

IAR Embedded Workbench[®]

IAR C/C++ Development Guide

Compiling and Linking

for Advanced RISC Machines Ltd's

ARM[®] Cores



COPYRIGHT NOTICE

© 1999–2016 IAR Systems AB.

No part of this document may be reproduced without the prior written consent of IAR Systems AB. The software described in this document is furnished under a license and may only be used or copied in accordance with the terms of such a license.

DISCLAIMER

The information in this document is subject to change without notice and does not represent a commitment on any part of IAR Systems. While the information contained herein is assumed to be accurate, IAR Systems assumes no responsibility for any errors or omissions.

In no event shall IAR Systems, its employees, its contractors, or the authors of this document be liable for special, direct, indirect, or consequential damage, losses, costs, charges, claims, demands, claim for lost profits, fees, or expenses of any nature or kind.

TRADEMARKS

IAR Systems, IAR Embedded Workbench, C-SPY, C-RUN, C-STAT, visualSTATE, Focus on Your Code, IAR KickStart Kit, IAR Experiment!, I-jet, I-jet Trace, I-scope, IAR Academy, IAR, and the logotype of IAR Systems are trademarks or registered trademarks owned by IAR Systems AB.

Microsoft and Windows are registered trademarks of Microsoft Corporation.

ARM, Thumb, and Cortex are registered trademarks of Advanced RISC Machines Ltd.

Adobe and Acrobat Reader are registered trademarks of Adobe Systems Incorporated.

All other product names are trademarks or registered trademarks of their respective owners.

EDITION NOTICE

Nineteenth edition: March 2016

Part number: DARM-19

This guide applies to version 7.60.x of IAR Embedded Workbench® for ARM.

Internal reference: M20, csrct2010.1, V_110411, IMAE.

Brief contents

Tables	35
Preface	37
Part 1. Using the build tools	45
Introduction to the IAR build tools	47
Developing embedded applications	53
Data storage	67
Functions	71
Linking using ILINK	83
Linking your application	101
The DLIB runtime environment	115
Assembler language interface	157
Using C	181
Using C++	191
Application-related considerations	203
Efficient coding for embedded applications	221
Part 2. Reference information	241
External interface details	243
Compiler options	253
Linker options	297
Data representation	329
Extended keywords	343

Pragma directives	359
Intrinsic functions	383
The preprocessor	419
C/C++ standard library functions	427
The linker configuration file	437
Section reference	463
The stack usage control file	469
IAR utilities	477
Implementation-defined behavior for Standard C	515
Implementation-defined behavior for C89	531
Index	543

Contents

Tables	35
Preface	37
Who should read this guide	37
Required knowledge	37
How to use this guide	37
What this guide contains	38
Part 1. Using the build tools	38
Part 2. Reference information	38
Other documentation	39
User and reference guides	40
The online help system	40
Further reading	41
Web sites	42
Document conventions	42
Typographic conventions	42
Naming conventions	43
 Part I. Using the build tools	45
Introduction to the IAR build tools	47
The IAR build tools—an overview	47
IAR C/C++ Compiler	47
IAR Assembler	48
The IAR ILINK Linker	48
Specific ELF tools	48
External tools	48
IAR language overview	49
Device support	49
Supported ARM devices	50
Preconfigured support files	50
Examples for getting started	50

Special support for embedded systems	51
Extended keywords	51
Pragma directives	51
Predefined symbols	51
Accessing low-level features	51
Developing embedded applications	53
Developing embedded software using IAR build tools	53
Mapping of memory	53
Communication with peripheral units	54
Event handling	54
System startup	54
Real-time operating systems	54
Interoperability with other build tools	55
The build process—an overview	55
The translation process	55
The linking process	56
After linking	58
Application execution—an overview	58
The initialization phase	59
The execution phase	62
The termination phase	62
Building applications—an overview	63
Basic project configuration	63
Processor configuration	64
Optimization for speed and size	65
Data storage	67
Introduction	67
Different ways to store data	67
Storage of auto variables and parameters	68
The stack	68
Dynamic memory on the heap	69
Potential problems	69

Functions	71
Function-related extensions	71
ARM and Thumb code	71
Execution in RAM	72
Interrupt functions for Cortex-M devices	73
Interrupts for Cortex-M	73
Interrupt functions for ARM7/9/11, Cortex-A, and Cortex-R devices	74
Interrupt functions	74
Installing exception functions	75
Interrupts and fast interrupts	76
Nested interrupts	77
Software interrupts	78
Interrupt operations	79
Inlining functions	79
C versus C++ semantics	80
Features controlling function inlining	80
Linking using ILINK	83
Linking—an overview	83
Modules and sections	84
The linking process in detail	85
Placing code and data—the linker configuration file	87
A simple example of a configuration file	88
Initialization at system startup	90
The initialization process	91
C++ dynamic initialization	92
Stack usage analysis	92
Introduction to stack usage analysis	92
Performing a stack usage analysis	93
Result of an analysis—the map file contents	94
Specifying additional stack usage information	95
Limitations	97
Situations where warnings are issued	97

Call graph log	98
Call graph XML output	99
Linking your application	101
Linking considerations	101
Choosing a linker configuration file	101
Defining your own memory areas	102
Placing sections	103
Reserving space in RAM	104
Keeping modules	105
Keeping symbols and sections	105
Application startup	105
Setting up stack memory	105
Setting up heap memory	106
Setting up the atexit limit	106
Changing the default initialization	106
Interaction between ILINK and the application	110
Standard library handling	110
Producing other output formats than ELF/DWARF	111
Veneers	111
Hints for troubleshooting	111
Relocation errors	111
Checking module consistency	113
Runtime model attributes	113
Using runtime model attributes	114
The DLIB runtime environment	115
Introduction to the runtime environment	115
Runtime environment functionality	115
Briefly about input and output (I/O)	116
Briefly about C-SPY emulated I/O	117
Briefly about retargeting	118
Setting up the runtime environment	119
Setting up your runtime environment	119
Retargeting—Adapting for your target system	120

Overriding library modules	122
Customizing and building your own runtime library	123
Additional information on the runtime environment	124
Runtime library configurations	125
Prebuilt runtime libraries	126
Formatters for printf	128
Formatters for scanf	129
The C-SPY emulated I/O mechanism	130
The semihosting mechanism	130
Math functions	131
System startup and termination	133
System initialization	137
The DLIB low-level I/O interface	138
abort	139
__acabi_assert	139
clock	140
__close	140
__exit	140
getenv	141
__getzone	141
__lseek	142
__open	142
raise	142
__read	143
remove	144
rename	144
signal	144
system	145
__time32, __time64	145
__write	145
Configuration symbols for printf and scanf	147
Configuration symbols for file input and output	148
Locale	148

Managing a multithreaded environment	150
Multithread support in the DLIB runtime environment	150
Enabling multithread support	151
TLS in the linker configuration file	155
Assembler language interface	157
Mixing C and assembler	157
Intrinsic functions	157
Mixing C and assembler modules	158
Inline assembler	158
Reference information for inline assembler	160
An example of how to use clobbered memory	166
Calling assembler routines from C	166
Creating skeleton code	167
Compiling the skeleton code	167
Calling assembler routines from C++	168
Calling convention	169
Function declarations	170
Using C linkage in C++ source code	170
Preserved versus scratch registers	171
Function entrance	172
Function exit	173
Examples	175
Call frame information	176
CFI directives	176
Creating assembler source with CFI support	177
Using C	181
C language overview	181
Extensions overview	182
Enabling language extensions	183
IAR C language extensions	183
Extensions for embedded systems programming	184
Relaxations to Standard C	186

Using C++	191
Overview—EC++ and EEC++	191
Embedded C++	191
Extended Embedded C++	192
Overview—Standard C++	193
Modes for exceptions and RTTI support	193
Exception handling	194
Enabling support for C++ and variants	196
C++ and EC++ feature descriptions	196
Using IAR attributes with Classes	196
Function types	197
Using static class objects in interrupts	198
Using New handlers	198
Templates	199
Debug support in C-SPY	199
EEC++ feature description	199
Templates	199
Variants of cast operators	199
Mutable	199
Namespace	199
The STD namespace	200
EC++ and C++ language extensions	200
Application-related considerations	203
Output format considerations	203
Stack considerations	204
Stack size considerations	204
Stack alignment	204
Exception stack	204
Heap considerations	205
Advanced versus basic heap	205
Heap size and standard I/O	206
Interaction between the tools and your application	206

Checksum calculation for verifying image integrity	208
Briefly about checksum calculation	208
Calculating and verifying a checksum	210
Troubleshooting checksum calculation	216
Linker optimizations	217
Virtual function elimination	217
AEABI compliance	217
Linking AEABI-compliant modules using the IAR ILINK linker ..	218
Linking AEABI-compliant modules using a third-party linker	219
Enabling AEABI compliance in the compiler	219
CMSIS integration	219
CMSIS DSP library	220
Customizing the CMSIS DSP library	220
Building with CMSIS on the command line	220
Building with CMSIS in the IDE	220
Efficient coding for embedded applications	221
Selecting data types	221
Using efficient data types	221
Floating-point types	222
Alignment of elements in a structure	222
Anonymous structs and unions	223
Controlling data and function placement in memory	224
Data placement at an absolute location	225
Data and function placement in sections	226
Data placement in registers	227
Controlling compiler optimizations	228
Scope for performed optimizations	229
Multi-file compilation units	229
Optimization levels	230
Speed versus size	231
Fine-tuning enabled transformations	231
Facilitating good code generation	234
Writing optimization-friendly source code	234

Saving stack space and RAM memory	235
Function prototypes	235
Integer types and bit negation	236
Protecting simultaneously accessed variables	236
Accessing special function registers	237
Passing values between C and assembler objects	239
Non-initialized variables	239
Part 2. Reference information	241
External interface details	243
Invocation syntax	243
Compiler invocation syntax	243
ILINK invocation syntax	244
Passing options	244
Environment variables	245
Include file search procedure	245
Compiler output	246
Error return codes	247
ILINK output	248
Diagnostics	248
Message format for the compiler	248
Message format for the linker	249
Severity levels	249
Setting the severity level	250
Internal error	250
Compiler options	253
Options syntax	253
Types of options	253
Rules for specifying parameters	253
Summary of compiler options	255
Descriptions of compiler options	260
--aapcs	260

--aeabi	260
--align_sp_on_irq	261
--arm	261
--c89	261
--char_is_signed	262
--char_is_unsigned	262
--cpu	262
--cpu_mode	263
--c++	264
-D	264
--debug, -r	265
--dependencies	265
--diag_error	266
--diag_remark	267
--diag_suppress	267
--diag_warning	267
--diagnostics_tables	268
--discard_unused_publics	268
--dlib_config	269
-e	269
--ec++	270
--eec++	270
--enable_hardware_workaround	270
--enable_multibytes	271
--enable_restrict	271
--endian	271
--enum_is_int	272
--error_limit	272
-f	273
--fpu	273
--guard_calls	274
--header_context	274
-I	275
--interwork	275

-l	275
--legacy	276
--lock_regs	277
--macro_positions_in_diagnostics	277
--make_all_definitions_weak	277
--mfc	278
--no_alignment_reduction	278
--no_call_frame_info	278
--no_clustering	279
--no_code_motion	279
--no_const_align	279
--no_cse	280
--no_exceptions	280
--no_fragments	281
--no_inline	281
--no_literal_pool	281
--no_loop_align	282
--no_mem_idioms	282
--no_path_in_file_macros	282
--no_rtti	283
--no_rw_dynamic_init	283
--no_scheduling	283
--no_size_constraints	284
--no_static_destruction	284
--no_system_include	284
--no_tbaa	285
--no_typedefs_in_diagnostics	285
--no_unaligned_access	285
--no_unroll	286
--no_var_align	286
--no_warnings	287
--no_wrap_diagnostics	287
-O	287
--only_stdout	288

--output, -o	288
--predef_macros	288
--preinclude	289
--preprocess	289
--public_equ	289
--relaxed_fp	290
--remarks	290
--require_prototypes	291
--ropi	291
--rwpi	292
--section	292
--separate_cluster_for_initialized_variables	293
--silent	293
--strict	293
--system_include_dir	294
--thumb	294
--use_c++_inline	294
--use_unix_directory_separators	295
--vectorize	295
--vla	295
--warn_about_c_style_casts	296
--warnings_affect_exit_code	296
--warnings_are_errors	296
Linker options	297
Summary of linker options	297
Descriptions of linker options	300
--advanced_heap	301
--basic_heap	301
--BE8	301
--BE32	302
--call_graph	302
--config	302
--config_def	303

--config_search	303
--cpp_init_routine	304
--cpu	304
--default_to_complex_ranges	304
--define_symbol	305
--dependencies	305
--diag_error	306
--diag_remark	306
--diag_suppress	307
--diag_warning	307
--diagnostics_tables	307
--do_segment_pad	308
--enable_hardware_workaround	308
--enable_stack_usage	309
--entry	309
--error_limit	309
--exception_tables	310
--export_builtin_config	310
--extra_init	311
-f	311
--force_exceptions	312
--force_output	312
--image_input	312
--inline	313
--keep	313
--log	314
--log_file	315
--mangled_names_in_messages	315
--map	315
--merge_duplicate_sections	316
--no_dynamic_rtti_elimination	316
--no_entry	317
--no_exceptions	317
--no_fragments	317

--no_library_search	318
--no_literal_pool	318
--no_locals	319
--no_range_reservations	319
--no_remove	319
--no_veneers	320
--no_vfe	320
--no_warnings	320
--no_wrap_diagnostics	321
--only_stdout	321
--output, -o	321
--pi_veneers	321
--place_holder	322
--redirect	322
--remarks	323
--search, -L	323
--semihosting	323
--silent	324
--skip_dynamic_initialization	324
--stack_usage_control	325
--strip	325
--threaded_lib	325
--use_full_std_template_names	326
--vfe	326
--warnings_affect_exit_code	327
--warnings_are_errors	327
--whole_archive	327
Data representation	329
Alignment	329
Alignment on the ARM core	330
Byte order	330
Basic data types—integer types	331
Integer types—an overview	331

Bool	331
The enum type	331
The char type	332
The wchar_t type	332
Bitfields	332
Basic data types—floating-point types	336
Floating-point environment	336
32-bit floating-point format	336
64-bit floating-point format	337
Representation of special floating-point numbers	337
Pointer types	337
Function pointers	337
Data pointers	338
Casting	338
Structure types	339
Alignment of structure types	339
General layout	339
Packed structure types	339
Type qualifiers	340
Declaring objects volatile	341
Declaring objects volatile and const	342
Declaring objects const	342
Data types in C++	342
Extended keywords	343
General syntax rules for extended keywords	343
Type attributes	343
Object attributes	345
Summary of extended keywords	346
Descriptions of extended keywords	347
__absolute	347
__arm	347
__big_endian	347
__fiq	348

__interwork	348
__intrinsic	348
__irq	348
__little_endian	349
__nested	349
__no_alloc, __no_alloc16	350
__no_alloc_str, __no_alloc_str16	350
__no_init	351
__noreturn	351
__packed	352
__ramfunc	353
__ro_placement	353
__root	354
__stackless	354
__swi	354
__task	356
__thumb	356
__weak	356
Pragma directives	359
Summary of pragma directives	359
Descriptions of pragma directives	361
bitfields	361
calls	362
call_graph_root	363
data_alignment	363
default_function_attributes	364
default_variable_attributes	365
diag_default	366
diag_error	366
diag_remark	366
diag_suppress	367
diag_warning	367
error	367

include_alias	368
inline	368
language	369
location	370
message	371
object_attribute	371
optimize	372
pack	373
__printf_args	374
public_equ	374
required	375
rtmodel	375
__scanf_args	376
section	376
STDC CX_LIMITED_RANGE	377
STDC FENV_ACCESS	377
STDC FP_CONTRACT	378
swi_number	378
type_attribute	378
unroll	379
vectorize	380
weak	380
Intrinsic functions	383
Summary of intrinsic functions	383
Intrinsic functions for Neon instructions	389
Descriptions of intrinsic functions	390
__CLREX	390
__CLZ	390
__disable_fiq	390
__disable_interrupt	391
__disable_irq	391
__DMB	391
__DSB	391

__enable_fiq	391
__enable_interrupt	392
__enable_irq	392
__get_BASEPRI	392
__get_CONTROL	392
__get_CPSR	392
__get_FAULTMASK	393
__get_FPSCR	393
__get_interrupt_state	393
__get_IPSR	394
__get_LR	394
__get_MSP	394
__get_PRIMASK	394
__get_PSP	394
__get_PSR	395
__get_SB	395
__get_SP	395
__ISB	395
__LDC	395
__LDCL	395
__LDC2	395
__LDC2L	395
__LDC_noidx	396
__LDCL_noidx	396
__LDC2_noidx	396
__LDC2L_noidx	396
__LDREX	397
__LDREXB	397
__LDREXD	397
__LDREXH	397
__MCR	397
__MCR2	397
__MRC	398
__MRC2	398

__no_operation	399
__PKHBT	399
__PKHTB	399
__PLD	400
__PLDW	400
__PLI	400
__QADD	400
__QDADD	400
__QDSUB	400
__QSUB	400
__QADD8	401
__QADD16	401
__QASX	401
__QSAX	401
__QSUB8	401
__QSUB16	401
__QCFlag	401
__QDOUBLE	401
__QFlag	402
__RBIT	402
__reset_Q_flag	402
__reset_QC_flag	402
__REV	403
__REV16	403
__REVSH	403
__SADD8	403
__SADD16	403
__SASX	403
__SSAX	403
__SSUB8	403
__SSUB16	403
__SEL	403
__set_BASEPRI	404
__set_CONTROL	404

