described in this document is subject to continuous developments and improvements. All particulars of the product and its use contained in this document are given by ARM in good faith. However, all warranties implied or expressed, including but not limited to implied warranties of merchantability, or fitness for purpose, are excluded.

This document is intended only to assist the reader in the use of the product. ARM shall not be liable for any loss or damage arising from the use of any information in this document, or any error or omission in such information, or any incorrect use of the product.

Where the term ARM is used it means “ARM or any of its subsidiaries as appropriate”.

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by ARM and the party that ARM delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

Web Address

<http://www.arm.com>

Contents

ARM Compiler toolchain v5.02 for μ Vision Using the Linker

Chapter 1	Conventions and feedback	
Chapter 2	Overview of the linker	
	2.1 About the linker	2-2
	2.2 Linker command-line syntax	2-3
	2.3 Linker command-line options listed in groups	2-4
	2.4 What the linker can accept as input	2-8
	2.5 What the linker outputs	2-9
	2.6 What the linker does when constructing an executable image	2-10
Chapter 3	Linking models supported by armlink	
	3.1 Overview of linking models	3-2
	3.2 Bare-metal linking model	3-3
	3.3 Partial linking model	3-4
Chapter 4	Image structure and generation	
	4.1 The image structure	4-3
	4.2 Input sections, output sections, regions, and Program Segments	4-5
	4.3 Load view and execution view of an image	4-6
	4.4 Methods of specifying an image memory map with the linker	4-8
	4.5 Types of simple image	4-10
	4.6 Type 1 image, one load region and contiguous execution regions	4-11
	4.7 Type 2 image, one load region and non-contiguous execution regions	4-13
	4.8 Type 3 image, two load regions and non-contiguous execution regions	4-15
	4.9 Image entry points	4-17
	4.10 About specifying an initial entry point	4-18

4.11	Section placement with the linker	4-19
4.12	Placing sections with FIRST and LAST attributes	4-21
4.13	Section alignment with the linker	4-22
4.14	Demand paging	4-23
4.15	About ordering execution regions containing Thumb code	4-25
4.16	Overview of veneers	4-26
4.17	Veneer sharing	4-27
4.18	Veneer types	4-28
4.19	Generation of position independent to absolute veneers	4-29
4.20	Reuse of veneers when scatter-loading	4-30
4.21	Using command-line options to control the generation of C++ exception tables ...	4-31
4.22	About weak references and definitions	4-32
4.23	How the linker performs library searching, selection, and scanning	4-35
4.24	Controlling how the linker searches for the ARM standard libraries	4-36
4.25	Specifying user libraries when linking	4-38
4.26	How the linker resolves references	4-39
4.27	Use of the strict family of options in the linker	4-41
Chapter 5	Using linker optimizations	
5.1	Elimination of common debug sections	5-2
5.2	Elimination of common groups or sections	5-3
5.3	Elimination of unused sections	5-4
5.4	Elimination of unused virtual functions	5-6
5.5	About linker feedback	5-7
5.6	Example of using linker feedback	5-9
5.7	About link-time code generation	5-11
5.8	Optimization with RW data compression	5-13
5.9	How the linker chooses a compressor	5-14
5.10	Overriding the compression algorithm used by the linker	5-15
5.11	How compression is applied	5-16
5.12	Working with RW data compression	5-17
5.13	Inlining functions with the linker	5-18
5.14	Factors that influence function inlining	5-19
5.15	Handling branches that optimize to a NOP	5-21
5.16	About reordering of tail calling sections	5-22
5.17	Restrictions on reordering of tail calling sections	5-23
5.18	About merging comment sections	5-24
Chapter 6	Getting information about images	
6.1	Linker options for getting information about images	6-2
6.2	Identifying the source of some link errors	6-3
6.3	Example of using the --info linker option	6-4
6.4	How to find where a symbol is placed when linking	6-6
Chapter 7	Accessing and managing symbols with armlink	
7.1	About mapping symbols	7-3
7.2	Accessing linker-defined symbols	7-4
7.3	Region-related symbols	7-5
7.4	Image\$\$ execution region symbols	7-6
7.5	Load\$\$ execution region symbols	7-7
7.6	Load\$\$LR\$\$ load region symbols	7-9
7.7	Region name values when not scatter-loading	7-10
7.8	Using scatter files	7-11
7.9	Importing linker-defined symbols in C and C++	7-12
7.10	Importing linker-defined symbols in ARM assembler	7-13
7.11	Section-related symbols	7-14
7.12	Image symbols	7-15
7.13	Input section symbols	7-16
7.14	Accessing symbols in another image	7-17
7.15	Creating a symdefs file	7-18

7.16	Outputting a subset of the global symbols	7-19
7.17	Reading a symdefs file	7-20
7.18	Symdefs file format	7-21
7.19	What is a steering file?	7-23
7.20	Specifying steering files on the linker command-line	7-24
7.21	Steering file command summary	7-25
7.22	Steering file format	7-26
7.23	Hiding and renaming global symbols with a steering file	7-27
7.24	Using \$Super\$\$ and \$Sub\$\$ to patch symbol definitions	7-28

Chapter 8

Using scatter files

8.1	About scatter-loading	8-3
8.2	When to use scatter-loading	8-4
8.3	Scatter-loading command-line option	8-5
8.4	Images with a simple memory map	8-6
8.5	Images with a complex memory map	8-8
8.6	Linker-defined symbols that are not defined when scatter-loading	8-10
8.7	Specifying stack and heap using the scatter file	8-11
8.8	What is a root region?	8-12
8.9	Creating root execution regions	8-13
8.10	Using the FIXED attribute to create root regions	8-16
8.11	Placing functions and data at specific addresses	8-17
8.12	Placing a named section explicitly using scatter-loading	8-22
8.13	Placing unassigned sections with the .ANY module selector	8-23
8.14	Examples of using placement algorithms for .ANY sections	8-26
8.15	Example of next_fit algorithm showing behavior of full regions, selectors, and priority .. 8-28	
8.16	Examples of using sorting algorithms for .ANY sections	8-30
8.17	Selecting veneer input sections in scatter-loading descriptions	8-32
8.18	Using __attribute__((section("name"))) to place code and data	8-33
8.19	Using __at sections to place sections at a specific address	8-35
8.20	Restrictions on placing __at sections	8-36
8.21	Automatic placement of __at sections	8-37
8.22	Manual placement of __at sections	8-39
8.23	Placing a key in flash memory using __at	8-41
8.24	Placing a structure over a peripheral register using __at	8-43
8.25	Placement of sections with overlays	8-44
8.26	About placing ARM C and C++ library code	8-47
8.27	Example of placing code in a root region	8-48
8.28	Example of placing ARM C library code	8-49
8.29	Example of placing ARM C++ library code	8-50
8.30	Example of placing ARM library helper functions	8-51
8.31	Reserving an empty region	8-52
8.32	About creating regions on page boundaries	8-54
8.33	Overalignment of execution regions and input sections	8-56
8.34	Using preprocessing commands in a scatter file	8-57
8.35	Expression evaluation in scatter files	8-58
8.36	Using expression evaluation in a scatter file to avoid padding	8-59
8.37	Equivalent scatter-loading descriptions for simple images	8-60
8.38	Type 1 image, one load region and contiguous execution regions	8-61
8.39	Type 2 image, one load region and non-contiguous execution regions	8-63
8.40	Type 3 image, two load regions and non-contiguous execution regions	8-65
8.41	Scatter file to ELF mapping	8-67

Chapter 1

Conventions and feedback

The following describes the typographical conventions and how to give feedback:

Typographical conventions

The following typographical conventions are used:

`monospace` Denotes text that can be entered at the keyboard, such as commands, file and program names, and source code.

monospace Denotes a permitted abbreviation for a command or option. The underlined text can be entered instead of the full command or option name.

monospace italic

Denotes arguments to commands and functions where the argument is to be replaced by a specific value.

`monospace bold`

Denotes language keywords when used outside example code.

italic Highlights important notes, introduces special terminology, denotes internal cross-references, and citations.

bold Highlights interface elements, such as menu names. Also used for emphasis in descriptive lists, where appropriate, and for ARM[®] processor signal names.

Feedback on this product

If you have any comments and suggestions about this product, contact your supplier and give:

- your name and company

- the serial number of the product
- details of the release you are using
- details of the platform you are using, such as the hardware platform, operating system type and version
- a small standalone sample of code that reproduces the problem
- a clear explanation of what you expected to happen, and what actually happened
- the commands you used, including any command-line options
- sample output illustrating the problem
- the version string of the tools, including the version number and build numbers.

Feedback on content

If you have comments on content then send an e-mail to errata@arm.com. Give:

- the title
- the number, ARM DUI 0377D
- if viewing online, the topic names to which your comments apply
- if viewing a PDF version of a document, the page numbers to which your comments apply
- a concise explanation of your comments.

ARM also welcomes general suggestions for additions and improvements.

ARM periodically provides updates and corrections to its documentation on the ARM Information Center, together with knowledge articles and *Frequently Asked Questions* (FAQs).

Other information

- ARM Product Manuals, http://www.keil.com/support/man_arm.htm
- Keil Support Knowledgebase, <http://www.keil.com/support/knowledgebase.asp>
- Keil Product Support, <http://www.keil.com/support/>
- ARM Glossary, <http://infocenter.arm.com/help/topic/com.arm.doc.aeg0014-/index.html>.

Chapter 2

Overview of the linker

The following topics give an overview of the ARM linker, arm1ink:

Concepts

- *About the linker on page 2-2*
- *What the linker can accept as input on page 2-8*
- *What the linker outputs on page 2-9*
- *What the linker does when constructing an executable image on page 2-10.*

Reference

- *Linker command-line syntax on page 2-3*
- *Linker command-line options listed in groups on page 2-4.*

2.1 About the linker

The linker, `arm1ink`, combines the contents of one or more object files with selected parts of one or more object libraries to produce:

- an ARM ELF image
- a partially linked ELF object that can be used as input in a subsequent link step
- ELF files that can be demand-paged efficiently.

The linker can:

- link ARM code and 16-bit Thumb[®] and 32-bit Thumb code
- generate interworking veneers to switch processor state when required
- generate inline veneers or long branch veneers, where required, to extend the range of branch instructions
- automatically select the appropriate standard C or C++ library variants to link with, based on the build attributes of the objects it is linking
- enable you to specify the locations of code and data within the system memory map, using either a command-line option or a scatter file
- perform Read/Write data compression to minimize ROM size
- perform unused section elimination to reduce the size of your output image
- control the generation of debug information in the output file
- generate a static callgraph and list the stack usage
- control the contents of the symbol table in output images
- show the sizes of code and data in the output
- use linker feedback to remove individual unused functions.

2.1.1 See also

Concepts

- [Demand paging on page 4-23](#)
- [About linker feedback on page 5-7](#)
- [Chapter 3 Linking models supported by armlink](#)
- [Chapter 4 Image structure and generation](#)
- [Chapter 5 Using linker optimizations](#)
- [Chapter 6 Getting information about images](#)
- [Chapter 7 Accessing and managing symbols with armlink](#)
- [Chapter 8 Using scatter files.](#)

2.2 Linker command-line syntax

The command for invoking the linker is:

```
armlink [options] [input-file-list]
```

options Linker command-line options.

input-file-list

A space-separated list of objects, libraries, or *symbol definitions* (symdefs) files.

2.2.1 See also

Reference

- [Linker command-line options listed in groups on page 2-4.](#)

Linker Reference:

- [input-file-list on page 2-65](#)
- [Chapter 2 Linker command-line options.](#)

2.3 Linker command-line options listed in groups

See the following command-line options in the *Linker Reference*:

Controlling library files and paths

- `--libpath=pathlist` on page 2-74
- `--library_type=lib` on page 2-75
- `--reduce_paths, --no_reduce_paths` on page 2-98
- `--scanlib, --no_scanlib` on page 2-108
- `--thumb2_library, --no_thumb2_library` on page 2-130
- `--userlibpath=pathlist` on page 2-136.

Controlling the linking of object files

- `--match=crossmangled` on page 2-84
- `--strict` on page 2-119
- `--strict_ph, --no_strict_ph` on page 2-122
- `--strict_relocations, --no_strict_relocations` on page 2-123
- `--unresolved=symbol` on page 2-134.

Controlling the output

- `--combreloc, --no_combreloc` on page 2-23
- `--ldpartial` on page 2-72
- `--output=filename` on page 2-89
- `--partial` on page 2-94
- `--reloc` on page 2-100.

Specifying the image memory map

- `--autoat, --no_autoat` on page 2-11
- `--predefine="string"` on page 2-96
- `--ro_base=address` on page 2-103
- `--ropi` on page 2-104
- `--rosplit` on page 2-105
- `--rw_base=address` on page 2-106
- `--rwp_i` on page 2-107
- `--scatter=filename` on page 2-109
- `--split` on page 2-117
- `--zi_base=address` on page 2-149.

Controlling debug information in an image

- `--bestdebug, --no_bestdebug` on page 2-14
- `--compress_debug, --no_compress_debug` on page 2-25
- `--debug, --no_debug` on page 2-31
- `--emit_debug_overlay_relocs` on page 2-39
- `--emit_debug_overlay_section` on page 2-40
- `--emit_non_debug_relocs` on page 2-41

Controlling the content of an image

- `--any_contingency` on page 2-5
- `--any_placement=algorithm` on page 2-6
- `--any_sort_order=order` on page 2-8
- `--api, --no_api` on page 2-9

- [--arm_only](#) on page 2-10
- [--branchnop](#), [--no_branchnop](#) on page 2-15
- [--comment_section](#), [--no_comment_section](#) on page 2-24
- [--cppinit](#), [--no_cppinit](#) on page 2-26
- [--cpu=name](#) on page 2-28
- [--datacompressor=opt](#) on page 2-30
- [--edit=file_list](#) on page 2-38
- [--emit_relocs](#) on page 2-42
- [--entry=location](#) on page 2-43
- [--exceptions](#), [--no_exceptions](#) on page 2-46
- [--exceptions_tables=action](#) on page 2-47
- [--filtercomment](#), [--no_filtercomment](#) on page 2-51
- [--fini=symbol](#) on page 2-52
- [--first=section_id](#) on page 2-53
- [--force_explicit_attr](#) on page 2-54
- [--fpu=name](#) on page 2-56
- [--init=symbol](#) on page 2-62
- [--inline](#), [--no_inline](#) on page 2-63
- [--keep=section_id](#) on page 2-67
- [--keep_protected_symbols](#) on page 2-69
- [--largeregions](#), [--no_largeregions](#) on page 2-70
- [--last=section_id](#) on page 2-71
- [--locals](#), [--no_locals](#) on page 2-80
- [--ltcg](#) on page 2-81
- [--max_visibility=type](#) on page 2-86
- [--merge](#), [--no_merge](#) on page 2-87
- [--muldefweak](#), [--no_muldefweak](#) on page 2-88
- [--override_visibility](#) on page 2-90
- [--pad=num](#) on page 2-91
- [--paged](#) on page 2-92
- [--pagesize=pagesize](#) on page 2-93
- [--ref_cpp_init](#), [--no_ref_cpp_init](#) on page 2-99
- [--remove](#), [--no_remove](#) on page 2-102
- [--sort=algorithm](#) on page 2-115
- [--startup=symbol](#), [--no_startup](#) on page 2-118
- [--strict_flags](#), [--no_strict_flags](#) on page 2-121
- [--tailreorder](#), [--no_tailreorder](#) on page 2-129
- [--tiebreaker=option](#) on page 2-131
- [--undefined=symbol](#) on page 2-132
- [--undefined_and_export=symbol](#) on page 2-133
- [--use_definition_visibility](#) on page 2-135
- [--vfemode=mode](#) on page 2-143.

Controlling veneer generation

- [--crosser_veneershare](#), [--no_crosser_veneershare](#) on page 2-29
- [--inlineveneer](#), [--no_inlineveneer](#) on page 2-64
- [--max_veneer_passes=value](#) on page 2-85
- [--piveneer](#), [--no_piveneer](#) on page 2-95

- `--veneerinject, --no_veneerinject` on page 2-137
- `--veneer_inject_type=type` on page 2-138
- `--veneer_pool_size=size` on page 2-139
- `--veneershare, --no_veneershare` on page 2-140.

Controlling byte addressing mode

- `--be8` on page 2-12
- `--be32` on page 2-13.

Controlling the extraction and presentation of image information

- `--callgraph, --no_callgraph` on page 2-16
- `--callgraph_file=filename` on page 2-18
- `--callgraph_output=fmt` on page 2-19
- `--cgfile=type` on page 2-20
- `--cgsymbol=type` on page 2-21
- `--cgundefined=type` on page 2-22
- `--feedback=filename` on page 2-48
- `--feedback_image=option` on page 2-49
- `--feedback_type=type` on page 2-50
- `--info=topic[,topic,...]` on page 2-58
- `--info_lib_prefix=opt` on page 2-61
- `--list_mapping_symbols, --no_list_mapping_symbols` on page 2-78
- `--load_addr_map_info, --no_load_addr_map_info` on page 2-79
- `--mangled, --unmangled` on page 2-82
- `--map, --no_map` on page 2-83
- `--section_index_display=type` on page 2-110
- `--symbols, --no_symbols` on page 2-127
- `--symdefs=filename` on page 2-128
- `--xref, --no_xref` on page 2-146
- `--xrefdbg, --no_xrefdbg` on page 2-147
- `--xref{from|to}=object(section)` on page 2-148.

Note

With the exception of `--callgraph`, the linker prints the information you request on the standard output stream, `stdout`, by default. You can redirect the information to a text file using the `--list` command-line option.

Controlling diagnostic messages

- `--diag_error=tag[,tag,...]` on page 2-32
- `--diag_remark=tag[,tag,...]` on page 2-33
- `--diag_style=arm|ide|gnu` on page 2-34
- `--diag_suppress=tag[,tag,...]` on page 2-35
- `--diag_warning=tag[,tag,...]` on page 2-36
- `--errors=filename` on page 2-45
- `--list=filename` on page 2-77
- `--remarks` on page 2-101
- `--show_full_path` on page 2-112
- `--show_parent_lib` on page 2-113
- `--show_sec_idx` on page 2-114

- `--strict_enum_size`, `--no_strict_enum_size` on page 2-120
- `--strict_symbols`, `--no_strict_symbols` on page 2-124
- `--strict_visibility`, `--no_strict_visibility` on page 2-125
- `--strict_wchar_size`, `--no_strict_wchar_size` on page 2-126
- `--verbose` on page 2-141.

Controlling alignment in legacy images

- `--legacyalign`, `--no_legacyalign` on page 2-73.

Miscellaneous

- `--cpu=list` on page 2-27
- `--eager_load_debug`, `--no_eager_load_debug` on page 2-37
- `--fpu=list` on page 2-55
- `--licretry` on page 2-76
- `--show_cmdline` on page 2-111
- `--version_number` on page 2-142
- `--via=filename` on page 2-144
- `--vsr` on page 2-145.

2.4 What the linker can accept as input

Input to `arm1ink` consists of one or more object files in ARM ELF. This format is described in the *ELF for the ARM Architecture (ARM IHI 0044)*.

Optionally, the following files can be used as input to `arm1ink`:

- one or more libraries created by the librarian, `armar`
- a symbol definitions file
- a scatter file
- a steering file.

2.4.1 See also

Tasks

- [Chapter 8 Using scatter files.](#)

Creating Static Software Libraries with `armar`:

- [Creating a new object library on page 3-2.](#)

Reference

- [Accessing symbols in another image on page 7-17.](#)

Linker Reference:

- [Chapter 3 Linker steering file command reference](#)
- [Chapter 4 Formal syntax of the scatter file.](#)

Other information

- *ELF for the ARM Architecture (ARM IHI 0044)*,
<http://infocenter.arm.com/help/topic/com.arm.doc.ih0044-/index.html>.

2.5 What the linker outputs

Output from `arm1ink` can be:

- an ELF executable image
- a partially-linked ELF object
- a relocatable ELF object.

You can use `fromelf` to convert an ELF executable image to other file formats, or to display and process the content of an ELF executable image.

2.5.1 See also

Concepts

- [Partial linking model](#) on page 3-4
- [Section placement with the linker](#) on page 4-19
- [The image structure](#) on page 4-3.

Using the *fromelf Image Converter*:

- [Chapter 2 Overview of the fromelf image converter](#).

2.6 What the linker does when constructing an executable image

When you use the linker to construct an executable image, it:

- resolves symbolic references between the input object files
- extracts object modules from libraries to satisfy otherwise unsatisfied symbolic references
- sorts input sections according to their attributes and names, and merges sections with similar attributes and names into contiguous chunks
- removes unused sections
- eliminates duplicate common groups and common code, data, and debug sections
- organizes object fragments into memory regions according to the grouping and placement information provided
- assigns addresses to relocatable values
- generates an executable image.

2.6.1 See also

Tasks

- [Elimination of common debug sections on page 5-2](#)
- [Elimination of unused sections on page 5-4.](#)

Concepts

- [The image structure on page 4-3.](#)

Chapter 3

Linking models supported by armlink

The following topics describe the linking models supported by the ARM linker, armlink:

Concepts

- *Overview of linking models on page 3-2*
- *Bare-metal linking model on page 3-3*
- *Partial linking model on page 3-4.*

3.1 Overview of linking models

A linking model is a group of command-line options and memory maps that control the behavior of the linker.

Bare-metal This model does not target any specific platform. It enables you to create an image with your own custom operating system, memory map, and, application code if required. Some limited dynamic linking support is available. You can specify additional options depending on whether or not a scatter file is in use.

Partial linking

This model produces a platform-independent object suitable for input to the linker in a subsequent link step. It can be used as an intermediate step in the development process and performs limited processing of input objects to produce a single output object.

Related options in each model can be combined to tighten control over the output.

3.1.1 See also

Concepts

- [Bare-metal linking model on page 3-3](#)
- [Partial linking model on page 3-4.](#)

3.2 Bare-metal linking model

The bare-metal model focuses on the conventional embedded market where the whole program, possibly including a *Real-Time Operating System* (RTOS), is linked in one pass. Very few assumptions can be made by the linker about the memory map of a bare metal system. Therefore, you must use the scatter-loading mechanism if you want more precise control.

By default, the linker attempts to resolve all the relocations statically. However, it is also possible to create a position-independent or relocatable image. Such an image can be executed from different addresses and have its relocations resolved at load or run-time. This can be achieved using a dynamic model.

With this type of model, you can:

- identify the regions that can be relocated or are position-independent using a scatter file or command-line options.
- identify the symbols that can be imported and exported using a steering file
- identify the shared libraries that are required by the ELF file using a steering file.

You can use the following options with this model:

- `--edit=`*file_list*
- `--scatter=`*file*.

You can use the following options when scatter-loading is not used:

- `--reloc`
- `--ro_base=`*address*
- `--ropi`
- `--rosplit`
- `--rw_base=`*address*
- `--rwpi`
- `--split`
- `--zi_base`.

3.2.1 See also

Concepts

- [Methods of specifying an image memory map with the linker on page 4-8.](#)

Reference

Linker Reference:

- [--edit=](#)*file_list* on page 2-38
- [--reloc](#) on page 2-100
- [--ro_base=](#)*address* on page 2-103
- [--ropi](#) on page 2-104
- [--rosplit](#) on page 2-105
- [--rw_base=](#)*address* on page 2-106
- [--rwpi](#) on page 2-107
- [--scatter=](#)*filename* on page 2-109
- [--split](#) on page 2-117
- [--zi_base=](#)*address* on page 2-149
- [Chapter 3 Linker steering file command reference.](#)

3.3 Partial linking model

Partial linking:

- eliminates duplicate copies of debug sections
- merges the symbol tables into one
- leaves unresolved references unresolved
- merges common data (COMDAT) groups
- generates an object that can be used as input to a subsequent link step.

A single output file is produced that can be used as input to a subsequent link step. If the linker finds multiple entry points in the input files it generates an error because the output file can have only one entry point.

To link with this model, use the `--partial` command-line option. Other linker command-line options supported by this model are:

- `--edit=`*file_list*
- `--exceptions_tables=`*action*.

Note

If you use partial linking, you cannot refer to the component objects by name in a scatter file. Therefore, you might have to update your scatter file.

3.3.1 See also

Concepts

- [What is a steering file? on page 7-23.](#)

Reference

- [Steering file format on page 7-26.](#)

Linker Reference:

- [--edit=](#)*file_list* on page 2-38
- [--exceptions_tables=](#)*action* on page 2-47
- [--partial](#) on page 2-94
- [Chapter 3 Linker steering file command reference.](#)

Chapter 4

Image structure and generation

The following topics describe the image structure and the functionality available in the ARM linker, `arm1ink`, to generate images:

Tasks

- *Using command-line options to control the generation of C++ exception tables on page 4-31*
- *Controlling how the linker searches for the ARM standard libraries on page 4-36*
- *Specifying user libraries when linking on page 4-38.*

Concepts

- *The image structure on page 4-3*
- *Input sections, output sections, regions, and Program Segments on page 4-5*
- *Load view and execution view of an image on page 4-6*
- *Methods of specifying an image memory map with the linker on page 4-8*
- *Types of simple image on page 4-10*
- *Type 1 image, one load region and contiguous execution regions on page 4-11*
- *Type 2 image, one load region and non-contiguous execution regions on page 4-13*
- *Type 3 image, two load regions and non-contiguous execution regions on page 4-15*
- *Image entry points on page 4-17*
- *About specifying an initial entry point on page 4-18*
- *Section placement with the linker on page 4-19*
- *Placing sections with FIRST and LAST attributes on page 4-21*
- *Section alignment with the linker on page 4-22*

- *Demand paging* on page 4-23
- *About ordering execution regions containing Thumb code* on page 4-25
- *Overview of veneers* on page 4-26
- *Veneer sharing* on page 4-27
- *Veneer types* on page 4-28
- *Generation of position independent to absolute veneers* on page 4-29
- *Reuse of veneers when scatter-loading* on page 4-30
- *About weak references and definitions* on page 4-32
- *How the linker performs library searching, selection, and scanning* on page 4-35
- *How the linker resolves references* on page 4-39
- *Use of the strict family of options in the linker* on page 4-41.

4.1 The image structure

The structure of an image is defined by the:

- number of its constituent regions and output sections
- positions in memory of these regions and sections when the image is loaded
- positions in memory of these regions and sections when the image executes.

Each link stage has a different view of the image:

ELF object file view (linker input)

The ELF object file view comprises input sections. The ELF object file can be:

- a relocatable file that holds code and data suitable for linking with other object files to create an executable file
- an executable file that holds a program suitable for execution.

Linker view The linker has two views for the address space of a program that become distinct in the presence of overlaid, position-independent, and relocatable program fragments (code or data):

- The load address of a program fragment is the target address that the linker expects an external agent such as a program loader, dynamic linker, or debugger to copy the fragment from the ELF file. This might not be the address at which the fragment executes.
- The execution address of a program fragment is the target address where the linker expects the fragment to reside whenever it participates in the execution of the program.

If a fragment is position-independent or relocatable, its execution address can vary during execution.

ELF image file view (linker output)

The ELF image file view comprises Program Segments and output sections:

- a load region corresponds to a Program Segment
- an execution region corresponds to up to three output sections:
 - RO section
 - RW section
 - ZI section.

One or more execution regions make up a load region.

Note

With `armlink`, the maximum size of a Program Segment is 2GB.

When describing a memory view:

- the term *root region* is used to describe a region that has the same load and execution addresses
- load regions are equivalent to ELF segments.