__set_CPSR	404
__set_FAULTMASK	404
__set_FPSCR	404
__set_interrupt_state	405
__set_LR	405
__set_MSP	405
__set_PRIMASK	405
__set_PSP	405
__set_SB	406
__set_SP	406
__SEV	406
__SHADD8	406
__SHADD16	406
__SHASX	406
__SHSAX	406
__SHSUB8	406
__SHSUB16	406
__SMLABB	407
__SMLABT	407
__SMLATB	407
__SMLATT	407
__SMLAWB	407
__SMLAWT	407
__SMLAD	407
__SMLADX	407
__SMLSD	407
__SMLSDX	407
__SMLALBB	408
__SMLALBT	408
__SMLALTB	408
__SMLALTT	408
__SMLALD	408
__SMLALDX	408
__SMLSLD	408

__SMLSLDX	408
__SMMLA	409
__SMMLAR	409
__SMMLS	409
__SMMLSR	409
__SMMUL	409
__SMMULR	409
__SMUAD	409
__SMUADX	409
__SMUSD	409
__SMUSDX	409
__SMUL	410
__SMULBB	410
__SMULBT	410
__SMULTB	410
__SMULTT	410
__SMULWB	410
__SMULWT	410
__SSAT	410
__SSAT16	411
__STC	411
__STCL	411
__STC2	411
__STC2L	411
__STC_noidx	412
__STCL_noidx	412
__STC2_noidx	412
__STC2L_noidx	412
__STREX	413
__STREXB	413
__STREXD	413
__STREXH	413
__SWP	413
__SWPB	413

__SXTAB	414
__SXTAB16	414
__SXTAH	414
__SXTB16	414
__UADD8	414
__UADD16	414
__UASX	414
__USAX	414
__USUB8	414
__USUB16	414
__UHADD8	415
__UHADD16	415
__UHASX	415
__UHSAX	415
__UHSUB8	415
__UHSUB16	415
__UMAAL	415
__UQADD8	416
__UQADD16	416
__UQASX	416
__UQSAX	416
__UQSUB8	416
__UQSUB16	416
__USAD8	416
__USADA8	416
__USAT	416
__USAT16	417
__UXTAB	417
__UXTAB16	417
__UXTAH	417
__UXTB16	417
__WFE	417
__WFI	417
__YIELD	417

The preprocessor	419
Overview of the preprocessor	419
Description of predefined preprocessor symbols	420
__AAPCS__	420
__AAPCS_VFP__	420
__ARM_ADVANCED_SIMD__	420
__ARM_MEDIA__	420
__ARM_PROFILE_M__	421
__ARMVFP__	421
__ARMVFP_D16__	421
__ARMVFP_FP16__	421
__ARMVFP_SP__	421
__BASE_FILE__	422
__BUILD_NUMBER__	422
__CORE__	422
__COUNTER__	422
__cplusplus	422
__CPU_MODE__	422
__DATE__	423
__embedded_cplusplus	423
__FILE__	423
__func__	423
__FUNCTION__	423
__IAR_SYSTEMS_ICC__	424
__ICCARM__	424
__LINE__	424
__LITTLE_ENDIAN__	424
__PRETTY_FUNCTION__	424
__ROPI__	424
__RWPI__	425
__STDC__	425
__STDC_VERSION__	425
__TIME__	425

__TIMESTAMP__	425
__VER__	425
Descriptions of miscellaneous preprocessor extensions	426
NDEBUG	426
#warning message	426
C/C++ standard library functions	427
C/C++ standard library overview	427
Header files	427
Library object files	428
Alternative more accurate library functions	428
Reentrancy	428
The longjmp function	429
DLIB runtime environment—implementation details	429
C header files	429
C++ header files	430
Library functions as intrinsic functions	433
Added C functionality	433
Symbols used internally by the library	435
The linker configuration file	437
Overview	437
Defining memories and regions	438
define memory directive	439
define region directive	439
Regions	440
Region literal	440
Region expression	441
Empty region	442
Section handling	443
define block directive	444
define overlay directive	445
initialize directive	446
do not initialize directive	449
keep directive	449

place at directive	450
place in directive	451
use init table directive	451
Section selection	452
section-selectors	452
extended-selectors	455
Using symbols, expressions, and numbers	456
check that directive	457
define symbol directive	457
export directive	458
expressions	459
numbers	460
Structural configuration	460
error directive	460
if directive	461
include directive	461
Section reference	463
Summary of sections	463
Descriptions of sections and blocks	464
.bss	464
CSTACK	464
.data	465
.data_init	465
__DLIB_PERTHREAD	465
.exc.text	465
HEAP	465
.iar.dynexit	466
.init_array	466
.intvec	466
IRQ_STACK	466
.noinit	467
.preinit_array	467
.prepreinit_array	467

.rodata	467
.text	467
.textrw	468
.textrw_init	468
The stack usage control file	469
Overview	469
C++ names	469
Stack usage control directives	469
call graph root directive	470
exclude directive	470
function directive	470
max recursion depth directive	471
no calls from directive	471
possible calls directive	472
Syntactic components	472
<i>category</i>	473
<i>func-spec</i>	473
<i>module-spec</i>	473
<i>name</i>	474
<i>call-info</i>	474
<i>stack-size</i>	474
<i>size</i>	475
IAR utilities	477
The IAR Archive Tool—iarchive	477
Invocation syntax	477
Summary of iarchive commands	478
Summary of iarchive options	479
Diagnostic messages	479
The IAR ELF Tool—ielftool	480
Invocation syntax	481
Summary of ielftool options	481
The IAR ELF Dumper—ielfdump	482
Invocation syntax	482

Summary of ielfdump options	483
The IAR ELF Object Tool—iobjmanip	484
Invocation syntax	484
Summary of iobjmanip options	484
Diagnostic messages	485
The IAR Absolute Symbol Exporter—isymexport	486
Invocation syntax	487
Summary of isymexport options	488
Steering files	488
Show directive	489
Hide directive	489
Rename directive	489
Diagnostic messages	490
Descriptions of options	492
--a	492
--all	492
--bin	492
--checksum	493
--code	496
--create	496
--delete, -d	497
--edit	497
--extract, -x	497
-f	498
--fill	499
--generate_vfe_header	499
--ihex	500
--no_header	500
--no_rel_section	500
--no_strtab	501
--offset	501
--output, -o	502
--parity	502
--ram_reserve_ranges	503

--range	504
--raw	504
--remove_file_path	505
--remove_section	505
--rename_section	505
--rename_symbol	506
--replace, -r	506
--reserve_ranges	507
--section, -s	507
--segment, -g	508
--self_reloc	508
--show_entry_as	509
--silent	509
--simple	509
--simple-ne	510
--source	510
--srec	510
--srec-len	511
--srec-s3only	511
--strip	511
--symbols	512
--titxt	512
--toc, -t	512
--use_full_std_template_names	513
--verbose, -V	513
--version	513
Implementation-defined behavior for Standard C	515
Descriptions of implementation-defined behavior	515
J.3.1 Translation	515
J.3.2 Environment	516
J.3.3 Identifiers	517
J.3.4 Characters	517
J.3.5 Integers	519

J.3.6 Floating point	519
J.3.7 Arrays and pointers	520
J.3.8 Hints	520
J.3.9 Structures, unions, enumerations, and bitfields	521
J.3.10 Qualifiers	521
J.3.11 Preprocessing directives	522
J.3.12 Library functions	524
J.3.13 Architecture	528
J.4 Locale	529
Implementation-defined behavior for C89	531
Descriptions of implementation-defined behavior	531
Translation	531
Environment	531
Identifiers	532
Characters	532
Integers	533
Floating point	534
Arrays and pointers	535
Registers	535
Structures, unions, enumerations, and bitfields	535
Qualifiers	536
Declarators	536
Statements	536
Preprocessing directives	536
Library functions for the IAR DLIB Runtime Environment	538
Index	543

Tables

1: Typographic conventions used in this guide	42
2: Naming conventions used in this guide	43
3: Sections holding initialized data	90
4: Description of a relocation error	112
5: Example of runtime model attributes	114
6: Library configurations	125
7: Formatters for printf	128
8: Formatters for scanf	129
9: DLIB low-level I/O interface functions	138
10: Descriptions of printf configuration symbols	147
11: Descriptions of scanf configuration symbols	147
12: Library objects using TLS	151
13: Macros for implementing TLS allocation	154
14: Inline assembler operand constraints	161
15: Supported constraint modifiers	162
16: Operand modifiers and transformations	163
17: List of valid clobbers	165
18: Registers used for passing parameters	172
19: Registers used for returning values	174
20: Call frame information resources defined in a names block	177
21: Language extensions	183
22: Section operators and their symbols	186
23: Exception stacks for ARM7/9/11, Cortex-A, and Cortex-R	204
24: Compiler optimization levels	230
25: Compiler environment variables	245
26: ILINK environment variables	245
27: Error return codes	247
28: Compiler options summary	255
29: Linker options summary	297
30: Integer types	331
31: Floating-point types	336

32: Extended keywords summary	346
33: Pragma directives summary	359
34: Intrinsic functions summary	383
35: Traditional Standard C header files—DLIB	430
36: C++ header files	431
37: Standard template library header files	432
38: New Standard C header files—DLIB	432
39: Examples of section selector specifications	454
40: Section summary	463
41: iarchive parameters	478
42: iarchive commands summary	478
43: iarchive options summary	479
44: ielftool parameters	481
45: ielftool options summary	481
46: ielfdumparm parameters	483
47: ielfdumparm options summary	483
48: iobjmanip parameters	484
49: iobjmanip options summary	484
50: isymexport parameters	487
51: isymexport options summary	488
52: Message returned by strerror()—DLIB runtime environment	530
53: Message returned by strerror()—DLIB runtime environment	541

Preface

Welcome to the *IAR C/C++ Development Guide for ARM*. The purpose of this guide is to provide you with detailed reference information that can help you to use the build tools to best suit your application requirements. This guide also gives you suggestions on coding techniques so that you can develop applications with maximum efficiency.

Who should read this guide

Read this guide if you plan to develop an application using the C or C++ language for the ARM core and need detailed reference information on how to use the build tools.

REQUIRED KNOWLEDGE

To use the tools in IAR Embedded Workbench, you should have working knowledge of:

- The architecture and instruction set of the ARM core (refer to the chip manufacturer's documentation)
- The C or C++ programming language
- Application development for embedded systems
- The operating system of your host computer.

For more information about the other development tools incorporated in the IDE, refer to their respective documentation, see *Other documentation*, page 39.

How to use this guide

When you start using the IAR C/C++ compiler and linker for ARM, you should read *Part 1. Using the build tools* in this guide.

When you are familiar with the compiler and linker and have already configured your project, you can focus more on *Part 2. Reference information*.

If you are new to using this product, we suggest that you first read the guide *Getting Started with IAR Embedded Workbench®* for an overview of the tools and the features that the IDE offers. The tutorials, which you can find in IAR Information Center, will help you get started using IAR Embedded Workbench.

What this guide contains

Below is a brief outline and summary of the chapters in this guide.

PART 1. USING THE BUILD TOOLS

- *Introduction to the IAR build tools* gives an introduction to the IAR build tools, which includes an overview of the tools, the programming languages, the available device support, and extensions provided for supporting specific features of the ARM core.
- *Developing embedded applications* gives the information you need to get started developing your embedded software using the IAR build tools.
- *Data storage* describes how to store data in memory.
- *Functions* gives a brief overview of function-related extensions—mechanisms for controlling functions—and describes some of these mechanisms in more detail.
- *Linking using ILINK* describes the linking process using the IAR ILINK Linker and the related concepts.
- *Linking your application* lists aspects that you must consider when linking your application, including using ILINK options and tailoring the linker configuration file.
- *The DLIB runtime environment* describes the DLIB runtime environment in which an application executes. It covers how you can modify it by setting options, overriding default library modules, or building your own library. The chapter also describes system initialization introducing the file `cstartup`, how to use modules for locale, and file I/O.
- *Assembler language interface* contains information required when parts of an application are written in assembler language. This includes the calling convention.
- *Using C* gives an overview of the two supported variants of the C language and an overview of the compiler extensions, such as extensions to Standard C.
- *Using C++* gives an overview of the two levels of C++ support: The industry-standard EC++ and IAR Extended EC++.
- *Application-related considerations* discusses a selected range of application issues related to using the compiler and linker.
- *Efficient coding for embedded applications* gives hints about how to write code that compiles to efficient code for an embedded application.

PART 2. REFERENCE INFORMATION

- *External interface details* provides reference information about how the compiler and linker interact with their environment—the invocation syntax, methods for passing options to the compiler and linker, environment variables, the include file

search procedure, and the different types of compiler and linker output. The chapter also describes how the diagnostic system works.

- *Compiler options* explains how to set options, gives a summary of the options, and contains detailed reference information for each compiler option.
- *Linker options* gives a summary of the options, and contains detailed reference information for each linker option.
- *Data representation* describes the available data types, pointers, and structure types. This chapter also gives information about type and object attributes.
- *Extended keywords* gives reference information about each of the ARM-specific keywords that are extensions to the standard C/C++ language.
- *Pragma directives* gives reference information about the pragma directives.
- *Intrinsic functions* gives reference information about functions to use for accessing ARM-specific low-level features.
- *The preprocessor* gives a brief overview of the preprocessor, including reference information about the different preprocessor directives, symbols, and other related information.
- *C/C++ standard library functions* gives an introduction to the C or C++ library functions, and summarizes the header files.
- *The linker configuration file* describes the purpose of the linker configuration file and describes its contents.
- *Section reference* gives reference information about the use of sections.
- *The stack usage control file* describes the syntax and semantics of stack usage control files.
- *IAR utilities* describes the IAR utilities that handle the ELF and DWARF object formats.
- *Implementation-defined behavior for Standard C* describes how the compiler handles the implementation-defined areas of Standard C.
- *Implementation-defined behavior for C89* describes how the compiler handles the implementation-defined areas of the C language standard C89.

Other documentation

User documentation is available as hypertext PDFs and as a context-sensitive online help system in HTML format. You can access the documentation from the Information Center or from the **Help** menu in the IAR Embedded Workbench IDE. The online help system is also available via the F1 key.

USER AND REFERENCE GUIDES

The complete set of IAR Systems development tools is described in a series of guides. Information about:

- System requirements and information about how to install and register the IAR Systems products, is available in the booklet Quick Reference (available in the product box) and the *Installation and Licensing Guide*.
- Getting started using IAR Embedded Workbench and the tools it provides, is available in the guide *Getting Started with IAR Embedded Workbench®*.
- Using the IDE for project management and building, is available in the *IDE Project Management and Building Guide for ARM*.
- Using the IAR C-SPY® Debugger, is available in the *C-SPY® Debugging Guide for ARM*.
- Programming for the IAR C/C++ Compiler for ARM and linking using the IAR ILINK Linker, is available in the *IAR C/C++ Development Guide for ARM*.
- Programming for the IAR Assembler for ARM, is available in the *IAR Assembler User Guide for ARM*.
- Using the IAR DLIB C/C++ standard library functions, is available in the *C/C++ standard library reference*, available in the online help system.
- Performing a static analysis using C-STAT and the required checks, is available in the *C-STAT® Static Analysis Guide*.
- Developing safety-critical applications using the MISRA C guidelines, is available in the *IAR Embedded Workbench® MISRA C:2004 Reference Guide* or the *IAR Embedded Workbench® MISRA C:1998 Reference Guide*.
- Using I-jet, refer to the *IAR Debug probes User Guide for I-jet®, I-jet Trace, and I-scope*.
- Using JTAGjet-Trace, refer to the *JTAGjet-Trace User Guide for ARM*.
- Using IAR J-Link and IAR J-Trace, refer to the *IAR J-Link and IAR J-Trace User Guide for JTAG Emulators for ARM Cores*.
- Porting application code and projects created with a previous version of the IAR Embedded Workbench for ARM, is available in the *IAR Embedded Workbench® Migration Guide*.

Note: Additional documentation might be available depending on your product installation.

THE ONLINE HELP SYSTEM

The context-sensitive online help contains:

- Information about project management, editing, and building in the IDE

- Information about debugging using the IAR C-SPY® Debugger
- Reference information about the menus, windows, and dialog boxes in the IDE
- Compiler reference information
- Keyword reference information for the DLIB library functions. To obtain reference information for a function, select the function name in the editor window and press F1.

FURTHER READING

These books might be of interest to you when using the IAR Systems development tools:

- Seal, David, and David Jagger. ARM Architecture Reference Manual. Addison-Wesley
- Barr, Michael, and Andy Oram, ed. *Programming Embedded Systems in C and C++*. O'Reilly & Associates.
- Furber, Steve, ARM System-on-Chip Architecture. Addison-Wesley.
- Harbison, Samuel P. and Guy L. Steele (contributor). *C: A Reference Manual*. Prentice Hall.
- Josuttis, Nicolai M. *The C++ Standard Library: A Tutorial and Reference*. Addison-Wesley.
- Kernighan, Brian W. and Dennis M. Ritchie. *The C Programming Language*. Prentice Hall.
- Labrosse, Jean J. *Embedded Systems Building Blocks: Complete and Ready-To-Use Modules in C*. R&D Books.
- Lippman, Stanley B. and Josée Lajoie. *C++ Primer*. Addison-Wesley.
- Mann, Bernhard. C für Mikrocontroller. Franzis-Verlag. [Written in German.]
- Meyers, Scott. *Effective C++: 50 Specific Ways to Improve Your Programs and Designs*. Addison-Wesley.
- Meyers, Scott. *More Effective C++*. Addison-Wesley.
- Meyers, Scott. *Effective STL*. Addison-Wesley.
- Sloss, Andrew N. et al, ARM System Developer's Guide: Designing and Optimizing System Software. Morgan Kaufmann.
- Stroustrup, Bjarne. *The C++ Programming Language*. Addison-Wesley.
- Stroustrup, Bjarne. *Programming Principles and Practice Using C++*. Addison-Wesley.
- Sutter, Herb. *Exceptional C++: 47 Engineering Puzzles, Programming Problems, and Solutions*. Addison-Wesley.

WEB SITES

Recommended web sites:

- The Advanced RISC Machines Ltd web site, **www.arm.com**, that contains information and news about the ARM cores.
- The IAR Systems web site, **www.iar.com**, that holds application notes and other product information.
- The web site of the C standardization working group, **www.open-std.org/jtc1/sc22/wg14**.
- The web site of the C++ Standards Committee, **www.open-std.org/jtc1/sc22/wg21**.
- Finally, the Embedded C++ Technical Committee web site, **www.caravan.net/ec2plus**, that contains information about the Embedded C++ standard.

Document conventions

When, in the IAR Systems documentation, we refer to the programming language C, the text also applies to C++, unless otherwise stated.

When referring to a directory in your product installation, for example `arm\doc`, the full path to the location is assumed, for example `c:\Program Files\IAR Systems\Embedded Workbench N.n\arm\doc`, where the initial digit of the version number reflects the initial digit of the version number of the IAR Embedded Workbench shared components.

TYPOGRAPHIC CONVENTIONS

The IAR Systems documentation set uses the following typographic conventions:

Style	Used for
<code>computer</code>	<ul style="list-style-type: none">• Source code examples and file paths.• Text on the command line.• Binary, hexadecimal, and octal numbers.
<i>parameter</i>	A placeholder for an actual value used as a parameter, for example <i>filename.h</i> where <i>filename</i> represents the name of the file.
[option]	An optional part of a directive, where [and] are not part of the actual directive, but any [,], {, or } are part of the directive syntax.
{option}	A mandatory part of a directive, where { and } are not part of the actual directive, but any [,], {, or } are part of the directive syntax.
[option]	An optional part of a command.

Table 1: Typographic conventions used in this guide





Style	Used for
[a b c]	An optional part of a command with alternatives.
{a b c}	A mandatory part of a command with alternatives.
bold	Names of menus, menu commands, buttons, and dialog boxes that appear on the screen.
<i>italic</i>	<ul style="list-style-type: none">• A cross-reference within this guide or to another guide.• Emphasis.
...	An ellipsis indicates that the previous item can be repeated an arbitrary number of times.
	Identifies instructions specific to the IAR Embedded Workbench® IDE interface.
	Identifies instructions specific to the command line interface.
	Identifies helpful tips and programming hints.
	Identifies warnings.

Table 1: Typographic conventions used in this guide (Continued)

NAMING CONVENTIONS

The following naming conventions are used for the products and tools from IAR Systems®, when referred to in the documentation:


Brand name	Generic term
IAR Embedded Workbench® for ARM	IAR Embedded Workbench®
IAR Embedded Workbench® IDE for ARM	the IDE
IAR C-SPY® Debugger for ARM	C-SPY, the debugger
IAR C-SPY® Simulator	the simulator
IAR C/C++ Compiler™ for ARM	the compiler
IAR Assembler™ for ARM	the assembler
IAR ILINK Linker™	ILINK, the linker
IAR DLIB Runtime Environment™	the DLIB runtime environment

Table 2: Naming conventions used in this guide

Part I. Using the build tools

This part of the *IAR C/C++ Development Guide for ARM* includes these chapters:

- Introduction to the IAR build tools
- Developing embedded applications
- Data storage
- Functions
- Linking using ILINK
- Linking your application
- The DLIB runtime environment
- Assembler language interface
- Using C
- Using C++
- Application-related considerations
- Efficient coding for embedded applications.





Introduction to the IAR build tools

- The IAR build tools—an overview
- IAR language overview
- Device support
- Special support for embedded systems

The IAR build tools—an overview

In the IAR product installation you can find a set of tools, code examples, and user documentation, all suitable for developing software for ARM-based embedded applications. The tools allow you to develop your application in C, C++, or in assembler language.



IAR Embedded Workbench® is a very powerful Integrated Development Environment (IDE) that allows you to develop and manage complete embedded application projects. It provides an easy-to-learn and highly efficient development environment with maximum code inheritance capabilities, comprehensive and specific target support. IAR Embedded Workbench promotes a useful working methodology, and thus a significant reduction of the development time.

For information about the IDE, see the *IDE Project Management and Building Guide for ARM*.



The compiler, assembler, and linker can also be run from a command line environment, if you want to use them as external tools in an already established project environment.

IAR C/C++ COMPILER

The IAR C/C++ Compiler for ARM is a state-of-the-art compiler that offers the standard features of the C and C++ languages, plus extensions designed to take advantage of the ARM-specific facilities.

IAR ASSEMBLER

The IAR Assembler for ARM is a powerful relocating macro assembler with a versatile set of directives and expression operators. The assembler features a built-in C language preprocessor and supports conditional assembly.

The IAR Assembler for ARM uses the same mnemonics and operand syntax as the Advanced RISC Machines Ltd ARM Assembler, which simplifies the migration of existing code. For more information, see the *IAR Assembler User Guide for ARM*.

THE IAR ILINK LINKER

The IAR ILINK Linker for ARM is a powerful, flexible software tool for use in the development of embedded controller applications. It is equally well suited for linking small, single-file, absolute assembler programs as it is for linking large, relocatable input, multi-module, C/C++, or mixed C/C++ and assembler programs.

SPECIFIC ELF TOOLS

ILINK both uses and produces industry-standard ELF and DWARF as object format, additional IAR utilities that handle these formats are provided:

- The IAR Archive Tool—`iarchive`—creates and manipulates a library (archive) of several ELF object files
- The IAR ELF Tool—`ielftool`—performs various transformations on an ELF executable image (such as, fill, checksum, format conversion etc)
- The IAR ELF Dumper for ARM—`ielfdumparm`—creates a text representation of the contents of an ELF relocatable or executable image
- The IAR ELF Object Tool—`iobjmanip`—is used for performing low-level manipulation of ELF object files
- The IAR Absolute Symbol Exporter—`isymexport`—exports absolute symbols from a ROM image file, so that they can be used when linking an add-on application.

Note: These ELF utilities are well-suited for object files produced by the tools from IAR Systems. Thus, we recommend using them instead of the GNU binary utilities.

EXTERNAL TOOLS

For information about how to extend the tool chain in the IDE, see the *IDE Project Management and Building Guide for ARM*.

IAR language overview

The IAR C/C++ Compiler for ARM supports:

- C, the most widely used high-level programming language in the embedded systems industry. You can build freestanding applications that follow these standards:
 - Standard C—also known as C99. Hereafter, this standard is referred to as *Standard C* in this guide.
 - C89—also known as C94, C90, C89, and ANSI C. This standard is required when MISRA C is enabled.
- C++, a modern object-oriented programming language with a full-featured library well suited for modular programming. Any of these standards can be used:
 - Standard C++—can be used with different levels of support for exceptions and runtime type information (RTTI).
 - Embedded C++ (EC++)—a subset of the C++ programming standard, which is intended for embedded systems programming. It is defined by an industry consortium, the Embedded C++ Technical committee. See the chapter *Using C++*.
 - IAR Extended Embedded C++ (EEC++)—EC++ with additional features such as full template support, multiple inheritance, namespace support, the new cast operators, as well as the Standard Template Library (STL).

Each of the supported languages can be used in *strict* or *relaxed* mode, or relaxed with IAR extensions enabled. The strict mode adheres to the standard, whereas the relaxed mode allows some common deviations from the standard.

For more information about C, see the chapter *Using C*.

For more information about C++, Embedded C++, and Extended Embedded C++, see the chapter *Using C++*.

For information about how the compiler handles the implementation-defined areas of the languages, see the chapter *Implementation-defined behavior for Standard C*.

It is also possible to implement parts of the application, or the whole application, in assembler language. See the *IAR Assembler User Guide for ARM*.

Device support

To get a smooth start with your product development, the IAR product installation comes with a wide range of device-specific support.

SUPPORTED ARM DEVICES

The IAR C/C++ Compiler for ARM supports several different ARM cores and devices based on the instruction sets version 4, 5, 6, 6M, and 7. The object code that the compiler generates is not always binary compatible between the cores. Therefore it is crucial to specify a processor option to the compiler. The default core is ARM7TDMI

PRECONFIGURED SUPPORT FILES

The IAR product installation contains preconfigured files for supporting different devices. If you need additional files for device support, they can be created using one of the provided ones as a template.

Header files for I/O

Standard peripheral units are defined in device-specific I/O header files with the filename extension `h`. The product package supplies I/O files for all devices that are available at the time of the product release. You can find these files in the `arm\inc\<vendor>` directory. Make sure to include the appropriate include file in your application source files. If you need additional I/O header files, they can be created using one of the provided ones as a template. For detailed information about the header file format, see `EWARM_HeaderFormat.pdf` located in the `arm\doc` directory.

Linker configuration files

The `arm\config` directory contains ready-made linker configuration files for all supported devices. The files have the filename extension `icf` and contain the information required by the linker. For more information about the linker configuration file, see *Placing code and data—the linker configuration file*, page 87, and for reference information, the chapter *The linker configuration file*.

Device description files

The debugger handles several of the device-specific requirements, such as definitions of available memory areas, peripheral registers and groups of these, by using device description files. These files are located in the `arm\config` directory and they have the filename extension `ddf`. The peripheral registers and groups of these can be defined in separate files (filename extension `sfr`), which in that case are included in the `ddf` file. For more information about these files, see the *C-SPY® Debugging Guide for ARM* and `EWARM_DDFFORMAT.pdf` located in the `arm\doc` directory.

EXAMPLES FOR GETTING STARTED

The `arm\examples` directory contains several hundreds of examples of working applications to give you a smooth start with your development. The complexity of the

examples ranges from simple LED blink to USB mass storage controllers. Examples are provided for most of the supported devices.

Special support for embedded systems

This section briefly describes the extensions provided by the compiler to support specific features of the ARM core.

EXTENDED KEYWORDS

The compiler provides a set of keywords that can be used for configuring how the code is generated. For example, there are keywords for controlling how to access and store data objects, as well as for controlling how a function should work internally and how it should be called/returned.

By default, language extensions are enabled in the IDE.

The command line option `-e` makes the extended keywords available, and reserves them so that they cannot be used as variable names. See, `-e`, page 269 for additional information.

For more information about the extended keywords, see the chapter *Extended keywords*. See also, *Data storage*, page 67 and *Functions*, page 71.

PRAGMA DIRECTIVES

The pragma directives control the behavior of the compiler, for example how it allocates memory, whether it allows extended keywords, and whether it issues warning messages.

The pragma directives are always enabled in the compiler. They are consistent with standard C, and are very useful when you want to make sure that the source code is portable.

For more information about the pragma directives, see the chapter *Pragma directives*.

PREDEFINED SYMBOLS

With the predefined preprocessor symbols, you can inspect your compile-time environment, for example time of compilation or the build number of the compiler.

For more information about the predefined symbols, see the chapter *The preprocessor*.

ACCESSING LOW-LEVEL FEATURES

For hardware-related parts of your application, accessing low-level features is essential. The compiler supports several ways of doing this: intrinsic functions, mixing C and

assembler modules, and inline assembler. For information about the different methods, see *Mixing C and assembler*, page 157.

Developing embedded applications

- Developing embedded software using IAR build tools
- The build process—an overview
- Application execution—an overview
- Building applications—an overview
- Basic project configuration

Developing embedded software using IAR build tools

Typically, embedded software written for a dedicated microcontroller is designed as an endless loop waiting for some external events to happen. The software is located in ROM and executes on reset. You must consider several hardware and software factors when you write this kind of software. To your help, you have compiler options, extended keywords, pragma directives, etc.

MAPPING OF MEMORY

Embedded systems typically contain various types of memory, such as on-chip RAM, external DRAM or SRAM, ROM, EEPROM, or flash memory.

As an embedded software developer, you must understand the features of the different types of memory. For example, on-chip RAM is often faster than other types of memories, and variables that are accessed often would in time-critical applications benefit from being placed here. Conversely, some configuration data might be accessed seldom but must maintain their value after power off, so they should be saved in EEPROM or flash memory.

For efficient memory usage, the compiler provides several mechanisms for controlling placement of functions and data objects in memory. For more information, see *Controlling data and function placement in memory*, page 224. The linker places sections of code and data in memory according to the directives you specify in the linker configuration file, see *Placing code and data—the linker configuration file*, page 87.

COMMUNICATION WITH PERIPHERAL UNITS

If external devices are connected to the microcontroller, you might need to initialize and control the signaling interface, for example by using chip select pins, and detect and handle external interrupt signals. Typically, this must be initialized and controlled at runtime. The normal way to do this is to use special function registers (SFR). These are typically available at dedicated addresses, containing bits that control the chip configuration.

Standard peripheral units are defined in device-specific I/O header files with the filename extension `.h`. See *Device support*, page 49. For an example, see *Accessing special function registers*, page 237.

EVENT HANDLING

In embedded systems, using *interrupts* is a method for handling external events immediately; for example, detecting that a button was pressed. In general, when an interrupt occurs in the code, the core immediately stops executing the code it runs, and starts executing an interrupt routine instead.

The compiler provides various primitives for managing hardware and software interrupts, which means that you can write your interrupt routines in C, see *Interrupt functions for Cortex-M devices*, page 73 and *Interrupt functions for ARM7/9/11, Cortex-A, and Cortex-R devices*, page 74.

SYSTEM STARTUP

In all embedded systems, system startup code is executed to initialize the system—both the hardware and the software system—before the `main` function of the application is called. The CPU imposes this by starting execution from a fixed memory address.

As an embedded software developer, you must ensure that the startup code is located at the dedicated memory addresses, or can be accessed using a pointer from the vector table. This means that startup code and the initial vector table must be placed in non-volatile memory, such as ROM, EPROM, or flash.

A C/C++ application further needs to initialize all global variables. This initialization is handled by the linker and the system startup code in conjunction. For more information, see *Application execution—an overview*, page 58.

REAL-TIME OPERATING SYSTEMS

In many cases, the embedded application is the only software running in the system. However, using an RTOS has some advantages.

For example, the timing of high-priority tasks is not affected by other parts of the program which are executed in lower priority tasks. This typically makes a program

more deterministic and can reduce power consumption by using the CPU efficiently and putting the CPU in a lower-power state when idle.

Using an RTOS can make your program easier to read and maintain, and in many cases smaller as well. Application code can be cleanly separated in tasks which are truly independent of each other. This makes teamwork easier, as the development work can be easily split into separate tasks which are handled by one developer or a group of developers.

Finally, using an RTOS reduces the hardware dependence and creates a clean interface to the application, making it easier to port the program to different target hardware.

INTEROPERABILITY WITH OTHER BUILD TOOLS

The IAR compiler and linker provide support for AEABI, the Embedded Application Binary Interface for ARM. For more information about this interface specification, see the www.arm.com web site.

The advantage of this interface is the interoperability between vendors supporting it; an application can be built up of libraries of object files produced by different vendors and linked with a linker from any vendor, as long as they adhere to the AEABI standard.

AEABI specifies full compatibility for C and C++ object code, and for the C library. The AEABI does not include specifications for the C++ library.

For more information about the AEABI support in the IAR build tools, see *AEABI compliance*, page 217.

The ARM IAR build tools with version numbers from 6.xx and up are not fully compatible with earlier versions of the product, see the *IAR Embedded Workbench® Migration Guide for ARM®* for more information.

For more information, see *Linker optimizations*, page 217.

The build process—an overview

This section gives an overview of the build process; how the various build tools—compiler, assembler, and linker—fit together, going from source code to an executable image.

To get familiar with the process in practice, you should run one or more of the tutorials available from the IAR Information Center.


THE TRANSLATION PROCESS

There are two tools in the IDE that translate application source files to intermediary object files. The IAR C/C++ Compiler and the IAR Assembler. Both produce

relocatable object files in the industry-standard format ELF, including the DWARF format for debug information.

Note: The compiler can also be used for translating C/C++ source code into assembler source code. If required, you can modify the assembler source code which then can be assembled into object code. For more information about the IAR Assembler, see the *IAR Assembler User Guide for ARM*.

This illustration shows the translation process:



After the translation, you can choose to pack any number of modules into an archive, or in other words, a library. The important reason you should use libraries is that each module in a library is conditionally linked in the application, or in other words, is only included in the application if the module is used directly or indirectly by a module supplied as an object file. Optionally, you can create a library; then use the IAR utility `iarchive`.

THE LINKING PROCESS

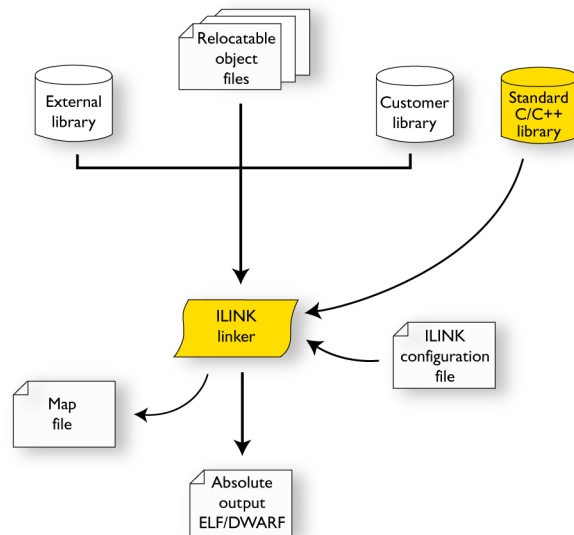
The relocatable modules, in object files and libraries, produced by the IAR compiler and assembler cannot be executed as is. To become an executable application, they must be *linked*.

Note: Modules produced by a toolset from another vendor can be included in the build as well. Be aware that this also might require a compiler utility library from the same vendor.

The IAR ILINK Linker (`ilinkarm.exe`) is used for building the final application. Normally, the linker requires the following information as input:

- Several object files and possibly certain libraries
- A program start label (set by default)
- The linker configuration file that describes placement of code and data in the memory of the target system

This illustration shows the linking process:



Note: The Standard C/C++ library contains support routines for the compiler, and the implementation of the C/C++ standard library functions.