The following figure shows the relationship between the views at each link stage:

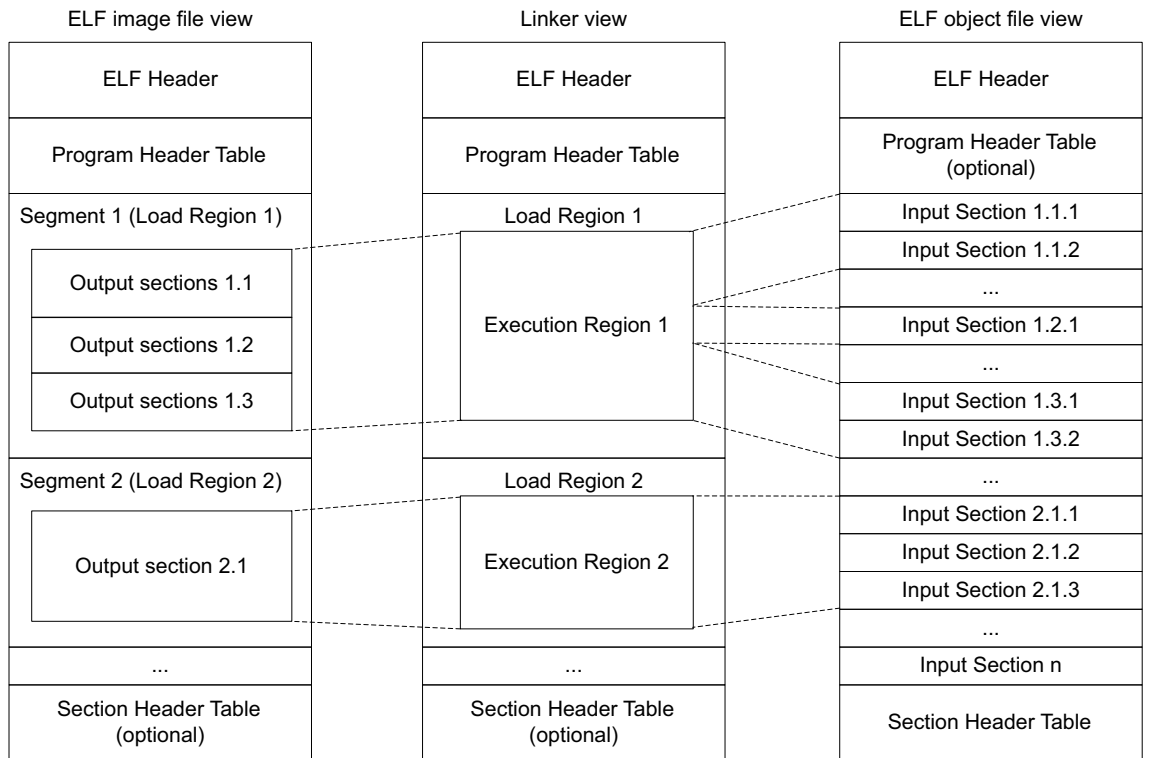


Figure 4-1 Relationship between sections, regions, and segments

4.1.1 See also

Concepts

- [Input sections, output sections, regions, and Program Segments](#) on page 4-5
- [Load view and execution view of an image](#) on page 4-6.

4.2 Input sections, output sections, regions, and Program Segments

An object or image file is constructed from a hierarchy of input sections, output sections, regions, and Program Segments:

Input section

An input section is an individual section from an input object file. It contains code, initialized data, or describes a fragment of memory that is not initialized or that must be set to zero before the image can execute. These properties are represented by attributes such as RO, RW and ZI. These attributes are used by `arm1ink` to group input sections into bigger building blocks called output sections and regions.

Output section

An output section is a group of input sections that have the same RO, RW, or ZI attribute, and that are placed contiguously in memory by the linker. An output section has the same attributes as its constituent input sections. Within an output section, the input sections are sorted according to the section placement rules.

Region

A region is a contiguous sequence of one, two, or three output sections depending on the contents of the number of sections with different attributes. The output sections in a region are sorted according to their attributes. The RO output section is first, then the RW output section, and finally the ZI output section. A region typically maps onto a physical memory device, such as ROM, RAM, or peripheral.

Program Segment

A Program Segment corresponds to a load region and contains output sections. Program Segments hold information such as text and data.

———— **Note** —————

With `arm1ink`, the maximum size of a Program Segment is 2GB.

4.2.1 See also

Concepts

- [The image structure on page 4-3](#)
- [Methods of specifying an image memory map with the linker on page 4-8](#)
- [Section placement with the linker on page 4-19.](#)

4.3 Load view and execution view of an image

Image regions are placed in the system memory map at load time. Before you can execute the image, you might have to move some of its regions to their execution addresses and create the ZI output sections. For example, initialized RW data might have to be copied from its load address in ROM to its execution address in RAM.

The memory map of an image has the following distinct views:

Load view Describes each image region and section in terms of the address where it is located when the image is loaded into memory, that is, the location before image execution starts.

Execution view Describes each image region and section in terms of the address where it is located during image execution.

The following figure shows these views:




Figure 4-2 Load and execution memory maps

The following table compares the load and execution views:

Table 4-1 Comparing load and execution views

Load	Description	Execution	Description
Load address	The address where a section or region is loaded into memory before the image containing it starts executing. The load address of a section or a non-root region can differ from its execution address	Execution address	The address where a section or region is located while the image containing it is being executed
Load region	A region in the load address space	Execution region	A region in the execution address space

4.3.1 See also

Concepts

- [The image structure on page 4-3](#)
- [Input sections, output sections, regions, and Program Segments on page 4-5](#)
- [Methods of specifying an image memory map with the linker on page 4-8](#)

- *Section placement with the linker on page 4-19.*

4.4 Methods of specifying an image memory map with the linker

An image can consist of any number of regions and output sections. Regions can have different load and execution addresses. To construct the memory map of an image, `armlink` must have information about:

- how input sections are grouped into output sections and regions
- where regions are to be located in the memory maps.

Depending on the complexity of the memory maps of the image, there are two ways to pass this information to `armlink`:

Using command-line options

The following options can be used for simple cases where an image has only one or two load regions and up to three execution regions:

- `--first`
- `--last`
- `--ro_base`
- `--rw_base`
- `--ropi`
- `--rwp_i`
- `--split`
- `--rosplit`
- `--zi_base`.

These options provide a simplified notation that gives the same settings as a scatter-loading description for a simple image.

Using a scatter file

A scatter file is a textual description of the memory layout and code and data placement. It is used for more complex cases where you require complete control over the grouping and placement of image components. To use a scatter file, specify `--scatter=filename` at the command-line.

———— Note —————

You cannot use `--scatter` with the other memory map related command-line options.

Table 4-2 Comparison of scatter file and equivalent command-line options

Scatter file	Equivalent command-line options
<pre>LR1 0x0000 0x20000 { ER1 0x0 0x20000 { init.o (INIT, +FIRST) *(+RO) } RAM 0x400000 { *(+RW) } RAM 0x405000 { *(+ZI) } }</pre>	<pre>--ro_base=0x0 --first=init.o(init) --rw_base=0x400000 --zi_base=0x405000</pre>

4.4.1 See also

Tasks

- [Chapter 8 Using scatter files.](#)

Concepts

- [The image structure](#) on page 4-3
- [Input sections, output sections, regions, and Program Segments](#) on page 4-5
- [Load view and execution view of an image](#) on page 4-6
- [Types of simple image](#) on page 4-10.

Reference

Linker Reference:

- [--first=section_id](#) on page 2-53
- [--last=section_id](#) on page 2-71
- [--ro_base=address](#) on page 2-103
- [--ropi](#) on page 2-104
- [--rosplit](#) on page 2-105
- [--rw_base=address](#) on page 2-106
- [--rwpi](#) on page 2-107
- [--scatter=filename](#) on page 2-109
- [--split](#) on page 2-117
- [--zi_base=address](#) on page 2-149.

4.5 Types of simple image

A simple image consists of a number of input sections of type RO, RW, and ZI. These input sections are collated to form the RO, RW, and ZI output sections. Depending on how the output sections are arranged within load and execution regions, there are three basic types of simple image:

- Type 1** One region in load view, three contiguous regions in execution view. Use the `--ro_base` option to create this type of image.
- Type 2** One region in load view, three non-contiguous regions in execution view. Use the `--ro_base` and `--rw_base` options to create this type of image.
- Type 3** Two regions in load view, three non-contiguous regions in execution view. Use the `--ro_base`, `--rw_base`, and `--split` options to create this type of image.

In all the simple image types:

- the first execution region contains the RO output section
- the second execution region contains the RW output section (if present)
- the third execution region contains the ZI output section (if present).

These execution regions are referred to as the RO, the RW, and the ZI execution region.

However, you can also use the `--rosplit` option for a Type 3 image. This option splits the default load region into two RO output sections, one for code and one for data.

You can also use the `--zi_base` command-line option to specify the base address of a ZI execution region for Type 1 and Type 2 images. This option is ignored if you also use the `--split` command-line option that is required for Type 3 images.

You can also create simple images with scatter files.

4.5.1 See also

Concepts

- [Type 1 image, one load region and contiguous execution regions](#) on page 4-11
- [Type 2 image, one load region and non-contiguous execution regions](#) on page 4-13
- [Type 3 image, two load regions and non-contiguous execution regions](#) on page 4-15
- [Equivalent scatter-loading descriptions for simple images](#) on page 8-60.

Reference

Linker Reference:

- [--ro_base=address](#) on page 2-103
- [--rosplit](#) on page 2-105
- [--rw_base=address](#) on page 2-106
- [--scatter=filename](#) on page 2-109
- [--split](#) on page 2-117
- [--zi_base=address](#) on page 2-149.

4.6 Type 1 image, one load region and contiguous execution regions

A Type 1 image consists of a single load region in the load view and three execution regions placed contiguously in the memory map. This approach is suitable for systems that load programs into RAM, for example, an OS bootloader or a desktop system.

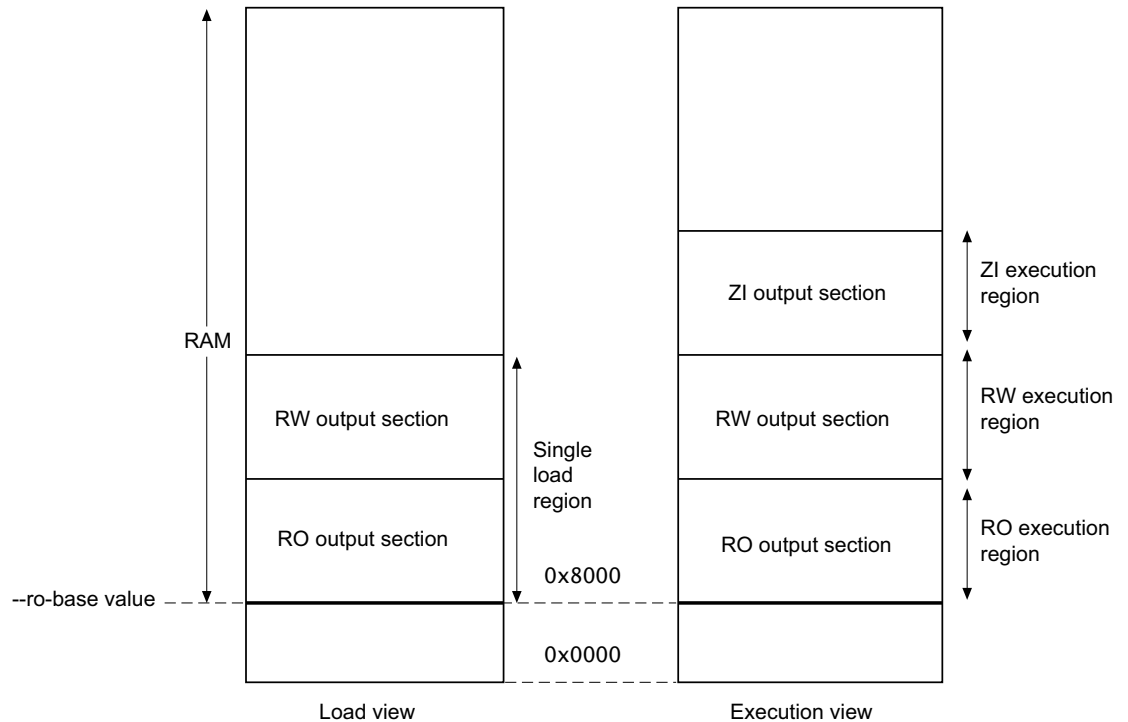


Figure 4-3 Simple type 1 image

Use the following command for images of this type:

```
armlink --ro_base 0x8000
```

———— **Note** ————

`0x8000` is the default address, so you do not have to specify `--ro_base` for the example.

4.6.1 Load view

The single load region consists of the RO and RW output sections, placed consecutively. The RO and RW execution regions are both root regions. The ZI output section does not exist at load time. It is created before execution, using the output section description in the image file.

4.6.2 Execution view

The three execution regions containing the RO, RW, and ZI output sections are arranged contiguously. The execution addresses of the RO and RW execution regions are the same as their load addresses, so nothing has to be moved from its load address to its execution address. However, the ZI execution region that contains the ZI output section is created at run-time.

Use `armlink` option `--ro_base address` to specify the load and execution address of the region containing the RO output. The default address is `0x8000`.

Use the `--zi_base` command-line option to specify the base address of a ZI execution region.

4.6.3 See also

Concepts

- [The image structure](#) on page 4-3
- [Input sections, output sections, regions, and Program Segments](#) on page 4-5
- [Load view and execution view of an image](#) on page 4-6.

Reference

Linker Reference:

- [--ro_base=address](#) on page 2-103
- [--zi_base=address](#) on page 2-149.

4.7 Type 2 image, one load region and non-contiguous execution regions

A Type 2 image consists of a single load region, and three execution regions in execution view. The RW execution region is not contiguous with the RO execution region. This approach is used, for example, for ROM-based embedded systems, where RW data is copied from ROM to RAM at startup:




Figure 4-4 Simple type 2 image

Use the following command for images of this type:

```
armlink --ro_base 0x0 --rw_base 0xA000
```

4.7.1 Load view

In the load view, the single load region consists of the RO and RW output sections placed consecutively, for example, in ROM. Here, the RO region is a root region, and the RW region is non-root. The ZI output section does not exist at load time. It is created at runtime.

4.7.2 Execution view

In the execution view, the first execution region contains the RO output section and the second execution region contains the RW and ZI output sections.

The execution address of the region containing the RO output section is the same as its load address, so the RO output section does not have to be moved. That is, it is a root region.

The execution address of the region containing the RW output section is different from its load address, so the RW output section is moved from its load address (from the single load region) to its execution address (into the second execution region). The ZI execution region, and its output section, is placed contiguously with the RW execution region.

Use `arm1ink` options `--ro_base address` to specify the load and execution address for the RO output section, and `--rw_base exec_address` to specify the execution address of the RW output section. If you do not use the `--ro_base` option to specify the address, the default value of `0x8000` is used by `arm1ink`. For an embedded system, `0x0` is typical for the `--ro_base` value. If you do not use the `--rw_base` option to specify the address, the default is to place RW directly above RO (as in a Type 1 image).

Use the `--zi_base` command-line option to specify the base address of a ZI execution region.

———— **Note** —————

The execution region for the RW and ZI output sections cannot overlap any of the load regions.

4.7.3 See also

Concepts

- [The image structure](#) on page 4-3
- [Input sections, output sections, regions, and Program Segments](#) on page 4-5
- [Load view and execution view of an image](#) on page 4-6
- [Type 1 image, one load region and contiguous execution regions](#) on page 4-11.

Reference

Linker Reference:

- [--ro_base=address](#) on page 2-103
- [--rw_base=address](#) on page 2-106
- [--zi_base=address](#) on page 2-149.

4.8 Type 3 image, two load regions and non-contiguous execution regions

A Type 3 image is similar to a Type 2 image except that the single load region is split into two root load regions.

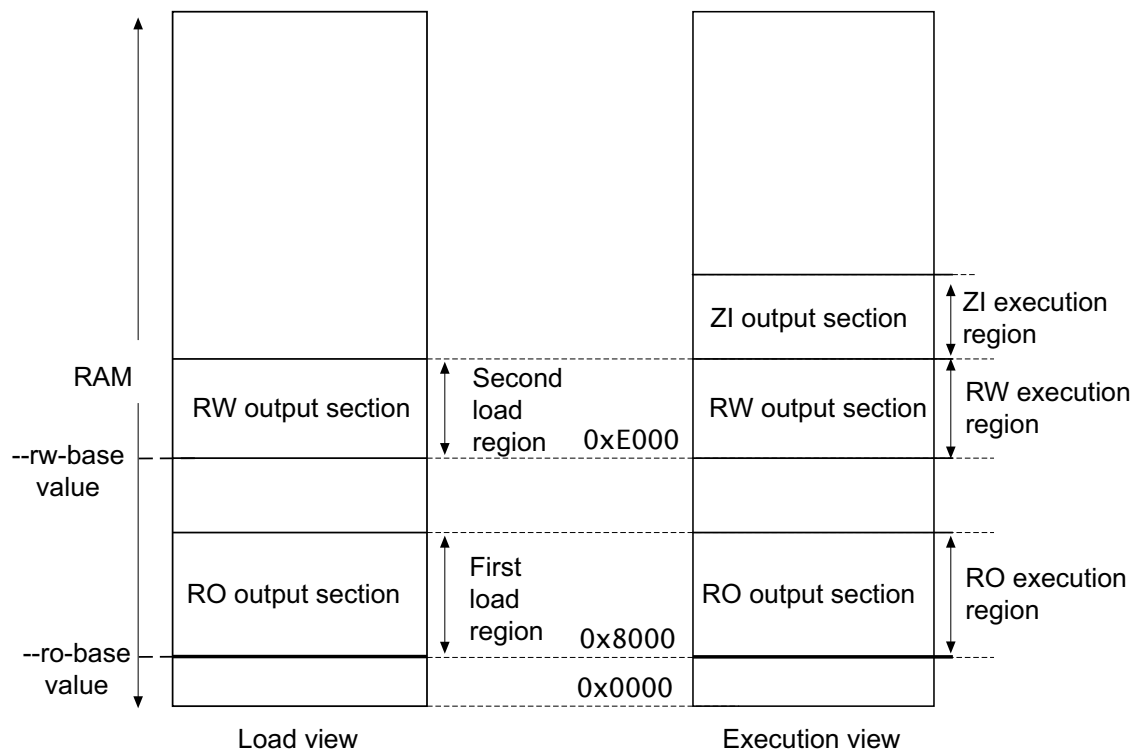


Figure 4-5 Simple type 3 image

Use the following command for images of this type:

```
armlink --split --ro_base 0x8000 --rw_base 0xE000
```

4.8.1 Load view

In the load view, the first load region consists of the RO output section, and the second load region consists of the RW output section. The ZI output section does not exist at load time. It is created before execution, using the description of the output section contained in the image file.

4.8.2 Execution view

In the execution view, the first execution region contains the RO output section, and the second execution region contains the RW and ZI output sections.

The execution address of the RO region is the same as its load address, so the contents of the RO output section do not have to be moved or copied from their load address to their execution address. Both RO and RW are root regions.

The execution address of the RW region is also the same as its load address, so the contents of the RW output section are not moved from their load address to their execution address. However, the ZI output section is created at run-time and is placed contiguously with the RW region.

Specify the load and execution address using the following linker options:

`--ro_base address`

Instructs `armlink` to set the load and execution address of the region containing the RO section at a four-byte aligned *address*, for example, the address of the first location in ROM. If you do not use the `--ro_base` option to specify the address, the default value of `0x8000` is used by `armlink`.

`--rw_base address`

Instructs `armlink` to set the execution address of the region containing the RW output section at a four-byte aligned *address*. If this option is used with `--split`, this specifies both the load and execution addresses of the RW region, for example, a root region.

`--split`

Splits the default single load region, that contains both the RO and RW output sections, into two root load regions:

- one containing the RO output section
- one containing the RW output section.

You can then place them separately using `--ro_base` and `--rw_base`.

4.8.3 See also

Concepts

- [The image structure on page 4-3](#)
- [Input sections, output sections, regions, and Program Segments on page 4-5](#)
- [Load view and execution view of an image on page 4-6](#)
- [Type 2 image, one load region and non-contiguous execution regions on page 4-13.](#)

Reference

Linker Reference:

- [--ro_base=address on page 2-103](#)
- [--rw_base=address on page 2-106](#)
- [--split on page 2-117.](#)

4.9 Image entry points

An entry point in an image is a location where program execution can start. There are two distinct types of entry point:

Initial entry point

The *initial* entry point for an image is a single value that is stored in the ELF header file. For programs loaded into RAM by an operating system or boot loader, the loader starts the image execution by transferring control to the initial entry point in the image.

An image can have only one initial entry point. The initial entry point can be, but is not required to be, one of the entry points set by the `ENTRY` directive.

Entry points set by the `ENTRY` directive

You can select one of many possible entry points for an image. An image can have only one entry point.

You create entry points in objects with the `ENTRY` directive in an assembler file. In embedded systems, this directive is typically used to mark code that is entered through the processor exception vectors, such as `RESET`, `IRQ`, and `FIQ`.

The directive marks the output code section with an `ENTRY` keyword that instructs the linker not to remove the section when it performs unused section elimination. For C and C++ programs, the `__main()` function in the C library is also an entry point.

If an embedded image is to be used by a loader, it must have a single initial entry point specified in the header. Use the `--entry` command-line option to select the entry point.

4.9.1 See also

Tasks

- [About specifying an initial entry point on page 4-18.](#)

Reference

Linker Reference:

- [--entry=location on page 2-43.](#)

Assembler Reference:

- [ENTRY on page 5-35.](#)

4.10 About specifying an initial entry point

You must specify at least one initial entry point for a program otherwise the linker produces a warning. Not every source file has to have an entry point. Multiple entry points in a single source file are not permitted.

For embedded applications with ROM at zero use `--entry 0x0`, or optionally `0xFFFF0000` for CPUs that are using high vectors.

The initial entry point must meet the following conditions:

- the image entry point must always lie within an execution region
- the execution region must not overlay another execution region, and must be a root execution region (the load address is the same as the execution address).

If you do not use the `--entry` option to specify the initial entry point then:

- if the input objects contain only one entry point set by the `ENTRY` directive, the linker uses that entry point as the initial entry point for the image
- the linker generates an image that does not contain an initial entry point when either:
 - more than one entry point has been specified by using the `ENTRY` directive
 - no entry point has been specified by using the `ENTRY` directive.

4.10.1 See also

Concepts

- [What is a root region?](#) on page 8-12.

Reference

Linker Reference:

- [--entry=location](#) on page 2-43.

Assembler Reference:

- [ENTRY](#) on page 5-35.

4.11 Section placement with the linker

By default, the linker places input sections in the following order when generating an image:

1. By attribute as follows:
 - a. read-only code
 - b. read-only data
 - c. read-write code
 - d. read-write data
 - e. zero-initialized data.
2. By input section name if they have the same attributes. Names are considered to be case-sensitive and are compared in alphabetical order using the ASCII collation sequence for characters.
3. By their relative positions in the input file if they have the same attributes and section names, except where overridden by `FIRST` or `LAST`.

Portions of the image are collected together into a minimum number of contiguous regions.

———— **Note** ————

The sorting order is unaffected by ordering within scatter files or object file names.

These rules mean that the positions of input sections with identical attributes and names included from libraries is not predictable. If you require more precise positioning, specify the individual modules explicitly in a scatter file, and include the modules in the input file list for the `armlink` command.

The base address of each input section is determined by the sorting order defined by the linker, and is correctly aligned within the output section that contains it.

By default, the linker creates an image consisting of an RO output section, an RW output section, and optionally a ZI output section. The RO output section can be protected at run-time on systems that have memory management hardware. RO sections can also be placed into ROM in the target.

Alternative sorting orders are available with the `--sort=algorithm` command-line option. The linker might change the *algorithm* to minimise the amount of veneers generated if no algorithm is chosen.

4.11.1 Handling unassigned sections

The linker might not be able to place some input sections in any execution region. When this happens, the linker generates an error message. This might occur because your current scatter file does not permit all possible module select patterns and input section selectors. How you fix this depends on the importance of placing these sections correctly:

- If the sections must be placed at specific locations, then modify your scatter file to include specific module selectors and input section selectors as required.
- If the placement of the unassigned sections is not important, you can use one or more `,ANY` module selectors with optional input section selectors.

When the linker creates an input section for each ELF section in the input objects, it increments a global counter. The value of this counter is stored in the section as the creation index. The creation index is used when placing any unassigned sections that have identical properties.

4.11.2 Example

The following scatter file shows how the linker places sections:

```
LoadRegion 0x8000
{
  ExecRegion1 0x0000 0x4000
  {
    *(sections)
    *(moresections)
  }
  ExecRegion2 0x4000 0x2000
  {
    *(evenmoresections)
  }
}
```

The order of execution regions within the load region is not altered by the linker.

4.11.3 See also

Tasks

- [Placing sections with FIRST and LAST attributes](#) on page 4-21
- [Placing functions and data at specific addresses](#) on page 8-17
- [Placing a named section explicitly using scatter-loading](#) on page 8-22
- [Placing unassigned sections with the .ANY module selector](#) on page 8-23.

Concepts

- [The image structure](#) on page 4-3
- [Input sections, output sections, regions, and Program Segments](#) on page 4-5
- [Load view and execution view of an image](#) on page 4-6
- [About ordering execution regions containing Thumb code](#) on page 4-25
- [Overview of veneers](#) on page 4-26
- [Section alignment with the linker](#) on page 4-22.

Reference

Linker Reference:

- [--sort=algorithm](#) on page 2-115
- [Syntax of an input section description](#) on page 4-23.

4.12 Placing sections with FIRST and LAST attributes

You can make sure that a section is placed either first or last in its execution region. For example, you might want to make sure the section containing the vector table is placed first in the image. To do this, use one of the following methods:

- If you are not using scatter-loading, use the `--first` and `--last` linker command-line options to place input sections.
- If you are using scatter-loading, use the attributes `FIRST` and `LAST` in the file to mark the first and last input sections in an execution region if the placement order is important. However, `FIRST` and `LAST` must not violate the basic attribute sorting order. For example, `FIRST RW` is placed after any read-only code or read-only data.

4.12.1 See also

Concepts

- [The image structure on page 4-3](#)
- [Input sections, output sections, regions, and Program Segments on page 4-5](#)
- [Load view and execution view of an image on page 4-6](#)
- [Section placement with the linker on page 4-19](#)
- [About scatter-loading on page 8-3.](#)

Reference

Linker Reference:

- [--first=section_id on page 2-53](#)
- [--last=section_id on page 2-71](#)
- [Syntax of an input section description on page 4-23.](#)

4.13 Section alignment with the linker

When input sections have been ordered and before the base addresses are fixed, `arm1ink` inserts padding, if required, to force each input section to start at an address that is a multiple of the input section alignment.

The linker permits ELF program headers and output sections to be aligned on a four-byte boundary regardless of the maximum alignment of the input sections. This enables `arm1ink` to minimize the amount of padding that it inserts into the image.

If you require strict conformance with the ELF specification then use the `--no_legacyalign` option. The linker faults the base address of a region if it is not aligned so padding might be inserted to ensure compliance. When `--no_legacyalign` is used the region alignment is the maximum alignment of any input section contained by the region.

If you are using scatter-loading, you can increase the alignment of a load region or execution region with the `ALIGN` attribute. For example, you can change an execution region that is normally four-byte aligned to be eight-byte aligned. However, you cannot reduce the natural alignment. For example, you cannot force two-byte alignment on a region that is normally four-byte aligned.

4.13.1 See also

Tasks

- [About creating regions on page boundaries on page 8-54.](#)

Reference

Linker Reference:

- [--legacyalign, --no_legacyalign on page 2-73](#)
- [Load region attributes on page 4-8](#)
- [Execution region attributes on page 4-12](#)
- [Example of aligning a base address in execution space but still tightly packed in load space on page 4-42.](#)

4.14 Demand paging

In operating systems that support virtual memory an ELF file can be loaded by mapping the ELF files into the address space of the process loading the file. When a virtual address in a page that is mapped to the file is accessed, the operating system loads that page from disk. ELF files that are to be used this way must conform to a certain format.

Use the `--paged` command-line option to enable demand paging mode. This helps produce ELF files that can be demand paged efficiently.

The basic constraints on the ELF file are:

- there is no difference between the load and execution address for any Output Section
- all `PT_LOAD` Program Headers have a minimum alignment, `pt_align`, of the page size for the operating system
- all `PT_LOAD` Program Headers have a file offset, `pt_offset`, that is congruent to the virtual address (`pt_addr`) modulo `pt_align`.

When `--paged` is on:

- The linker automatically generates the Program Headers from the execution region base addresses. The usual situation where one load region generates one Program Header no longer holds.
- The operating system page size is controlled by the `--pagesize` command-line option.
- The linker attempts to place the ELF Header and Program Header in the first `PT_LOAD` program header, if space is available.

4.14.1 Example

This is an example of a demand paged scatter file:

```
LR1 GetPageSize() + SizeOfHeaders()
{
  ER_RO +0
  {
    *(+RO)
  }
  ER_RW +GetPageSize()
  {
    *(+RW)
  }
  ER_ZI +0
  {
    *(+ZI)
  }
}
```

4.14.2 See also

Concepts

- [About scatter-loading on page 8-3](#)

Reference

Linker Reference:

- [--paged on page 2-92](#)
- [--pagesize=pagesize on page 2-93](#)

- *--pagesize=pagesize* on page 2-93
- *--scatter=filename* on page 2-109
- *GetPageSize()* function on page 4-44
- *SizeOfHeaders()* function on page 4-45.

4.15 About ordering execution regions containing Thumb code

The 16-bit Thumb branch range is 4MB. When an execution region contains Thumb code that exceeds 4MB, `arm1ink` attempts to order sections that are at a similar average call depth and to place the most commonly called sections centrally. This helps to minimize the number of veneers generated.

The 32-bit Thumb branch range is 16MB. Section re-ordering is only required if that limit is exceeded.

To disable section re-ordering, use the `--no_largeregions` command-line option.

4.15.1 See also

Concepts

- [Section placement with the linker on page 4-19](#)
- [Overview of veneers on page 4-26](#).

Reference

Linker Reference:

- [--largeregions, --no_largeregions on page 2-70](#).