During the linking, the linker might produce error messages and logging messages on `stdout` and `stderr`. The log messages are useful for understanding why an application was linked the way it was, for example, why a module was included or a section removed.

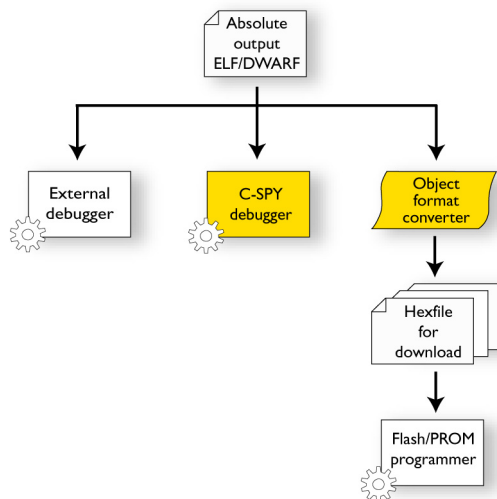
For more information about the procedure performed by the linker, see *The linking process in detail*, page 85.

AFTER LINKING

The IAR ILINK Linker produces an absolute object file in ELF format that contains the executable image. After linking, the produced absolute executable image can be used for:

- Loading into the IAR C-SPY Debugger or any other compatible external debugger that reads ELF and DWARF.
- Programming to a flash/PROM using a flash/PROM programmer. Before this is possible, the actual bytes in the image must be converted into the standard Motorola 32-bit S-record format or the Intel Hex-32 format. For this, use `ielftool`, see *The IAR ELF Tool—ielftool*, page 480.

This illustration shows the possible uses of the absolute output ELF/DWARF file:



Application execution—an overview

This section gives an overview of the execution of an embedded application divided into three phases, the:

- Initialization phase
- Execution phase
- Termination phase.

THE INITIALIZATION PHASE


Initialization is executed when an application is started (the CPU is reset) but before the `main` function is entered. The initialization phase can for simplicity be divided into:

- Hardware initialization, which generally at least initializes the stack pointer.
The hardware initialization is typically performed in the system startup code `cstartup.s` and if required, by an extra low-level routine that you provide. It might include resetting/starting the rest of the hardware, setting up the CPU, etc, in preparation for the software C/C++ system initialization.
- Software C/C++ system initialization
Typically, this includes assuring that every global (statically linked) C/C++ symbol receives its proper initialization value before the `main` function is called.
- Application initialization
This depends entirely on your application. It can include setting up an RTOS kernel and starting initial tasks for an RTOS-driven application. For a bare-bone application, it can include setting up various interrupts, initializing communication, initializing devices, etc.


For a ROM/flash-based system, constants and functions are already placed in ROM. All symbols placed in RAM must be initialized before the `main` function is called. The linker has already divided the available RAM into different areas for variables, stack, heap, etc.

The following sequence of illustrations gives a simplified overview of the different stages of the initialization.

- I When an application is started, the system startup code first performs hardware initialization, such as initialization of the stack pointer to point at the end of the predefined stack area:




- 2 Then, memories that should be zero-initialized are cleared, in other words, filled with zeros:

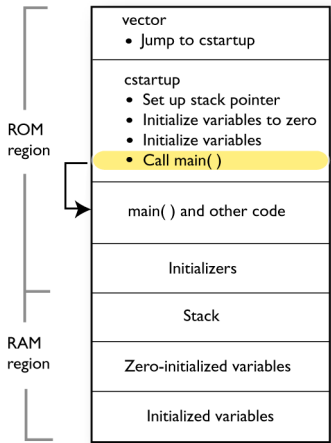


Typically, this is data referred to as *zero-initialized data*; variables declared as, for example, `int i = 0;`

- 3 For *initialized data*, data declared, for example, like `int i = 6;` the initializers are copied from ROM to RAM:



4 Finally, the `main` function is called:



For more information about each stage, see *System startup and termination*, page 133. For more information about initialization of data, see *Initialization at system startup*, page 90.

THE EXECUTION PHASE

The software of an embedded application is typically implemented as a loop which is either interrupt-driven or uses polling for controlling external interaction or internal events. For an interrupt-driven system, the interrupts are typically initialized at the beginning of the `main` function.

In a system with real-time behavior and where responsiveness is critical, a multi-task system might be required. This means that your application software should be complemented with a real-time operating system. In this case, the RTOS and the different tasks must also be initialized at the beginning of the `main` function.

THE TERMINATION PHASE

Typically, the execution of an embedded application should never end. If it does, you must define a proper end behavior.

To terminate an application in a controlled way, either call one of the Standard C library functions `exit`, `_Exit`, or `abort`, or return from `main`. If you return from `main`, the `exit` function is executed, which means that C++ destructors for static and global variables are called (C++ only) and all open files are closed.

Of course, in case of incorrect program logic, the application might terminate in an uncontrolled and abnormal way—a system crash.

For more information about this, see *System termination*, page 136.

Building applications—an overview

In the command line interface, this line compiles the source file `myfile.c` into the object file `myfile.o` using the default settings:

```
iccarm myfile.c
```

You must also specify some critical options, see *Basic project configuration*, page 63.

On the command line, this line can be used for starting the linker:

```
ilinkarm myfile.o myfile2.o -o a.out --config my_configfile.icf
```

In this example, `myfile.o` and `myfile2.o` are object files, and `my_configfile.icf` is the linker configuration file. The option `-o` specifies the name of the output file.

Note: By default, the label where the application starts is `__iar_program_start`. You can use the `--entry` command line option to change this.



When building a project, the IAR Embedded Workbench IDE can produce extensive build information in the **Build** messages window. This information can be useful, for example, as a base for producing batch files for building on the command line. You can copy the information and paste it in a text file. To activate extensive build information, choose **Tools>Options> Messages** and select the option **Show build messages: All**.

Basic project configuration

This section gives an overview of the basic settings for the project setup that are needed to make the compiler and linker generate the best code for the ARM device you are using. You can specify the options either from the command line interface or in the IDE.

You need to make settings for:

- Processor configuration, that is processor variant, CPU mode, interworking, VFP and floating-point arithmetic, and byte order
- Optimization settings
- Runtime environment, see *Setting up the runtime environment*, page 119
- Customizing the ILINK configuration, see the chapter *Linking your application*.

In addition to these settings, many other options and settings can fine-tune the result even further. For information about how to set options and for a list of all available

options, see the chapters *Compiler options*, *Linker options*, and the *IDE Project Management and Building Guide for ARM*, respectively.

PROCESSOR CONFIGURATION

To make the compiler generate optimum code, you should configure it for the ARM core you are using.

Processor variant

The IAR C/C++ Compiler for ARM supports several different ARM cores and devices based on the instruction sets version 4, 5, 6, and 7. All supported cores support Thumb instructions and 64-bit multiply instructions. The object code that the compiler generates is not always binary compatible between the cores. Therefore it is crucial to specify a processor option to the compiler. The default core is ARM7TDMI.



See the *IDE Project Management and Building Guide for ARM*, for information about setting the **Processor variant** option in the IDE.



Use the `--cpu` option to specify the ARM core; see `--arm`, page 261 and `--thumb`, page 294, for syntax information.

Interworking

When code is compiled with the `--interwork` option, ARM and Thumb code can be freely mixed. Interworking functions can be called from both ARM and Thumb code. Interworking is default for devices based on the instruction sets version 5, 6, and 7, or when using the `--aeabi` compiler option. See also *Veneers*, page 111.



See the *IDE Project Management and Building Guide for ARM* for information about setting the **Generate interwork code** option in the IDE.



Use the `--interwork` option to specify interworking capabilities for your project; see `--interwork`, page 275, for syntax information.

VFP and floating-point arithmetic

If you are using an ARM core that contains a Vector Floating Point (VFP) coprocessor, you can use the `--fpu` option to generate code that carries out floating-point operations utilizing the coprocessor, instead of using the software floating-point library routines.



See the *IDE Project Management and Building Guide for ARM*, for information about setting the **FPU** option in the IDE.



Use the `--fpu` option to specify the ARM core; see `--fpu`, page 273 for syntax information.

Byte order

The compiler supports the big-endian and little-endian byte order. All user and library modules in your application must use the same byte order.



See the *IDE Project Management and Building Guide for ARM* for information about setting the **Endian mode** option in the IDE.



Use the `--endian` option to specify the byte order for your project; see `--endian`, page 271, for syntax information.

OPTIMIZATION FOR SPEED AND SIZE

The compiler's optimizer performs, among other things, dead-code elimination, constant propagation, inlining, common sub-expression elimination, static clustering, instruction scheduling, and precision reduction. It also performs loop optimizations, such as unrolling and induction variable elimination.

You can choose between several optimization levels, and for the highest level you can choose between different optimization goals—*size*, *speed*, or *balanced*. Most optimizations will make the application both smaller and faster. However, when this is not the case, the compiler uses the selected optimization goal to decide how to perform the optimization.

The optimization level and goal can be specified for the entire application, for individual files, and for individual functions. In addition, some individual optimizations, such as function inlining, can be disabled.

For information about compiler optimizations and for more information about efficient coding techniques, see the chapter *Efficient coding for embedded applications*.

Data storage

- Introduction
- Storage of auto variables and parameters
- Dynamic memory on the heap

Introduction

An ARM core can address 4 Gbytes of continuous memory, ranging from 0x00000000 to 0xFFFFFFFF. Different types of physical memory can be placed in the memory range. A typical application will have both read-only memory (ROM) and read/write memory (RAM). In addition, some parts of the memory range contain processor control registers and peripheral units.

DIFFERENT WAYS TO STORE DATA

In a typical application, data can be stored in memory in three different ways:

- Auto variables

All variables that are local to a function, except those declared static, are stored either in registers or on the stack. These variables can be used as long as the function executes. When the function returns to its caller, the memory space is no longer valid. For more information, see *Storage of auto variables and parameters*, page 68.
- Global variables, module-static variables, and local variables declared `static`

In this case, the memory is allocated once and for all. The word static in this context means that the amount of memory allocated for this kind of variables does not change while the application is running. The ARM core has one single address space and the compiler supports full memory addressing.
- Dynamically allocated data.

An application can allocate data on the *heap*, where the data remains valid until it is explicitly released back to the system by the application. This type of memory is useful when the number of objects is not known until the application executes. Note that there are potential risks connected with using dynamically allocated data in systems with a limited amount of memory, or systems that are expected to run for a long time. For more information, see *Dynamic memory on the heap*, page 69.

Storage of auto variables and parameters

Variables that are defined inside a function—and not declared static—are named auto variables by the C standard. A few of these variables are placed in processor registers; the rest are placed on the stack. From a semantic point of view, this is equivalent. The main differences are that accessing registers is faster, and that less memory is required compared to when variables are located on the stack.

Auto variables can only live as long as the function executes; when the function returns, the memory allocated on the stack is released.

THE STACK

The stack can contain:

- Local variables and parameters not stored in registers
- Temporary results of expressions
- The return value of a function (unless it is passed in registers)
- Processor state during interrupts
- Processor registers that should be restored before the function returns (callee-save registers).

The stack is a fixed block of memory, divided into two parts. The first part contains allocated memory used by the function that called the current function, and the function that called it, etc. The second part contains free memory that can be allocated. The borderline between the two areas is called the top of stack and is represented by the stack pointer, which is a dedicated processor register. Memory is allocated on the stack by moving the stack pointer.

A function should never refer to the memory in the area of the stack that contains free memory. The reason is that if an interrupt occurs, the called interrupt function can allocate, modify, and—of course—deallocate memory on the stack.

See also *Stack considerations*, page 204 and *Setting up stack memory*, page 105.

Advantages

The main advantage of the stack is that functions in different parts of the program can use the same memory space to store their data. Unlike a heap, a stack will never become fragmented or suffer from memory leaks.

It is possible for a function to call itself either directly or indirectly—a recursive function—and each invocation can store its own data on the stack.

Potential problems

The way the stack works makes it impossible to store data that is supposed to live after the function returns. The following function demonstrates a common programming mistake. It returns a pointer to the variable `x`, a variable that ceases to exist when the function returns.

```
int *MyFunction()
{
    int x;
    /* Do something here. */
    return &x; /* Incorrect */
}
```

Another problem is the risk of running out of stack. This will happen when one function calls another, which in turn calls a third, etc., and the sum of the stack usage of each function is larger than the size of the stack. The risk is higher if large data objects are stored on the stack, or when recursive functions are used.

Dynamic memory on the heap

Memory for objects allocated on the heap will live until the objects are explicitly released. This type of memory storage is very useful for applications where the amount of data is not known until runtime.

In C, memory is allocated using the standard library function `malloc`, or one of the related functions `calloc` and `realloc`. The memory is released again using `free`.

In C++, a special keyword, `new`, allocates memory and runs constructors. Memory allocated with `new` must be released using the keyword `delete`.

For information about how to set up the size for heap memory, see *Setting up heap memory*, page 106.

POTENTIAL PROBLEMS

Applications that are using heap-allocated objects must be designed very carefully, because it is easy to end up in a situation where it is not possible to allocate objects on the heap.

The heap can become exhausted if your application uses too much memory. It can also become full if memory that no longer is in use was not released.

For each allocated memory block, a few bytes of data for administrative purposes is required. For applications that allocate a large number of small blocks, this administrative overhead can be substantial.

There is also the matter of fragmentation; this means a heap where small sections of free memory is separated by memory used by allocated objects. It is not possible to allocate a new object if no piece of free memory is large enough for the object, even though the sum of the sizes of the free memory exceeds the size of the object.

Unfortunately, fragmentation tends to increase as memory is allocated and released. For this reason, applications that are designed to run for a long time should try to avoid using memory allocated on the heap.

Functions

- Function-related extensions
- ARM and Thumb code
- Execution in RAM
- Interrupt functions for Cortex-M devices
- Interrupt functions for ARM7/9/11, Cortex-A, and Cortex-R devices
- Inlining functions

Function-related extensions

In addition to supporting Standard C, the compiler provides several extensions for writing functions in C. Using these, you can:

- Generate code for the different CPU modes ARM and Thumb Typically, this refers to target-specific keywords, such as `__arm`, `__thumb`.
- Execute functions in RAM
- Write interrupt functions for the different devices
- Control function inlining
- Facilitate function optimization
- Access hardware features.

The compiler uses compiler options, extended keywords, pragma directives, and intrinsic functions to support this.

For more information about optimizations, see *Efficient coding for embedded applications*, page 221. For information about the available intrinsic functions for accessing hardware operations, see the chapter *Intrinsic functions*.

ARM and Thumb code

The IAR C/C++ Compiler for ARM can generate code for either the 32-bit ARM, or the 16-bit Thumb or Thumb2 instruction set. Use the `--cpu_mode` option, alternatively the `--arm` or `--thumb` options, to specify which instruction set should be used for your project. For individual functions, it is possible to override the project setting by using

the extended keywords `__arm` and `__thumb`. You can freely mix ARM and thumb code in the same application, as long as the code is interworking.

When performing function calls, the compiler always attempts to generate the most efficient assembler language instruction or instruction sequence available. As a result, 4 Gbytes of continuous memory in the range `0x0-0xFFFFFFFF` can be used for placing code. There is a limit of 4 Mbytes per code module.

The size of all code pointers is 4 bytes. There are restrictions to implicit and explicit casts from code pointers to data pointers or integer types or vice versa. For further information about the restrictions, see *Pointer types*, page 337.

In the chapter *Assembler language interface*, the generated code is studied in more detail in the description of calling C functions from assembler language and vice versa.

Execution in RAM

The `__ramfunc` keyword makes a function execute in RAM. In other words it places the function in a section that has read/write attributes. The function is copied from ROM to RAM at system startup just like any initialized variable, see *System startup and termination*, page 133.

The keyword is specified before the return type:

```
__ramfunc void foo(void);
```

If a function declared `__ramfunc` tries to access ROM, the compiler will issue a warning.

If the whole memory area used for code and constants is disabled—for example, when the whole flash memory is being erased—only functions and data stored in RAM may be used. Interrupts must be disabled unless the interrupt vector and the interrupt service routines are also stored in RAM.

String literals and other constants can be avoided by using initialized variables. For example, the following lines:

```
__ramfunc void test()
{
    /* myc: initializer in ROM */
    const int myc[] = { 10, 20 };

    /* string literal in ROM */
    msg("Hello");
}
```


can be rewritten to:

```
__ramfunc void test()
{
    /* myc: initialized by cstartup */
    static int myc[] = { 10, 20 };

    /* hello: initialized by cstartup */
    static char hello[] = "Hello";

    msg(hello);
}
```

For more information, see *Initializing code—copying ROM to RAM*, page 109.

Interrupt functions for Cortex-M devices

Cortex-M has a different interrupt mechanism than previous ARM architectures, which means the primitives provided by the compiler are also different.

INTERRUPTS FOR CORTEX-M

On Cortex-M, an interrupt service routine enters and returns in the same way as a normal function, which means no special keywords are required. Thus, the keywords `__irq`, `__fiq`, and `__nested` are not available when you compile for Cortex-M.

These exception function names are defined in `cstartup_M.c` and `cstartup_M.s`. They are referred to by the library exception vector code:

```
NMI_Handler
HardFault_Handler
MemManage_Handler
BusFault_Handler
UsageFault_Handler
SVC_Handler
DebugMon_Handler
PendSV_Handler
SysTick_Handler
```

The vector table is implemented as an array. It should always have the name `__vector_table`, because the C-SPY debugger looks for that symbol when determining where the vector table is located.

The predefined exception functions are defined as weak symbols. A weak symbol is only included by the linker as long as no duplicate symbol is found. If another symbol is defined with the same name, it will take precedence. Your application can therefore simply define its own exception function by just defining it using the correct name from

the list above. If you need other interrupts or other exception handlers, you must make a copy of the `cstartup_M.c` or `cstartup_M.s` file and make the proper addition to the vector table.