4.16 Overview of veneers

Veneers are small sections of code generated by the linker and inserted into your program. The BL instruction is PC-relative and has a limited branch range. Therefore, `arm1ink` must generate veneers when a branch involves a destination beyond the branching range of the BL instruction.

The range of a BL instruction is 32MB for ARM, 16MB for 32-bit Thumb, and 4MB for 16-bit Thumb. A veneer extends the range of the branch by becoming the intermediate target of the branch instruction. The veneer then sets the PC to the destination address. This enables the veneer to branch anywhere in the 4 GB address space. If ARM and Thumb are mixed, the veneer also handles processor state changes.

The linker can generate the following veneer types depending on what is required:

- inline veneers
- short branch veneers
- long branch veneers.

`arm1ink` creates one input section called `Veneer$$Code` for each veneer. A veneer is generated only if no other existing veneer can satisfy the requirements. If two input sections contain a long branch to the same destination, only one veneer is generated that is shared by both branch instructions. A veneer is only shared in this way if it can be reached by both sections.

If you are using ARMv4T, `arm1ink` generates veneers when a branch involves change of state between ARM and Thumb. You still get interworking veneers for ARMv5TE and later when using conditional branches, because there is no conditional BL instruction. Veneers for state changes are also required for B instructions in ARMv5 and later.

4.16.1 See also

Concepts

- [Veneer sharing on page 4-27](#)
- [Veneer types on page 4-28](#)
- [Generation of position independent to absolute veneers on page 4-29](#)
- [Reuse of veneers when scatter-loading on page 4-30.](#)

4.17 Veneer sharing

If multiple objects result in the same veneer being created, the linker creates a single instance of that veneer. The veneer is then shared by those objects.

You can use the command-line option `--no_veneershare` to specify that veneers are not shared. This assigns ownership of the created veneer section to the object that created the veneer and so enables you to select veneers from a particular object in a scatter file, for example:

```
LR 0x8000
{
  ER_ROOT +0
  {
    object1.o(Veneer$$Code)
  }
}
```

Be aware that veneer sharing makes it impossible to assign an owning object. Using `--no_veneershare` provides a more consistent image layout. However, this comes at the cost of a significant increase in code size, because of the extra veneers generated by the linker.

4.17.1 See also

Concepts

- [Overview of veneers](#) on page 4-26
- [About scatter-loading](#) on page 8-3.

Reference

Linker Reference:

- [--veneershare, --no_veneershare](#) on page 2-140
- [Chapter 4 Formal syntax of the scatter file.](#)

4.18 Veneer types

Veneers have different capabilities and use different code pieces. The linker selects the most appropriate, smallest, and fastest depending on the branching requirements:

- Inline veneer:
 - used to perform only a state change
 - the veneer must be inserted just before the target section to be in range
 - an ARM-Thumb interworking veneer has a range of 256 bytes so the function entry point must appear within 256 bytes of the veneer
 - a Thumb-ARM interworking veneer has a range of zero bytes so the function entry point must appear immediately after the veneer
 - an inline veneer is always position-independent.
- Short branch veneer:
 - an interworking Thumb to ARM short branch veneer has a range of 32MB, the range for an ARM instruction
 - a short branch veneer is always position-independent.
- Long branch veneer:
 - can branch anywhere in the 4GB address space
 - all long branch veneers are also interworking veneers
 - there are different long branch veneers for absolute or position-independent code.

When you are using veneers be aware of the following:

- The inline veneer limitations mean that you cannot move inline veneers out of an execution region using a scatter file. Use the command-line option `--no_inlineveneer` to prevent the generation of inline veneers.
- All veneers cannot be collected into one input section because the resulting veneer input section might not be within range of other input sections. If the sections are not within addressing range, long branching is not possible.
- The linker generates position-independent variants of the veneers automatically. However, because such veneers are larger than non position-independent variants, the linker only does this where necessary, that is, where the source and destination execution regions are both position-independent and are rigidly related.

Veneers are generated to optimize code size. `arm1ink`, therefore, chooses the variant in order of preference:

1. Inline veneer.
2. Short branch veneer.
3. Long veneer.

4.18.1 See also

Concepts

- [Overview of veneers on page 4-26.](#)

Reference

Linker Reference:

- [--inlineveneer; --no_inlineveneer on page 2-64](#)
- [--max_veneer_passes=value on page 2-85.](#)

4.19 Generation of position independent to absolute veneers

The normal call instruction encodes the address of the target as an offset from the calling address. When calling from *position independent* (PI) code to absolute code the offset cannot be calculated at link time, so the linker must insert a long-branch veneer.

The generation of PI to absolute veneers can be controlled using the `--piveneer` option, that is set by default. When this option is turned off using `--no_piveneer`, the linker generates an error when a call from PI code to absolute code is detected.

4.19.1 See also

Concepts

- [Overview of veneers on page 4-26.](#)

Reference

Linker Reference:

- [--max_veneer_passes=value on page 2-85](#)
- [--piveneer, --no_piveneer on page 2-95.](#)

4.20 Reuse of veneers when scatter-loading

The linker reuses veneers whenever possible, but there are some limitations on the reuse of veneers in protected load regions and overlaid execution regions.

A scatter file enables you to create regions that share the same area of RAM:

- If you use the PROTECTED keyword for a load region it prevents:
 - overlapping of load regions
 - veneer sharing
 - string sharing with the --merge option.
- If you use the OVERLAY keyword for a region, no other execution region can reuse a veneer placed in an overlay execution region.

If it is not possible to reuse a veneer, new veneers are created instead. Unless you have instructed the linker to place veneers somewhere specific using scatter-loading, a veneer is always placed in the execution region that contains the call requiring the veneer.

4.20.1 See also

Concepts

- [Overview of veneers on page 4-26.](#)

Reference

- [Load region attributes on page 4-8](#)
- [Address attributes for load and execution regions on page 4-15.](#)

4.21 Using command-line options to control the generation of C++ exception tables

By default, or if the option `--exceptions` is specified, the image can contain exception tables. Exception tables are discarded silently if no code throws an exception. However, if the option `--no_exceptions` is specified, the linker generates an error if any exceptions sections are present after unused sections have been eliminated.

You can use the `--no_exceptions` option to ensure that your code is exceptions free. The linker generates an error message to highlight that exceptions have been found and does not produce a final image.

However, you can use the `--no_exceptions` option with the `--diag_warning` option to downgrade the error message to a warning. The linker produces a final image but also generates a message to warn you that exceptions have been found.

The linker can create exception tables for legacy objects that contain debug frame information. The linker can do this safely for C and assembly language objects. By default, the linker does not create exception tables. This is the same as using the linker option `--exceptions_tables=nocreate`.

The linker option `--exceptions_tables=unwind` enables the linker to use the `.debug_frame` information to create a register-restoring unwinding table for each section in your image that does not already have an exception table. If this is not possible, the linker creates a nounwind table instead.

Use the linker option `--exceptions_tables=cantunwind` to create a nounwind table for each section in your image that does not already have an exception table.

———— Note ————

Be aware of the following:

- With the default settings, that is, `--exceptions` `--exceptions_tables=nocreate`, it is not safe to throw an exception through C or assembly code, unless the C code is compiled with the option `--exceptions`.
- The linker can generate frame unwinding instructions from objects with `.debug_frame` information. Frame unwinding is sufficient for C and assembler code. It is not sufficient for C++ code, because it does not call the destructors for the objects on the stack that is being unwound.

The cleanup code for C++ must be generated by the compiler with the `--exceptions` option.

4.21.1 See also

Reference

Linker Reference:

- [--diag_warning=tag\[tag,...\]](#) on page 2-36
- [--exceptions, --no_exceptions](#) on page 2-46
- [--exceptions_tables=action](#) on page 2-47.

Compiler Reference:

- [--exceptions, --no_exceptions](#) on page 3-65.

4.22 About weak references and definitions

Weak references and definitions provide additional flexibility in the way the linker includes various functions and variables in a build. These references are typically to library functions.

Weak references

If the linker cannot resolve normal, non-weak, references to symbols included in the link, it attempts to do so by finding the symbol in a library:

- If it is unable to find such a reference, the linker reports an error.
- If such a reference is resolved, the section it is resolved to is marked as used. This ensures the section is not removed by the linker as an unused section. Each non-weak reference must be resolved by exactly one definition. If there are multiple definitions, the linker reports an error.

Function or variable declarations in C source files can be marked with the `__weak` qualifier. As with `extern`, this qualifier tells the compiler that a function or variable is declared in another source file. Because the definition of this function or variable might not be available to the compiler, it creates a weak reference to be resolved by the linker.

The linker does not load an object from a library to resolve a weak reference. It is able to resolve the weak reference only if the definition is included in the image for other reasons. The weak reference does not cause the linker to mark the section containing the definition as used, so it might be removed by the linker as unused. The definition might already exist in the image for several reasons:

- The symbol is strongly referenced somewhere else in the code.
- The symbol definition exists in the same ELF section as a symbol definition that is included for any of these reasons.
- The symbol definition is in a section that has been specified using `--keep`, or contains an ENTRY point.
- The symbol definition is in an object file included in the link and the `--no_remove` option is used. The object file is not referenced from a library unless that object file within the library is explicitly included on the linker command-line.

In summary, a weak reference is resolved if the definition is already included in the image, but it does not determine if that definition is included.

An unresolved weak function call is replaced with either:

- A no-operation instruction, NOP.
- A branch with link instruction, BL, to the following instruction. That is, the function call just does not happen.

Weak definitions

A function definition, or an exported label in assembler, can also be marked as weak, as can a variable definition. In this case, a weak symbol definition is created in the object file.

A weak definition can be used to resolve any reference to that symbol in the same way as a normal definition. However, if another non-weak definition of that symbol exists in the build, the linker uses that definition instead of the weak definition, and does not produce an error due to multiply-defined symbols.

4.22.1 Example of a weak reference

A library contains a function `foo()`, that is called in some builds of an application but not in others. If it is used, `init_foo()` must be called first. Weak references can be used to automate the call to `init_foo()`.

The library can define `init_foo()` and `foo()` in the same ELF section. The application initialization code must call `init_foo()` weakly. If the application includes `foo()` for any reason, it also includes `init_foo()` and this is called from the initialization code. In any builds that do not include `foo()`, the call to `init_foo()` is removed by the linker.

Typically, the code for multiple functions defined within a single source file is placed into a single ELF section by the compiler. However, certain build options might alter this behavior, so you must use them with caution if your build is relying on the grouping of files into ELF sections:

- The compiler command-line option `--split_sections` results in each function being placed in its own section. In this example, compiling the library with this option results in `foo()` and `init_foo()` being placed in separate sections. Therefore `init_foo()` is not automatically included in the build due to a call to `foo()`.
- The linker feedback mechanism, `--feedback`, records `init_foo()` as being unused during the link step. This causes the compiler to place `init_foo()` into its own section during subsequent compilations, so that it can be removed.
- The compiler directive `#pragma arm section` also instructs the compiler to generate a separate ELF section for some functions.

In this example, there is no need to rebuild the initialization code between builds that include `foo()` and do not include `foo()`. There is also no possibility of accidentally building an application with a version of the initialization code that does not call `init_foo()`, and other parts of the application that call `foo()`.

An example of `foo.c` source code that is typically built into a library is:

```
void init_foo()
{
    // Some initialization code
}

void foo()
{
    // A function that is included in some builds
    // and requires init_foo() to be called first.
}
```

An example of `init.c` is:

```
__weak void init_foo(void);
int main(void)
{
    init_foo();
    // Rest of code that may make calls to foo() directly or indirectly.
}
```

An example of a weak reference generated by the assembler is:

```
init.s:

    IMPORT init_foo WEAK
    AREA init, CODE, readonly
```

```

    BL init_foo
    ;Rest of code
END

```

4.22.2 Example of a weak definition

A simple or dummy implementation of a function can be provided as a weak definition. This enables you to build software with defined behavior without having to provide a full implementation of the function. It also enables you to provide a full implementation for some builds if required.

4.22.3 See also

Concepts

- [How the linker performs library searching, selection, and scanning on page 4-35](#)
- [How the linker resolves references on page 4-39.](#)

Reference

Linker Reference:

- [--feedback=filename on page 2-48](#)
- [--keep=section_id on page 2-67](#)
- [--remove, --no_remove on page 2-102.](#)

Compiler Reference:

- [--split_sections on page 3-156](#)
- [__weak on page 5-27](#)
- [__attribute__\(\(weak\)\) function attribute on page 5-60](#)
- [__attribute__\(\(weakref\("target"\)\)\) function attribute on page 5-61](#)
- [__attribute__\(\(weak\)\) variable attribute on page 5-80](#)
- [__attribute__\(\(weakref\("target"\)\)\) variable attribute on page 5-81](#)
- [#pragma arm section \[section_type_list\] on page 5-86.](#)

Assembler Reference:

- [NOP on page 3-132](#)
- [B, BL, BX, BLX, and BXJ on page 3-45](#)
- [ENTRY on page 5-35](#)
- [EXPORT or GLOBAL on page 5-37.](#)

4.23 How the linker performs library searching, selection, and scanning

The linker always searches user libraries before the ARM libraries. If you specify the `--no_scanlib` command-line option, the linker does not search for the default ARM libraries and uses only those libraries that are specified in the input file list to resolve references.

The linker creates an internal list of libraries as follows:

1. Any libraries explicitly specified in the input file list are added to the list.
2. The user-specified search path is examined to identify ARM standard libraries to satisfy requests embedded in the input objects.

The best-suited library variants are chosen from the searched directories and their subdirectories. Libraries supplied by ARM have multiple variants that are named according to the attributes of their members.

Be aware of the following differences between the way the linker adds object files to the image and the way it adds libraries to the image:

- Each object file in the input list is added to the output image unconditionally, whether or not anything refers to it. At least one object must be specified.
- A member from a library is included in the output only if:
 - an object file or an already-included library member makes a non-weak reference to it
 - the linker is explicitly instructed to add it.

————— **Note** —————

If a library member is explicitly requested in the input file list, the member is loaded even if it does not resolve any current references. In this case, an explicitly requested member is treated as if it is an ordinary object.

Unresolved references to weak symbols do not cause library members to be loaded.

4.23.1 See also

Tasks

- [About weak references and definitions on page 4-32](#)
- [Controlling how the linker searches for the ARM standard libraries on page 4-36.](#)

Reference

Linker Reference:

- [--keep=section_id on page 2-67](#)
- [--remove, --no_remove on page 2-102](#)
- [--scanlib, --no_scanlib on page 2-108.](#)

4.24 Controlling how the linker searches for the ARM standard libraries

By default, the linker searches for the ARM standard libraries in `./lib`, relative to the location of the `armlink` executable. You can also control how the linker searches for the ARM standard libraries with either the `--libpath` command-line option or the `ARMLIB` or `ARMCCnLIB` environment variables.

Some libraries are stored in subdirectories. If the compiler requires a library from a particular subdirectory, it adds an import of a special symbol to identify the subdirectory to the linker. The names of subdirectories are placed in each compiled object by using a symbol of the form `Lib$$Request$$sub_dir_name`.

4.24.1 Using the `--libpath` command-line option

Use the `--libpath` command-line option with a comma-separated list of parent directories. This list must end with the parent directory of the ARM library directories `armlib` and `cpplib`.

The linker searches subdirectories given by the symbol `Lib$$Request$$sub_dir_name`, if you include the path separator character on the end of the library path:

- `\` on Windows
- `/` on Red Hat Linux.

For example, for `--libpath=mylibs\` and the symbol `Lib$$Request$$armlib` the linker searches the directories:

```
mylibs
mylibs\armlib
```

———— **Note** —————

When the linker command-line option `--libpath` is used, any paths specified by the `ARMCCnLIB` variable are not searched.

The sequential nature of the search ensures that the linker chooses the library that appears earlier in the list if two or more libraries define the same symbol.

4.24.2 Using the `ARMCCnLIB` or `ARMLIB` environment variable

You can use either of the `ARMLIB` or `ARMCCnLIB` environment variables to specify a library path.

The linker searches subdirectories given by the symbol `Lib$$Request$$sub_dir_name`, if you include the path separator character on the end of the path specified in `ARMCCnLIB`:

- `\` on Windows
- `/` on Red Hat Linux.

For example, if `ARMCC5LIB` is set to `install_directory\lib\`, the linker searches the directories:

```
lib
lib\armlib
lib\cpplib
```

4.24.3 Library search order

The linker searches for libraries in the following order:

1. Relative to the current path.
2. At the location specified with the command-line option `--libpath`.

3. At the location specified in ARMCCnLIB.
4. At the location specified in ARMLIB.
5. At the location specified in ../lib.

4.24.4 How the linker selects ARM library variants

The ARM Compiler toolchain includes a number of variants of each of the libraries, that are built using different build options. For example, architecture versions, endianness, and instruction set. The variant of the ARM library is coded into the library name. The linker must select the best-suited variant from each of the directories identified during the library search.

The linker accumulates the attributes of each input object and then selects the library variant best suited to those attributes. If more than one of the selected libraries are equally suited, the linker retains the first library selected and rejects all others.

4.24.5 See also

Concepts

- [How the linker performs library searching, selection, and scanning on page 4-35.](#)

Reference

Linker Reference:

- [--libpath=pathlist on page 2-74.](#)

Using ARM® C and C++ Libraries and Floating-Point Support:

- [C and C++ library naming conventions on page 2-120.](#)

ARM® C and C++ Libraries and Floating-Point Support Reference:

- [Chapter 2 The C and C++ libraries.](#)

Introducing the ARM Compiler toolchain:

- [Toolchain environment variables on page 2-12.](#)

4.25 Specifying user libraries when linking

To specify user libraries:

- include them with path information explicitly in the input file list
- add the `--userlibpath` option to the `armlink` command line with a comma-separated list of directories, and then specify the names of the libraries as input files.

If you do not specify a full path name to a library on the command line, the linker tries to locate the library in the directories specified by the `--userlibpath` option. For example, if the directory `/mylib` contains `my_lib.a` and `other_lib.a`, add `/mylib/my_lib.a` to the input file list with the command:

```
armlink --userlibpath /mylib my_lib.a *.o
```

If you add a particular member from a library this does not add the library to the list of searchable libraries used by the linker. To load a specific member *and* add the library to the list of searchable libraries include the library *filename* on its own as well as specifying *library(member)*. For example, to load `strcmp.o` and place `mystring.lib` on the searchable library list add the following to the input file list:

```
mystring.lib(strcmp.o) mystring.lib
```

———— Note —————

Any search paths used for the ARM standard libraries specified by either the linker command-line option `--libpath` or the `ARMLIB` or `ARMCCnLIB` environment variables are not searched for user libraries.

4.25.1 See also

Tasks

- [Controlling how the linker searches for the ARM standard libraries on page 4-36.](#)

Reference

Linker Reference:

- [--libpath=pathlist on page 2-74](#)
- [--userlibpath=pathlist on page 2-136.](#)

ARM® C and C++ Libraries and Floating-Point Support Reference:

- [Chapter 2 The C and C++ libraries.](#)

Introducing the ARM Compiler toolchain:

- [Toolchain environment variables on page 2-12.](#)

4.26 How the linker resolves references

When the linker has constructed the list of libraries, it repeatedly scans each library in the list to resolve references. There are two separate lists of files that are maintained. The lists are scanned in the following order to resolve all dependencies:

1. List of system libraries found in `../lib`, or the directories specified by `--libpath`, `ARMCCnLIB`, or `ARMLIB`. These might also be specified by the `-Jdir[,dir,...]` compiler option.
2. The list of all other files that have been loaded. These might be specified by the `-Idir[,dir,...]` compiler option.

Each list is scanned using the following process:

1. Search all specified directories to select the most compatible library variants.
2. Add the variants to the list of libraries.
3. Scan each of the libraries to load the required members:
 - a. For each currently unsatisfied non-weak reference, search sequentially through the list of libraries for a matching definition. The first definition found is marked for step **b**.
 The sequential nature of the search ensures that the linker chooses the library that appears earlier in the list if two or more libraries define the same symbol. This enables you to override function definitions from other libraries, for example, the ARM C libraries, by adding your libraries to the input file list. However you must be careful to consistently override all the symbols in a library member. If you do not, you risk the objects from both libraries being loaded when there is a reference to an overridden symbol and a reference to a symbol that was not overridden. This results in a multiple symbol definition error L6200E for each overridden symbol.
 - b. Load the library members marked in step **a**. As each member is loaded it might satisfy some unresolved references, possibly including weak ones. Loading a library member might also create new unresolved weak and non-weak references.
 - c. Repeat these stages until all non-weak references are either resolved or cannot be resolved by any library.
4. If any non-weak reference remains unsatisfied at the end of the scanning operation, generate an error message.

4.26.1 See also

Concepts

- [About weak references and definitions on page 4-32](#)
- [How the linker performs library searching, selection, and scanning on page 4-35](#)
- [Controlling how the linker searches for the ARM standard libraries on page 4-36](#)
- [Specifying user libraries when linking on page 4-38.](#)

Reference

Linker Reference:

- [--libpath=pathlist on page 2-74.](#)

Compiler Reference:

- [-Idir\[,dir,...\] on page 3-83](#)
- [-Jdir\[,dir,...\] on page 3-93.](#)

Introducing the ARM Compiler toolchain:

- [Toolchain environment variables](#) on page 2-12.

Errors and Warnings Reference:

- [List of the armlink error and warning messages](#) on page 4-3.

4.27 Use of the strict family of options in the linker

The strict family of options are not directly related to error severity. Usually, you add a strict option because the standard linker checks are not precise enough or are potentially noisy with legacy objects.

The strict family of options are:

- `--strict`
- `--[no_]strict_enum_size`
- `--[no_]strict_flags`
- `--[no_]strict_ph`
- `--[no_]strict_relocations`
- `--[no_]strict_symbols`
- `--[no_]strict_visibility`
- `--[no_]strict_wchar_size`.

4.27.1 See also

Reference

Linker Reference:

- [--strict](#) on page 2-119
- [--strict_enum_size, --no_strict_enum_size](#) on page 2-120
- [--strict_flags, --no_strict_flags](#) on page 2-121
- [--strict_ph, --no_strict_ph](#) on page 2-122
- [--strict_relocations, --no_strict_relocations](#) on page 2-123
- [--strict_symbols, --no_strict_symbols](#) on page 2-124
- [--strict_visibility, --no_strict_visibility](#) on page 2-125
- [--strict_wchar_size, --no_strict_wchar_size](#) on page 2-126.

Chapter 5

Using linker optimizations

The following topics describe the optimizations available in the linker, `arm1ink`:

Tasks

- *Overriding the compression algorithm used by the linker on page 5-15*
- *Working with RW data compression on page 5-17*
- *Inlining functions with the linker on page 5-18*
- *Handling branches that optimize to a NOP on page 5-21*

Concepts

- *Elimination of common debug sections on page 5-2*
- *Elimination of common groups or sections on page 5-3*
- *Elimination of unused sections on page 5-4*
- *Elimination of unused virtual functions on page 5-6*
- *About linker feedback on page 5-7*
- *Example of using linker feedback on page 5-9*
- *About link-time code generation on page 5-11*
- *Optimization with RW data compression on page 5-13*
- *How the linker chooses a compressor on page 5-14*
- *How compression is applied on page 5-16*
- *Factors that influence function inlining on page 5-19*
- *About reordering of tail calling sections on page 5-22*
- *Restrictions on reordering of tail calling sections on page 5-23*
- *About merging comment sections on page 5-24.*

5.1 Elimination of common debug sections

In DWARF 2, the compiler and assembler generate one set of debug sections for each source file that contributes to a compilation unit. `arm1ink` can detect multiple copies of a debug section for a particular source file and discard all but one copy in the final image. This can result in a considerable reduction in image debug size.

In DWARF 3, common debug sections are placed in common groups. `arm1ink` discards all but one copy of each group with the same signature.

5.1.1 See also

Concepts

- [Input sections, output sections, regions, and Program Segments](#) on page 4-5.
- [Elimination of common groups or sections](#) on page 5-3
- [Elimination of unused sections](#) on page 5-4
- [Elimination of unused virtual functions](#) on page 5-6.

Reference

Compiler Reference:

- [--debug, --no_debug](#) on page 3-39.

Assembler Reference:

- [--debug](#) on page 2-20.

Other information

- The DWARF Debugging Standard web site, <http://www.dwarfstd.org/>.

5.2 Elimination of common groups or sections

The ARM compiler generates complete objects for linking. Therefore:

- If there are inline functions in C and C++ sources, each object contains the out-of-line copies of the inline functions that the object requires.
- If templates are used in C++ sources, each object contains the template functions that the object requires.

When these functions are declared in a common header file, the functions might be defined many times in separate objects that are subsequently linked together. To eliminate duplicates, the compiler compiles these functions into separate instances of common code sections or groups.

It is possible that the separate instances of common code sections, or groups, are not identical. Some of the copies, for example, might be found in a library that has been built with different, but compatible, build options, different optimization, or debug options.

If the copies are not identical, `arm1ink` retains the best available variant of each common code section, or group, based on the attributes of the input objects. `arm1ink` discards the rest.

If the copies are identical, `arm1ink` retains the first section or group located.

You control this optimization with the following linker options:

- Use the `--bestdebug` option to use the largest common data (COMDAT) group (likely to give the best debug view).
- Use the `--no_bestdebug` option to use the smallest COMDAT group (likely to give the smallest code size). This is the default.

Because `--no_bestdebug` is the default, the final image is the same regardless of whether you generate debug tables during compilation with `--debug`.

5.2.1 See also

Concepts

- [Input sections, output sections, regions, and Program Segments on page 4-5.](#)
- [Elimination of common debug sections on page 5-2](#)
- [Elimination of unused sections on page 5-4](#)
- [Elimination of unused virtual functions on page 5-6.](#)

Using the Compiler:

- [Inline functions on page 5-29.](#)

Reference

Linker Reference:

- [--bestdebug, --no_bestdebug on page 2-14.](#)

Compiler Reference:

- [--debug, --no_debug on page 3-39.](#)

5.3 Elimination of unused sections

Unused section elimination is the most significant optimization on image size that is performed by the linker. It removes unreachable code and data from the final image.

Unused section elimination is suppressed in cases that might result in the removal of all sections.

To control this optimization use the `--remove`, `--no_remove`, `--first`, `--last`, and `--keep` linker options.

Unused section elimination requires an entry point. Therefore, if there is no entry point specified for an image, use the `--entry` linker option to specify an entry point and permit unused section elimination to work, if it is enabled.

Note

By default, unused section elimination is disabled if you are building DLLs with `--dll`, or shared libraries with `--shared`. Therefore, you must explicitly include `--remove` to re-enable unused section elimination.

Use the `--info unused` linker option to instruct the linker to generate a list of the unused sections that it eliminates.

An input section is retained in the final image under the following conditions:

- if it contains an entry point
- if it is referred to, directly or indirectly, by a non-weak reference from an input section containing an entry point
- if it is specified as the first or last input section by the `--first` or `--last` option (or a scatter-loading equivalent)
- if it is marked as unremovable by the `--keep` option.

Note

Compilers usually collect functions and data together and emit one section for each category. The linker can only eliminate a section if it is entirely unused.

You can mark a function or variable in source code with the `__attribute__((used))` attribute. This causes `armcc` to generate the symbol `__tagsym$$used` for each of the functions and variables, and ensures that the function or variable is not removed by the linker.

You can also use the `--split_sections` compiler command-line option to instruct the compiler to generate one ELF section for each function in the source file.

5.3.1 See also

Concepts

- [Input sections, output sections, regions, and Program Segments on page 4-5](#)
- [About weak references and definitions on page 4-32](#)
- [Elimination of common debug sections on page 5-2](#)
- [Elimination of common groups or sections on page 5-3](#)
- [Elimination of unused virtual functions on page 5-6.](#)

Reference

Linker Reference:

- *--entry=location* on page 2-43
- *--first=section_id* on page 2-53
- *--info=topic[,topic,...]* on page 2-58
- *--keep=section_id* on page 2-67
- *--last=section_id* on page 2-71
- *--remove, --no_remove* on page 2-102.

Compiler Reference:

- *--split_sections* on page 3-156
- *__attribute__((used)) function attribute* on page 5-58
- *__attribute__((used)) variable attribute* on page 5-78.

5.4 Elimination of unused virtual functions

Unused section elimination efficiently removes unused functions from C code. In C++ applications, virtual functions and *RunTime Type Information* (RTTI) objects are referenced by pointer tables, known as vtables. Without extra information, the linker cannot determine which vtable entries are accessed at runtime. This means that the standard unused section elimination algorithm used by the linker cannot guarantee to remove unused virtual functions and RTTI objects.

Virtual Function Elimination (VFE) is a refinement of unused section elimination to reduce ROM size in images generated from C++ code. This optimization can be used to eliminate unused virtual functions and RTTI objects from your code.

An input section that contains more than one function can only be eliminated if all the functions are unused. The linker cannot remove unused functions from within a section.

VFE is a collaboration between the ARM compiler and the linker whereby the compiler supplies extra information about unused virtual functions that is then used by the linker. Based on this analysis, the linker is able to remove unused virtual functions and RTTI objects.

———— **Note** —————

For VFE to work, the assembler requires all objects using C++ to have VFE annotations. If the linker finds a C++ mangled symbol name in the symbol table of an object and VFE information is not present, it turns off the optimization.

The compiler places the extra information in sections with names beginning `.arm_vfe`. These sections are ignored by the linker when it is not VFE-aware.

5.4.1 See also

Concepts

- [Elimination of common debug sections on page 5-2](#)
- [Elimination of common groups or sections on page 5-3](#)
- [Elimination of unused sections on page 5-4](#).

Reference

Linker Reference:

- [--vfemode=mode on page 2-143](#).

Compiler Reference:

- [--rtti, --no_rtti on page 3-149](#).

5.5 About linker feedback

Linker feedback is a collaboration between the compiler and linker that can increase the amount of unused code that can be removed from an ELF image.

The feedback option produces a text file containing a list of unused functions, and functions that have been inlined by the linker. This information can be fed back to the compiler, which can rebuild the objects, placing these functions in their own sections. These sections can then be removed by the linker during usual unused section elimination.

The feedback file has the following format:

```
;<FEEDBACK># ARM Linker, N.nn [Build num]: Last Updated: day mmm dd hh:mm:ss yyyy
;VERSION 0.2
;FILE filename.o
unused_function <= USED 0
inlined_function <= LINKER_INLINED
...
```

The feedback file contains an entry for each object file. Each entry contains:

- the object filename specified as a comment:
;FILE *filename.o*
- a list of the functions in that file that are not used:
unused_function <= USED 0
- a list of the functions in that file that are inlined by the linker:
inlined_function <= LINKER_INLINED.

To use linker feedback, specify `--feedback file` on the linker and compiler command lines.

Note

The compiler issues a warning message if no feedback file exists. Therefore, you might want to leave the `--feedback file` option off the first invocation of the compiler.

Additional feedback options are provided by the linker:

- If you are using scatter-loading then an executable ELF image cannot be created if your code does not fit into the region limits described in your scatter file. In this case use the `--feedback_image=option` command-line option.
- To control the information that the linker puts into the feedback file, use the `--feedback_type=type` command-line option. You can control whether or not to list functions that require interworking or unused functions.

5.5.1 See also

Tasks

- [Inlining functions with the linker on page 5-18](#)
- [Chapter 8 Using scatter files.](#)

Concepts

- [Example of using linker feedback on page 5-9](#)

Reference

Linker Reference:

- [--feedback=*filename* on page 2-48](#)
- [--feedback_image=*option* on page 2-49](#)

- *--feedback_type=type* on page 2-50
- *--inline, --no_inline* on page 2-63
- *--scatter=filename* on page 2-109.

Compiler Reference:

- *--feedback=filename* on page 3-68.

5.6 Example of using linker feedback

To see how linker feedback works:

1. Create a file `fb.c` containing the code shown in this example:

Example 5-1 Feedback example

```
#include <stdio.h>

void legacy()
{
    printf("This is a legacy function that is no longer used.\n");
}

int cubed(int i)
{
    return i*i*i;
}

void main()
{
    int n = 3;
    printf("%d cubed = %d\n",n,cubed(n));
}
```

2. Compile the program, and ignore the warning that the feedback file does not exist:

```
armcc --asm -c --feedback fb.txt fb.c
```

This inlines the `cubed()` function by default, and creates an assembler file `fb.s` and an object file `fb.o`. In the assembler file, the code for `legacy()` and `cubed()` is still present. Because of the inlining, there is no call to `cubed()` from `main`.

An out-of-line copy of `cubed()` is kept because it is not declared as **static**.

3. Link the object file to create the linker feedback file with the command line:

```
armlink --info sizes --list fbout1.txt --feedback fb.txt fb.o -o fb.axf
```

Linker diagnostics are output to the file `fbout1.txt`.

The linker feedback file identifies the source file that contains the unused functions in a comment (not used by the compiler) and includes entries for the `legacy()` and `cubed()` functions:

```
;<FEEDBACK># ARM Linker, 5.02 [Build num]: Last Updated: Date
;VERSION 0.2
;FILE fb.o
cubed <= USED 0
legacy <= USED 0
```

This shows that the functions are not used.

4. Repeat the compile and link stages with a different diagnostics file:

```
armcc --asm -c --feedback fb.txt fb.c
```

```
armlink --info sizes --list fbout2.txt fb.o -o fb.axf
```

5. Compare the two diagnostics files, `fbout1.txt` and `fbout2.txt`, to see the sizes of the image components (for example, Code, RO Data, RW Data, and ZI Data). The Code component is smaller.

In the assembler file, `fb.s`, the `legacy()` and `cubed()` functions are no longer in the same area as the `main()` function. They are compiled into their own ELF sections. Therefore, `arm1link` can remove the `legacy()` and `cubed()` functions from the final image.

Note

To get the maximum benefit from linker feedback you have to do a full compile and link at least twice. However, a single compile and link using feedback from a previous build is usually sufficient.

5.6.1 See also

Concepts

- [About linker feedback on page 5-7.](#)

Reference

Linker Reference:

- [--feedback=filename on page 2-48](#)
- [--feedback_image=option on page 2-49](#)
- [--feedback_type=type on page 2-50](#)
- [--info=topic\[,topic,...\] on page 2-58](#)
- [--list=filename on page 2-77](#)
- [--scatter=filename on page 2-109.](#)

Compiler Reference:

- [--asm on page 3-16](#)
- [-c on page 3-23](#)
- [--feedback=filename on page 3-68.](#)
- [--inline, --no_inline on page 3-90.](#)

5.7 About link-time code generation

Link-time code generation (LTCG) enables cross source-file optimization by delaying code generation until the link stage. This can significantly reduce code size. To enable LTCG, compile your source with `-c --ltcg` to create objects in an intermediate format. You must then link these object files with `--ltcg` to instruct the linker to perform code generation.

Both `armcc` and `armlink` have a `--ltcg` command-line option to enable LTCG.

Note

The LTCG feature is deprecated. As an alternative ARM recommends you use the `--multifile` compiler option.

5.7.1 Considerations when using LTCG

Be aware of the following when using LTCG:

- If no input objects are compiled with `--ltcg`, the final object is the same as that produced without LTCG.
- Debug information is not preserved in any object files compiled with `--ltcg`. However, debug information is still preserved for any object files that are not compiled with `--ltcg`.
- You cannot use specific object file names in scatter files because LTCG causes a temporary object file to be created and used for linking. If you are using a scatter file, then you must match files using the wildcard `*`. Otherwise no match is made on the temporary file.
- If there is more than one entry point in the input object files provided to the linker, the linker clears all entry points. Therefore, you must specify an entry point on the linker command-line with:

```
--entry=symbol
```

Any input section containing an entry point is not removed. For example, where an object is built without `--ltcg`.

Note

The linker option `--entry=object(symbol)` is not supported when using LTCG.

5.7.2 Example

The following example shows a typical use of LTCG:

1. Create ELF object files `one.o` and `two.o` with `--ltcg`:

```
armcc -c --ltcg one.c -o one.o
```

```
armcc -c --ltcg two.c -o two.o
```

The compiler generates intermediate code into the object files.

2. Link them using the command:

```
armlink --ltcg one.o two.o -o ltcg_image.axf
```

The linker:

1. Combines all the immediate code together, losing the link to original object file names.
2. Performs code generation for the intermediate code.

- Creates the output image.

5.7.3 About the intermediate object files generated by link-time code generation

An intermediate ELF object file generated by *link-time code generation* (LTCG) includes a section called `.il`. This section is marked with the flags `SHF_EXECINSTR` and `SHF_ALLOC`.

Using the intermediate object file `one.o` from the example, you can use the `fromelf` command to view these sections, for example:

```
fromelf --text -v one.o
...
=====

** Section #1

    Name      : .il
    Type      : SHT_PROGBITS (0x00000001)
    Flags     : SHF_ALLOC + SHF_EXECINSTR (0x00000006)
    ...
    Link      : SHN_UNDEF
    ...
=====

** Section #3

    Name      : .rel.il
    Type      : SHT_REL (0x00000009)
    Flags     : None (0x00000000)
    ...
    Link      : Section 2 (.symtab)
    Info      : Section 1 (.il)
    ...
=====
...
```

5.7.4 See also

Reference

Linker Reference:

- [--entry=location](#) on page 2-43
- [--ltcg](#) on page 2-81.

Compiler Reference:

- [-c](#) on page 3-23
- [--ltcg](#) on page 3-110
- [--multifile, --no_multifile](#) on page 3-117.

Using the fromelf Image Converter:

- [--text](#) on page 4-53.

5.8 Optimization with RW data compression

RW data areas typically contain a large number of repeated values, such as zeros, that makes them suitable for compression. RW data compression is enabled by default to minimize ROM size.

The linker compresses the data. This data is then decompressed on the target at run time.

The ARM libraries contain some decompression algorithms and the linker chooses the optimal one to add to your image to decompress the data areas when the image is executed. You can override the algorithm chosen by the linker.

5.8.1 See also

Concepts

- [How compression is applied on page 5-16.](#)

Tasks

- [Overriding the compression algorithm used by the linker on page 5-15](#)
- [Working with RW data compression on page 5-17.](#)

Concepts

- [How the linker chooses a compressor on page 5-14](#)
- [How compression is applied on page 5-16.](#)

5.9 How the linker chooses a compressor

armlink gathers information about the content of data sections before choosing the most appropriate compression algorithm to generate the smallest image. If compression is appropriate, the linker can only use one data compressor for all the compressible data sections in the image. Different compression algorithms might be tried on these sections to produce the best overall size. Compression is applied automatically if:

Compressed data size + Size of decompressor < Uncompressed data size

When a compressor has been chosen, armlink adds the decompressor to the code area of your image. If the final image does not contain any compressed data, no decompressor is added.

5.9.1 See also

Concepts

- [Optimization with RW data compression](#) on page 5-13
- [How compression is applied](#) on page 5-16.

Tasks

- [Overriding the compression algorithm used by the linker](#) on page 5-15
- [Working with RW data compression](#) on page 5-17.

5.10 Overriding the compression algorithm used by the linker

You can override the compression algorithm used by the linker by either:

- using the `--datacompressor off` option to turn off compression
- specifying a compression algorithm.

To specify a compression algorithm, use the number of the required compressor on the linker command line, for example:

```
armlink --datacompressor 2 ...
```

Use the command-line option `--datacompressor list` to get a list of compression algorithms available in the linker:

```
armlink --datacompressor list
...
Num      Compression algorithm
=====
0        Run-length encoding
1        Run-length encoding, with LZ77 on small-repeats
2        Complex LZ77 compression
```

When choosing a compression algorithm be aware that:

- compressor 0 performs well on data with large areas of zero-bytes but few nonzero bytes
- compressor 1 performs well on data where the nonzero bytes are repeating
- compressor 2 performs well on data that contains repeated values.

The linker prefers compressor 0 or 1 where the data contains mostly zero-bytes (>75%). Compressor 2 is chosen where the data contains few zero-bytes (<10%). If the image is made up only of ARM code, then ARM decompressors are used automatically. If the image contains any Thumb code, Thumb decompressors are used. If there is no clear preference, all compressors are tested to produce the best overall size.

Note

It is not possible to add your own compressors into the linker. The algorithms that are available, and how the linker chooses to use them, might change in the future.

5.10.1 See also

Concepts

- [Optimization with RW data compression](#) on page 5-13
- [How the linker chooses a compressor](#) on page 5-14
- [How compression is applied](#) on page 5-16.

Tasks

- [Working with RW data compression](#) on page 5-17.

Reference

Linker Reference:

- `--datacompressor=opt` on page 2-30.

5.11 How compression is applied

Run-length compression encodes data as non-repeated bytes and repeated zero-bytes. Non-repeated bytes are output unchanged, followed by a count of zero-bytes. Lempel-Ziv 1977 (LZ77) compression keeps track of the last n bytes of data seen and, when a phrase is encountered that has already been seen, it outputs a pair of values corresponding to the position of the phrase in the previously-seen buffer of data, and the length of the phrase.

5.11.1 See also

Concepts

- [Optimization with RW data compression on page 5-13](#)
- [How the linker chooses a compressor on page 5-14.](#)

Tasks

- [Overriding the compression algorithm used by the linker on page 5-15](#)
- [Working with RW data compression on page 5-17.](#)

Reference

Linker Reference:

- [--datacompressor=opt on page 2-30.](#)

5.12 Working with RW data compression

When working with RW data compression:

- Use the linker option `--map` to see where compression has been applied to regions in your code.
- The linker in *RealView Compilation Tools* (RVCT) v4.0 and later turns off RW compression if there is a reference from a compressed region to a linker-defined symbol that uses a load address.
- If you are using an ARM processor with on-chip cache, enable the cache after decompression to avoid code coherency problems.

Compressed data sections are automatically decompressed at run time, providing `__main` is executed, using code from the ARM libraries. This code must be placed in a root region. This is best done using `InRoot$$Sections` in a scatter file.

If you are using a scatter file, you can specify that a load or execution region is not to be compressed by adding the `NOCOMPRESS` attribute.

5.12.1 See also

Concepts

- [Optimization with RW data compression](#) on page 5-13
- [How the linker chooses a compressor](#) on page 5-14
- [How compression is applied](#) on page 5-16
- [Load\\$\\$ execution region symbols](#) on page 7-7
- [Chapter 8 Using scatter files.](#)

Tasks

- [Overriding the compression algorithm used by the linker](#) on page 5-15.

Reference

Linker Reference:

- [--map, --no_map](#) on page 2-83
- [Chapter 4 Formal syntax of the scatter file.](#)

5.13 Inlining functions with the linker

The linker can inline small functions in place of a branch instruction to that function. For the linker to be able to do this, the function (without the return instruction) must fit in the four bytes of the branch instruction.

Use the `--inline` and `--no_inline` command-line options to control branch inlining.

If branch inlining optimization is enabled, the linker scans each function call in the image and then inlines as appropriate. When the linker finds a suitable function to inline, it replaces the function call with the instruction from the function that is being called.

The linker applies branch inlining optimization before any unused sections are eliminated so that inlined sections can also be removed if they are no longer called.

———— **Note** ————

The linker can inline two 16-bit instructions in place of the 32-bit Thumb BL instruction.

Use the `--info=inline` command-line option to list all the inlined functions.

5.13.1 See also

Reference

- [Factors that influence function inlining](#) on page 5-19
- [Elimination of unused sections](#) on page 5-4.

Linker Reference:

- `--info=topic[,topic,...]` on page 2-58
- `--inline, --no_inline` on page 2-63.

5.14 Factors that influence function inlining

The following factors influence the way functions are inlined:

- The linker handles only the simplest cases and does not inline any instructions that read or write to the PC because this depends on the location of the function.
- If your image contains both ARM and Thumb code, functions that are called from the opposite state must be built for interworking. The linker can inline functions containing up to two 16-bit Thumb instructions. However, an ARM calling function can only inline functions containing a single 16-bit Thumb instruction or 32-bit Thumb instruction.
- The action that the linker takes depends on the size of the function being called. The following table shows the state of both the calling function and the function being called:

Table 5-1 Inlining small functions

Calling function state	Called function state	Called function size
ARM	ARM	4 to 8 bytes
ARM	Thumb	2 to 6 bytes
Thumb	Thumb	2 to 6 bytes

The linker can inline in different states if there is an equivalent instruction available. For example, if a Thumb instruction is `adds r0, r0` then the linker can inline the equivalent ARM instruction. It is not possible to inline from ARM to Thumb because there is less chance of Thumb equivalent to an ARM instruction.

- For a function to be inlined, the last instruction of the function must be either:
`MOV pc, lr`
or
`BX lr`
A function that consists only of a return sequence can be inlined as a NOP.
- A conditional ARM instruction can only be inlined if either:
 - The condition on the BL matches the condition on the instruction being inlined. For example, `BLEQ` can only inline an instruction with a matching condition like `ADDEQ`.
 - The BL instruction or the instruction to be inlined is unconditional. An unconditional ARM BL can inline any conditional or unconditional instruction that satisfies all the other criteria. An instruction that cannot be conditionally executed cannot be inlined if the BL instruction is conditional.
- A BL that is the last instruction of a Thumb *If-Then* (IT) block cannot inline a 16-bit Thumb instruction or a 32-bit MRS, MSR, or CPS instruction. This is because the IT block changes the behavior of the instructions within its scope so inlining the instruction changes the behavior of the program.

5.14.1 See also

Concepts

- [Handling branches that optimize to a NOP on page 5-21.](#)

Using the Assembler:

- [Conditional instructions on page 6-2.](#)

Reference*Assembler Reference:*

- *ADD, SUB, RSB, ADC, SBC, and RSC* on page 3-30
- *B, BL, BX, BLX, and BXJ* on page 3-45
- *CPS* on page 3-57
- *MOV and MVN* on page 3-97
- *MRS (PSR to general-purpose register)* on page 3-104
- *MSR (general-purpose register to PSR)* on page 3-107
- *IT* on page 3-65.

5.15 Handling branches that optimize to a NOP

By default, the linker replaces any branch with a relocation that resolves to the next instruction with a NOP instruction. This optimization can also be applied if the linker reorders tail calling sections.

However, there are cases where you might want to disable the option, for example, when performing verification or pipeline flushes.

To control this optimization, use the `--branchnop` and `--no_branchnop` command-line options.

5.15.1 See also

Concepts

- [About reordering of tail calling sections on page 5-22.](#)

Reference

Linker Reference:

- [--branchnop, --no_branchnop on page 2-15.](#)

5.16 About reordering of tail calling sections

A tail calling section is a section that contains a branch instruction at the end of the section. If the branch instruction has a relocation that targets a function at the start of another section, the linker can place the tail calling section immediately before the called section. The linker can then optimize the branch instruction at the end of the tail calling section to a NOP instruction.

You can take advantage of this behavior by using the command-line option `--tailreorder` to move tail calling sections immediately before their target.

Use the `--info=tailreorder` command-line option to display information about any tail call optimizations performed by the linker.

5.16.1 See also

Concepts

- [Veneer types](#) on page 4-28
- [Handling branches that optimize to a NOP](#) on page 5-21
- [Restrictions on reordering of tail calling sections](#) on page 5-23.

Reference

Linker Reference:

- [--info=topic\[,topic,...\]](#) on page 2-58
- [--tailreorder, --no_tailreorder](#) on page 2-129.

5.17 Restrictions on reordering of tail calling sections

The linker:

- Can only move one tail calling section for each tail call target. If there are multiple tail calls to a single section, the tail calling section with an identical section name is moved before the target. If no section name is found in the tail calling section that has a matching name, then the linker moves the first section it encounters.
- Cannot move a tail calling section out of its execution region.
- Does not move tail calling sections before inline veneers.

5.17.1 See also

Concepts

- [About reordering of tail calling sections on page 5-22.](#)

5.18 About merging comment sections

If input object files have any `.comment` sections that are identical, then the linker merges them to produce the smallest `.common` section while retaining all useful information.

The linker associates each input `.comment` section with the filename of the corresponding input object. If it merges identical `.comment` sections, then all the filenames that contain the common section are listed before the section contents, for example:

```
file1.o  
file2.o  
.comment section contents.
```

The linker merges these sections by default. To prevent the merging of identical `.comment` sections, use the `--no_filtercomment` command-line option.

Note

If you do not want to retain the information in a `.comment` section, then you can use the `--no_comment_section` option to strip this section from the image.

5.18.1 See also

Reference

Linker Reference:

- [--comment_section, --no_comment_section](#) on page 2-24
- [--filtercomment, --no_filtercomment](#) on page 2-51

Chapter 6

Getting information about images

The following topics describe how to get image information from armLink:

Tasks

- *Identifying the source of some link errors on page 6-3*
- *How to find where a symbol is placed when linking on page 6-6.*

Concepts

- *Linker options for getting information about images on page 6-2*
- *Example of using the --info linker option on page 6-4.*

6.1 Linker options for getting information about images

You can use following options to get information about how your image is generated by the linker:

`--info=topic[,topic,...]`

Displays information about various topics.

`--map`

Displays the image memory map, and contains the address and the size of each load region, execution region, and input section in the image, including linker-generated input sections. It also shows how RW data compression is applied.

`--section_index_display`

Use with `--map` to change the display of the index column.

`--show_cmdline`

Outputs the command-line used by the linker.

`--symbols`

Displays a list of each local and global symbol used in the link step, and its value.

`--verbose`

Displays detailed information about the link operation, including the objects that are included and the libraries that contain them.

`--xref`

Displays a list of all cross-references between input sections.

`--xrefdbg`

Displays a list of all cross-references between input debug sections.

The information can be written to a file using the `--list=filename` option.

6.1.1 See also

Concepts

- [Section alignment with the linker on page 4-22](#)
- [Optimization with RW data compression on page 5-13.](#)
- [Identifying the source of some link errors on page 6-3](#)
- [Example of using the `--info` linker option on page 6-4](#)

Reference

Linker Reference:

- [--info=topic\[,topic,...\] on page 2-58](#)
- [--list=filename on page 2-77](#)
- [--map, --no_map on page 2-83](#)
- [--section_index_display=type on page 2-110](#)
- [--show_cmdline on page 2-111](#)
- [--symbols, --no_symbols on page 2-127](#)
- [--verbose on page 2-141](#)
- [--xref, --no_xref on page 2-146](#)
- [--xrefdbg, --no_xrefdbg on page 2-147.](#)

6.2 Identifying the source of some link errors

You can use `--info` inputs to identify the source of some link errors. For example, you can search the output to locate undefined references from library objects or multiply defined symbols caused by retargeting some library functions and not others. Search backwards from the end of this output to find and resolve link errors.

You can also use the `--verbose` option to output similar text with additional information on the linker operations.

6.2.1 See also

Concepts

- [Linker options for getting information about images](#) on page 6-2.

Reference

Linker Reference:

- `--info=topic[,topic,...]` on page 2-58
- `--verbose` on page 2-141.

(incl. Generated)

armlink might generate image contents, for example, interworking veneers, and input sections such as region tables. If the Object Totals row includes this type of data, it is shown in this row.

In the example, there are 19 bytes of RO data in total, of which 16 bytes is linker-generated RO data.

Library Totals

Shows how many bytes are occupied by library members that have been extracted and added to the image as individual objects.

(incl. Padding)

armlink inserts padding, if required, to force section alignment. If the Object Totals row includes this type of data, it is shown in the associated (incl. Padding) row. Similarly, if the Library Totals row includes this type of data, it is shown in its associated row.

In the example, there are 19 bytes of RO data in the object total, of which 3 bytes is linker-generated padding, and 805 bytes of RO data in the library total, with 6 bytes of padding.

Grand Totals

Shows the true size of the image. In the example, there are 10200 bytes of ZI data (in Object Totals) and 300 of ZI data (in Library Totals) giving a total of 10500 bytes.

ELF Image Totals

If you are using RW data compression (the default) to optimize ROM size, the size of the final image changes and this is reflected in the output from `--info`. Compare the number of bytes under Grand Totals and ELF Image Totals to see the effect of compression.

In the example, RW data compression is not enabled. If data is compressed, the RW value changes.

ROM Totals

Shows the minimum size of ROM required to contain the image. This does not include ZI data and debug information which is not stored in the ROM.

6.3.1 See also**Concepts**

- [Linker options for getting information about images on page 6-2.](#)

Reference

Linker Reference:

- [--info=topic\[,topic,...\] on page 2-58.](#)

6.4 How to find where a symbol is placed when linking

To find where a symbol is placed in an ELF image file when linking, use the `--keep=section_id` and `--map` options to view the image memory map. For example, if `object(section)` is the section containing the symbol, enter:

```
armlink --keep=object(section) --map s.o --output=s.axf
```

The memory map shows where the section containing the symbol is placed.

6.4.1 Example

Do the following:

1. Create the file `s.c` containing the following source code:


```
long long altstack[10] __attribute__((section("STACK"), zero_init));

int main()
{
    return sizeof(altstack);
}
```
2. Compile the source:


```
armcc -c s.c -o s.o
```
3. Link the object `s.o`, keeping the `STACK` symbol and displaying the image memory map:


```
armlink --keep=s.o(STACK) --map s.o --output=s.axf
```
4. Locate the `STACK` symbol in the output, for example:

```
...
Execution Region ER_RW (Base: 0x000081c8, Size: 0x00000000, Max: 0xffffffff, ABSOLUTE)
**** No section assigned to this execution region ****
```

```
Execution Region ER_ZI (Base: 0x000081c8, Size: 0x000000b0, Max: 0xffffffff, ABSOLUTE)
```

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
0x000081c8	0x00000060	Zero	RW	42	.bss	libspace.o(c_4.1)
0x00008228	0x00000050	Zero	RW	2	STACK	s.o

This shows that the stack is placed in execution region `ER_ZI`.

6.4.2 See also

Tasks

- [Using fromelf to find where a symbol is placed in an executable ELF image on page 3-8.](#)

Reference

Compiler Reference:

- [-c on page 3-23](#)
- [-o filename on page 3-120.](#)

Linker Reference:

- [--keep=section_id on page 2-67](#)
- [--map, --no_map on page 2-83](#)
- [--output=filename on page 2-89.](#)

Chapter 7

Accessing and managing symbols with armlink

The following topics describe how to access and manage symbols with the linker, armlink:

Tasks

- *About mapping symbols on page 7-3*
- *Accessing linker-defined symbols on page 7-4*
- *Using scatter files on page 7-11*
- *Importing linker-defined symbols in C and C++ on page 7-12*
- *Importing linker-defined symbols in ARM assembler on page 7-13*
- *Accessing symbols in another image on page 7-17*
- *Creating a symdefs file on page 7-18*
- *Outputting a subset of the global symbols on page 7-19*
- *Reading a symdefs file on page 7-20*
- *Specifying steering files on the linker command-line on page 7-24*
- *Hiding and renaming global symbols with a steering file on page 7-27*
- *Using `$Super$$` and `$Sub$$` to patch symbol definitions on page 7-28.*

Concepts

- *Region-related symbols on page 7-5*
- *Region name values when not scatter-loading on page 7-10*
- *Section-related symbols on page 7-14*
- *What is a steering file? on page 7-23.*

Reference

- *Image\$\$ execution region symbols on page 7-6*
- *Load\$\$ execution region symbols on page 7-7*

- *Load\$\$LR\$\$ load region symbols on page 7-9*
- *Image symbols on page 7-15*
- *Input section symbols on page 7-16*
- *Symdefs file format on page 7-21*
- *Steering file command summary on page 7-25*
- *Steering file format on page 7-26.*

7.1 About mapping symbols

Mapping symbols are generated by armcc and armasm to identify inline transitions between:

- code and data at literal pool boundaries
- ARM code and Thumb code, such as ARM/Thumb interworking veneers.

The mapping symbols are:

\$a	start of a sequence of ARM instructions
\$t	start of a sequence of Thumb instructions
\$t.x	start of a sequence of ThumbEE instructions
\$d	start of a sequence of data items, such as a literal pool.

armlink generates the \$d.realdata mapping symbol to communicate to fromelf that the data is from a non-executable section. Therefore, the code and data sizes output by fromelf -z are the same as the output from armlink --info sizes, for example:

```
Code (inc. data)  R0 Data
      x           y       z
```

In this example, the y is marked with \$d, and R0 Data is marked with \$d.realdata.

Note

Symbols beginning with the characters \$v are mapping symbols related to VFP and might be output when building for a target with VFP. Avoid using symbols beginning with \$v in your source code.

7.1.1 See also

Concepts

Using the Assembler:

- [Symbol naming rules](#) on page 8-3.

Reference

Linker Reference:

- [--list_mapping_symbols, --no_list_mapping_symbols](#) on page 2-78
- [--strict_symbols, --no_strict_symbols](#) on page 2-124.

Using the fromelf Image Converter:

- [--text](#) on page 4-53.

Other information

- [ELF for the ARM Architecture](#),
<http://infocenter.arm.com/help/topic/com.arm.doc.ih0044-/index.html>.

7.2 Accessing linker-defined symbols

The linker defines some symbols that contain the character sequence `$$`. These symbols, and all other external names containing the sequence `$$`, are names reserved by ARM.

You can import these symbolic addresses and use them as relocatable addresses by your assembly language programs, or refer to them as **extern** symbols from your C or C++ source code.

Be aware that:

- If you use the `--strict` compiler command-line option, the compiler does not accept symbol names containing dollar symbols. To re-enable support, include the `--dollar` option on the compiler command line.
- Linker-defined symbols are only generated when your code references them.

7.2.1 See also

Concepts

- [Importing linker-defined symbols in C and C++ on page 7-12](#)
- [Importing linker-defined symbols in ARM assembler on page 7-13.](#)

Reference

Compiler Reference:

- [--dollar, --no_dollar on page 3-57](#)
- [--strict, --no_strict on page 3-157.](#)

7.3 Region-related symbols

The linker generates the following types of region-related symbols for each region in the image:

- Image\$\$
- Load\$\$
- Load\$\$LR\$\$.