The intrinsic functions `__get_CPSR` and `__set_CPSR` are not available when you compile for Cortex-M. Instead, if you need to get or set values of these or other registers, you can use inline assembler. For more information, see *Passing values between C and assembler objects*, page 239.

Interrupt functions for ARM7/9/11, Cortex-A, and Cortex-R devices

The IAR C/C++ Compiler for ARM provides the following primitives related to writing interrupt functions for ARM7/9/11, Cortex-A, and Cortex-R devices:

- The extended keywords: `__irq`, `__fiq`, `__swi`, `__nested`,
- The intrinsic functions: `__enable_interrupt`, `__disable_interrupt`, `__get_interrupt_state`, `__set_interrupt_state`.

Note: Cortex-M has a different interrupt mechanism than other ARM devices, and for these devices a different set of primitives is available. For more information, see *Interrupt functions for Cortex-M devices*, page 73.

INTERRUPT FUNCTIONS

In embedded systems, using interrupts is a method for handling external events immediately; for example, detecting that a button was pressed.

Interrupt service routines

In general, when an interrupt occurs in the code, the core immediately stops executing the code it runs, and starts executing an interrupt routine instead. It is important that the environment of the interrupted function is restored after the interrupt is handled (this includes the values of processor registers and the processor status register). This makes it possible to continue the execution of the original code after the code that handled the interrupt was executed.

The compiler supports interrupts, software interrupts, and fast interrupts. For each interrupt type, an interrupt routine can be written.

All interrupt functions must be compiled in ARM mode; if you are using Thumb mode, use the `__arm` extended keyword or the `#pragma type_attribute=__arm` directive to override the default behavior. This is not applicable for Cortex-M devices.

Interrupt vectors and the interrupt vector table

Each interrupt routine is associated with a vector address/instruction in the exception vector table, which is specified in the ARM cores documentation. The interrupt vector is the address in the exception vector table. For the ARM cores, the exception vector table starts at address 0x0.

By default, the vector table is populated with a *default interrupt handler* which loops indefinitely. For each interrupt source that has no explicit interrupt service routine, the default interrupt handler will be called. If you write your own service routine for a specific vector, that routine will override the default interrupt handler.

Defining an interrupt function—an example

To define an interrupt function, the `__irq` or the `__fiq` keyword can be used. For example:

```
__irq __arm void IRQ_Handler(void)
{
    /* Do something */
}
```

See the ARM cores documentation for more information about the interrupt vector table.

Note: An interrupt function must have the return type `void`, and it cannot specify any parameters.

Interrupt and C++ member functions

Only `static` member functions can be interrupt functions. When a non-static member function is called, it must be applied to an object. When an interrupt occurs and the interrupt function is called, there is no object available to apply the member function to.

INSTALLING EXCEPTION FUNCTIONS

All interrupt functions and software interrupt handlers must be installed in the vector table. This is done in assembler language in the system startup file `cstartup.s`.

The default implementation of the ARM exception vector table in the standard runtime library jumps to predefined functions that implement an infinite loop. Any exception that occurs for an event not handled by your application will therefore be caught in the infinite loop (B.).

The predefined functions are defined as weak symbols. A weak symbol is only included by the linker as long as no duplicate symbol is found. If another symbol is defined with the same name, it will take precedence. Your application can therefore simply define its own exception function by just defining it using the correct name.

These exception function names are defined in `cstartup.s` and referred to by the library exception vector code:

```
Undefined_Handler
SWI_Handler
Prefetch_Handler
Abort_Handler
IRQ_Handler
FIQ_Handler
```

To implement your own exception handler, define a function using the appropriate exception function name from the list above.

For example to add an interrupt function in C, it is sufficient to define an interrupt function named `IRQ_Handler`:

```
__irq __arm void IRQ_Handler()
{
}
```

An interrupt function must have C linkage, read more in *Calling convention*, page 169.

If you use C++, an interrupt function could look, for example, like this:

```
extern "C"
{
    __irq __arm void IRQ_Handler(void);
}
__irq __arm void IRQ_Handler(void)
{
}
```

No other changes are needed.

INTERRUPTS AND FAST INTERRUPTS

The interrupt and fast interrupt functions are easy to handle as they do not accept parameters or have a return value. Use any of these keywords:

- To declare an interrupt function, use the `__irq` extended keyword or the `#pragma type_attribute=__irq` directive. For syntax information, see `__irq`, page 348 and `type_attribute`, page 378, respectively.
- To declare a fast interrupt function, use the `__fiq` extended keyword or the `#pragma type_attribute=__fiq` directive. For syntax information, see `__fiq`, page 348, and `type_attribute`, page 378, respectively.

Note: An interrupt function (`irq`) and a fast interrupt function (`fiq`) must have a return type of `void` and cannot have any parameters. A software interrupt function (`swi`) may

have parameters and return values. By default, only four registers, R0–R3, can be used for parameters and only the registers R0–R1 can be used for return values.

NESTED INTERRUPTS

Interrupts are automatically disabled by the ARM core prior to entering an interrupt handler. If an interrupt handler re-enables interrupts, calls functions, and another interrupt occurs, then the return address of the interrupted function—stored in LR—is overwritten when the second IRQ is taken. In addition, the contents of SPSR will be destroyed when the second interrupt occurs. The `__irq` keyword itself does not save and restore LR and SPSR. To make an interrupt handler perform the necessary steps needed when handling nested interrupts, the keyword `__nested` must be used in addition to `__irq`. The function prolog—function entrance sequence—that the compiler generates for nested interrupt handlers will switch from IRQ mode to system mode. Make sure that both the IRQ stack and system stack is set up. If you use the default `cstartup.s` file, both stacks are correctly set up.

Compiler-generated interrupt handlers that allow nested interrupts are supported for IRQ interrupts only. The FIQ interrupts are designed to be serviced quickly, which in most cases mean that the overhead of nested interrupts would be too high.

This example shows how to use nested interrupts with the ARM vectored interrupt controller (VIC):

```
__irq __nested __arm void interrupt_handler(void)
{
    void (*interrupt_task)();
    unsigned int vector;

    /* Get interrupt vector. */
    vector = VICVectAddr;

    interrupt_task = (void(*)()) vector;

    /* Allow other IRQ interrupts to be serviced. */
    __enable_interrupt();

    /* Execute the task associated with this interrupt. */

    (*interrupt_task)();
}
```

Note: The `__nested` keyword requires the processor mode to be in either User or System mode.

SOFTWARE INTERRUPTS

Software interrupt functions are slightly more complex than other interrupt functions, in the way that they need a software interrupt handler (a dispatcher), are invoked (called) from running application software, and that they accept arguments and have return values. The mechanisms for calling a software interrupt function and how the software interrupt handler dispatches the call to the actual software interrupt function is described here.

Calling a software interrupt function

To call a software interrupt function from your application source code, the assembler instruction `SVC #immed` is used, where `immed` is an integer value that is referred to as the software interrupt number—or `swi_number`—in this guide. The compiler provides an easy way to implicitly generate this instruction from C/C++ source code, by using the `__swi` keyword and the `#pragma swi_number` directive when declaring the function.

A `__swi` function can for example be declared like this:

```
#pragma swi_number=0x23
__swi int swi_function(int a, int b);
```

In this case, the assembler instruction `SVC 0x23` will be generated where the function is called.

Software interrupt functions follow the same calling convention regarding parameters and return values as an ordinary function, except for the stack usage, see *Calling convention*, page 169.

For more information, see `__swi`, page 354, and `swi_number`, page 378, respectively.

The software interrupt handler and functions

The interrupt handler, for example `SWI_Handler` works as a dispatcher for software interrupt functions. It is invoked from the interrupt vector and is responsible for retrieving the software interrupt number and then calling the proper software interrupt function. The `SWI_Handler` must be written in assembler as there is no way to retrieve the software interrupt number from C/C++ source code.

The software interrupt functions

The software interrupt functions can be written in C or C++. Use the `__swi` keyword in a function definition to make the compiler generate a return sequence suited for a specific software interrupt function. The `#pragma swi_number` directive is not needed in the interrupt function definition.

For more information, see `__swi`, page 354.

Setting up the software interrupt stack pointer

If software interrupts will be used in your application, then the software interrupt stack pointer (`SVC_STACK`) must be set up and some space must be allocated for the stack. The `SVC_STACK` pointer can be set up together with the other stacks in the `cstartup.s` file. As an example, see the set up of the interrupt stack pointer. Relevant space for the `SVC_STACK` pointer is set up in the linker configuration file, see *Setting up stack memory*, page 105.

INTERRUPT OPERATIONS

An interrupt function is called when an external event occurs. Normally it is called immediately while another function is executing. When the interrupt function has finished executing, it returns to the original function. It is imperative that the environment of the interrupted function is restored; this includes the value of processor registers and the processor status register.

When an interrupt occurs, the following actions are performed:

- The operating mode is changed corresponding to the particular exception
- The address of the instruction following the exception entry instruction is saved in `R14` of the new mode
- The old value of the `CPSR` is saved in the `SPSR` of the new mode
- Interrupt requests are disabled by setting bit 7 of the `CPSR` and, if the exception is a fast interrupt, further fast interrupts are disabled by setting bit 6 of the `CPSR`
- The `PC` is forced to begin executing at the relevant vector address.

For example, if an interrupt for vector `0x18` occurs, the processor will start to execute code at address `0x18`. The memory area that is used as start location for interrupts is called the interrupt vector table. The content of the interrupt vector is normally a branch instruction jumping to the interrupt routine.

Note: If the interrupt function enables interrupts, the special processor registers needed to return from the interrupt routine must be assumed to be destroyed. For this reason they must be stored by the interrupt routine to be restored before it returns. This is handled automatically if the `__nested` keyword is used.

Inlining functions

Function inlining means that a function, whose definition is known at compile time, is integrated into the body of its caller to eliminate the overhead of the function call. This optimization, which is performed at optimization level High, normally reduces execution time, but might increase the code size. The resulting code might become more

difficult to debug. Whether the inlining actually occurs is subject to the compiler's heuristics.

The compiler heuristically decides which functions to inline. Different heuristics are used when optimizing for speed, size, or when balancing between size and speed. Normally, code size does not increase when optimizing for size.

C VERSUS C++ SEMANTICS

In C++, all definitions of a specific inline function in separate translation units must be exactly the same. If the function is not inlined in one or more of the translation units, then one of the definitions from these translation units will be used as the function implementation.

In C, you must manually select one translation unit that includes the non-inlined version of an inline function. You do this by explicitly declaring the function as `extern` in that translation unit. If you declare the function as `extern` in more than one translation unit, the linker will issue a *multiple definition* error. In addition, in C, inline functions cannot refer to static variables or functions.

For example:

```
// In a header file.
static int sX;
inline void F(void)
{
    //static int sY; // Cannot refer to statics.
    //sX;           // Cannot refer to statics.
}

// In one source file.
// Declare this F as the non-inlined version to use.
extern inline void F();
```

FEATURES CONTROLLING FUNCTION INLINING

There are several mechanisms for controlling function inlining:

- The `inline` keyword advises the compiler that the function defined immediately after the directive should be inlined.

If you compile your function in C or C++ mode, the keyword will be interpreted according to its definition in Standard C or Standard C++, respectively.

The main difference in semantics is that in Standard C you cannot (in general) simply supply an inline definition in a header file. You must supply an external definition in one of the compilation units, by designating the inline definition as being external in that compilation unit.

- `#pragma inline` is similar to the `inline` keyword, but with the difference that the compiler always uses C++ inline semantics.

By using the `#pragma inline` directive you can also disable the compiler's heuristics to either force inlining or completely disable inlining. For more information, see *inline*, page 368.

- `--use_cplusplus_inline` forces the compiler to use C++ semantics when compiling a Standard C source code file.
- `--no_inline`, `#pragma optimize=no_inline`, and `#pragma inline=never` all disable function inlining. By default, function inlining is enabled at optimization level High.

The compiler can only inline a function if the definition is known. Normally, this is restricted to the current translation unit. However, when the `--mfc` compiler option for multi-file compilation is used, the compiler can inline definitions from all translation units in the multi-file compilation unit. For more information, see *Multi-file compilation units*, page 229.

For more information about the function inlining optimization, see *Function inlining*, page 232.

Linking using ILINK

- Linking—an overview
- Modules and sections
- The linking process in detail
- Placing code and data—the linker configuration file
- Initialization at system startup
- Stack usage analysis

Linking—an overview

The IAR ILINK Linker is a powerful, flexible software tool for use in the development of embedded applications. It is equally well suited for linking small, single-file, absolute assembler programs as it is for linking large, relocatable, multi-module, C/C++, or mixed C/C++ and assembler programs.

The linker combines one or more relocatable object files—produced by the IAR Systems compiler or assembler—with selected parts of one or more object libraries to produce an executable image in the industry-standard format *Executable and Linking Format* (ELF).

The linker will automatically load only those library modules—user libraries and Standard C or C++ library variants—that are actually needed by the application you are linking. Further, the linker eliminates duplicate sections and sections that are not required.

ILINK can link both ARM and Thumb code, as well as a combination of them. By automatically inserting additional instructions (veneers), ILINK will assure that the destination will be reached for any calls and branches, and that the processor state is switched when required. For more details about how to generate veneers, see *Veneers*, page 111.

The linker uses a *configuration file* where you can specify separate locations for code and data areas of your target system memory map. This file also supports automatic handling of the application's initialization phase, which means initializing global variable areas and code areas by copying initializers and possibly decompressing them as well.

The final output produced by ILINK is an absolute object file containing the executable image in the ELF (including DWARF for debug information) format. The file can be downloaded to C-SPY or any other compatible debugger that supports ELF/DWARF, or it can be stored in EPROM or flash.

To handle ELF files, various tools are included. For information about included utilities, see *Specific ELF tools*, page 48.

Modules and sections

Each relocatable object file contains one module, which consists of:

- Several sections of code or data
- Runtime attributes specifying various types of information, for example the version of the runtime environment
- Optionally, debug information in DWARF format
- A symbol table of all global symbols and all external symbols used.

A *section* is a logical entity containing a piece of data or code that should be placed at a physical location in memory. A section can consist of several *section fragments*, typically one for each variable or function (symbols). A section can be placed either in RAM or in ROM. In a normal embedded application, sections that are placed in RAM do not have any content, they only occupy space.

Each section has a name and a type attribute that determines the content. The type attribute is used (together with the name) for selecting sections for the ILINK configuration. The most commonly used attributes are:

code	Executable code
readonly	Constant variables
readwrite	Initialized variables
zeroinit	Zero-initialized variables

Note: In addition to these section types—sections that contain the code and data that are part of your application—a final object file will contain many other types of sections, for example sections that contain debugging information or other type of meta information.

A section is the smallest linkable unit; but if possible, ILINK can exclude smaller units—section fragments—from the final application. For more information, see *Keeping modules*, page 105, and *Keeping symbols and sections*, page 105.

At compile time, data and functions are placed in different sections. At link time, one of the most important functions of the linker is to assign addresses to the various sections used by the application.

The IAR build tools have many predefined section names. See the chapter *Section reference* for more information about each section.

You can group sections together for placement by using blocks. See *define block directive*, page 444.

The linking process in detail

The relocatable modules in object files and libraries, produced by the IAR compiler and assembler, cannot be executed as is. To become an executable application, they must be *linked*.


Note: Modules produced by a toolset from another vendor can be included in the build as well, as long as the module is AEABI (ARM Embedded Application Binary Interface) compliant. Be aware that this also might require a compiler utility library from the same vendor.

The linker is used for the link process. It normally performs the following procedure (note that some of the steps can be turned off by command line options or by directives in the linker configuration file):

- Determine which modules to include in the application. Modules provided in object files are always included. A module in a library file is only included if it provides a definition for a global symbol that is referenced from an included module.
- Select which standard library files to use. The selection is based on attributes of the included modules. These libraries are then used for satisfying any still outstanding undefined symbols.
- Handle symbols with more than one definition. If there is more than one non-weak definition, an error is emitted. Otherwise, one of the definitions is picked (the non-weak one, if there is one) and the others are suppressed. Weak definitions are typically used for inline and template functions. If you need to override some of the non-weak definitions from a library module, you must ensure that the library module is not included (typically by providing alternate definitions for all the symbols your application uses in that library module).
- Determine which sections/section fragments from the included modules to include in the application. Only those sections/section fragments that are actually needed by the application are included. There are several ways to determine of which sections/section fragments that are needed, for example, the `__root` object attribute, the `#pragma required` directive, and the `keep` linker directive. In case of duplicate sections, only one is included.

- Where appropriate, arrange for the initialization of initialized variables and code in RAM. The `initialize` directive causes the linker to create extra sections to enable copying from ROM to RAM. Each section that will be initialized by copying is divided into two sections, one for the ROM part and one for the RAM part. If manual initialization is not used, the linker also arranges for the startup code to perform the initialization.
- Determine where to place each section according to the section placement directives in the *linker configuration file*. Sections that are to be initialized by copying appear twice in the matching against placement directives, once for the ROM part and once for the RAM part, with different attributes. During the placement, the linker also adds any required veneers to make a code reference reach its destination or to switch CPU modes.
- Produce an absolute file that contains the executable image and any debug information provided. The contents of each needed section in the relocatable input files is calculated using the relocation information supplied in its file and the addresses determined when placing sections. This process can result in one or more relocation failures if some of the requirements for a particular section are not met, for instance if placement resulted in the destination address for a PC-relative jump instruction being out of range for that instruction.
- Optionally, produce a map file that lists the result of the section placement, the address of each global symbol, and finally, a summary of memory usage for each module and library.

This illustration shows the linking process:



During the linking, ILINK might produce error messages and logging messages on `stdout` and `stderr`. The log messages are useful for understanding why an application was linked as it was. For example, why a module or section (or section fragment) was included.