If you are using a scatter file these symbols are generated for each region in the scatter file.

If you are not using scatter-loading, the symbols are generated for the default region names. That is, the region names are fixed and the same types of symbol are supplied.

7.3.1 See also

Concepts

- [Image\\$\\$ execution region symbols on page 7-6](#)
- [Load\\$\\$ execution region symbols on page 7-7](#)
- [Load\\$\\$LR\\$\\$ load region symbols on page 7-9](#)
- [Region name values when not scatter-loading on page 7-10.](#)

7.4 Image\$\$ execution region symbols

The following table shows the symbols that the linker generates for every execution region present in the image. All the symbols refer to execution addresses after the C library is initialized.

Table 7-1 Image\$\$ execution region symbols

Symbol	Description
Image\$\$region_name\$\$Base	Execution address of the region.
Image\$\$region_name\$\$Length	Execution region length in bytes excluding ZI length.
Image\$\$region_name\$\$Limit	Address of the byte beyond the end of the non-ZI part of the execution region.
Image\$\$region_name\$\$RO\$\$Base	Execution address of the RO output section in this region.
Image\$\$region_name\$\$RO\$\$Length	Length of the RO output section in bytes.
Image\$\$region_name\$\$RO\$\$Limit	Address of the byte beyond the end of the RO output section in the execution region.
Image\$\$region_name\$\$RW\$\$Base	Execution address of the RW output section in this region.
Image\$\$region_name\$\$RW\$\$Length	Length of the RW output section in bytes.
Image\$\$region_name\$\$RW\$\$Limit	Address of the byte beyond the end of the RW output section in the execution region.
Image\$\$region_name\$\$ZI\$\$Base	Execution address of the ZI output section in this region.
Image\$\$region_name\$\$ZI\$\$Length	Length of the ZI output section in bytes.
Image\$\$region_name\$\$ZI\$\$Limit	Address of the byte beyond the end of the ZI output section in the execution region.

7.4.1 See also

Tasks

- [Importing linker-defined symbols in C and C++ on page 7-12](#)
- [Importing linker-defined symbols in ARM assembler on page 7-13.](#)

Concepts

- [Region-related symbols on page 7-5](#)
- [Region name values when not scatter-loading on page 7-10.](#)

7.5 Load execution region symbols

The linker performs an extra address assignment and relocation pass for relocations that refer to load addresses after RW compression. This delayed relocation permits more information about load addresses to be used in linker-defined symbols.

———— **Note** ————

Load $_{region_name}$ symbols apply only to execution regions, and Load $_{LR}_{load_region_name}$ symbols apply only to load regions.

The following table shows the symbols that the linker generates for every Load execution region present in the image. All the symbols refer to execution addresses after the C library is initialized.

Table 7-2 Load execution region symbols

Symbol	Description
Load $_{region_name}$ Base	Load address of the region.
Load $_{region_name}$ Length	Region length in bytes.
Load $_{region_name}$ Limit	Address of the byte beyond the end of the execution region.
Load $_{region_name}$ ROBase	Address of the RO output section in this execution region.
Load $_{region_name}$ ROLength	Length of the RO output section in bytes.
Load $_{region_name}$ ROLimit	Address of the byte beyond the end of the RO output section in the execution region.
Load $_{region_name}$ RWBase	Address of the RW output section in this execution region.
Load $_{region_name}$ RWLength	Length of the RW output section in bytes.
Load $_{region_name}$ RWLimit	Address of the byte beyond the end of the RW output section in the execution region.
Load $_{region_name}$ ZIBase	Load address of the ZI output section in this execution region.
Load $_{region_name}$ ZILength	Load length of the ZI output section in bytes. The Load Length of ZI is zero unless <i>region_name</i> has the ZEROPAD scatter-loading keyword set. If ZEROPAD is set then: Load Length = Image $_{region_name}$ ZILength
Load $_{region_name}$ ZILimit	Load address of the byte beyond the end of the ZI output section in the execution region.

All symbols in this table refer to load addresses before the C library is initialized. Be aware of the following:

- The symbols are absolute because section-relative symbols can only have execution addresses.
- The symbols take into account RW compression.

- The symbols do not include ZI output section because it does not exist before the C library is initialized.
- All relocations from RW compressed execution regions must be performed before compression, because the linker cannot resolve a delayed relocation on compressed data.
- If the linker detects a relocation from an RW-compressed region to a linker-defined symbol that depends on RW compression, then the linker disables compression for that region.
- Any zero bytes written to the file are visible. Therefore, the Limit and Length values must take into account the zero bytes written into the file.

7.5.1 See also

Concepts

- [Optimization with RW data compression](#) on page 5-13
- [Region-related symbols](#) on page 7-5
- [Image\\$\\$ execution region symbols](#) on page 7-6
- [Load\\$\\$LR\\$\\$ load region symbols](#) on page 7-9
- [Region name values when not scatter-loading](#) on page 7-10.

Reference:

Linker Reference:

- [Execution region attributes](#) on page 4-12

7.6 Load region symbols

A Load region can contain many execution regions, so there are no separate `$_R0` and `$_RW` components.

———— **Note** —————

`Load $_{region_name}$` symbols apply only to load regions, and `Load $_{region_name}$` symbols apply only to execution regions.

The following table shows the symbols that the linker generates for every Load region present in the image.

Table 7-3 Load region symbols

Symbol	Description
<code>Load$_{region_name}$\$_Base</code>	address of the load region
<code>Load$_{region_name}$\$_Length</code>	length of the load region
<code>Load$_{region_name}$\$_Limit</code>	address of the byte beyond the end of the load region

7.6.1 See also

Concepts

- [The image structure on page 4-3](#)
- [Input sections, output sections, regions, and Program Segments on page 4-5](#)
- [Load view and execution view of an image on page 4-6](#)
- [Region-related symbols on page 7-5.](#)

7.7 Region name values when not scatter-loading

If you are not using scatter-loading, the linker uses region name values of:

- ER_R0, for the read-only execution region
- ER_RW, for the read-write execution region
- ER_ZI, for the zero-initialized execution region.

You can insert these names into the following symbols to obtain the required address:

- Image\$\$ execution region symbols
- Load\$\$ execution region symbols.

For example, Load\$\$ER_R0\$\$Base.

Note

- The ZI output sections of an image are not created statically, but are automatically created dynamically at runtime. Therefore, there is no load address symbol for ZI output sections.
 - It is recommended that you use region-related symbols in preference to section-related symbols.
-

7.7.1 See also

Concepts

- [Region-related symbols on page 7-5](#)
- [Image\\$\\$ execution region symbols on page 7-6](#)
- [Load\\$\\$ execution region symbols on page 7-7](#)
- [Section-related symbols on page 7-14.](#)

7.8 Using scatter files

If you are using scatter-loading, the names from a scatter file are used in the linker defined symbols. The scatter file:

- Names all the execution regions in the image, and provides their load and execution addresses.
- Defines both stack and heap. The linker also generates special stack and heap symbols.

7.8.1 See also

Tasks

- [Chapter 8 Using scatter files.](#)

Reference

Linker Reference:

- [--scatter=filename on page 2-109.](#)

7.9 Importing linker-defined symbols in C and C++

You can import linker-defined symbols into your C or C++ source code either by value or by reference:

Import by value

```
extern unsigned int symbol_name;
```

Import by reference

```
extern void *symbol_name;
```

If you declare a symbol as an int, then you must use the address-of operator (&) to obtain the correct value as shown in these examples:

Example 7-1 Importing a linker-defined symbol

```
extern unsigned int Image$$ZI$$Limit;
config.heap_base = (unsigned int) &Image$$ZI$$Limit;
```

Example 7-2 Importing symbols that define a ZI output section

```
extern unsigned int Image$$ZI$$Length;
extern char Image$$ZI$$Base[];
memset(Image$$ZI$$Base,0,(unsigned int)&Image$$ZI$$Length);
```

7.9.1 See also

Concepts

- [Image\\$\\$ execution region symbols on page 7-6.](#)

7.10 Importing linker-defined symbols in ARM assembler

To import linker-defined symbols into your assembler source code, use the `IMPORT` directive and create a 32-bit data word to hold the value of the symbol, for example:

```
IMPORT |Image$$ZI$$Limit|  
...  
zi_limit DCD |Image$$ZI$$Limit|
```

To load the value into a register, such as `r1`, use the `LDR` instruction:

```
LDR r1, zi_limit
```

The `LDR` instruction must be able to reach the 32-bit data word. The accessible memory range varies between ARM and Thumb, and the architecture you are using.

7.10.1 See also

Concepts

- [Image\\$\\$ execution region symbols on page 7-6.](#)

Reference

Assembler Reference:

- [Memory access instructions on page 3-9](#)
- [IMPORT and EXTERN on page 5-59.](#)

7.11 Section-related symbols

Section-related symbols are symbols generated by the linker when it creates an image without scatter-loading.

The linker generates the following types of section-related symbols:

- Image symbols, if you use command-line options to create a simple image. A simple image has three output sections (RO, RW, and ZI) that produce the three execution regions.
- Input section symbols, for every input section present in the image.

The linker sorts sections within an execution region first by attribute RO, RW, or ZI, then by name. So, for example, all `.text` sections are placed in one contiguous block. A contiguous block of sections with the same attribute and name is known as a *consolidated section*.

7.11.1 See also

Concepts

- [Image symbols on page 7-15](#)
- [Input section symbols on page 7-16](#).

7.12 Image symbols

Image symbols are generated by the linker when you use a command-line option to create a simple image.

The following table shows the image symbols:

Table 7-4 Image symbols

Symbol	Section type	Description
Image\$\$RO\$\$Base	Output	Address of the start of the RO output section.
Image\$\$RO\$\$Limit	Output	Address of the first byte beyond the end of the RO output section.
Image\$\$RW\$\$Base	Output	Address of the start of the RW output section.
Image\$\$RW\$\$Limit	Output	Address of the byte beyond the end of the ZI output section. (The choice of the end of the ZI region rather than the end of the RW region is to maintain compatibility with legacy code.)
Image\$\$ZI\$\$Base	Output	Address of the start of the ZI output section.
Image\$\$ZI\$\$Limit	Output	Address of the byte beyond the end of the ZI output section.

If you are using a scatter file, the image symbols are undefined. If your code accesses any of these symbols, you must treat them as a weak reference.

The standard implementation of `__user_setup_stackheap()` uses the value in `Image$$ZI$$Limit`. Therefore, if you are using a scatter file you must manually place the stack and heap. You can do this either:

- in a scatter file using one of the following methods:
 - define separate stack and heap regions called `ARM_LIB_STACK` and `ARM_LIB_HEAP`
 - define a combined region containing both stack and heap called `ARM_LIB_STACKHEAP`.
- by re-implementing `__user_setup_stackheap()` to set the heap and stack boundaries.

7.12.1 See also

Tasks

- [Linker-defined symbols that are not defined when scatter-loading on page 8-10](#)
- [Specifying stack and heap using the scatter file on page 8-11.](#)

Concepts

- [Types of simple image on page 4-10](#)
- [About weak references and definitions on page 4-32.](#)

Using the Compiler:

- [Stack use in C and C++ on page 5-17.](#)

Reference

Using ARM® C and C++ Libraries and Floating-Point Support:

- [__user_setup_stackheap\(\) on page 2-60.](#)

7.13 Input section symbols

Input section symbols are generated by the linker for every input section present in the image.

The following table shows the input section symbols:

Table 7-5 Section-related symbols

Symbol	Section type	Description
<i>SectionName\$\$Base</i>	Input	Address of the start of the consolidated section called <i>SectionName</i> .
<i>SectionName\$\$Length</i>	Input	Length of the consolidated section called <i>SectionName</i> (in bytes).
<i>SectionName\$\$Limit</i>	Input	Address of the byte beyond the end of the consolidated section called <i>SectionName</i> .

If your code refers to the input-section symbols, it is assumed that you expect all the input sections in the image with the same name to be placed contiguously in the image memory map.

If your scatter file places input sections non-contiguously, the linker issues an error. This is because the use of the base and limit symbols over non-contiguous memory is ambiguous.

7.13.1 See also

Tasks

- [Chapter 8 Using scatter files.](#)

Concepts

- [Input sections, output sections, regions, and Program Segments](#) on page 4-5.

7.14 Accessing symbols in another image

If you want one image to know the global symbol values of another image, you can use a *symbol definitions* (symdefs) file.

You can use this, for example, if you have one image that always resides in ROM and multiple images that are loaded into RAM. The images loaded into RAM can access global functions and data from the image located in ROM.

7.14.1 See also

Tasks

- [Creating a symdefs file on page 7-18](#)
- [Reading a symdefs file on page 7-20.](#)

Reference

- [Symdefs file format on page 7-21.](#)

7.15 Creating a symdefs file

Use the armlink option `--symdefs=filename` to generate a symdefs file.

The linker produces a symdefs file during a successful final link stage. It is not produced for partial linking or for unsuccessful final linking.

Note

If *filename* does not exist, the file is created containing all the global symbols. If *filename* exists, the existing contents of *filename* are used to select the symbols that are output when the linker rewrites the file. This means that only the existing symbols in the filename are updated, and no new symbols (if any) are added at all. If you do not want this behavior, ensure that any existing symdefs file is deleted before the link step.

7.15.1 See also

Concepts

- [Accessing symbols in another image on page 7-17.](#)

Reference

- [Symdefs file format on page 7-21.](#)

Linker Reference:

- [--symdefs=filename on page 2-128.](#)

7.16 Outputting a subset of the global symbols

By default, all global symbols are written to the symdefs file. When a symdefs file exists, the linker uses its contents to restrict the output to a subset of the global symbols.

For an application `image1` containing symbols that you want to expose to another application using a symdefs file:

1. Specify `--symdefs=filename` when you are doing a final link for `image1`. The linker creates a symdefs file `filename`.
2. Open `filename` in a text editor, remove any symbol entries you do not want in the final list, and save the file.
3. Specify `--symdefs=filename` when you are doing a final link for `image1`.
You can edit `filename` at any time to add comments and link `image1` again, for example, to update the symbol definitions after one or more objects used to create `image1` have changed.

You can now use the symdefs file to link additional applications.

7.16.1 See also

Concepts

- [Accessing symbols in another image on page 7-17.](#)

Reference

- [Symdefs file format on page 7-21.](#)

Linker Reference:

- [--symdefs=filename on page 2-128.](#)

7.17 Reading a symdefs file

A symdefs file can be considered as an object file with symbol information but no code or data. To read a symdefs file, add it to your file list as you do for any object file. The linker reads the file and adds the symbols and their values to the output symbol table. The added symbols have ABSOLUTE and GLOBAL attributes.

If a partial link is being performed, the symbols are added to the output object symbol table. If a full link is being performed, the symbols are added to the image symbol table.

The linker generates error messages for invalid rows in the file. A row is invalid if:

- any of the columns are missing
- any of the columns have invalid values.

The symbols extracted from a symdefs file are treated in exactly the same way as symbols extracted from an object symbol table. The same restrictions apply regarding multiple symbol definitions.

———— **Note** —————

The same function name or symbol name cannot be defined in both ARM code and in Thumb code.

7.17.1 See also

Reference

- [Symdefs file format on page 7-21.](#)

7.18 Symdefs file format

The symdefs file defines symbols and their values. The file consists of:

Identification line

The identification line in a symdefs file comprises:

- an identifying string, #<SYMDEFS>#, which must be the first 11 characters in the file for the linker to recognize it as a symdefs file
- linker version information, in the format:
ARM Linker, *N.nn* [*Build num*]:
- date and time of the most recent update of the symdefs file, in the format:
Last Updated: *Date*

The version and update information are not part of the identifying string.

Comments You can insert comments manually with a text editor. Comments have the following properties:

- The first line must start with the special identifying comment #<SYMDEFS>#. This comment is inserted by the linker when the file is produced and must not be manually deleted.
- Any line where the first non-whitespace character is a semicolon (;) or hash (#) is a comment.
- A semicolon (;) or hash (#) after the first non-whitespace character does not start a comment.
- Blank lines are ignored and can be inserted to improve readability.

Symbol information

The symbol information is provided on a single line, and comprises:

Symbol value	The linker writes the absolute address of the symbol in fixed hexadecimal format, for example, <code>0x00008000</code> . If you edit the file, you can use either hexadecimal or decimal formats for the address value.
Type flag	A single letter to show symbol type:
	A ARM code
	T Thumb code
	D Data
	N Number.
Symbol name	The symbol name.

7.18.1 Example symdefs file

This example shows a typical symdefs file format:

Example 7-3 Symdefs file format

```
#<SYMDEFS># ARM Linker, 5.01 [Build num]: Last Updated: Date
;value type name, this is an added comment
0x00008000 A __main
0x00008004 A __scatterload
0x000080E0 T main
0x0000814D T _main_arg
```

```
0x0000814D T __argv_alloc
0x00008199 T __rt_get_argv
...
# This is also a comment, blank lines are ignored
...
0x0000A4FC D __stdin
0x0000A540 D __stdout
0x0000A584 D __stderr
0xFFFFFFFF N __SIG_IGN
```

7.18.2 See also

Concepts

- [Accessing symbols in another image on page 7-17](#)
- [Creating a symdefs file on page 7-18](#)
- [Outputting a subset of the global symbols on page 7-19](#)
- [Reading a symdefs file on page 7-20.](#)

7.19 What is a steering file?

A steering file is a text file that contains a set of commands to edit the symbol tables of output objects and the dynamic sections of images. Steering file commands enable you to:

- manage symbols in the symbol table
- control the copying of symbols from the static symbol table to the dynamic symbol table
- store information about the libraries that a link unit depends on.

For example, you can use steering files to protect intellectual property, or avoid namespace clashes.

7.19.1 See also

Tasks

- [Specifying steering files on the linker command-line](#) on page 7-24.

Reference

- [Steering file command summary](#) on page 7-25
- [Steering file format](#) on page 7-26.

Linker Reference:

- [--edit=file_list](#) on page 2-38.

7.20 Specifying steering files on the linker command-line

Use the option `--edit file-list` to specify one or more steering files on the linker command-line.

When you specify more than one steering file, you can use either of the following command-line formats:

```
armlink --edit file1 --edit file2 --edit file3
```

```
armlink --edit file1,file2,file3
```

Do not include spaces between the comma and the filenames when using a comma-separated list.

7.20.1 See also

Concepts

- [What is a steering file? on page 7-23.](#)

Reference

- [Steering file command summary on page 7-25](#)
- [Steering file format on page 7-26.](#)

Linker Reference:

- [EXPORT on page 3-2](#)
- [HIDE on page 3-3](#)
- [IMPORT on page 3-4](#)
- [RENAME on page 3-5](#)
- [REQUIRE on page 3-7](#)
- [RESOLVE on page 3-8](#)
- [SHOW on page 3-10.](#)

7.21 Steering file command summary

The steering file commands are:

Table 7-6 Steering file command summary

Command	Description
EXPORT	Specifies that a symbol can be accessed by other shared objects or executables.
HIDE	Makes defined global symbols in the symbol table anonymous.
IMPORT	Specifies that a symbol is defined in a shared object at runtime.
RENAME	Renames defined and undefined global symbol names.
REQUIRE	Creates a DT_NEEDED tag in the dynamic array. DT_NEEDED tags specify dependencies to other shared objects used by the application, for example, a shared library.
RESOLVE	Matches specific undefined references to a defined global symbol.
SHOW	Makes global symbols visible. This command is useful if you want to make a specific symbol visible that is hidden using a HIDE command with a wildcard.

Note

The steering file commands control only global symbols. Local symbols are not affected by any of these commands.

7.21.1 See also

Tasks

- [Specifying steering files on the linker command-line](#) on page 7-24.

Concepts

- [What is a steering file?](#) on page 7-23.

Reference

- [Steering file format](#) on page 7-26.

Linker Reference:

- [EXPORT](#) on page 3-2
- [HIDE](#) on page 3-3
- [IMPORT](#) on page 3-4
- [RENAME](#) on page 3-5
- [REQUIRE](#) on page 3-7
- [RESOLVE](#) on page 3-8
- [SHOW](#) on page 3-10.

7.22 Steering file format

A steering file is a plain text file of the following format:

- Lines with a semicolon (;) or hash (#) character as the first non-whitespace character are interpreted as comments. A comment is treated as a blank line.
- Blank lines are ignored.
- Each non-blank, non-comment line is either a command, or part of a command that is split over consecutive non-blank lines.
- Command lines that end with a comma (,) as the last non-whitespace character is continued on the next non-blank line.

Each command line consists of a command, followed by one or more comma-separated operand groups. Each operand group comprises either one or two operands, depending on the command. The command is applied to each operand group in the command. The following rules apply:

- Commands are case-insensitive, but are conventionally shown in uppercase.
- Operands are case-sensitive because they must be matched against case-sensitive symbol names. You can use wildcard characters in operands.

Commands are applied to global symbols only. Other symbols, such as local symbols, are not affected.

The following example shows a sample steering file:

Example 7-4 Example steering file

```
; Import my_func1 as func1
IMPORT my_func1 AS func1

# Rename a very long function name to a shorter name
RENAME a_very_long_function_name AS,
      short_func_name
```

7.22.1 See also

Tasks

- [Specifying steering files on the linker command-line](#) on page 7-24.

Concepts

- [What is a steering file?](#) on page 7-23.

Reference

- [Steering file command summary](#) on page 7-25.

Linker Reference:

- [EXPORT](#) on page 3-2
- [HIDE](#) on page 3-3
- [IMPORT](#) on page 3-4
- [RENAME](#) on page 3-5
- [REQUIRE](#) on page 3-7
- [RESOLVE](#) on page 3-8
- [SHOW](#) on page 3-10.

7.23 Hiding and renaming global symbols with a steering file

You can use a steering file to hide and rename global symbol names in output files. You use the HIDE and RENAME commands accordingly.

For example, you can use steering files to protect intellectual property, or avoid namespace clashes.

Example of renaming a symbol:

Example 7-5 RENAME steering command example

```
RENAME func1 AS my_func1
```

Example of hiding symbols:

Example 7-6 HIDE steering command example

```
; Hides all global symbols with the 'internal' prefix
HIDE internal*
```

7.23.1 See also

Tasks

- [Specifying steering files on the linker command-line](#) on page 7-24.

Concepts

- [What is a steering file?](#) on page 7-23
- [Steering file command summary](#) on page 7-25.

Reference

- [Steering file format](#) on page 7-26.

Linker Reference:

- [--edit=file_list](#) on page 2-38
- [HIDE](#) on page 3-3
- [RENAME](#) on page 3-5.

7.24 Using \$\$Super\$\$ and \$\$Sub\$\$ to patch symbol definitions

There are situations where an existing symbol cannot be modified because, for example, it is located in an external library or in ROM code. In such cases you can use the \$\$Super\$\$ and \$\$Sub\$\$ patterns to patch an existing symbol.

To patch the definition of the function foo():

`$$Super$$foo` Identifies the original unpatched function foo(). Use this to call the original function directly.

`$$Sub$$foo` Identifies the new function that is called instead of the original function foo(). Use this to add processing before or after the original function.

Note

The \$\$Sub\$\$ and \$\$Super\$\$ mechanism only works at static link time, \$\$Super\$\$ references cannot be imported or exported into the dynamic symbol table.

The following example shows how to insert a call to the function ExtraFunc() before the call to the legacy function foo().

Example 7-7 Using \$\$Super\$\$ and \$\$Sub\$\$

```
extern void ExtraFunc(void); extern void $$Super$$foo(void):

/* this function is called instead of the original foo() */
void $$Sub$$foo(void)
{
    ExtraFunc(); /* does some extra setup work */
    $$Super$$foo(); /* calls the original foo() function */
    /* To avoid calling the original foo() function
     * omit the $$Super$$foo(); function call.
     */
}

```

7.24.1 See also

Other information

- *ELF for the ARM Architecture*,
<http://infocenter.arm.com/help/topic/com.arm.doc.ih0044-/index.html>

Chapter 8

Using scatter files

The following topics describe how you use scatter files with `armlink` to create complex images:

Tasks

- *Specifying stack and heap using the scatter file* on page 8-11
- *Creating root execution regions* on page 8-13
- *Using the FIXED attribute to create root regions* on page 8-16
- *Placing functions and data at specific addresses* on page 8-17
- *Placing a named section explicitly using scatter-loading* on page 8-22
- *Placing unassigned sections with the .ANY module selector* on page 8-23
- *Selecting veneer input sections in scatter-loading descriptions* on page 8-32
- *Using `__attribute__((section("name")))` to place code and data* on page 8-33
- *Using `__at` sections to place sections at a specific address* on page 8-35
- *Placing a key in flash memory using `__at`* on page 8-41
- *Placing a structure over a peripheral register using `__at`* on page 8-43
- *Reserving an empty region* on page 8-52
- *Using preprocessing commands in a scatter file* on page 8-57
- *Using expression evaluation in a scatter file to avoid padding* on page 8-59.

Concepts

- *About scatter-loading* on page 8-3
- *When to use scatter-loading* on page 8-4
- *Scatter-loading command-line option* on page 8-5

- *Images with a simple memory map on page 8-6*
- *Images with a complex memory map on page 8-8*
- *Linker-defined symbols that are not defined when scatter-loading on page 8-10*
- *What is a root region? on page 8-12*
- *Examples of using placement algorithms for .ANY sections on page 8-26*
- *Example of next_fit algorithm showing behavior of full regions, selectors, and priority on page 8-28*
- *Examples of using sorting algorithms for .ANY sections on page 8-30*
- *Restrictions on placing __at sections on page 8-36*
- *Automatic placement of __at sections on page 8-37*
- *Manual placement of __at sections on page 8-39*
- *Placement of sections with overlays on page 8-44*
- *About placing ARM C and C++ library code on page 8-47*
- *Example of placing code in a root region on page 8-48*
- *Example of placing ARM C library code on page 8-49*
- *Example of placing ARM C++ library code on page 8-50*
- *Example of placing ARM library helper functions on page 8-51*
- *About creating regions on page boundaries on page 8-54*
- *Overalignment of execution regions and input sections on page 8-56*
- *Expression evaluation in scatter files on page 8-58*
- *Equivalent scatter-loading descriptions for simple images on page 8-60*
- *Type 1 image, one load region and contiguous execution regions on page 8-61*
- *Type 2 image, one load region and non-contiguous execution regions on page 8-63*
- *Type 3 image, two load regions and non-contiguous execution regions on page 8-65*
- *Scatter file to ELF mapping on page 8-67.*

8.1 About scatter-loading

The scatter-loading mechanism enables you to specify the memory map of an image to the linker using a description in a text file. Scatter-loading gives you complete control over the grouping and placement of image components. You can use scatter-loading to create simple images, but it is generally only used for images that have a complex memory map. That is, where multiple memory regions are scattered in the memory map at load and execution time.

An image memory map is made up of regions and output sections. Every region in the memory map can have a different load and execution address.

To construct the memory map of an image, the linker must have:

- grouping information that describes how input sections are grouped into output sections and regions
- placement information that describes the addresses where regions are to be located in the memory maps.

When the linker creates an image using a scatter file, it creates some region-related symbols. The linker creates these special symbols only if your code references them.

8.1.1 See also

Concepts

- [The image structure on page 4-3](#)
- [Region-related symbols on page 7-5](#)
- [When to use scatter-loading on page 8-4](#)
- [Scatter file to ELF mapping on page 8-67](#).

8.2 When to use scatter-loading

The command-line options to the linker give some control over the placement of data and code, but complete control of placement requires more detailed instructions than can be entered on the command line.

Situations where scatter-loading is either required or very useful:

Complex memory maps

Code and data that must be placed into many distinct areas of memory require detailed instructions on where to place the sections in the memory space.

Different types of memory

Many systems contain a variety of physical memory devices such as flash, ROM, SDRAM, and fast SRAM. A scatter-loading description can match the code and data with the most appropriate type of memory. For example, interrupt code might be placed into fast SRAM to improve interrupt response time but infrequently-used configuration information might be placed into slower flash memory.

Memory-mapped peripherals

The scatter-loading description can place a data section at a precise address in the memory map so that memory mapped peripherals can be accessed.

Functions at a constant location

A function can be placed at the same location in memory even though the surrounding application has been modified and recompiled. This is useful for jump table implementation.

Using symbols to identify the heap and stack

Symbols can be defined for the heap and stack location when the application is linked.

Scatter-loading is usually required for implementing embedded systems because these use ROM, RAM, and memory-mapped peripherals.

8.2.1 See also

Concepts

- [About scatter-loading on page 8-3.](#)

8.3 Scatter-loading command-line option

The `arm1ink` command-line option for using scatter-loading is:

```
--scatter=description_file
```

This instructs the linker to construct the image memory map as described in *description_file*.

Be aware that you cannot use `--scatter` with the following memory map related command-line options:

- `--partial`
- `--ro_base`
- `--rw_base`
- `--ropi`
- `--rwpi`
- `--rosplit`
- `--split`
- `--reloc`
- `--startup`
- `--zi_base`.

8.3.1 See also

Concepts

- [About scatter-loading on page 8-3](#)
- [When to use scatter-loading on page 8-4](#)
- [Equivalent scatter-loading descriptions for simple images on page 8-60.](#)

Reference

Linker Reference:

- [--partial on page 2-94](#)
- [--reloc on page 2-100](#)
- [--ro_base=address on page 2-103](#)
- [--ropi on page 2-104](#)
- [--rosplit on page 2-105](#)
- [--rw_base=address on page 2-106](#)
- [--rwpi on page 2-107](#)
- [--scatter=filename on page 2-109](#)
- [--split on page 2-117](#)
- [--startup=symbol, --no_startup on page 2-118](#)
- [--zi_base=address on page 2-149](#)
- [Chapter 4 Formal syntax of the scatter file.](#)

8.4 Images with a simple memory map

If an image has a simple memory map, you can either:

- use a scatter file
- specify the memory map using basic linker command-line options.

The following figure shows a simple memory map:

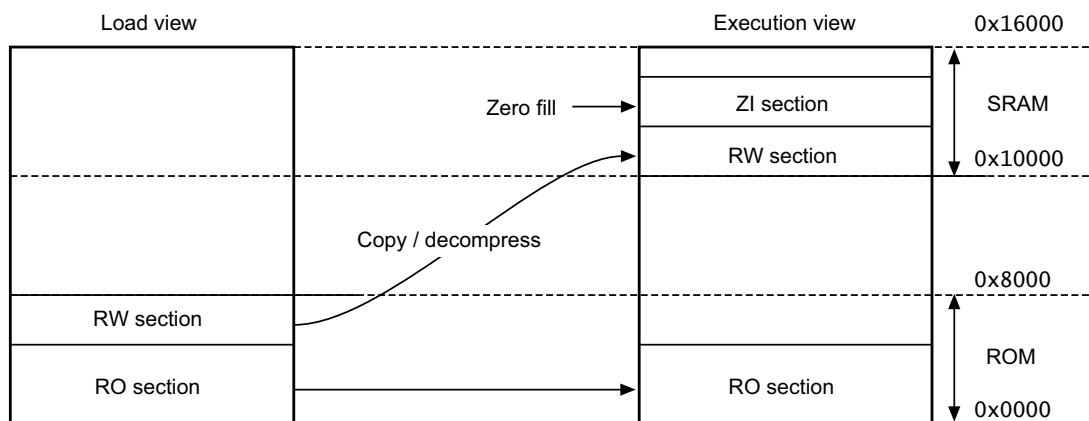


Figure 8-1 Simple scatter-loaded memory map

The following example shows the corresponding scatter-loading description that loads the segments from the object file into memory:

Example 8-1 Simple memory map in a scatter file

```
LOAD_ROM 0x0000 0x8000 ; Name of load region (LOAD_ROM),
                       ; Start address for load region (0x0000),
                       ; Maximum size of load region (0x8000)
{
  EXEC_ROM 0x0000 0x8000 ; Name of first exec region (EXEC_ROM),
                          ; Start address for exec region (0x0000),
                          ; Maximum size of first exec region (0x8000)
  {
    * (+RO) ; Place all code and RO data into
            ; this exec region
  }
  SRAM 0x10000 0x6000 ; Name of second exec region (RAM),
                      ; Start address of second exec region (0x10000),
                      ; Maximum size of second exec region (0x6000)
  {
    * (+RW, +ZI) ; Place all RW and ZI data into
                 ; this exec region
  }
}
```

The maximum size specifications for the regions are optional. However, if you include them, they enable the linker to check that a region does not overflow its boundary.

In this example, you can achieve the same result, apart from the limit checking with the following linker command-line:

```
armlink --ro_base 0x0 --rw_base 0x10000
```

8.4.1 See also

Concepts

- [About scatter-loading](#) on page 8-3
- [When to use scatter-loading](#) on page 8-4
- [Scatter file to ELF mapping](#) on page 8-67.

Reference

Linker Reference:

- [--ro_base=address](#) on page 2-103
- [--rw_base=address](#) on page 2-106.

8.5 Images with a complex memory map

For images with a complex memory map, you cannot specify the memory map using basic linker command-line options. Such images require the use of a scatter file.

The following figure shows a complex memory map:




Figure 8-2 Complex memory map

The following example shows the corresponding scatter-loading description that loads the segments from the program1.o and program2.o files into memory:

Example 8-2 Complex memory map in a scatter file

```

LOAD_ROM_1 0x0000          ; Start address for first load region (0x0000)
{
  EXEC_ROM_1 0x0000        ; Start address for first exec region (0x0000)
  {
    program1.o (+R0)       ; Place all code and R0 data from
                          ; program1.o into this exec region
  }

  DRAM 0x18000 0x8000     ; Start address for this exec region (0x18000),
                          ; Maximum size of this exec region (0x8000)
  {
    program1.o (+RW, +ZI) ; Place all RW and ZI data from
                          ; program1.o into this exec region
  }
}

LOAD_ROM_2 0x4000          ; Start address for second load region (0x4000)
{
  EXEC_ROM_2 0x4000        ; Start address for second exec region (0x4000)
  {
    program2.o (+R0)       ; Place all code and R0 data from
                          ; program2.o into this exec region
  }
}

```

```

SRAM 0x8000 0x8000
{
    program2.o (+RW, +ZI) ; Place all RW and ZI data from
                          ; program2.o into this exec region
}
}

```

Caution

The scatter-loading description in this example specifies the location for code and data for program1.o and program2.o only. If you link an additional module, for example, program3.o, and use this description file, the location of the code and data for program3.o is not specified.

Unless you want to be very rigorous in the placement of code and data, it is advisable to use the * or .ANY specifier to place leftover code and data.

8.5.1 See also

Tasks

- [Creating root execution regions on page 8-13](#)
- [Using the FIXED attribute to create root regions on page 8-16.](#)

Concepts

- [About scatter-loading on page 8-3](#)
- [When to use scatter-loading on page 8-4](#)
- [Scatter file to ELF mapping on page 8-67.](#)

Linker Reference:

- [Scatter files containing relative base address load regions and a ZI execution region on page 4-37.](#)

8.6 Linker-defined symbols that are not defined when scatter-loading

Be aware that the following symbols are undefined when a scatter file is used:

- Image\$\$RW\$\$Base
- Image\$\$RW\$\$Limit
- Image\$\$RO\$\$Base
- Image\$\$RO\$\$Limit
- Image\$\$ZI\$\$Base
- Image\$\$ZI\$\$Limit

If you use a scatter file but do not use the special region names for stack and heap, or do not re-implement `__user_setup_stackheap()`, an error message is generated.

8.6.1 See also

Tasks

- [Accessing linker-defined symbols on page 7-4](#)
- [Specifying stack and heap using the scatter file on page 8-11.](#)

8.7 Specifying stack and heap using the scatter file

The ARM C library provides multiple implementations of the function `__user_setup_stackheap()`, and can select the correct one for you automatically from information given in a scatter file.

To select the two region memory model, define two special execution regions in your scatter file named `ARM_LIB_HEAP` and `ARM_LIB_STACK`. Both regions have the `EMPTY` attribute. This causes the library to select the non-default implementation of `__user_setup_stackheap()` that uses the value of the symbols:

- `Image$$ARM_LIB_STACK$$Base`
- `Image$$ARM_LIB_STACK$$ZI$$Limit`
- `Image$$ARM_LIB_HEAP$$Base`
- `Image$$ARM_LIB_HEAP$$ZI$$Limit`

Only one `ARM_LIB_STACK` or `ARM_LIB_HEAP` region can be specified, and you must allocate a size, for example:

```
LOAD_FLASH ...
{
    ...
    ARM_LIB_STACK 0x40000 EMPTY -0x20000 ; Stack region growing down
    { }
    ARM_LIB_HEAP 0x28000000 EMPTY 0x80000 ; Heap region growing up
    { }
    ...
}
```

You can use a combined stack and heap region by defining a single execution region named `ARM_LIB_STACKHEAP`, with the `EMPTY` attribute. This causes `__user_setup_stackheap()` to use the value of the symbols `Image$$ARM_LIB_STACKHEAP$$Base` and `Image$$ARM_LIB_STACKHEAP$$ZI$$Limit`.

———— Note —————

If you re-implement `__user_setup_stackheap()`, this overrides all library implementations.

8.7.1 See also

Reference

- [Region-related symbols on page 7-5.](#)

Using ARM® C and C++ Libraries and Floating-Point Support:

- [__user_setup_stackheap\(\) on page 2-60](#)
- [Legacy function __user_initial_stackheap\(\) on page 2-71.](#)

8.8 What is a root region?

A root region is a region with the same load and execution address. The initial entry point of the image must be in a root region. If the initial entry point is not in a root region, the link fails and the linker gives an error message.

Example 8-3 Root region with the same load and execution address

```

LR_1 0x040000      ; load region starts at 0x40000
{
  ER_RO 0x040000   ; load address = execution address
  {
    * (+RO)        ; all RO sections (must include section with
                  ; initial entry point)
  }
  ...              ; rest of scatter-loading description
}

```

8.8.1 See also

Tasks

- [Creating root execution regions on page 8-13](#)
- [Using the FIXED attribute to create root regions on page 8-16](#)
- [About placing ARM C and C++ library code on page 8-47.](#)

Concepts

- [The image structure on page 4-3.](#)

8.9 Creating root execution regions

To specify a region as a root region in a scatter file you can:

- Specify ABSOLUTE as the attribute for the execution region, either explicitly or by permitting it to default, and use the same address for the first execution region and the enclosing load region. To make the execution region address the same as the load region address, either:
 - Specify the same numeric value for both the base address for the execution region and the base address for the load region.
 - Specify a +0 offset for the first execution region in the load region.
 If an offset of zero (+0) is specified for all subsequent execution regions in the load region, then all execution regions not following an execution region containing ZI are also root regions.

The following example shows an implicitly defined root region:

Example 8-4 Implicit root region with the same load and execution address

```

LR_1 0x040000      ; load region starts at 0x40000
{
  ; start of execution region descriptions
  ER_RO 0x040000   ; load address = execution address
  {
    * (+RO)        ; all RO sections (must include section with
                  ; initial entry point)
  }
  ...             ; rest of scatter-loading description
}

```

- Use the FIXED execution region attribute to ensure that the load address and execution address of a specific region are the same. You can use the FIXED attribute to place any execution region at a specific address in ROM. For example, the following memory map shows fixed execution regions:

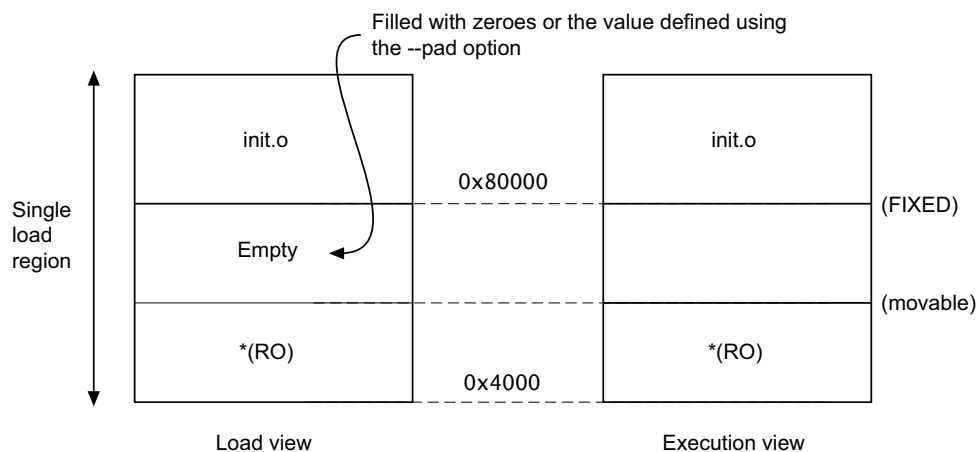


Figure 8-3 Memory map for fixed execution regions

The following example shows the corresponding scatter-loading description:

Example 8-5 Using the FIXED attribute

```

LR_1 0x040000          ; load region starts at 0x40000
{
    ER_RO 0x040000      ; start of execution region descriptions
    {
        ER_RO 0x040000 ; load address = execution address
        {
            * (+RO)     ; RO sections other than those in init.o
        }
    }
    ER_INIT 0x080000 FIXED ; load address and execution address of this
    {
        init.o(+RO)    ; execution region are fixed at 0x80000
    }
    ...                ; all RO sections from init.o
    ...                ; rest of scatter-loading description
}

```

8.9.1 Examples of misusing the FIXED attribute

The following example shows common cases where the FIXED execution region attribute is misused:

Example 8-6 Misuse of the FIXED attribute

```

LR1 0x8000
{
    ER_LOW +0 0x1000
    {
        * (+RO)
    }
    ; At this point the next available Load and Execution address is 0x8000 + size of
    ; contents of ER_LOW. The maximum size is limited to 0x1000 so the next available Load
    ; and Execution address is at most 0x9000
    ER_HIGH 0xF0000000 FIXED
    {
        * (+RW+ZI)
    }
    ; The required execution address and load address is 0xF0000000. The linker inserts
    ; 0xF0000000 - (0x8000 + size of(ER_LOW)) bytes of padding so that load address matches
    ; execution address
}

; The other common misuse of FIXED is to give a lower execution address than the next
; available load address.

LR_HIGH 0x100000000
{
    ER_LOW 0x1000 FIXED
    {
        * (+RO)
    }
    ; The next available load address in LR_HIGH is 0x100000000. The required Execution
    ; address is 0x1000. Because the next available load address in LR_HIGH must increase
    ; monotonically the linker cannot give ER_LOW a Load Address lower than 0x100000000
}

```

8.9.2 See also

Tasks

- [Using the FIXED attribute to create root regions on page 8-16.](#)

Concepts

- [What is a root region? on page 8-12.](#)

Linker Reference:

- [About load region descriptions on page 4-5](#)
- [About execution region descriptions on page 4-9](#)
- [Considerations when using a relative address +offset for load regions on page 4-17](#)
- [Considerations when using a relative address +offset for execution regions on page 4-18.](#)

Reference

Linker Reference:

- [Load region attributes on page 4-8](#)
- [Execution region attributes on page 4-12](#)
- [Address attributes for load and execution regions on page 4-15.](#)

Assembler Reference:

- [ENTRY on page 5-35.](#)

8.10 Using the FIXED attribute to create root regions

You can use the `FIXED` attribute in an execution region scatter file to create root regions that load and execute at fixed addresses.

`FIXED` is used to create multiple root regions within a single load region and therefore typically a single ROM device. For example, you can use this to place a function or a block of data, such as a constant table or a checksum, at a fixed address in ROM so that it can be accessed easily through pointers.

If you specify, for example, that some initialization code is to be placed at start of ROM and a checksum at the end of ROM, some of the memory contents might be unused. Use the `*` or `.ANY` module selector to flood fill the region between the end of the initialization block and the start of the data block.

To make your code easier to maintain and debug, it is suggested that you use the minimum amount of placement specifications in scatter files and leave the detailed placement of functions and data to the linker.

You cannot specify component objects that have been partially linked. For example, if you partially link the objects `obj1.o`, `obj2.o`, and `obj3.o` together to produce `obj_all.o`, the component object names are discarded in the resulting object. Therefore, you cannot refer to one of the objects by name, for example, `obj1.o`. You can refer only to the combined object `obj_all.o`.

Note

There are some situations where using `FIXED` and a single load region are not appropriate. Other techniques for specifying fixed locations are:

- If your loader can handle multiple load regions, place the RO code or data in its own load region.
 - If you do not require the function or data to be at a fixed location in ROM, use `ABSOLUTE` instead of `FIXED`. The loader then copies the data from the load region to the specified address in RAM. `ABSOLUTE` is the default attribute.
 - To place a data structure at the location of memory-mapped I/O, use two load regions and specify `UNINIT`. `UNINIT` ensures that the memory locations are not initialized to zero.
-

8.10.1 See also

Concepts

Linker Reference:

- [About execution region descriptions on page 4-9.](#)

Reference

Linker Reference:

- [Load region attributes on page 4-8](#)
- [Execution region attributes on page 4-12](#)
- [Address attributes for load and execution regions on page 4-15.](#)

8.11 Placing functions and data at specific addresses

Normally, the compiler produces RO, RW and ZI sections from a single source file. These regions contain all the code and data from the source file. To place a single function or data item at a fixed address, you must enable the linker to process the function or data separately from the rest of the input files.

The linker has two methods that enable you to place a section at a specific address:

- You can create a scatter file that defines an execution region at the required address with a section description that selects only one section.
- For a specially-named section the linker can get the placement address from the section name. These specially-named sections are called `__at` sections.

To place a function or variable at a specific address it must be placed in its own section. There are several ways to do this:

- Place the function or data item in its own source file.
- Use `__attribute__((at(address)))` to place variables in a separate section at a specific address.
- Use `__attribute__((section("name")))` to place functions and variables in a named section.
- Use the AREA directive from assembly language. In assembly code, the smallest locatable unit is an AREA.
- Use the `--split_sections` compiler option to generate one ELF section for each function in the source file.

This option results in a small increase in code size for some functions because it reduces the potential for sharing addresses, data, and string literals between functions. However, this can help to reduce the final image size overall by enabling the linker to remove unused functions when you specify `armlink --remove`.

8.11.1 Example of placing a variable at a specific address without scatter-loading

This example shows how to modify your source code to place code and data at specific addresses, and does not require a scatter file:

1. Create the source file `main.c` containing the following code:

```
#include <stdio.h>

extern int sqr(int n1);
int gSquared __attribute__((at(0x5000))); // Place at 0x5000

int main()
{
    gSquared=sqr(3);
    printf("Value squared is: %d\n", gSquared);
}
```

2. Create the source file `function.c` containing the following code:

```
int sqr(int n1)
{
    return n1*n1;
}
```

3. Compile and link the sources:

```
armcc -c -g function.c
armcc -c -g main.c
armlink --map function.o main.o -o squared.axf
```

The `--map` option displays the memory map of the image. Also, `--autoat` is the default.

In this example, `__attribute__((at(0x5000)))` specifies that the global variable `gSquared` is to be placed at the absolute address `0x5000`. `gSquared` is placed in the execution region `ER$.ARM.__AT_0x00005000` and load region `LR$.ARM.__AT_0x00005000`.

———— **Note** —————

Although the address is specified as `0x5000` in the source file, the region and section name addresses are normalized to eight hexadecimal digits.

The memory map shows:

```
...
Load Region LR$.ARM.__AT_0x00005000 (Base: 0x00005000, Size: 0x00000004, Max: 0x00000004, ABSOLUTE)

Execution Region ER$.ARM.__AT_0x00005000 (Base: 0x00005000, Size: 0x00000004, Max: 0x00000004, ABSOLUTE,
UNINIT)

Base Addr   Size       Type  Attr   Idx   E Section Name      Object
0x00005000  0x00000004  Zero  RW           13  .ARM.__AT_0x00005000  main.o
```

8.11.2 Example of placing a variable in a named section with scatter-loading

This example shows how to modify your source code to place code and data in a specific section using a scatter file:

1. Create the source file `main.c` containing the following code:

```
#include <stdio.h>

extern int sqr(int n1);
int gSquared __attribute__((section("foo"))); // Place in section foo

int main()
{
    gSquared=sqr(3);
    printf("Value squared is: %d\n", gSquared);
}
```

2. Create the source file `function.c` containing the following code:

```
int sqr(int n1)
{
    return n1*n1;
}
```

3. Create the scatter file `scatter.sc` containing the following load region:

```
LR1 0x0000 0x20000
{
    ER1 0x0 0x2000
    {
        *(+R0) ; rest of code and read-only data
    }
    ER2 0x8000 0x2000
    {
        main.o
    }
}
```

```

ER3 0x10000 0x2000
{
    function.o
    *(foo)                ; Place gSquared in ER3
}
RAM 0x200000 (0x1FF00-0x2000) ; RW & ZI data to be placed at 0x200000
{
    *(+RW, +ZI)
}
ARM_LIB_STACK 0x800000 EMPTY -0x10000
{
}
ARM_LIB_HEAP +0 EMPTY 0x10000
{
}
}

```

The ARM_LIB_STACK and ARM_LIB_HEAP regions are required because the program is being linked with the semihosting libraries.

4. Compile and link the sources:

```

armcc -c -g function.c
armcc -c -g main.c
armmlink --map --scatter=scatter.scat function.o main.o -o squared.axf

```

The --map option displays the memory map of the image. Also, --autoat is the default.

In this example, `__attribute__((section("foo")))` specifies that the global variable `gSquared` is to be placed in a section called `foo`. The scatter file specifies that the section `foo` is to be placed in the ER3 execution region.

The memory map shows:

```

Load Region LR1 (Base: 0x00000000, Size: 0x000015d0, Max: 0x00020000, ABSOLUTE)
...
Execution Region ER3 (Base: 0x00010000, Size: 0x00000010, Max: 0x00002000, ABSOLUTE)

Base Addr    Size           Type  Attr    Idx  E Section Name    Object
0x00010000   0x0000000c    Code  RO      3    .text             function.o
0x0001000c   0x00000004    Data  RW      15    foo                main.o
...

```

———— Note ————

If you omit `*(foo)` from the scatter file, the section is placed in the region of the same type. That is RAM in this example.

8.11.3 Example of placing a variable at a specific address with scatter-loading

This example shows how to modify your source code to place code and data at a specific address using a scatter file:

1. Create the source file `main.c` containing the following code:

```

#include <stdio.h>

extern int sqr(int n1);

// Place at address 0x10000
const int gValue __attribute__((section(".ARM.__at_0x10000"))) = 3;

int main()

```

```

{
    int squared;
    squared=sqr(gValue);
    printf("Value squared is: %d\n", squared);
}

```

2. Create the source file `function.c` containing the following code:

```

int sqr(int n1)
{
    return n1*n1;
}

```

3. Create the scatter file `scatter.scat` containing the following load region:

```

LR1 0x0
{
    ER1 0x0
    {
        *(+RO)                ; rest of code and read-only data
    }
    ER2 +0
    {
        function.o
        *(.ARM.__at_0x10000)    ; Place gValue at 0x10000
    }
    RAM 0x200000 (0x1FF00-0x2000) ; RW & ZI data to be placed at 0x200000
    {
        *(+RW, +ZI)
    }
    ARM_LIB_STACK 0x800000 EMPTY -0x10000
    {
    }
    ARM_LIB_HEAP +0 EMPTY 0x10000
    {
    }
}

```

The `ARM_LIB_STACK` and `ARM_LIB_HEAP` regions are required because the program is being linked with the semihosting libraries.

4. Compile and link the sources:

```

armcc -c -g function.c
armcc -c -g main.c
armlink --no_autoat --scatter=scatter.scat --map function.o main.o -o squared.axf

```

The `--map` option displays the memory map of the image.

The memory map shows that the variable is placed in the ER2 execution region at address `0x10000`:

```

...
Execution Region ER2 (Base: 0x00001578, Size: 0x0000ea8c, Max: 0xffffffff, ABSOLUTE)

Base Addr   Size           Type  Attr   Idx   E Section Name      Object
0x00001578  0x0000000c    Code  RO     3     .text               function.o
0x00001584  0x0000ea7c    PAD
0x00010000  0x00000004    Data  RO     15    .ARM.__at_0x10000  main.o
...

```

In this example, the size of ER1 is unknown. Therefore, gValue might be placed in ER1 or ER2. To make sure that gValue is placed in ER2, you must include the corresponding selector in ER2 and link with the `--no_autoat` command-line option. If you omit `--no_autoat`, gValue is placed in a separate load region `LR$$.ARM. __at_0x10000` that contains the execution region `ER$$.ARM. __at_0x10000`.

8.11.4 See also

Tasks

- [Placing a named section explicitly using scatter-loading](#) on page 8-22
- [Using `__at` sections to place sections at a specific address](#) on page 8-35.

Concept

- [Restrictions on placing `__at` sections](#) on page 8-36.

Reference

Compiler Reference:

- [--split_sections](#) on page 3-156
- [__attribute__\(\(section\("name"\)\)\)](#) function attribute on page 5-56
- [__attribute__\(\(at\(address\)\)\)](#) variable attribute on page 5-70
- [__attribute__\(\(section\("name"\)\)\)](#) variable attribute on page 5-75
- [#pragma arm section \[section_type_list\]](#) on page 5-86.

Linker Reference:

- [--autoat, --no_autoat](#) on page 2-11
- [--map, --no_map](#) on page 2-83
- [--scatter=filename](#) on page 2-109.

Assembler Reference:

- [AREA](#) on page 5-13.

8.12 Placing a named section explicitly using scatter-loading

The following example shows how to place a named section explicitly using scatter-loading:

Example 8-7 Explicit section placement

```

LR1 0x0 0x10000
{
  ER1 0x0 0x2000          ; Root Region, containing init code
  {
    init.o (INIT, +FIRST) ; place init code at exactly 0x0
    *(+RO)                ; rest of code and read-only data
  }
  RAM_RW 0x400000 (0x1FF00-0x2000) ; RW & ZI data to be placed at 0x400000
  {
    *(+RW)
  }
  RAM_ZI +0
  {
    *(+ZI)
  }
  DATABLOCK 0x1FF00 0xFF          ; execution region at 0x1FF00
  {
    data.o(+RO-DATA)             ; maximum space available for table is 0xFF
    ; place RO data between 0x1FF00 and 0x1FFFF
  }
}

```

In this example, the scatter-loading description places:

- The initialization code is placed in the INIT section in the `init.o` file. This example shows that the code from the INIT section is placed first, at address `0x0`, followed by the remainder of the RO code and all of the RO data except for the RO data in the object `data.o`.
- All global RW variables in RAM at `0x400000`.
- A table of RO-DATA from `data.o` at address `0x1FF00`.

8.12.1 See also

Tasks

- [Using the FIXED attribute to create root regions on page 8-16.](#)

Concepts

Linker Reference:

- [About load region descriptions on page 4-5](#)
- [About execution region descriptions on page 4-9.](#)

Reference

Linker Reference:

- [Load region attributes on page 4-8](#)
- [Execution region attributes on page 4-12](#)
- [Address attributes for load and execution regions on page 4-15.](#)

Assembler Reference:

- [ENTRY on page 5-35.](#)

8.13 Placing unassigned sections with the .ANY module selector

The linker attempts to place input sections into specific execution regions. For any input sections that cannot be resolved, and where the placement of those sections is not important, you can use the .ANY module selector in the scatter file.

In most cases, using a single .ANY selector is equivalent to using the * module selector. However, unlike *, you can specify .ANY in multiple execution regions.

8.13.1 Default rules for placing unassigned sections

By default, the linker places unassigned sections using the following criteria:

- Place an unassigned section in the execution region that currently has the most free space. You can specify a maximum amount of space to use for unassigned sections with the execution region attribute ANY_SIZE.
- Sort sections in descending size order.

8.13.2 Placement rules when using multiple .ANY selectors

If more than one .ANY selector is present in a scatter file, the linker takes the unassigned section with the largest size and assigns the section to the most specific .ANY execution region that has enough free space. For example, .ANY(.text) is judged to be more specific than .ANY(+R0).

If several execution regions are equally specific, then the section is assigned to the execution region with the most available remaining space.

For example:

- If you have two equally specific execution regions where one has a size limit of 0x2000 and the other has no limit, then all the sections are assigned to the second unbounded .ANY region.
- If you have two equally specific execution regions where one has a size limit of 0x2000 and the other has a size limit of 0x3000, then the first sections to be placed are assigned to the second .ANY region of size limit 0x3000 until the remaining size of the second .ANY is reduced to 0x2000. From this point, sections are assigned alternately between both .ANY execution regions.

8.13.3 Prioritizing .ANY sections

You can give a priority ordering if you have multiple .ANY sections with the .ANYnum selector, where num is a positive integer from zero upwards. The highest priority is given to the selector with the highest integer.

The following example shows how to use .ANYnum:

```

lr1 0x8000 1024
{
  er1 +0 512
  {
    .ANY1(+R0) ; evenly distributed with er3
  }
  er2 +0 256
  {
    .ANY2(+R0) ; Highest priority, so filled first
  }
  er3 +0 256
  {

```

```

        .ANY1(+R0) ; evenly distributed with er1
    }
}

```

8.13.4 Controlling the placement of input sections for multiple .ANY selectors

You can modify how the linker places unassigned input sections when using multiple .ANY selectors by using a different placement algorithm or a different sort order. The following command-line options are available:

- `--any_placement=algorithm`, where *algorithm* is one of `first_fit`, `worst_fit`, `best_fit`, or `next_fit`
- `--any_sort_order=order`, where *order* is one of `cmdline` or `descending_size`.

Use `first_fit` when you want to fill regions in order.