Note: To see the actual content of an ELF object file, use `iefldumparm`. See *The IAR ELF Dumper—iefldump*, page 482.

Placing code and data—the linker configuration file

The placement of sections in memory is performed by the IAR ILINK Linker. It uses the *linker configuration file* where you can define how ILINK should treat each section and how they should be placed into the available memories.

A typical linker configuration file contains definitions of:

- Available addressable memories
- Populated regions of those memories
- How to treat input sections
- Created sections

- How to place sections into the available regions.

The file consists of a sequence of declarative directives. This means that the linking process will be governed by all directives at the same time.

To use the same source code with different derivatives, just rebuild the code with the appropriate configuration file.

A SIMPLE EXAMPLE OF A CONFIGURATION FILE

Assume a simple 32-bit architecture that has these memory prerequisites:

- There are 4 Gbytes of addressable memory.
- There is ROM memory in the address range 0x0000–0x10000.
- There is RAM memory in the range 0x20000–0x30000.
- The stack has an alignment of 8.
- The system startup code must be located at a fixed address.

A simple configuration file for this assumed architecture can look like this:

```
/* The memory space denoting the maximum possible amount
   of addressable memory */
define memory Mem with size = 4G;

/* Memory regions in an address space */
define region ROM = Mem:[from 0x00000 size 0x10000];
define region RAM = Mem:[from 0x20000 size 0x10000];

/* Create a stack */
define block STACK with size = 0x1000, alignment = 8 { };

/* Handle initialization */
initialize by copy { readwrite }; /* Initialize RW sections */

/* Place startup code at a fixed address */
place at start of ROM { readonly section .cstartup };

/* Place code and data */
place in ROM { readonly }; /* Place constants and initializers in
                             ROM: .rodata and .data_init */
place in RAM { readwrite, /* Place .data, .bss, and .noinit */
               block STACK }; /* and STACK */
```

This configuration file defines one addressable memory `Mem` with the maximum of 4 Gbytes of memory. Further, it defines a ROM region and a RAM region in `Mem`, namely ROM and RAM. Each region has the size of 64 Kbytes.


The file then creates an empty block called `STACK` with a size of 4 Kbytes in which the application stack will reside. To create a *block* is the basic method which you can use to get detailed control of placement, size, etc. It can be used for grouping sections, but also as in this example, to specify the size and placement of an area of memory.

Next, the file defines how to handle the initialization of variables, read/write type (`readwrite`) sections. In this example, the initializers are placed in ROM and copied at startup of the application to the RAM area. By default, ILINK may compress the initializers if this appears to be advantageous.

The last part of the configuration file handles the actual placement of all the sections into the available regions. First, the startup code—defined to reside in the read-only (`readonly`) section `.cstartup`—is placed at the start of the ROM region, that is at address 0x10000. Note that the part within `{ }` is referred to as *section selection* and it selects the sections for which the directive should be applied to. Then the rest of the read-only sections are placed in the ROM region. Note that the section selection `{ readonly section .cstartup }` takes precedence over the more generic section selection `{ readonly }`.

Finally, the read/write (`readwrite`) sections and the `STACK` block are placed in the RAM region.

This illustration gives a schematic overview of how the application is placed in memory:



In addition to these standard directives, a configuration file can contain directives that define how to:

- Map a memory that can be addressed in multiple ways
- Handle conditional directives
- Create symbols with values that can be used in the application
- More in detail, select the sections a directive should be applied to
- More in detail, initialize code and data.

For more details and examples about customizing the linker configuration file, see the chapter *Linking your application*.

For more information about the linker configuration file, see the chapter *The linker configuration file*.

Initialization at system startup

In Standard C, all static variables—variables that are allocated at a fixed memory address—must be initialized by the runtime system to a known value at application startup. This value is either an explicit value assigned to the variable, or if no value is given, it is cleared to zero. In the compiler, there are exceptions to this rule, for example variables declared `__no_init`, which are not initialized at all.

The compiler generates a specific type of section for each type of variable initialization:

Categories of declared data	Source	Section type	Section name	Section content
Zero-initialized data	<code>int i;</code>	Read/write data, zero-init	<code>.bss</code>	None
Zero-initialized data	<code>int i = 0;</code>	Read/write data, zero-init	<code>.bss</code>	None
Initialized data (non-zero)	<code>int i = 6;</code>	Read/write data	<code>.data</code>	The initializer
Non-initialized data	<code>__no_init int i;</code>	Read/write data, zero-init	<code>.noinit</code>	None
Constants	<code>const int i = 6;</code>	Read-only data	<code>.rodata</code>	The constant
Code	<code>__ramfunc void myfunc() {}</code>	Read/write code	<code>.textrw</code>	The code

Table 3: Sections holding initialized data

Note: Clustering of static variables might group zero-initialized variables together with initialized data in `.data`. The compiler can decide to place constants in the `.text` section to avoid loading the address of a constant from a constant table.

For information about all supported sections, see the chapter *Section reference*.

THE INITIALIZATION PROCESS

Initialization of data is handled by ILINK and the system startup code in conjunction.

To configure the initialization of variables, you must consider these issues:

- Sections that should be zero-initialized, or not initialized at all (`__no_init`) are handled automatically by ILINK.
- Sections that should be initialized, except for zero-initialized sections, should be listed in an `initialize` directive.

Normally during linking, a section that should be initialized is split into two sections, where the original initialized section will keep the name. The contents are placed in the new initializer section, which will get the original name suffixed with `_init`. The initializers should be placed in ROM and the initialized sections in RAM, by means of placement directives. The most common example is the `.data` section which the linker splits into `.data` and `.data_init`.

- Sections that contains constants should not be initialized; they should only be placed in flash/ROM.

In the linker configuration file, it can look like this:

```
/* Handle initialization */
initialize by copy { readwrite }; /* Initialize RW sections */

/* Place startup code at a fixed address */
place at start of ROM { readonly section .cstartup };

/* Place code and data */
place in ROM { readonly }; /* Place constants and initializers in
                             ROM: .rodata and .data_init */
place in RAM { readwrite, /* Place .data, .bss, and .noinit */
               block STACK }; /* and STACK */
```

Note: When compressed initializers are used (see *initialize directive*, page 446), the contents sections (that is, sections with the `_init` suffix) are not listed as separate sections in the map file. Instead, they are combined into aggregates of “initializer bytes”. You can place the contents sections the usual way in the linker configuration file; however, this affects the placement (and possibly the number) of the “initializer bytes” aggregates.

For more information about and examples of how to configure the initialization, see *Linking considerations*, page 101.

C++ DYNAMIC INITIALIZATION

The compiler places subroutine pointers for performing C++ dynamic initialization into sections of the ELF section types `SHT_PREINIT_ARRAY` and `SHT_INIT_ARRAY`. By default, the linker will place these into a linker-created block, ensuring that all sections of the section type `SHT_PREINIT_ARRAY` are placed before those of the type `SHT_INIT_ARRAY`. If any such sections were included, code to call the routines will also be included.

The linker-created blocks are only generated if the linker configuration does not contain section selector patterns for the `preinit_array` and `init_array` section types. The effect of the linker-created blocks will be very similar to what happens if the linker configuration file contains this:

```
define block SHT$$PREINIT_ARRAY { preinit_array };
define block SHT$$INIT_ARRAY { init_array };
define block CPP_INIT with fixed order { block
                                SHT$$PREINIT_ARRAY,
                                block SHT$$INIT_ARRAY };
```

If you put this into your linker configuration file, you must also mention the `CPP_INIT` block in one of the section placement directives. If you wish to select where the linker-created block is placed, you can use a section selector with the name `".init_array"`.

See also *section-selectors*, page 452.

Stack usage analysis

This section describes how to perform a stack usage analysis using the linker.

In the `ARM\src` directory, you can find an example project that demonstrates stack usage analysis.

INTRODUCTION TO STACK USAGE ANALYSIS

Under the right circumstances, the linker can accurately calculate the maximum stack usage for each call graph, starting from the program start, interrupt functions, tasks etc. (each function that is not called from another function, in other words, the root).

If you enable stack usage analysis, a stack usage chapter will be added to the linker map file, listing for each call graph root the particular call chain which results in the maximum stack depth.

The analysis is only accurate if there is accurate stack usage information for each function in the application.

In general, the compiler will generate this information for each C function, but if there are indirect calls (calls using function pointers) in your application, you must supply a list of possible functions that can be called from each calling function.

If you use a stack usage control file, you can also supply stack usage information for functions in modules that do not have stack usage information.

You can use the `check` directive in your linker configuration file to check that the stack usage calculated by the linker does not exceed the stack space you have allocated.

PERFORMING A STACK USAGE ANALYSIS

- 1 Enable stack usage analysis:



In the IDE, choose **Project>Options>Linker>Advanced>Enable stack usage analysis**



On the command line, use the linker option `--enable_stack_usage`

See `--enable_stack_usage`, page 309.

- 2 Enable the linker map file:



In the IDE, choose **Project>Options>Linker>List>Generate linker map file**



On the command line, use the linker option `--map`

- 3 Link your project. Note that the linker will issue warnings related to stack usage under certain circumstances, see *Situations where warnings are issued*, page 97.
- 4 Review the linker map file, which now contains a stack usage chapter with a summary of the stack usage for each call graph root. For more information, see *Result of an analysis—the map file contents*, page 94.
- 5 For more details, analyze the call graph log, see *Call graph log*, page 98.

Note that there are limitations and sources of inaccuracy in the analysis, see *Limitations*, page 97.

You might need to specify more information to the linker to get a more representative result. See *Specifying additional stack usage information*, page 95



In the IDE, choose **Project>Options>Linker>Advanced>Enable stack usage analysis>Control file**



On the command line, use the linker option `--stack_usage_control`
See `--stack_usage_control`, page 325.

- 6 To add an automatic check that you have allocated memory enough for the stack, use the `check that` directive in your linker configuration file. For example, assuming a stack block named `MY_STACK`, you can write like this:

```
check that size(block MY_STACK) >=maxstack("Program entry")
+ totalstack("interrupt") + 100;
```

When linking, the linker emits an error if the expression is false (zero). In this example, an error will be emitted if the sum of the following exceeds the size of the `MY_STACK` block:

- The maximum stack usage in the category `Program entry` (the main program).
- The sum of each individual maximum stack usage in the category `interrupt` (assuming that all interrupt routines need space at the same time).
- A safety margin of 100 bytes (to account for stack usage not visible to the analysis).

See `check that directive`, page 457.

RESULT OF AN ANALYSIS—THE MAP FILE CONTENTS

When stack usage analysis is enabled, the linker map file contains a stack usage chapter with a summary of the stack usage for each call graph root category, and lists the call chain that results in the maximum stack depth for each call graph root. This is an example of what the stack usage chapter in the map file might look like:

```
*****
***  STACK USAGE
***

Call Graph Root Category  Max Use  Total Use
-----
interrupt                 104       136
Program entry             168       168

Program entry
  "__iar_program_start": 0x000085ac
  Maximum call chain                                168 bytes
```

```

    "__iar_program_start"          0
    "__cmain"                      0
    "main"                          8
    "printf"                       24
    "_PrintfTiny"                  56
    "_Prout"                       16
    "putchar"                      16
    "__write"                      0
    "__dwrite"                     0
    "__iar_sh_stdout"              24
    "__iar_get_ttio"               24
    "__iar_lookup_ttioh"           0

interrupt
    "FaultHandler": 0x00008434

    Maximum call chain              32 bytes

    "FaultHandler"                  32

interrupt
    "IRQHandler": 0x00008424

    Maximum call chain              104 bytes

    "IRQHandler"                   24
    "do_something" in suexample.o [1] 80

```

The summary contains the depth of the deepest call chain in each category as well as the sum of the depths of the deepest call chains in that category.

Each call graph root belongs to a call graph root category to enable convenient calculations in `check that` directives.

SPECIFYING ADDITIONAL STACK USAGE INFORMATION

To specify additional stack usage information you can use either a stack usage control file (`suc`) where you specify stack usage control directives or annotate the source code.

You can:

- Specify complete stack usage information (call graph root category, stack usage, and possible calls) for a function, by using the stack usage control directive

function. Typically, you do this if stack usage information is missing, for example in an assembler module. In your `suc` file you can for example write like this:

```
function MyFunc: 32,  
    calls MyFunc2,  
    calls MyFunc3, MyFunc4: 16;
```

```
function [interrupt] MyInterruptHandler: 44;
```

See also *function directive*, page 470.

- Exclude certain functions from stack usage analysis, by using the stack usage control directive `exclude`. In your `suc` file you can for example write like this:

```
exclude MyFunc5, MyFunc6;
```

See also *exclude directive*, page 470.

- Specify a list of possible destinations for indirect calls in a function, by using the stack usage control directive `possible calls`. Use this for functions which are known to perform indirect calls and where you know exactly which functions that might be called in this particular application. In your `suc` file you can for example write like this:

```
possible calls MyFunc7: MyFunc8, MyFunc9;
```

If the information about which functions that might be called is available at compile time, consider using the `#pragma calls` directive instead.

See also *possible calls directive*, page 472 and *calls*, page 362.

- Specify that functions are call graph roots, including an optional call graph root category, by using the stack usage control directive `call graph root` or the `#pragma call_graph_root` directive. In your `suc` file you can for example write like this:

```
call graph root [task]: MyFunc10, MyFunc11;
```

If your interrupt functions have not already been designated as call graph roots by the compiler, you must do so manually. You can do this either by using the `#pragma call_graph_root` directive in your source code or by specifying a directive in your `suc` file, for example:

```
call graph root [interrupt]: Irq1Handler, Irq2Handler;
```

See also *call graph root directive*, page 470 and *call_graph_root*, page 363.

- Specify a maximum number of iterations through any of the cycles in the recursion nest of which the function is a member. In your `suc` file you can for example write like this:

```
max recursion depth MyFunc12: 10;
```

- Selectively suppress the warning about unmentioned functions referenced by a module for which you have supplied stack usage information in the stack usage

control file. Use the `no calls from` directive in your `suc` file, for example like this:

```
no calls from [file.o] to MyFunc13, MyFunc14;
```

For more information, see the chapter *Specifying additional stack usage information*, page 95.

LIMITATIONS

Apart from missing or incorrect stack usage information, there are also other sources of inaccuracy in the analysis:

- The linker cannot always identify all functions in object modules that lack stack usage information. In particular, this might be a problem with object modules written in assembly language or produced by non-IAR tools. You can provide stack usage information for such modules using a stack usage control file, and for assembly language modules you can also annotate the assembler source code with `CFI` directives to provide stack usage information. See the *IAR Assembler User Guide for ARM*.
- If you use inline assembler to change the frame size or to perform function calls, this will not be reflected in the analysis.
- Extra space consumed by other sources (the processor, an operating system, etc) is not accounted for.
- Source code that uses exceptions is not supported.
- If you use other forms of function calls, like software interrupts, they will not be reflected in the call graph.
- Using multi-file compilation (`--mfc`) can interfere with using a stack usage control file to specify properties of module-local functions in the involved files.

Note that stack usage analysis produces a worst case result. The program might not actually ever end up in the maximum call chain, by design, or by coincidence. In particular, the set of possible destinations for a virtual function call in C++ might sometimes include implementations of the function in question which cannot, in fact, be called from that point in the code.



Stack usage analysis is only a complement to actual measurement. If the result is important, you need to perform independent validation of the results of the analysis.

SITUATIONS WHERE WARNINGS ARE ISSUED

When stack usage analysis is enabled in the linker, warnings will be generated in the following circumstances:

- There is a function without stack usage information.

- There is an indirect call site in the application for which a list of possible called functions has not been supplied.
- There are no known indirect calls, but there is an uncalled function that is not known to be a call graph root.
- The application contains recursion (a cycle in the call graph) for which no maximum recursion depth has been supplied, or which is of a form for which the linker is unable to calculate a reliable estimate of stack usage.
- There are calls to a function declared as a call graph root.
- You have used the stack usage control file to supply stack usage information for functions in a module that does not have such information, and there are functions referenced by that module which have not been mentioned as being called in the stack usage control file.

CALL GRAPH LOG

To help you interpret the results of the stack usage analysis, there is a log output option that produces a simple text representation of the call graph (`--log call_graph`).

Example output:

```

Program entry:
0 __iar_program_start [168]
0 __cmain [168]
0 __iar_data_init3 [16]
  8 __iar_zero_init3 [8]
    16 - [0]
  8 __iar_copy_init3 [8]
    16 - [0]
0 __low_level_init [0]
0 main [168]
  8 printf [160]
    32 _PrintfTiny [136]
      88 _Prout [80]
        104 putchar [64]
          120 __write [48]
            120 __dwrite [48]
              120 __iar_sh_stdout [48]
                144 __iar_get_ttio [24]
                  168 __iar_lookup_ttioh [0]
                    120 __iar_sh_write [24]
                      144 - [0]
            88 __aeabi_uidiv [0]
              88 __aeabi_idiv0 [0]
                88 strlen [0]
          0 exit [8]
            0 _exit [8]
              0 __exit [8]
                0 __iar_close_ttio [8]
                  8 __iar_lookup_ttioh [0] ***
          0 __exit [8] ***

```

Each line consists of this information:

- The stack usage at the point of call of the function
- The name of the function, or a single '-' to indicate usage in a function at a point with no function call (typically in a leaf function)
- The stack usage along the deepest call chain from that point. If no such value could be calculated, "[---]" is output instead. "***" marks functions that have already been shown.

CALL GRAPH XML OUTPUT

The linker can also produce a call graph file in XML format. This file contains one node for each function in your application, with the stack usage and call information relevant

to that function. It is intended to be input for post-processing tools and is not particularly human-readable.

For more information about the XML format used, see the `callGraph.txt` file in your product installation.