Use `best_fit` when you want to fill regions to their maximum.

Use `worst_fit` when you want to fill regions evenly. With equal sized regions and sections `worst_fit` fills regions cyclically.

Use `next_fit` when you need a more deterministic fill pattern.

If the linker attempts to fill a region to its limit, as it does with `first_fit` and `best_fit`, it might overflow the region. This is because linker-generated content such as padding and veneers are not known until sections have been assigned to .ANY selectors. If this occurs you might see the following error:

```
Error: L6220E: Execution region regionname size (size bytes) exceeds limit (limit bytes).
```

The `--any_contingency` option prevents the linker from filling the region up to its maximum. It reserves a portion of the region's size for linker-generated content and fills this contingency area only if no other regions have space. It is enabled by default for the `first_fit` and `best_fit` algorithms, because they are most likely to exhibit this behavior.

8.13.5 Specifying the maximum size permitted for placing unassigned sections

The execution region attribute `ANY_SIZE max_size` enables you to specify the maximum size in a region that `armlink` can fill with unassigned sections.

Be aware of the following restrictions when using this keyword:

- *max_size* must be less than or equal to the region size
- you can use `ANY_SIZE` on a region without a .ANY selector but it is ignored by `armlink`.

When `ANY_SIZE` is present, `armlink`:

- Does not override a given .ANY size. That is, it does not reduce the priority then try to fit more sections in later.
- Never recalculates contingency.
- Never assigns sections in the contingency space.

`ANY_SIZE` does not require `--any_contingency` to be specified. However, when `--any_contingency` is specified and `ANY_SIZE` is not, `armlink` attempts to adjust contingencies. The aims are to:

- never overflow a .ANY region
- never refuse to place a section in a contingency reserved space.

If you specify `--any_contingency` on the command line, it is ignored for regions that have `ANY_SIZE` specified. It is used as normal for regions that do not have `ANY_SIZE` specified.

The following example shows how to use ANY_SIZE:

```
LOAD_REGION 0x0 0x3000
{
  ER_1 0x0 ANY_SIZE 0xF00 0x1000
  {
    .ANY
  }
  ER_2 0x0 ANY_SIZE 0xFB0 0x1000
  {
    .ANY
  }
  ER_3 0x0 ANY_SIZE 0x1000 0x1000
  {
    .ANY
  }
}
```

In this example:

- ER_1 has 0x100 reserved for linker-generated content.
- ER_2 has 0x50 reserved for linker-generated content. That is about the same as the automatic contingency of --any_contingency.
- ER_3 has no reserved space. Therefore 100% of the region is filled, with no contingency for veneers. Omitting the ANY_SIZE parameter causes 98% of the region to be filled, with a two percent contingency for veneers.

8.13.6 See also

Concepts

- [Examples of using placement algorithms for .ANY sections on page 8-26](#)
- [Example of next_fit algorithm showing behavior of full regions, selectors, and priority on page 8-28](#)
- [Examples of using sorting algorithms for .ANY sections on page 8-30.](#)

Linker Reference:

- [How the linker resolves multiple matches when processing scatter files on page 4-27](#)
- [Behavior when .ANY sections overflow because of linker-generated content on page 4-29.](#)

Reference

Linker Reference:

- [--any_contingency on page 2-5](#)
- [--any_placement=algorithm on page 2-6](#)
- [--any_sort_order=order on page 2-8](#)
- [--info=topic\[,topic,...\] on page 2-58](#)
- [--map, --no_map on page 2-83](#)
- [--section_index_display=type on page 2-110](#)
- [--tiebreaker=option on page 2-131](#)
- [Syntax of an input section description on page 4-23.](#)

8.14 Examples of using placement algorithms for .ANY sections

These examples show the operation of the placement algorithms for R0-CODE sections in sections.o.

The input section properties and ordering are shown in the following table:

Table 8-1 Input section properties

Name	Size
sec1	0x4
sec2	0x4
sec3	0x4
sec4	0x4
sec5	0x4
sec6	0x4

The scatter file used for the examples is:

```
LR 0x100
{
  ER_1 0x100 0x10
  {
    .ANY
  }

  ER_2 0x200 0x10
  {
    .ANY
  }
}
```

———— **Note** ————

These examples have `--any_contingency` disabled.

8.14.1 Example for first_fit, next_fit, and best_fit

This example shows the situation where several sections of equal size are assigned to two regions with one selector. The selectors are equally specific, equivalent to `.ANY(+R0)` and have no priority.

Execution Region ER_1 (Base: 0x00000100, Size: 0x00000010, Max: 0x00000010, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
0x00000100	0x00000004	Code	R0	1	sec1	sections.o
0x00000104	0x00000004	Code	R0	2	sec2	sections.o
0x00000108	0x00000004	Code	R0	3	sec3	sections.o
0x0000010c	0x00000004	Code	R0	4	sec4	sections.o

Execution Region ER_2 (Base: 0x00000200, Size: 0x00000008, Max: 0x00000010, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
-----------	------	------	------	-----	----------------	--------

0x00000200	0x00000004	Code	RO	5	sec5	sections.o
0x00000204	0x00000004	Code	RO	6	sec6	sections.o

In this example:

- For `first_fit` the linker first assigns all the sections it can to ER_1, then moves on to ER_2 because that is the next available region.
- For `next_fit` the linker does the same as `first_fit`. However, when ER_1 is full it is marked as FULL and is not considered again. In this example, ER_1 is completely full. ER_2 is then considered.
- For `best_fit` the linker assigns sec1 to ER_1. It then has two regions of equal priority and specificity, but ER_1 has less space remaining. Therefore, the linker assigns sec2 to ER_1, and continues assigning sections until ER_1 is full.

8.14.2 Example for `worst_fit`

This example shows the image memory map when using the `worst_fit` algorithm.

Execution Region ER_1 (Base: 0x00000100, Size: 0x0000000c, Max: 0x00000010, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
0x00000100	0x00000004	Code	RO	1	sec1	sections.o
0x00000104	0x00000004	Code	RO	3	sec3	sections.o
0x00000108	0x00000004	Code	RO	5	sec5	sections.o

Execution Region ER_2 (Base: 0x00000200, Size: 0x0000000c, Max: 0x00000010, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
0x00000200	0x00000004	Code	RO	2	sec2	sections.o
0x00000204	0x00000004	Code	RO	4	sec4	sections.o
0x00000208	0x00000004	Code	RO	6	sec6	sections.o

The linker first assigns sec1 to ER_1. It then has two equally specific and priority regions. It assigns sec2 to the one with the most free space, ER_2 in this example. The regions now have the same amount of space remaining, so the linker assigns sec3 to the first one that appears in the scatter file, that is ER_1.

———— **Note** ————

The behavior of `worst_fit` is the default behavior in this version of the linker, and it is the only algorithm available and earlier linker versions.

8.14.3 See also

Concepts

- [Placing unassigned sections with the .ANY module selector on page 8-23](#)
- [Example of next_fit algorithm showing behavior of full regions, selectors, and priority on page 8-28.](#)

Reference

Linker Reference:

- [--any_placement=algorithm on page 2-6](#)
- [--scatter=filename on page 2-109.](#)

8.15 Example of next_fit algorithm showing behavior of full regions, selectors, and priority

This example shows the operation of the next_fit placement algorithm for RO-CODE sections in sections.o.

The input section properties and ordering are shown in the following table:

Table 8-2 Input section properties

Name	Size
sec1	0x4
sec2	0x4
sec3	0x4
sec4	0x4
sec5	0x4
sec6	0x4

The scatter file used for the examples is:

```
LR 0x100
{
  ER_1 0x100 0x20
  {
    .ANY1(+RO-CODE)
  }

  ER_2 0x200 0x20
  {
    .ANY2(+RO)
  }

  ER_3 0x300 0x20
  {
    .ANY3(+RO)
  }
}
```

———— **Note** ————

This example has --any_contingency disabled.

The next_fit algorithm is different to the others in that it never revisits a region that is considered to be full. This example also shows the interaction between priority and specificity of selectors - this is the same for all the algorithms.

Execution Region ER_1 (Base: 0x00000100, Size: 0x00000014, Max: 0x00000020, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
0x00000100	0x00000014	Code	RO	1	sec1	sections.o

Execution Region ER_2 (Base: 0x00000200, Size: 0x0000001c, Max: 0x00000020, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E Section Name	Object
-----------	------	------	------	-----	----------------	--------

0x00000200	0x00000010	Code	RO	3	sec3	sections.o
0x00000210	0x00000004	Code	RO	4	sec4	sections.o
0x00000214	0x00000004	Code	RO	5	sec5	sections.o
0x00000218	0x00000004	Code	RO	6	sec6	sections.o

Execution Region ER_3 (Base: 0x00000300, Size: 0x00000014, Max: 0x00000020, ABSOLUTE)

Base Addr	Size	Type	Attr	Idx	E	Section Name	Object
0x00000300	0x00000014	Code	RO	2		sec2	sections.o

In this example:

- The linker places sec1 in ER_1 because ER_1 has the most specific selector. ER_1 now has 0x6 bytes remaining.
- The linker then tries to place sec2 in ER_1, because it has the most specific selector, but there is not enough space. Therefore, ER_1 is marked as full and is not considered in subsequent placement steps. The linker chooses ER_3 for sec2 because it has higher priority than ER_2.
- The linker then tries to place sec3 in ER_3. It does not fit, so ER_3 is marked as full and the linker places sec3 in ER_2.
- The linker now processes sec4. This is 0x4 bytes so it can fit in either ER_1 or ER_3. Because both of these sections have previously been marked as full, they are not considered. The linker places all remaining sections in ER_2.
- If another section sec7 of size 0x8 exists, and is processed after sec6 the example fails to link. The algorithm does not attempt to place the section in ER_1 or ER_3 because they have previously been marked as full.

8.15.1 See also

Concepts

- [Placing unassigned sections with the .ANY module selector on page 8-23](#)
- [Examples of using placement algorithms for .ANY sections on page 8-26.](#)

Linker Reference:

- [How the linker resolves multiple matches when processing scatter files on page 4-27](#)
- [Behavior when .ANY sections overflow because of linker-generated content on page 4-29.](#)

Reference

Linker Reference:

- [--any_placement=algorithm on page 2-6](#)
- [--scatter=filename on page 2-109.](#)

8.16 Examples of using sorting algorithms for .ANY sections

These examples show the operation of the sorting algorithms for R0-CODE sections in sections_a.o and sections_b.o.

The input section properties and ordering are shown in the following tables:

Table 8-3 Input section properties for sections_a.o

Name	Size
seca_1	0x4
seca_2	0x4
seca_3	0x10
seca_4	0x14

Table 8-4 Input section properties for sections_b.o

Name	Size
secb_1	0x4
secb_2	0x4
secb_3	0x10
secb_4	0x14

8.16.1 Descending size example

The following linker command-line options are used for this example:

```
--any_sort_order=descending_size sections_a.o sections_b.o --scatter scatter.txt
```

The order that the sections are processed by the .ANY assignment algorithm is:

Table 8-5 Sort order for descending_size algorithm

Name	Size
seca_4	0x14
secb_4	0x14
seca_3	0x10
secb_3	0x10
seca_1	0x4
seca_2	0x4
secb_1	0x4
secb_2	0x4

Sections of the same size use the tiebreak specified by --tiebreaker.

8.16.2 Command-line example

The following linker command-line options are used for this example:

```
--any_sort_order=cmdline sections_a.o sections_b.o --scatter scatter.txt
```

The order that the sections are processed by the `.ANY` assignment algorithm is:

Table 8-6 Sort order for `cmdline` algorithm

Name	Size
seca_1	0x4
secb_1	0x4
seca_2	0x4
secb_2	0x4
seca_3	0x10
secb_3	0x10
seca_4	0x14
secb_4	0x14

Sections with the same command-line index use the tiebreak specified by `--tiebreaker`.

8.16.3 See also

Concepts

- [Placing unassigned sections with the `.ANY` module selector on page 8-23.](#)

Reference

Linker Reference:

- [--any_sort_order=order on page 2-8](#)
- [--scatter=filename on page 2-109](#)
- [--tiebreaker=option on page 2-131.](#)

8.17 Selecting veneer input sections in scatter-loading descriptions

Veneers are used to switch between ARM and Thumb code or to perform a longer program jump than can be specified in a single instruction. You can place veneers at a specific location by including the linker-generated symbol `Veneer$$Code` in a scatter file. At most, one execution region in the scatter file can have the `*(Veneer$$Code)` section selector.

If it is safe to do so, the linker places veneer input sections into the region identified by the `*(Veneer$$Code)` section selector. It might not be possible for a veneer input section to be assigned to the region because of address range problems or execution region size limitations. If the veneer cannot be added to the specified region, it is added to the execution region containing the relocated input section that generated the veneer.

Note

Instances of `*(IWV$$Code)` in scatter files from earlier versions of ARM tools are automatically translated into `*(Veneer$$Code)`. Use `*(Veneer$$Code)` in new descriptions.

`*(Veneer$$Code)` is ignored when the amount of code in an execution region exceeds 4Mb of Thumb code, 16Mb of 32-bit Thumb code, and 32Mb of ARM code.

8.17.1 See also

Concepts

- [Overview of veneers on page 4-26.](#)

8.18 Using `__attribute__((section("name")))` to place code and data

You can place code and data by separating them into their own objects without having to use toolchain-specific pragmas or attributes. However, you can also use `__attribute__((section("name")))` to place an item in a separate ELF section. You can then use a scatter file to place the named sections at specific locations.

To use `__attribute__((section("name")))` to place a variable in a separate section:

1. Use `__attribute__((section("name")))` to specify the named section where the variable is to be placed, for example:

Example 8-8 Naming a section

```
int variable __attribute__((section("foo"))) = 10;
```

2. Use a scatter file to place the named section, for example:

Example 8-9 Placing a section

```
FLASH 0x24000000 0x4000000
{
    ... ; rest of code

    ADDER 0x08000000
    {
        file.o (foo) ; select section foo from file.o
    }
}
```

The following example shows the memory map for the FLASH load region:

```
...
Load Region FLASH (Base: 0x24000000, Size: 0x00000004, Max: 0x04000000, ABSOLUTE)
  Execution Region ADDER (Base: 0x08000000, Size: 0x00000004, Max: 0xffffffff, ABSOLUTE)
    Base Addr   Size       Type  Attr   Idx   E Section Name   Object
    0x08000000  0x00000004  Data  RW     16    foo              file.o
...
```

Be aware of the following:

- linking with `--autoat` or `--no_autoat` does not affect the placement
- if scatter-loading is not used, the section is placed in the default `ER_RW` execution region of the `LR_1` load region
- if you have a scatter file that does not include the `foo` selector, then the section is placed in the defined `RW` execution region.

You can also place a function at a specific address using `.ARM.__at_address` as the section name. For example, to place the function `sqr` at `0x20000`, specify:

```
int sqr(int n1) __attribute__((section(".ARM.__at_0x20000")));

int sqr(int n1)
```

```
{  
    return n1*n1;  
}
```

8.18.1 See also

Reference

- [Using __at sections to place sections at a specific address](#) on page 8-35
- [Restrictions on placing __at sections](#) on page 8-36.

Compiler Reference:

- [__attribute__\(\(section\("name"\)\)\)](#) function attribute on page 5-56
- [__attribute__\(\(section\("name"\)\)\)](#) variable attribute on page 5-75
- [#pragma arm section \[section_type_list\]](#) on page 5-86.

Linker Reference:

- [--autoat, --no_autoat](#) on page 2-11
- [--scatter=filename](#) on page 2-109.

8.19 Using `__at` sections to place sections at a specific address

You can give a section a special name that encodes the address where it must be placed. You specify the name as follows:

```
.ARM.__at_address
```

where:

`address` is the required address of the section. You can specify this in hexadecimal or decimal. Sections in the form of `.ARM.__at_address` are referred to by the abbreviation `__at`.

In the compiler, you can assign variables to `__at` sections by:

- explicitly naming the section using the `__attribute__((section("name")))`
- using the attribute `__at` that sets up the name of the section for you.

Example 8-10 Assigning variables to `__at` sections in C or C++ code

```
// place variable1 in a section called .ARM.__AT_0x00008000
int variable1 __attribute__((at(0x8000))) = 10;

// place variable2 in a section called .ARM.__at_0x8000
int variable2 __attribute__((section(".ARM.__at_0x8000"))) = 10;
```

———— Note —————

When using `__attribute__((at(address)))`, the part of the `__at` section name representing `address` is normalized to an 8 digit hexadecimal number. The name of the section is only significant if you are trying to match the section by name in a scatter file. The linker automatically assigns `__at` sections when you use the `--autoat` command-line option. This option is the default.

8.19.1 See also

Concepts

- [Placing functions and data at specific addresses](#) on page 8-17
- [Using `__attribute__\(\(section\("name"\)\)\)` to place code and data](#) on page 8-33
- [Restrictions on placing `__at` sections](#) on page 8-36
- [Automatic placement of `__at` sections](#) on page 8-37
- [Manual placement of `__at` sections](#) on page 8-39
- [Placing a key in flash memory using `__at`](#) on page 8-41
- [Placing a structure over a peripheral register using `__at`](#) on page 8-43.

Reference

Linker Reference:

- [--autoat, --no_autoat](#) on page 2-11.

Compiler Reference:

- [`__attribute__\(\(section\("name"\)\)\)` function attribute](#) on page 5-56
- [`__attribute__\(\(at\(address\)\)\)` variable attribute](#) on page 5-70
- [`__attribute__\(\(section\("name"\)\)\)` variable attribute](#) on page 5-75.

8.20 Restrictions on placing __at sections

The following restrictions apply when placing __at sections at specific addresses:

- __at section address ranges must not overlap, unless the overlapping sections are placed in different overlay regions
- __at sections are not permitted in position independent execution regions
- you must not reference the linker-defined symbols \$\$Base, \$\$Limit and \$\$Length of an __at section
- __at sections must have an address that is a multiple of their alignment
- __at sections ignore any +FIRST or +LAST ordering constraints.

8.20.1 See also

Concepts

- [Using __at sections to place sections at a specific address](#) on page 8-35.

8.21 Automatic placement of __at sections

The automatic placement of __at sections is enabled by default. This feature is controlled by the linker command-line option, --autoat.

———— **Note** —————

You cannot use __at section placement with position independent execution regions.

When linking with the --autoat option, the __at sections are not placed by the scatter-loading selectors. Instead, the linker places the __at section in a compatible region. If no compatible region is found, the linker creates a load and execution region for the __at section.

All linker --autoat created execution regions have the UNINIT scatter-loading attribute. If you require a ZI __at section to be zero-initialized then it must be placed within a compatible region. A linker --autoat created execution region must have a base address that is at least 4 byte-aligned. The linker produces an error message if any region is incorrectly aligned.

A compatible region is one where:

- The __at address lies within the execution region base and limit, where limit is the base address + maximum size of execution region. If no maximum size is set, the linker sets the limit for placing __at sections as the current size of the execution region without __at sections plus a constant, 10240 bytes.
- The execution region meets at least one of the following conditions:
 - it has a selector that matches the __at section by the standard scatter-loading rules
 - it has at least one section of the same type (RO, RW or ZI) as the __at section
 - it does not have the EMPTY attribute.

———— **Note** —————

The linker considers an __at section with type RW compatible with RO.

The following example shows the sections .ARM.__at_0x0000 type RO, .ARM.__at_0x2000 type RW, .ARM.__at_0x4000 type ZI, and .ARM.__at_0x8000 type ZI:

Example 8-11 Placement of the variables in C or C++ code

```
// place the RW variable in a section called .ARM.__at_0x2000
int foo __attribute__((section(".ARM.__at_0x2000"))) = 100;

// place the ZI variable in a section called .ARM.__at_0x4000
int bar __attribute__((section(".ARM.__at_0x4000"), zero_init));

// place the ZI variable in a section called .ARM.__at_0x8000
int variable __attribute__((section(".ARM.__at_0x8000"), zero_init));
```

Example 8-12 Automatic placement of __at sections

```
LR1 0x0
{
  ER_RO 0x0 0x2000
  {
    *(+RO) ; .ARM.__at_0x0 lies within the bounds of ER_RO
  }
}
```

```

ER_RW 0x2000 0x2000
{
    *(+RW)      ; .ARM.__at_0x2000 lies within the bounds of ER_RW
}
ER_ZI 0x4000 0x2000
{
    *(+ZI)      ; .ARM.__at_0x4000 lies within the bounds of ER_ZI
}
}

```

; the linker creates a load and execution region for the __at section
; .ARM.__at_0x8000 because it lies outside all candidate regions.

8.21.1 See also

Concepts

- [Using __attribute__\(\(section\("name"\)\)\) to place code and data on page 8-33](#)
- [Using __at sections to place sections at a specific address on page 8-35](#)
- [Restrictions on placing __at sections on page 8-36](#)
- [Manual placement of __at sections on page 8-39](#)
- [Placing a key in flash memory using __at on page 8-41](#)
- [Placing a structure over a peripheral register using __at on page 8-43.](#)

Linker Reference:

- [About execution region descriptions on page 4-9.](#)

Reference

Linker Reference:

- [--autoat, --no_autoat on page 2-11](#)
- [--ro_base=address on page 2-103](#)
- [--rw_base=address on page 2-106](#)
- [--zi_base=address on page 2-149](#)
- [Execution region attributes on page 4-12.](#)

8.22 Manual placement of `__at` sections

You can use the standard section placement rules to place `__at` sections when using the `--no_autoat` command-line option.

———— **Note** —————

You cannot use `__at` section placement with position independent execution regions.

The following example shows the placement of read-only sections `.ARM.__at_0x2000` and the read-write section `.ARM.__at_0x4000`. Load and execution regions are not created automatically in manual mode. An error is produced if an `__at` section cannot be placed in an execution region.

Example 8-13 Placement of the variables in C or C++ code

```
// place the RO variable in a section called .ARM.__at_0x2000
const int F00 __attribute__((section(".ARM.__at_0x2000"))) = 100;

// place the RW variable in a section called .ARM.__at_0x4000
int bar __attribute__((section(".ARM.__at_0x4000")));
```

Example 8-14 Manual placement of `__at` sections

```
LR1 0x0
{
  ER_RO 0x0 0x2000
  {
    *(+RO)           ; .ARM.__at_0x0 is selected by +RO
  }
  ER_R02 0x2000
  {
    *(.ARM.__at_0x2000) ; .ARM.__at_0x2000 is selected by the section named
                        ; .ARM.__at_0x2000
  }
  ER2 0x4000
  {
    *(+RW +ZI)       ; .ARM.__at_0x4000 is selected by +RW
  }
}
```

8.22.1 See also

Concepts

- [Using `__attribute__\(\(section\("name"\)\)\)` to place code and data](#) on page 8-33
- [Using `__at` sections to place sections at a specific address](#) on page 8-35
- [Restrictions on placing `__at` sections](#) on page 8-36
- [Automatic placement of `__at` sections](#) on page 8-37
- [Placing a key in flash memory using `__at`](#) on page 8-41
- [Placing a structure over a peripheral register using `__at`](#) on page 8-43.

Linker Reference:

- [About execution region descriptions](#) on page 4-9.

Reference

Linker Reference:

- [--autoat, --no_autoat](#) on page 2-11
- [Execution region attributes](#) on page 4-12.

8.23 Placing a key in flash memory using `__at`

Some flash devices require a key to be written to an address to activate certain features. An `__at` section provides a simple method of writing a value to a specific address.

Assuming a device has flash memory from `0x8000` to `0x10000` and a key is required in address `0x8000`. To do this with an `__at` section, you must declare a variable so that the compiler can generate a section called `.ARM.__at_0x8000`.

Example 8-15 Placement of the flash key variable in C or C++ code

```
// place flash_key in a section called .ARM.__at_0x8000
long flash_key __attribute__((section(".ARM.__at_0x8000")));
```

The following example shows a scatter file with manual placement of the flash execution region:

Example 8-16 Manual placement of flash execution regions

```
ER_FLASH 0x8000 0x2000
{
    *(+RO)
    *(.ARM.__at_0x8000) ; key
}
```

Use the linker command-line option `--no_autoat` to enable manual placement.

The following example shows a scatter file with automatic placement of the flash execution region. Use the linker command-line option `--autoat` to enable automatic placement.

Example 8-17 Automatic placement of flash execution regions

```
ER_FLASH 0x8000 0x2000
{
    *(+RO) ; other code and read-only data, the
           ; __at section is automatically selected
}
```

8.23.1 See also

Tasks

- [Placing sections with *FIRST* and *LAST* attributes on page 4-21](#)
- [Using `__at` sections to place sections at a specific address on page 8-35.](#)

Concepts

- [Automatic placement of `__at` sections on page 8-37](#)
- [Manual placement of `__at` sections on page 8-39.](#)

Linker Reference:

- [About execution region descriptions on page 4-9.](#)

Reference

Linker Reference:

- `--autoat, --no_autoat` on page 2-11.

Compiler Reference:

- `__attribute__((section("name")))` variable attribute on page 5-75.

8.24 Placing a structure over a peripheral register using `__at`

To place an uninitialized variable over a peripheral register, you can use a ZI `__at` section. Assuming a register is available for use at `0x10000000`, define a ZI `__at` section called `.ARM.__at_0x10000000`. For example:

```
int foo __attribute__((section(".ARM.__at_0x10000000"), zero_init));
```

The following example shows the a scatter file with the manual placement of the ZI `__at` section:

Example 8-18 Manual placement of ZI `__at` sections

```
ER_PERIPHERAL 0x10000000 UNINIT
{
  *(.ARM.__at_0x10000000)
}
```

Using automatic placement, and assuming that there is no other execution region near `0x10000000`, the linker automatically creates a region with the UNINIT attribute at `0x10000000`. The UNINIT attribute creates an execution region containing uninitialized data or memory-mapped I/O.

8.24.1 See also

Concepts

- [Using `__at` sections to place sections at a specific address](#) on page 8-35.

Linker Reference:

- [About execution region descriptions](#) on page 4-9.

Reference

Linker Reference:

- [Execution region attributes](#) on page 4-12.

Compiler Reference:

- [`__attribute__\(\(section\("name"\)\)\)` variable attribute](#) on page 5-75.

8.25 Placement of sections with overlays

You can use the `OVERLAY` attribute in a scatter file to place multiple execution regions at the same address. An overlay manager is required to make sure that only one execution region is instantiated at a time. The ARM Compiler toolchain does not provide an overlay manager.

The following example shows the definition of a static section in RAM followed by a series of overlays. Here, only one of these sections is instantiated at a time.

Example 8-19 Specifying a root region

```

EMB_APP 0x8000
{
    .
    .
    STATIC_RAM 0x0                ; contains most of the RW and ZI code/data
    {
        * (+RW,+ZI)
    }
    OVERLAY_A_RAM 0x1000 OVERLAY  ; start address of overlay...
    {
        module1.o (+RW,+ZI)
    }
    OVERLAY_B_RAM 0x1000 OVERLAY
    {
        module2.o (+RW,+ZI)
    }
    ...                            ; rest of scatter-loading description...
}

```

A region marked as `OVERLAY` is not initialized by the C library at startup. The contents of the memory used by the overlay region are the responsibility of an overlay manager. If the region contains initialized data, use the `NOCOMPRESS` attribute to prevent RW data compression.

The linker defined symbols can be used to obtain the addresses required to copy the code and data.

The `OVERLAY` attribute can be used on a single region that is not the same address as a different region. Therefore, an overlay region can be used as a method to prevent the initialization of particular regions by the C library startup code. As with any overlay region these must be manually initialized in your code.

An overlay region can have a relative base. The behavior of an overlay region with a `+offset` base address depends on the regions that precede it and the value of `+offset`. The linker places consecutive `+offset` regions at the same base address if they have the same `+offset` value.

When a `+offset` execution region ER follows a contiguous overlapping block of overlay execution regions the base address of ER is:

limit address of the overlapping block of overlay execution regions + `offset`

The following table shows the effect of *+offset* when used with the OVERLAY attribute. REGION1 appears immediately before REGION2 in the scatter file:

Table 8-7 Using relative offset in overlays

REGION1 is set with OVERLAY	+offset	REGION2 Base Address
NO	<offset>	REGION1 Limit + <offset>
YES	+0	REGION1 Base Address
YES	<none-zero offset>	REGION1 Limit + <none-zero offset>

The following example shows the use of relative offsets with overlays and the effect on execution region addresses:

Example 8-20 Example of relative offset in overlays

```

EMB_APP 0x8000{
  CODE 0x8000
  {
    *(+R0)
  }

  # REGION1 Base = CODE limit
  REGION1 +0 OVERLAY
  {
    module1.o(*)
  }

  # REGION2 Base = REGION1 Base
  REGION2 +0 OVERLAY
  {
    module2.o(*)
  }

  # REGION3 Base = REGION2 Base = REGION1 Base
  REGION3 +0 OVERLAY
  {
    module3.o(*)
  }

  # REGION4 Base = REGION3 Limit + 4
  Region4 +4 OVERLAY
  {
    module4.o(*)
  }
}

```

If the length of the non-overlay area is unknown, a zero relative offset can be used to specify the start address of an overlay so that it is placed immediately after the end of the static section.

You can use the following command-line options to add extra debug information to the image:

- `--emit_debug_overlay_relocs`
- `--emit_debug_overlay_section`.

These permit an overlay-aware debugger to track which overlay is currently active.

8.25.1 See also

Concepts

- [Using `__at` sections to place sections at a specific address](#) on page 8-35.

Linker Reference:

- [About load region descriptions](#) on page 4-5
- [About execution region descriptions](#) on page 4-9
- [Considerations when using a relative address `+offset` for load regions](#) on page 4-17
- [Considerations when using a relative address `+offset` for execution regions](#) on page 4-18.

Reference

- [Accessing linker-defined symbols](#) on page 7-4.

Linker Reference:

- [--emit_debug_overlay_relocs](#) on page 2-39
- [--emit_debug_overlay_section](#) on page 2-40
- [Load region attributes](#) on page 4-8
- [Execution region attributes](#) on page 4-12
- [Address attributes for load and execution regions](#) on page 4-15.

Compiler Reference:

- [__attribute__\(\(section\("name"\)\)\)](#) variable attribute on page 5-75.

Other information

- [ABI for the ARM Architecture: Support for Debugging Overlaid Programs](#), <http://infocenter.arm.com/help/topic/com.arm.doc.ih0049-/index.html>

8.26 About placing ARM C and C++ library code

You can place code from the ARM standard C and C++ libraries in a scatter file. Use `*armlib` or `*cpplib*` so that the linker can resolve library naming in your scatter file.

Some ARM C and C++ library sections must be placed in a root region, for example `__main.o`, `__scatter*.o`, `__dc*.o`, and `*Region$$Table`. This list can change between releases. The linker can place all these sections automatically in a future-proof way with `InRoot$$Sections`.

8.26.1 See also

Tasks

- [Creating root execution regions on page 8-13](#)
- [Using the *FIXED* attribute to create root regions on page 8-16.](#)

Concepts

- [What is a root region? on page 8-12](#)
- [Example of placing code in a root region on page 8-48](#)
- [Example of placing ARM C library code on page 8-49](#)
- [Example of placing ARM C++ library code on page 8-50.](#)

8.27 Example of placing code in a root region

Use a scatter file to specify a root section in the same way as a named section. The following example uses the section selector `InRoot$$Sections` to place all sections that must be in a root region:

Example 8-21 Specifying a root region

```

ROM_LOAD 0x0000 0x4000
{
  ROM_EXEC 0x0000 0x4000      ; root region at 0x0
  {
    vectors.o (Vect, +FIRST) ; Vector table
    * (InRoot$$Sections)     ; All library sections that must be in a
                             ; root region, for example, __main.o,
                             ; __scatter*.o, __dc*.o, and * Region$$Table
  }
  RAM 0x10000 0x8000
  {
    * (+RO, +RW, +ZI)       ; all other sections
  }
}

```

8.27.1 See also

Tasks

- [Creating root execution regions on page 8-13](#)
- [Using the `FIXED` attribute to create root regions on page 8-16.](#)

Concepts

- [What is a root region? on page 8-12](#)
- [About placing ARM C and C++ library code on page 8-47](#)
- [Example of placing ARM C library code on page 8-49](#)
- [Example of placing ARM C++ library code on page 8-50.](#)

8.28 Example of placing ARM C library code

The following example shows how to place C library code:

Example 8-22 Placing ARM C library code

```

ROM1 0
{
  * (InRoot$$Sections)
  * (+R0)
}
ROM2 0x1000
{
  *armlib/c_* (+R0) ; all ARM-supplied C library functions
}
ROM3 0x2000
{
  *armlib/h_* (+R0) ; just the ARM-supplied __ARM_*
  ; redistributable library functions
}
RAM1 0x3000
{
  *armlib* (+R0) ; all other ARM-supplied library code
  ; for example, floating-point libraries
}
RAM2 0x4000
{
  * (+RW, +ZI)
}

```

The name `armlib` is used to indicate the ARM C library files that are located in the `install_directory\lib\armlib` directory.

8.28.1 See also

Concepts

- [About placing ARM C and C++ library code on page 8-47](#)
- [Example of placing code in a root region on page 8-48](#)
- [Example of placing ARM C++ library code on page 8-50.](#)

Reference

Using ARM® C and C++ Libraries and Floating Point Support.

- [C and C++ library naming conventions on page 2-120.](#)

8.29 Example of placing ARM C++ library code

The following is a C++ program that is to be scatter-loaded:

```
#include <iostream>

using namespace std;

extern "C" int foo ()
{
    cout << "Hello" << endl;
    return 1;
}
```

To place the C++ library code, define the scatter file as follows:

```
LR 0x0
{
    ER1 0x0
    {
        *armlib*(+R0)
    }

    ER2 +0
    {
        *cplib*(+R0)
        *(.init_array) ; Section .init_array must be placed explicitly,
                        ; otherwise it is shared between two regions, and
                        ; the linker is unable to decide where to place it.
    }

    ER3 +0
    {
        *(+R0)
    }

    ER4 +0
    {
        *(+RW,+ZI)
    }
}
```

The name *install_directory\lib\armlib* is used to indicate the ARM C library files that are located in the *armlib* directory.

The name *install_directory\lib\cplib* is used to indicate the ARM C++ library files that are located in the *cplib* directory.

8.29.1 See also

Concepts

- [About placing ARM C and C++ library code on page 8-47](#)
- [Example of placing code in a root region on page 8-48](#)
- [Example of placing ARM C library code on page 8-49.](#)

Reference

Using ARM® C and C++ Libraries and Floating Point Support:

- [C and C++ library naming conventions on page 2-120.](#)

8.30 Example of placing ARM library helper functions

ARM library helper functions are generated by the compiler in the resulting object files. Therefore, you cannot use `armlib` and `cpplib` in a scatter file to place these functions.

To place the helper functions specify `*.* (i.__ARM_*)` in your scatter file. The `*.*` part is important if you have `*(+R0)` in your scatter file.

Be aware that if you use `*(i.__ARM_*)` the following error is generated:

Error: L6223E: Ambiguous selectors...

This is because of the scatter-loading rules for resolving multiple matches.

8.30.1 See also

Concepts

ARM Compiler toolchain Linker Reference:

- [How the linker resolves multiple matches when processing scatter files on page 4-27.](#)

Reference

Errors and Warnings Reference:

- [List of the armlink error and warning messages on page 4-3.](#)

8.31 Reserving an empty region

You can use the `EMPTY` attribute in an execution region scatter-loading description to reserve an empty block of memory for the stack.

The block of memory does not form part of the load region, but is assigned for use at execution time. Because it is created as a dummy ZI region, the linker uses the following symbols to access it:

- `Image$$region_name$$ZI$$Base`
- `Image$$region_name$$ZI$$Limit`
- `Image$$region_name$$ZI$$Length`.

If the length is given as a negative value, the address is taken to be the end address of the region. This must be an absolute address and not a relative one.

In the following example, the execution region definition `STACK 0x800000 EMPTY -0x10000` defines a region called `STACK` that starts at address `0x7F0000` and ends at address `0x800000`:

Example 8-23 Reserving a region for the stack

```

LR_1 0x800000                                ; load region starts at 0x800000
{
  STACK 0x800000 EMPTY -0x10000             ; region ends at 0x800000 because of the
                                             ; negative length. The start of the region
                                             ; is calculated using the length.
  {
                                             ; Empty region used to place stack
  }
  HEAP +0 EMPTY 0x10000                     ; region starts at the end of previous
                                             ; region. End of region calculated using
                                             ; positive length
  {
                                             ; Empty region used to place heap
  }
  ...                                       ; rest of scatter-loading description...
}

```

Note

The dummy ZI region that is created for an `EMPTY` execution region is not initialized to zero at runtime.

If the address is in relative (`+offset`) form and the length is negative, the linker generates an error.

The following figure shows a diagrammatic representation for this example.




Figure 8-4 Reserving a region for the stack

In this example, the linker generates the symbols:

```
Image$$STACK$$ZI$$Base      = 0x7f0000
Image$$STACK$$ZI$$Limit     = 0x800000
Image$$STACK$$ZI$$Length    = 0x10000
Image$$HEAP$$ZI$$Base       = 0x800000
Image$$HEAP$$ZI$$Limit      = 0x810000
Image$$HEAP$$ZI$$Length     = 0x10000
```

———— **Note** ————

The `EMPTY` attribute applies only to an execution region. The linker generates a warning and ignores an `EMPTY` attribute used in a load region definition.

The linker checks that the address space used for the `EMPTY` region does not coincide with any other execution region.

8.31.1 See also

Concepts

Linker Reference:

- [About execution region descriptions on page 4-9.](#)

Reference

- [Image\\$\\$ execution region symbols on page 7-6.](#)

Linker Reference:

- [Execution region attributes on page 4-12.](#)

8.32 About creating regions on page boundaries

You can produce an ELF file that can be loaded directly to a target with each execution region starting at a page boundary.

The linker provides the following built-in functions to help create load and execution regions on page boundaries:

- `AlignExpr`
- `GetPageSize`.

Note

Alignment on an execution region causes both the load address and execution address to be aligned.

The following example produces an ELF file with each execution region starting on a new page:

Example 8-24 Creating regions on page boundaries

```
LR1 GetPageSize() + SizeOfHeaders()
{
  ER_RO +0
  {
    *(+RO)
  }
  ER_RW +GetPageSize()
  {
    *(+RW)
  }
  ER_ZI +0
  {
    *(+ZI)
  }
}
```

The default page size `0x8000`, is used. You can change the page size with the `--pagesize` command-line option.

8.32.1 See also

Concepts

- [Demand paging on page 4-23](#)
- [Overalignment of execution regions and input sections on page 8-56](#)
- [Expression evaluation in scatter files on page 8-58](#)
- [Using expression evaluation in a scatter file to avoid padding on page 8-59.](#)

Reference

Linker Reference:

- [--pagesize=pagesize on page 2-93](#)
- [Load region attributes on page 4-8](#)
- [Execution region attributes on page 4-12](#)
- [Example of aligning a base address in execution space but still tightly packed in load space on page 4-42](#)
- [AlignExpr\(expr, align\) function on page 4-43](#)

- [GetPageSize\(\) function](#) on page 4-44.

8.33 Overallignment of execution regions and input sections

There are situations when you want to overallign code and data sections. How you deal with them depends on whether or not you have access to the source code:

- If you have access to the original source code, you can do this at compile time with the `__align(n)` keyword or the `--min_array_alignment` command-line option, for example.
- If you do not have access to the source code, then you must use the following alignment specifiers in a scatter file:

ALIGNALL Increases the section alignment of all the sections in an execution region, for example:

```
ER_DATA ... ALIGNALL 8
{
    ... ;selectors
}
```

OVERALIGN

Increases the alignment of a specific section, for example:

```
ER_DATA ...
{
    *.o(.bar, OVERALIGN 8)
    ... ;selectors
}
```

8.33.1 See also

Concepts

- [About creating regions on page boundaries on page 8-54](#)

Reference

Linker Reference:

- [Execution region attributes on page 4-12](#)
- [About input section descriptions on page 4-22.](#)

Compiler Reference:

- [__align on page 5-6](#)
- [--min_array_alignment=opt on page 3-114.](#)

8.34 Using preprocessing commands in a scatter file

You can pass a scatter file through a C preprocessor. This permits access to all the features of the C preprocessor.

Use the first line in the scatter file to specify a preprocessor command that the linker invokes to process the file. The command is of the form:

```
#! preprocessor [pre_processor_flags]
```

Most typically the command is `#! armcc -E`. This passes the scatter file through the `armcc` preprocessor.

You can:

- add preprocessing directives to the top of the scatter file
- use simple expression evaluation in the scatter file.

For example, a scatter file, `file.scat`, might contain:

```
#! armcc -E

#define ADDRESS 0x20000000
#include "include_file_1.h"

lr1 ADDRESS
{
    ...
}
```

The linker parses the preprocessed scatter file and treats the directives as comments.

You can also use preprocessing of a scatter file in conjunction with the `--predefine` command-line option. For this example:

1. Modify `file.scat` to delete the directive `#define ADDRESS 0x20000000`.
2. Specify the command:

```
armlink --predefine="-DADDRESS=0x20000000" --scatter=file.scat
```

8.34.1 See also

Concepts

- [Expression evaluation in scatter files on page 8-58](#).

Reference

Linker Reference:

- [--predefine="string" on page 2-96](#).

8.35 Expression evaluation in scatter files

The linker can carry out simple expression evaluation with a restricted set of operators. The operators are +, -, *, /, AND, OR, and parentheses. The implementation of OR and AND follows C operator precedence rules.

8.35.1 Example of using expression evaluation

Use the directives:

```
#define BASE_ADDRESS 0x8000
#define ALIAS_NUMBER 0x2
#define ALIAS_SIZE 0x400

#define AN_ADDRESS (BASE_ADDRESS+(ALIAS_NUMBER*ALIAS_SIZE))
```

The scatter file might contain:

```
LOAD_FLASH AN_ADDRESS ; start address
```

After preprocessing, this evaluates to:

```
LOAD_FLASH ( 0x8000 + ( 0x2 * 0x400 )) ; start address
```

After evaluation, the linker parses the scatter file to produce the load region:

```
LOAD_FLASH 0x8800 ; start address
```

8.35.2 See also

Concepts

- [Using preprocessing commands in a scatter file on page 8-57](#)
- [Using expression evaluation in a scatter file to avoid padding on page 8-59.](#)

8.36 Using expression evaluation in a scatter file to avoid padding

Using the `ALIGN`, `ALIGNALL`, or `FIXED` attributes in a scatter file can result in a large amount of padding in the image. To remove this padding, you can use expression evaluation to specify the start address of a load region and execution region. The built-in function `AlignExpr` is available to help you specify address expressions.

8.36.1 Example to avoid padding in scatter file

The following scatter file produces an image with padding:

```
LR1 0x4000
{
    ER1 +0 ALIGN 0x8000
    {
        ...
    }
}
```

Using the `ALIGN` keyword `ER1` is aligned to a `0x8000` boundary in both the load and the execution view. To align in the load view, the linker must insert `0x4000` bytes of padding.

The following scatter file produces an image without padding:

```
LR1 0x4000
{
    ER1 AlignExpr(+0, 0x8000)
    {
        ...
    }
}
```

Using `AlignExpr` the result of `+0` is aligned to a `0x8000` boundary. This creates an execution region with a load address of `0x4000` but an Execution Address of `0x8000`.

8.36.2 See also

Concepts

- [Expression evaluation in scatter files on page 8-58.](#)

Linker Reference:

- [Example of aligning a base address in execution space but still tightly packed in load space on page 4-42.](#)

Reference

Linker Reference:

- [Execution region attributes on page 4-12](#)
- [AlignExpr\(expr, align\) function on page 4-43.](#)

8.37 Equivalent scatter-loading descriptions for simple images

The command-line options `--reloc`, `--ro_base`, `--rw_base`, `--ropi`, `--rwpi`, and `--split` create the simple image types:

- Type 1 image, one load region and contiguous execution regions
- Type 2 image, one load region and non-contiguous execution regions
- Type 3 image, two load regions and non-contiguous execution regions.

You can create the same image types by using the `--scatter` command-line option and a file containing one of the corresponding scatter-loading descriptions.

8.37.1 See also

Concepts

- [Types of simple image on page 4-10](#)
- [Type 1 image, one load region and contiguous execution regions on page 8-61](#)
- [Type 2 image, one load region and non-contiguous execution regions on page 8-63](#)
- [Type 3 image, two load regions and non-contiguous execution regions on page 8-65.](#)

Linker Reference:

- [About load region descriptions on page 4-5.](#)

Reference

Linker Reference:

- [--reloc on page 2-100](#)
- [--ro_base=address on page 2-103](#)
- [--ropi on page 2-104](#)
- [--rw_base=address on page 2-106](#)
- [--rwpi on page 2-107](#)
- [--scatter=filename on page 2-109](#)
- [--split on page 2-117](#)
- [Load region attributes on page 4-8.](#)

8.38 Type 1 image, one load region and contiguous execution regions

An image of this type consists of a single load region in the load view and three execution regions in the execution view. The execution regions are placed contiguously in the memory map.

`--ro_base address` specifies the load and execution address of the region containing the RO output section. The following example shows the scatter-loading description equivalent to using `--ro_base 0x040000`:

Example 8-25 Single load region and contiguous execution regions

```
LR_1 0x040000    ; Define the load region name as LR_1, the region starts at 0x040000.
{
  ER_RO +0      ; First execution region is called ER_RO, region starts at end of previous region.
                ; However, since there is no previous region, the address is 0x040000.
  {
    * (+RO)     ; All RO sections go into this region, they are placed consecutively.
  }
  ER_RW +0      ; Second execution region is called ER_RW, the region starts at the end of the
                ; previous region. The address is 0x040000 + size of ER_RO region.
  {
    * (+RW)     ; All RW sections go into this region, they are placed consecutively.
  }
  ER_ZI +0      ; Last execution region is called ER_ZI, the region starts at the end of the
                ; previous region at 0x040000 + the size of the ER_RO regions + the size of
                ; the ER_RW regions.
  {
    * (+ZI)     ; All ZI sections are placed consecutively here.
  }
}
}
```

In this example:

- This description creates an image with one load region called LR_1 that has a load address of 0x040000.
- The image has three execution regions, named ER_RO, ER_RW, and ER_ZI, that contain the RO, RW, and ZI output sections respectively. RO, RW are root regions. ZI is created dynamically at runtime. The execution address of ER_RO is 0x040000. All three execution regions are placed contiguously in the memory map by using the *+offset* form of the base designator for the execution region description. This enables an execution region to be placed immediately following the end of the preceding execution region.

Use the `--reloc` option to make relocatable images. Used on its own, `--reloc` makes an image similar to simple type 1, but the single load region has the RELOC attribute.

8.38.1 ropi example variant

In this variant, the execution regions are placed contiguously in the memory map. However, `--ropi` marks the load and execution regions containing the RO output section as position-independent.

The following example shows the scatter-loading description equivalent to using `--ro_base 0x010000 --ropi`:

Example 8-26 Position-independent code

```

LR_1 0x010000 PI      ; The first load region is at 0x010000.
{
  ER_RO +0           ; The PI attribute is inherited from parent.
                    ; The default execution address is 0x010000, but the code can be moved.
  {
    * (+RO)         ; All the RO sections go here.
  }
  ER_RW +0 ABSOLUTE ; PI attribute is overridden by ABSOLUTE.
  {
    * (+RW)         ; The RW sections are placed next. They cannot be moved.
  }
  ER_ZI +0          ; ER_ZI region placed after ER_RW region.
  {
    * (+ZI)         ; All the ZI sections are placed consecutively here.
  }
}

```

ER_RO, the RO execution region, inherits the PI attribute from the load region LR_1. The next execution region, ER_RW, is marked as ABSOLUTE and uses the *+offset* form of base designator. This prevents ER_RW from inheriting the PI attribute from ER_RO. Also, because the ER_ZI region has an offset of +0, it inherits the ABSOLUTE attribute from the ER_RW region.

8.38.2 See also**Concepts**

Linker Reference:

- [About load region descriptions on page 4-5](#)
- [Considerations when using a relative address +offset for load regions on page 4-17](#)
- [Considerations when using a relative address +offset for execution regions on page 4-18.](#)

Reference

Linker Reference:

- [--reloc on page 2-100](#)
- [--ro_base=address on page 2-103](#)
- [--ropi on page 2-104](#)
- [Load region attributes on page 4-8.](#)

8.39 Type 2 image, one load region and non-contiguous execution regions

An image of this type consists of a single load region in the load view and three execution regions in the execution view. It is similar to images of type 1 except that the RW execution region is not contiguous with the RO execution region.

--ro_base=*address1* specifies the load and execution address of the region containing the RO output section. --rw_base=*address2* specifies the execution address for the RW execution region.

The following example shows the scatter-loading description equivalent to using --ro_base=0x010000 --rw_base=0x040000:

Example 8-27 Single load region and multiple execution regions

```

LR_1 0x010000      ; Defines the load region name as LR_1
{
  ER_RO +0        ; The first execution region is called ER_RO and starts at end of previous region.
                  ; Because there is no previous region, the address is 0x010000.
  {
    * (+RO)       ; All RO sections are placed consecutively into this region.
  }
  ER_RW 0x040000  ; Second execution region is called ER_RW and starts at 0x040000.
  {
    * (+RW)       ; All RW sections are placed consecutively into this region.
  }
  ER_ZI +0        ; The last execution region is called ER_ZI.
                  ; The address is 0x040000 + size of ER_RW region.
  {
    * (+ZI)       ; All ZI sections are placed consecutively here.
  }
}

```

In this example:

- This description creates an image with one load region, named LR_1, with a load address of 0x010000.
- The image has three execution regions, named ER_RO, ER_RW, and ER_ZI, that contain the RO, RW, and ZI output sections respectively. The RO region is a root region. The execution address of ER_RO is 0x010000.
- The ER_RW execution region is not contiguous with ER_RO. Its execution address is 0x040000.
- The ER_ZI execution region is placed immediately following the end of the preceding execution region, ER_RW.

8.39.1 rwp_i example variant

This is similar to images of type 2 with --rw_base where the RW execution region is separate from the RO execution region. However, --rwp_i marks the execution regions containing the RW output section as position-independent.

The following example shows the scatter-loading description equivalent to using --ro_base=0x010000 --rw_base=0x018000 --rwp_i:

Example 8-28 Position-independent data

```

LR_1 0x010000      ; The first load region is at 0x010000.
{
  ER_RO +0         ; Default ABSOLUTE attribute is inherited from parent. The execution address
                  ; is 0x010000. The code and RO data cannot be moved.
  {
    * (+RO)        ; All the RO sections go here.
  }
  ER_RW 0x018000 PI ; PI attribute overrides ABSOLUTE
  {
    * (+RW)        ; The RW sections are placed at 0x018000 and they can be moved.
  }
  ER_ZI +0         ; ER_ZI region placed after ER_RW region.
  {
    * (+ZI)        ; All the ZI sections are placed consecutively here.
  }
}

```

ER_RO, the RO execution region, inherits the ABSOLUTE attribute from the load region LR_1. The next execution region, ER_RW, is marked as PI. Also, because the ER_ZI region has an offset of +0, it inherits the PI attribute from the ER_RW region.

Similar scatter-loading descriptions can also be written to correspond to the usage of other combinations of `--ropi` and `--rwpi` with type 2 and type 3 images.

8.39.2 See also**Concepts**

Linker Reference:

- [About load region descriptions on page 4-5](#)
- [Considerations when using a relative address +offset for load regions on page 4-17](#)
- [Considerations when using a relative address +offset for execution regions on page 4-18.](#)

Reference

Linker Reference:

- [--ro_base=address on page 2-103](#)
- [--rw_base=address on page 2-106](#)
- [--rwpi on page 2-107](#)
- [Load region attributes on page 4-8.](#)

8.40 Type 3 image, two load regions and non-contiguous execution regions

Type 3 images consist of two load regions in load view and three execution regions in execution view. They are similar to images of type 2 except that the single load region in type 2 is now split into two load regions.

Relocate and split load regions using the following linker options:

`--reloc` The combination `--reloc --split` makes an image similar to simple type 3, but the two load regions now have the RELOC attribute.

`--ro_base=address1`
 Specifies the load and execution address of the region containing the RO output section.

`--rw_base=address2`
 Specifies the load and execution address for the region containing the RW output section.

`--split` Splits the default single load region (that contains the RO and RW output sections) into two load regions. One load region contains the RO output section and one contains the RW output section.

The following example shows the scatter-loading description equivalent to using

`--ro_base=0x010000 --rw_base=0x040000 --split`:

Example 8-29 Multiple load regions

```
LR_1 0x010000    ; The first load region is at 0x010000.
{
  ER_RO +0      ; The address is 0x010000.
  {
    * (+RO)
  }
}
LR_2 0x040000    ; The second load region is at 0x040000.
{
  ER_RW +0      ; The address is 0x040000.
  {
    * (+RW)     ; All RW sections are placed consecutively into this region.
  }
  ER_ZI +0      ; The address is 0x040000 + size of ER_RW region.
  {
    * (+ZI)     ; All ZI sections are placed consecutively into this region.
  }
}
}
```

In this example:

- This description creates an image with two load regions, named LR_1 and LR_2, that have load addresses 0x010000 and 0x040000.
- The image has three execution regions, named ER_RO, ER_RW and ER_ZI, that contain the RO, RW, and ZI output sections respectively. The execution address of ER_RO is 0x010000.
- The ER_RW execution region is not contiguous with ER_RO, because its execution address is 0x040000.

- The ER_ZI execution region is placed immediately following the end of the preceding execution region, ER_RW.

8.40.1 Relocatable load regions example variant

This type 3 image also consists of two load regions in load view and three execution regions in execution view. However, `--reloc` is used to specify that the two load regions now have the RELOC attribute.

The following example shows the scatter-loading description equivalent to using `--ro_base 0x010000 --rw_base 0x040000 --reloc --split`:

Example 8-30 Relocatable load regions

```
LR_1 0x010000 RELOC
{
  ER_RO + 0
  {
    * (+RO)
  }
}

LR2 0x040000 RELOC
{
  ER_RW + 0
  {
    * (+RW)
  }

  ER_ZI +0
  {
    * (+ZI)
  }
}
```

8.40.2 See also

Concepts

Linker Reference:

- [About load region descriptions on page 4-5](#)
- [Considerations when using a relative address +offset for load regions on page 4-17](#)
- [Considerations when using a relative address +offset for execution regions on page 4-18.](#)

Reference

Linker Reference:

- [--reloc on page 2-100](#)
- [--ro_base=address on page 2-103](#)
- [--rw_base=address on page 2-106](#)
- [--split on page 2-117](#)
- [Load region attributes on page 4-8](#)
- [Address attributes for load and execution regions on page 4-15.](#)

8.41 Scatter file to ELF mapping

For simple images, ELF executable files contain segments:

- a load region is represented by an ELF Program Segment with type `PT_LOAD`
- an execution region is represented by up to three ELF Sections:
 - one for RO
 - one for RW
 - one for ZI.

For example, you might have a scatter file similar to the following:

Example 8-31 Scatter file

```
LOAD 0x8000
{
  EXEC_ROM +0
  {
    *(+RO)
  }
  RAM +0
  {
    *(+RW,+ZI)
  }
  HEAP +0x100 EMPTY 0x100
  {
  }
  STACK +0 EMPTY 0x400
  {
  }
}
```

This scatter file creates a single Program Segment with type `PT_LOAD` for the load region with address `0x8000`.

A single Output Section with type `SHT_PROGBITS` is created to represent the contents of `EXEC_ROM`. Two Output Sections are created to represent RAM. The first has a type `SHT_PROGBITS` and contains the initialized read/write data. The second has a type of `SHT_NOBITS` and describes the zero-initialized data.

The heap and stack are described in the ELF file by `SHT_NOBITS` sections.

Enter the following `fromelf` command to see the scatter-loaded sections in the image:

```
fromelf --text -v my_image.axf
```

To display the symbol table, enter the command:

```
fromelf --text -s -v my_image.axf
```

The following is an example of the `fromelf` output showing the `LOAD`, `EXEC_ROM`, `RAM`, `HEAP`, and `STACK` sections:

Example 8-32 Scatter-loaded sections in the ELF image

```

...
=====
** Program header #0

    Type           : PT_LOAD (1)
    File Offset    : 52 (0x34)
    Virtual Addr   : 0x00008000
    Physical Addr  : 0x00008000
    Size in file   : 764 bytes (0x2fc)
    Size in memory: 2140 bytes (0x85c)
    Flags          : PF_X + PF_W + PF_R + PF_ARM_ENTRY (0x80000007)
    Alignment      : 4
=====

** Section #1

    Name           : EXEC_ROM
...
    Addr           : 0x00008000
    File Offset    : 52 (0x34)
    Size           : 740 bytes (0x2e4)
...
=====

** Section #2

    Name           : RAM
...
    Addr           : 0x000082e4
    File Offset    : 792 (0x318)
    Size           : 20 bytes (0x14)
...
=====

** Section #3

    Name           : RAM
...
    Addr           : 0x000082f8
    File Offset    : 812 (0x32c)
    Size           : 96 bytes (0x60)
...
=====

** Section #4

    Name           : HEAP
...
    Addr           : 0x00008458
    File Offset    : 812 (0x32c)
    Size           : 256 bytes (0x100)
...
=====

** Section #5

    Name           : STACK
...
    Addr           : 0x00008558

```

File Offset : 812 (0x32c)
Size : 1024 bytes (0x400)
...

8.41.1 See also

Concepts

- [About scatter-loading on page 8-3](#)
- [Images with a simple memory map on page 8-6.](#)