Linking your application

- Linking considerations
- Hints for troubleshooting
- Checking module consistency

Linking considerations

Before you can link your application, you must set up the configuration required by ILINK. Typically, you must consider:

- Choosing a linker configuration file
- Defining your own memory areas
- Placing sections
- Reserving space in RAM
- Keeping modules
- Keeping symbols and sections
- Application startup
- Setting up stack memory
- Setting up heap memory
- Setting up the atexit limit
- Changing the default initialization
- Interaction between ILINK and the application
- Standard library handling
- Producing other output formats than ELF/DWARF

CHOOSING A LINKER CONFIGURATION FILE

The `config` directory contains two ready-made templates for the linker configuration files:

- `generic.icf`, designed for all cores except for Cortex-M cores
- `generic_cortex.icf`, designed for all Cortex-M cores.

The files contain the information required by ILINK. The only change, if any, you will normally have to make to the supplied configuration file is to customize the start and end addresses of each region so they fit the target system memory map. If, for example, your

application uses additional external RAM, you must also add details about the external RAM memory area.

For some devices, device-specific configuration files are automatically selected.

To edit a linker configuration file, use the editor in the IDE, or any other suitable editor. Alternatively, choose **Project>Options>Linker** and click the **Edit** button on the **Config** page to open the dedicated linker configuration file editor.

Do not change the original template file. We recommend that you make a copy in the working directory, and modify the copy instead. If you are using the linker configuration file editor in the IDE, the IDE will make a copy for you.

Each project in the IDE should have a reference to one, and only one, linker configuration file. This file can be edited, but for the majority of all projects it is sufficient to configure the vital parameters in **Project>Options>Linker>Config**.

DEFINING YOUR OWN MEMORY AREAS

The default configuration file that you selected has predefined ROM and RAM regions. This example will be used as a starting-point for all further examples in this chapter:

```
/* Define the addressable memory */
define memory Mem with size = 4G;

/* Define a region named ROM with start address 0 and to be 64
Kbytes large */
define region ROM = Mem:[from 0 size 0x10000];

/* Define a region named RAM with start address 0x20000 and to be
64 Kbytes large */
define region RAM = Mem:[from 0x20000 size 0x10000];
```

Each region definition must be tailored for the actual hardware.

To find out how much of each memory that was filled with code and data after linking, inspect the memory summary in the map file (command line option `--map`).

Adding an additional region

To add an additional region, use the `define region` directive, for example:

```
/* Define a 2nd ROM region to start at address 0x80000 and to be
128 Kbytes large */
define region ROM2 = Mem:[from 0x80000 size 0x20000];
```

Merging different areas into one region

If the region is comprised of several areas, use a region expression to merge the different areas into one region, for example:

```
/* Define the 2nd ROM region to have two areas. The first with
the start address 0x80000 and 128 Kbytes large, and the 2nd with
the start address 0xC0000 and 32 Kbytes large */
define region ROM2 = Mem:[from 0x80000 size 0x20000]
                    | Mem:[from 0xC0000 size 0x08000];
```

or equivalently

```
define region ROM2 = Mem:[from 0x80000 to 0xC7FFF]
                    -Mem:[from 0xA0000 to 0xBFFFF];
```

PLACING SECTIONS

The default configuration file that you selected places all predefined sections in memory, but there are situations when you might want to modify this. For example, if you want to place the section that holds constant symbols in the `CONSTANT` region instead of in the default place. In this case, use the `place` in directive, for example:

```
/* Place sections with readonly content in the ROM region */
place in ROM {readonly};

/* Place the constant symbols in the CONSTANT region */
place in CONSTANT {readonly section .rodata};
```

Note: Placing a section—used by the IAR build tools—in a different memory which use a different way of referring to its content, will fail.

For the result of each placement directive after linking, inspect the placement summary in the map file (the command line option `--map`).

Placing a section at a specific address in memory

To place a section at a specific address in memory, use the `place at` directive, for example:

```
/* Place section .vectors at address 0 */
place at address Mem:0x0 {readonly section .vectors};
```

Placing a section first or last in a region

To place a section first or last in a region is similar, for example:

```
/* Place section .vectors at start of ROM */
place at start of ROM {readonly section .vectors};
```

Declare and place your own sections

To declare new sections—in addition to the ones used by the IAR build tools—to hold specific parts of your code or data, use mechanisms in the compiler and assembler. For example:

```
/* Place a variable in that section. */
const short MyVariable @ "MYOWNSECTION" = 0xF0F0;
```

This is the corresponding example in assembler language:

```
name      createSection
section MYOWNSECTION:CONST ; Create a section,
                           ; and fill it with
dc16      0xF0F0           ; constant bytes.
end
```

To place your new section, the original place in ROM {readonly}; directive is sufficient.

However, to place the section `MyOwnSection` explicitly, update the linker configuration file with a place in directive, for example:

```
/* Place MyOwnSection in the ROM region */
place in ROM {readonly section MyOwnSection};
```

RESERVING SPACE IN RAM

Often, an application must have an empty uninitialized memory area to be used for temporary storage, for example a heap or a stack. It is easiest to achieve this at link time. You must create a block with a specified size and then place it in a memory.

In the linker configuration file, it can look like this:

```
define block TempStorage with size = 0x1000, alignment = 4 { };
place in RAM { block TempStorage };
```

To retrieve the start of the allocated memory from the application, the source code could look like this:

```
/* Define a section for temporary storage. */
#pragma section = "TempStorage"
char *GetTempStorageStartAddress()
{
    /* Return start address of section TempStorage. */
    return __section_begin("TempStorage");
}
```


KEEPING MODULES

If a module is linked as an object file, it is always kept. That is, it will contribute to the linked application. However, if a module is part of a library, it is included only if it is symbolically referred to from other parts of the application. This is true, even if the library module contains a root symbol. To assure that such a library module is always included, use `iarchive` to extract the module from the library, see *The IAR Archive Tool—iarchive*, page 477.

For information about included and excluded modules, inspect the log file (the command line option `--log modules`).

For more information about modules, see *Modules and sections*, page 84.

KEEPING SYMBOLS AND SECTIONS

By default, ILINK removes any sections, section fragments, and global symbols that are not needed by the application. To retain a symbol that does not appear to be needed—or actually, the section fragment it is defined in—you can either use the root attribute on the symbol in your C/C++ or assembler source code, or use the ILINK option `--keep`. To retain sections based on attribute names or object names, use the directive `keep` in the linker configuration file.

To prevent ILINK from excluding sections and section fragments, use the command line options `--no_remove` or `--no_fragments`, respectively.

For information about included and excluded symbols and sections, inspect the log file (the command line option `--log sections`).

For more information about the linking procedure for keeping symbols and sections, see *The linking process*, page 56.

APPLICATION STARTUP

By default, the point where the application starts execution is defined by the `__iar_program_start` label, which is defined to point at the start of the `cstartup.s` file. The label is also communicated via ELF to any debugger that is used.

To change the start point of the application to another label, use the ILINK option `--entry`; see `--entry`, page 309.

SETTING UP STACK MEMORY

The size of the `CSTACK` block is defined in the linker configuration file. To change the allocated amount of memory, change the block definition for `CSTACK`:

```
define block CSTACK with size = 0x2000, alignment = 8{ };
```

Specify an appropriate size for your application.

For more information about the stack, see *Stack considerations*, page 204.

SETTING UP HEAP MEMORY

The size of the heap is defined in the linker configuration file as a block:

```
define block HEAP with size = 0x1000, alignment = 8{ };
place in RAM {block HEAP};
```

Specify the appropriate size for your application. If you use a heap, you must allocate at least 50 bytes for it.

SETTING UP THE ATEXTIT LIMIT

By default, the `atexit` function can be called a maximum of 32 times from your application. To either increase or decrease this number, add a line to your configuration file. For example, to reserve room for 10 calls instead, write:

```
define symbol __iar_maximum_atexit_calls = 10;
```

CHANGING THE DEFAULT INITIALIZATION

By default, memory initialization is performed during application startup. ILINK sets up the initialization process and chooses a suitable packing method. If the default initialization process does not suit your application and you want more precise control over the initialization process, these alternatives are available:

- Suppressing initialization
- Choosing the packing algorithm
- Manual initialization
- Initializing code—copying ROM to RAM.

For information about the performed initializations, inspect the log file (the command line option `--log initialization`).

Suppressing initialization

If you do not want the linker to arrange for initialization by copying, for some or all sections, make sure that those sections do not match a pattern in an `initialize by copy` directive (or use an `except` clause to exclude them from matching). If you do not want any initialization by copying at all, you can omit the `initialize by copy` directive entirely.

This can be useful if your application, or just your variables, are loaded into RAM by some other mechanism before application startup.

Choosing a packing algorithm

To override the default packing algorithm, write for example:

```
initialize by copy with packing = lz77 { readwrite };
```

For more information about the available packing algorithms, see *initialize directive*, page 446.

Manual initialization

In the usual case, the `initialize by copy` directive is used for making the linker arrange for initialization by copying (with or without packing) of sections with content at application startup. The linker achieves this by logically creating an initialization section for each such section, holding the content of the section, and turning the original section into a section without content. Then, the linker adds table elements to the initialization table so that the initialization will be performed at application startup. You can use `initialize manually` to suppress the creation of table elements to take control over when and how the elements are copied. This is useful for overlays, but also in a number of other circumstances.

For sections without content (zero-initialized sections), the situation is reversed. The linker arranges for zero initialization of all such sections at application startup, except for those that have been mentioned in a `do not initialize` directive. Usually, only `.noinit` sections are specified in a `do not initialize` directive, but you can add any zero-initialized sections you like, and take direct control over when and how these sections are initialized.

Simple copying example with an implicit block

Assume that you have some initialized variables in `MYSECTION`. If you add this directive to your linker configuration file:

```
initialize manually { section MYSECTION };
```

you can use this source code example to initialize the section:

```
#pragma section = "MYSECTION"
#pragma section = "MYSECTION_init"

void DoInit()
{
    char * from = __section_begin("MYSECTION_init");
    char * to   = __section_begin("MYSECTION");
    memcpy(to, from, __section_size("MYSECTION"));
}
```

This piece of source code takes advantage of the fact that if you use `__section_begin` (and related operators) with a section name, a synthetic block is created by the linker for those sections.

Example with explicit blocks

Assume that you instead of needing manual initialization for variables in a specific section, you need it for all initialized variables from a particular library. In that case, you must create explicit blocks for both the variables and the content. Like this:

```
initialize manually      { section .data      object mylib.a };
define block MYBLOCK     { section .data      object mylib.a };
define block MYBLOCK_init { section .data_init object mylib.a };
```

You must also place the two new blocks using one of the section placement directives, the block `MYBLOCK` in RAM and the block `MYBLOCK_init` in ROM.

Then you can initialize the sections using the same source code as in the previous example, only with `MYBLOCK` instead of `MYSECTION`.

Overlay example

This is a simple overlay example that takes advantage of automatic block creation:

```
initialize manually { section MYOVERLAY* };

define overlay MYOVERLAY { section MYOVERLAY1 };
define overlay MYOVERLAY { section MYOVERLAY2 };
```

You must also place overlay `MYOVERLAY` somewhere in RAM. The copying could look like this:

```
#pragma section = "MYOVERLAY"
#pragma section = "MYOVERLAY1_init"
#pragma section = "MYOVERLAY2_init"

void SwitchToOverlay1()
{
    char * from = __section_begin("MYOVERLAY1_init");
    char * to   = __section_begin("MYOVERLAY");
    memcpy(to, from, __section_size("MYOVERLAY1_init"));
}

void SwitchToOverlay2()
{
    char * from = __section_begin("MYOVERLAY2_init");
    char * to   = __section_begin("MYOVERLAY");
    memcpy(to, from, __section_size("MYOVERLAY2_init"));
}
```

Initializing code—copying ROM to RAM

Sometimes, an application copies pieces of code from flash/ROM to RAM. You can direct the linker to arrange for this to be done automatically at application startup, or do it yourself at some later time using the techniques described in *Manual initialization*, page 107.

You need to list the code sections that should be copied in an `initialize by copy` directive. The easiest way is usually to place the relevant functions in a particular section (for example, `RAMCODE`), and add section `RAMCODE` to your `initialize by copy` directive. For example:

```
initialize by copy { rw, section RAMCODE };
```

If you need to place the `RAMCODE` functions in some particular location, you must mention them in a placement directive, otherwise they will be placed together with other read/write sections.

If you need to control the manner and/or time of copying, you must use an `initialize manually` directive instead. See *Manual initialization*, page 107.

If the functions need to run without accessing the flash/ROM, you can use the `__ramfunc` keyword when compiling. See *Execution in RAM*, page 72.

Running all code from RAM

If you want to copy the entire application from ROM to RAM at program startup, use the `initilize by copy` directive, for example:

```
initialize by copy { readonly, readwrite };
```

The `readwrite` pattern will match all statically initialized variables and arrange for them to be initialized at startup. The `readonly` pattern will do the same for all read-only code and data, except for code and data needed for the initialization.

Because the function `__low_level_init`, if present, is called before initialization, it and anything it needs, will not be copied from ROM to RAM either. In some circumstances—for example, if the ROM contents are no longer available to the program after startup—you might need to avoid using the same functions during startup and in the rest of the code.

If anything else should not be copied, include it in an `except` clause. This can apply to, for example, the interrupt vector table.

It is also recommended to exclude the C++ dynamic initialization table from being copied to RAM, as it is typically only read once and then never referenced again. For example, like this:

```
initialize by copy { readonly, readwrite }
    except { section .intvec,          /* Don't copy
                                         interrupt table */
            section .init_array }; /* Don't copy
                                         C++ init table */
```

INTERACTION BETWEEN ILINK AND THE APPLICATION

ILINK provides the command line options `--config_def` and `--define_symbol` to define symbols which can be used for controlling the application. You can also use symbols to represent the start and end of a continuous memory area that is defined in the linker configuration file. For more information, see *Interaction between the tools and your application*, page 206.

To change a reference to one symbol to another symbol, use the ILINK command line option `--redirect`. This is useful, for example, to redirect a reference from a non-implemented function to a stub function, or to choose one of several different implementations of a certain function, for example, how to choose the DLIB formatter for the standard library functions `printf` and `scanf`.

The compiler generates mangled names to represent complex C/C++ symbols. If you want to refer to these symbols from assembler source code, you must use the mangled names.

For information about the addresses and sizes of all global (statically linked) symbols, inspect the entry list in the map file (the command line option `--map`).

For more information, see *Interaction between the tools and your application*, page 206.

STANDARD LIBRARY HANDLING

By default, ILINK determines automatically which variant of the standard library to include during linking. The decision is based on the sum of the runtime attributes available in each object file and the library options passed to ILINK.

To disable the automatic inclusion of the library, use the option `--no_library_search`. In this case, you must explicitly specify every library file to be included. For information about available library files, see *Prebuilt runtime libraries*, page 126.

PRODUCING OTHER OUTPUT FORMATS THAN ELF/DWARF

ILINK can only produce an output file in the ELF/DWARF format. To convert that format into a format suitable for programming PROM/flash, see *The IAR ELF Tool—ielftool*, page 480.

VENEERS

The ARM cores need to use veneers on two occasions:

- When calling an ARM function from Thumb mode or vice versa; the veneer then changes the state of the microprocessor. The called function must be an interworking function, see *Interworking*, page 64. If the core supports the BLX instruction, a veneer is not needed for changing modes.
- When calling a function that it cannot normally reach; the veneer introduces code which makes the call successfully reach the destination.

Code for veneers can be inserted between any caller and called function. As a result, the R12 register must be treated as a scratch register at function calls, including functions written in assembler. This also applies to jumps.

For more information, see `--no_veneers`, page 320.

Hints for troubleshooting

ILINK has several features that can help you manage code and data placement correctly, for example:

- Messages at link time, for examples when a relocation error occurs
- The `--log` option that makes ILINK log information to `stdout`, which can be useful to understand why an executable image became the way it is, see `--log`, page 314
- The `--map` option that makes ILINK produce a memory map file, which contains the result of the linker configuration file, see `--map`, page 315.

RELOCATION ERRORS

For each instruction that cannot be relocated correctly, ILINK will generate a *relocation error*. This can occur for instructions where the target is out of reach or is of an incompatible type, or for many other reasons.

A relocation error produced by ILINK can look like this:

```
Error[Lp002]: relocation failed: out of range or illegal value
  Kind      : R_XXX_YYY[0x1]
  Location  : 0x40000448
              "myfunc" + 0x2c
              Module: somecode.o
              Section: 7 (.text)
              Offset: 0x2c
  Destination: 0x9000000c
              "read"
              Module: read.o(iolib.a)
              Section: 6 (.text)
              Offset: 0x0
```

The message entries are described in this table:

Message entry	Description
Kind	The relocation directive that failed. The directive depends on the instruction used.
Location	<p>The location where the problem occurred, described with the following details:</p> <ul style="list-style-type: none">• The instruction address, expressed both as a hexadecimal value and as a label with an offset. In this example, 0x40000448 and "myfunc" + 0x2c.• The module, and the file. In this example, the module somecode.o.• The section number and section name. In this example, section number 7 with the name .text.• The offset, specified in number of bytes, in the section. In this example, 0x2c.
Destination	<p>The target of the instruction, described with the following details:</p> <ul style="list-style-type: none">• The instruction address, expressed both as a hexadecimal value and as a label with an offset. In this example, 0x9000000c and "read" (thus, no offset).• The module, and when applicable the library. In this example, the module read.o and the library iolib.a.• The section number and section name. In this example, section number 6 with the name .text.• The offset, specified in number of bytes, in the section. In this example, 0x0.

Table 4: Description of a relocation error

Possible solutions

In this case, the distance from the instruction in `myfunc` to `__read` is too long for the branch instruction.

Possible solutions include ensuring that the two `.text` sections are allocated closer to each other or using some other calling mechanism that can reach the required distance. It is also possible that the referring function tried to refer to the wrong target and that this caused the range error.

Different range errors have different solutions. Usually, the solution is a variant of the ones presented above, in other words modifying either the code or the section placement.

Checking module consistency

This section introduces the concept of runtime model attributes, a mechanism used by the tools provided by IAR Systems to ensure that modules that are linked into an application are compatible, in other words, are built using compatible settings. The tools use a set of predefined runtime model attributes. In addition to these, you can define your own that you can use to ensure that incompatible modules are not used together.

Note: In addition to the predefined attributes, compatibility is also checked against the AEABI runtime attributes. These attributes deal mainly with object code compatibility, etc. They reflect compilation settings and are not user-configurable.

RUNTIME MODEL ATTRIBUTES

A runtime attribute is a pair constituted of a named key and its corresponding value. In general, two modules can only be linked together if they have the same value for each key that they both define.

There is one exception: if the value of an attribute is `*`, then that attribute matches any value. The reason for this is that you can specify this in a module to show that you have considered a consistency property, and this ensures that the module does not rely on that property.

Note: For IAR predefined runtime model attributes, the linker checks them in several ways.

Example

In this table, the object files could (but do not have to) define the two runtime attributes `color` and `taste`:

Object file	Color	Taste
file1	blue	not defined
file2	red	not defined
file3	red	*
file4	red	spicy
file5	red	lean

Table 5: Example of runtime model attributes

In this case, `file1` cannot be linked with any of the other files, since the runtime attribute `color` does not match. Also, `file4` and `file5` cannot be linked together, because the `taste` runtime attribute does not match.

On the other hand, `file2` and `file3` can be linked with each other, and with either `file4` or `file5`, but not with both.

USING RUNTIME MODEL ATTRIBUTES

To ensure module consistency with other object files, use the `#pragma rtmodel` directive to specify runtime model attributes in your C/C++ source code. For example, if you have a UART that can run in two modes, you can specify a runtime model attribute, for example `uart`. For each mode, specify a value, for example `model` and `mode2`. Declare this in each module that assumes that the UART is in a particular mode. This is how it could look like in one of the modules:

```
#pragma rtmodel="uart", "model"
```

Alternatively, you can also use the `rtmodel` assembler directive to specify runtime model attributes in your assembler source code. For example:

```
rtmodel "uart", "model"
```

Note that key names that start with two underscores are reserved by the compiler. For more information about the syntax, see *rtmodel*, page 375 and the *IAR Assembler User Guide for ARM*, respectively.

At link time, the IAR ILINK Linker checks module consistency by ensuring that modules with conflicting runtime attributes will not be used together. If conflicts are detected, an error is issued.

The DLIB runtime environment

- Introduction to the runtime environment
- Setting up the runtime environment
- Additional information on the runtime environment
- Managing a multithreaded environment

Introduction to the runtime environment

A *runtime environment* is the environment in which your application executes.

This section contains information about:

- *Runtime environment functionality*, page 115
- *Briefly about input and output (I/O)*, page 116
- *Briefly about C-SPY emulated I/O*, page 117
- *Briefly about retargeting*, page 118

RUNTIME ENVIRONMENT FUNCTIONALITY

The *DLIB runtime environment* supports Standard C and C++ and consists of:

- The *C/C++ standard library*, both its interface (provided in the system header files) and its implementation.
- Startup and exit code.
- Low-level I/O interface for managing input and output (I/O).
- Special compiler support, for instance functions for switch handling or integer arithmetics.
- Support for hardware features:
 - Direct access to low-level processor operations by means of *intrinsic* functions, such as functions for interrupt mask handling
 - Peripheral unit registers and interrupt definitions in include files
 - The Vector Floating Point (VFP) coprocessor.

Runtime environment functions are provided in a *runtime library*.

The runtime library is delivered both as a prebuilt library and (depending on your product package) as source files. The prebuilt libraries are available in different *configurations* to meet various needs, see *Runtime library configurations*, page 125. You can find the libraries in the product subdirectories `arm\lib` and `arm\src\lib`, respectively.


For more information about the library, see the chapter *C/C++ standard library functions*.

BRIEFLY ABOUT INPUT AND OUTPUT (I/O)

Every application must communicate with its environment. The application might for example display information on an LCD, read a value from a sensor, get the current date from the operating system, etc. Typically, your application performs I/O via the C/C++ standard library or some third-party library.

There are many functions in the C/C++ standard library that deal with I/O, including functions for: standard character streams, file system access, time and date, miscellaneous system actions, and termination and assert. This set of functions is referred to as the *standard I/O interface*.

On a desktop computer or a server, the operating system is expected to provide I/O functionality to the application via the standard I/O interface in the runtime environment. However, in an embedded system, the runtime library cannot assume that such functionality is present, or even that there is an operating system at all. Therefore, the low-level part of the standard I/O interface is not completely implemented by default:




To make the standard I/O interface work, you can:

- Let the C-SPY debugger emulate I/O operations on the host computer, see *Briefly about C-SPY emulated I/O*, page 117
- *Retarget* the standard I/O interface to your target system by providing a suitable implementation of the interface, see *Briefly about retargeting*, page 118.

It is possible to mix these two approaches. You can, for example, let debug printouts and asserts be emulated by the C-SPY debugger, but implement your own file system. The debug printouts and asserts are useful during debugging, but no longer needed when running the application stand-alone (not connected to the C-SPY debugger).

BRIEFLY ABOUT C-SPY EMULATED I/O

C-SPY emulated I/O is a mechanism which lets the runtime environment interact with the C-SPY debugger to emulate I/O actions on the host computer:



For example, when C-SPY emulated I/O is enabled:

- Standard character streams are directed to the C-SPY **Terminal I/O** window
- File system operations are performed on the host computer
- Time and date functions return the time and date of the host computer
- Termination and failed asserts break execution and notify the C-SPY debugger.

This behavior can be valuable during the early development of an application, for example in an application that uses file I/O before any flash file system I/O drivers are implemented, or if you need to debug constructions in your application that use `stdin` and `stdout` without the actual hardware device for input and output being available.


See *Setting up your runtime environment*, page 119 and *The semihosting mechanism*, page 130.

BRIEFLY ABOUT RETARGETING

Retargeting is the process where you adapt the runtime environment so that your application can execute I/O operations on your target system.

The standard I/O interface is large and complex. To make retargeting easier, the DLIB runtime environment is designed so that it performs all I/O operations through a small set of simple functions, which is referred to as the *DLIB low-level I/O interface*. By default, the functions in the low-level interface lack usable implementations. Some are unimplemented, others have stub implementations that do not perform anything except returning error codes.

To retarget the standard I/O interface, all you have to do is to provide implementations for the functions in the DLIB low-level I/O interface.



For example, if your application calls the functions `printf` and `fputc` in the standard I/O interface, the implementations of those functions both call the low-level function `__write` to output individual characters. To make them work, you just need to provide an implementation of the `__write` function—either by implementing it yourself, or by using a third-party implementation.

For information about how to override library modules with your own implementations, see *Overriding library modules*, page 122. See also *The DLIB low-level I/O interface*, page 138 for information about the functions that are part of the interface.

Setting up the runtime environment

This section contains these tasks:

- *Setting up your runtime environment*, page 119
A runtime environment with basic project settings to be used during the initial phase of development.
- *Retargeting—Adapting for your target system*, page 120
- *Overriding library modules*, page 122
- *Customizing and building your own runtime library*, page 123

See also:

- *Managing a multithreaded environment*, page 150 for information about how to adapt the runtime environment to treat all library objects according to whether they are global or local to a thread.

SETTING UP YOUR RUNTIME ENVIRONMENT

You can set up the runtime environment based on some basic project settings. It is also often convenient to let the C-SPY debugger manage things like standard streams, file I/O, and various other system interactions. This basic runtime environment can be used for simulation before you have any target hardware.

To set up the runtime environment:

- 1 Before you build your project, choose **Project>Options>General Options** to open the **Options** dialog box.
- 2 On the **Library Configuration** page, verify the following settings:
 - **Library**: choose which *library configuration* to use. Typically, choose **Tiny**, **Normal**, or **Full**.
For information about the various library configurations, see *Runtime library configurations*, page 125.
- 3 On the **Library Options** page, select **Auto** for both **Printf formatter** and **Scanf formatter**. This means that the linker will automatically choose the appropriate formatters based on information from the compiler. For more information about the available formatters and how to choose one manually, see *Formatters for printf*, page 128 and *Formatters for scanf*, page 129, respectively.
- 4 To enable C-SPY emulated I/O, choose **Project>Options>General Options>Library Configuration** and choose **Semihosted** (`--semihosted`) or **IAR breakpoint** (`--semihosting=iar_breakpoint`).

Note that for some Cortex-M devices it is also possible to direct `stdout/stderr` via SWO. This can significantly improve `stdout/stderr` performance compared to semihosting. For hardware requirements, see the *C-SPY® Debugging Guide for ARM*.



To enable `stdout` via SWO on the command line, use the linker option `--redirect __iar_sh_stdout=__iar_sh_stdout_swo`.



To enable `stdout` via SWO in the IDE, select the **Semihosted** option and the **stdout/stderr via SWO** option.

See *Briefly about C-SPY emulated I/O*, page 117 and *The semihosting mechanism*, page 130.

- 5 On some systems, terminal output might be slow because the host computer and the target system must communicate for each character.

For this reason, a replacement for the `__write` function called `__write_buffered` is included in the runtime library. This module buffers the output and sends it to the debugger one line at a time, speeding up the output. Note that this function uses about 80 bytes of RAM memory.



To use this feature in the IDE, choose **Project>Options>General Options>Library Options>** and select the option **Buffered terminal output**.



To enable this function on the command line, add this to the linker command line:

```
--redirect __write=__write_buffered
```

- 6 Some math functions are available in different versions: default versions, smaller than the default versions, and larger but more accurate than default versions. Consider which versions you should use.

For more information, see *Math functions*, page 131.

- 7 When you build your project, a suitable prebuilt library and library configuration file are automatically used based on the project settings you made.

For information about which project settings affect the choice of library file, see *Runtime library configurations*, page 125.

You have now set up a runtime environment that can be used while developing your application source code.

RETARGETING—ADAPTING FOR YOUR TARGET SYSTEM

Before you can run your application on your target system, you must adapt some parts of the runtime environment, typically the system initialization and the DLIB low-level I/O interface functions.

To adapt your runtime environment for your target system:

1 Adapt system initialization.

It is likely that you must adapt the system initialization, for example, your application might need to initialize interrupt handling, I/O handling, watchdog timers, etc. You do this by implementing the routine `__low_level_init`, which is executed before the data sections are initialized. See *System startup and termination*, page 133 and *System initialization*, page 137. Note that you can find device-specific examples on this in the example projects provided in the product installation; see the Information Center.

2 Adapt the runtime library for your target system. To implement such functions, you need a good understanding of the DLIB low-level I/O interface, see *Briefly about retargeting*, page 118.

Typically, you must implement your own functions if your application uses:

- Standard streams for input and output

If any of these streams are used by your application, for example by the functions `printf` and `scanf`, you must implement your versions of the low-level functions `__read` and `__write`.

The low-level functions identify I/O streams, such as an open file, with a file handle that is a unique integer. The I/O streams normally associated with `stdin`, `stdout`, and `stderr` have the file handles 0, 1, and 2, respectively. When the handle is -1, all streams should be flushed. Streams are defined in `stdio.h`.

- File input and output

The library contains a large number of powerful functions for file I/O operations, such as `fopen`, `fclose`, `fprintf`, `fputs`, etc. All these functions call a small set of low-level functions, each designed to accomplish one particular task; for example, `__open` opens a file, and `__write` outputs characters. Implement your version of these low-level functions.

- signal and raise

If the default implementation of these functions does not provide the functionality you need, you can implement your own versions.

- Time and date

To make the time and date functions work, you must implement the functions `clock`, `__time32`, `__time64`, and `__getzone`. Whether you use `__time32` or `__time64` depends on which interface you use for `time_t`, see *time.h*, page 435.

- Assert, see `__aeabi_assert`, page 139.

- Environment interaction

If the default implementation of `system` or `getenv` does not provide the functionality you need, you can implement your own versions.

For more information about the functions, see *The DLIB low-level I/O interface*, page 138.

The library files that you can override with your own versions are located in the `arm\src\lib` directory.

- 3 When you have implemented your functions of the low-level I/O interface, you must add your version of these functions to your project. For information about this, see *Overriding library modules*, page 122.

Note: If you have implemented a DLIB low-level I/O interface function and added it to a project that you have built with support for C-SPY emulated I/O, your low-level function will be used and not the functions provided with C-SPY emulated I/O. For example, if you implement your own version of `__write`, output to the C-SPY **Terminal I/O** window will not be supported. See *Briefly about C-SPY emulated I/O*, page 117.

- 4 Before you can execute your application on your target system, you must rebuild your project with a Release build configuration. This means that the linker will not include the C-SPY emulated I/O mechanism and the low-level I/O functions it provides. If your application calls any of the low-level functions of the standard I/O interface, either directly or indirectly, and your project does not contain these, the linker will issue an error for every missing low-level function. Also, note that the `NDEBUG` symbol is defined in a Release build configuration, which means asserts will no longer be generated. For more information, see `__aeabi_assert`, page 139.

OVERRIDING LIBRARY MODULES

To override a library function and replace it with your own implementation:

- 1 Use a template source file—a library source file or another template—and place a copy of it in your project directory.

The library files that you can override with your own versions are located in the `arm\src\lib` directory.

- 2 Modify the file.

Note: To override the functions in a module, you must provide alternative implementations for all the needed symbols in the overridden module. Otherwise you will get error messages about duplicate definitions.

- 3 Add the modified file to your project, like any other source file.

Note: If you have implemented a DLIB low-level I/O interface function and added it to a project that you have built with support for C-SPY emulated I/O, your low-level function will be used and not the functions provided with C-SPY emulated I/O. For

example, if you implement your own version of `__write`, output to the C-SPY **Terminal I/O** window will not be supported. See *Briefly about C-SPY emulated I/O*, page 117.

You have now finished the process of overriding the library module with your version.

CUSTOMIZING AND BUILDING YOUR OWN RUNTIME LIBRARY

If the prebuilt library configurations do not meet your requirements, you can customize your own library configuration, but that requires that you *rebuild* the library.

Building a customized library is a complex process. Therefore, consider carefully whether it is really necessary. You must build your own runtime library when:

- You want to define your own library configuration with support for locale, file descriptors, multibyte characters, etc. This will include or exclude certain parts of the DLIB runtime environment.

In those cases, you must:

- Make sure that you have installed the library source code (`scr\lib`). If not already installed, you can install it using the IAR License Manager, see the *Installation and Licensing Guide*.
- Set up a library project
- Make the required library customizations
- Build your customized runtime library
- Finally, make sure your application project will use the customized runtime library.

To set up a library project:

- I In the IDE, choose **Project>Create New Project** and

Note that when you create a new library project from a template, the majority of the files included in the new project are the original installation files. If you are going to modify these files, make copies of them first and replace the original files in the project with these copies.

To customize the library functionality:

- I The library functionality is determined by a set of *configuration symbols*. The default values of these symbols are defined in the file `DLib_Defaults.h` which you can find in `arm\inc\c`. This read-only file describes the configuration possibilities. Note that you should not modify this file.

In addition, your custom library has its own *library configuration file* `dlarmCustom.h`—which you can find in `arm\config\template\project`—and which sets up that specific library with the required library configuration. Customize

this file by setting the values of the configuration symbols according to the application requirements. For example, to enable the `ll` qualifier in `printf` format strings, write in your library configuration file:

```
#define _DLIB_PRINTF_LONG_LONG 1
```

For information about configuration symbols that you might want to customize, see:

- *Configuration symbols for printf and scanf*, page 147
- *Configuration symbols for file input and output*, page 148
- *Locale*, page 148
- *Managing a multithreaded environment*, page 150

- 2 When you are finished, build your library project with the appropriate project options.

After you build your library, you must make sure to use it in your application project.



To build IAR Embedded Workbench projects from the command line, use the IAR Command Line Build Utility (`iarbuild.exe`). However, no make or batch files for building the library from the command line are provided. For information about the build process and the IAR Command Line Build Utility, see the *IDE Project Management and Building Guide for ARM*.

To use the customized runtime library in your application project:

- 1 In the IDE, choose **Project>Options>General Options** and click the **Library Configuration** tab.
- 2 From the **Library** drop-down menu, choose **Custom DLIB**.
- 3 In the **Configuration file** text box, locate your library configuration file.
- 4 Click the **Library** tab, also in the **Linker** category. Use the **Additional libraries** text box to locate your library file.

Additional information on the runtime environment

This section gives additional information on the runtime environment:

- *Runtime library configurations*, page 125
- *Prebuilt runtime libraries*, page 126
- *Formatters for printf*, page 128
- *Formatters for scanf*, page 129
- *The C-SPY emulated I/O mechanism*, page 130
- *The semihosting mechanism*, page 130
- *Math functions*, page 131

- *System startup and termination*, page 133
- *System initialization*, page 137
- *The DLIB low-level I/O interface*, page 138
- *Configuration symbols for printf and scanf*, page 147
- *Configuration symbols for file input and output*, page 148
- *Locale*, page 148

RUNTIME LIBRARY CONFIGURATIONS

The runtime library is provided with different *library configurations*, where each configuration is suitable for different application requirements.

The runtime library configuration is defined in the *library configuration file*. It contains information about what functionality is part of the runtime environment. The less functionality you need in the runtime environment, the smaller the environment becomes.

These predefined library configurations are available:




Library configuration	Description
Normal DLIB (default)	No locale interface, C locale, no file descriptor support, and no multibyte characters in printf and scanf.
Full DLIB	Full locale interface, C locale, file descriptor support, and multibyte characters in printf and scanf.

Table 6: Library configurations

Note: In addition to these predefined library configurations, you can provide your own configuration, see *Customizing and building your own runtime library*, page 123

If you do not specify a library configuration explicitly you will get the default configuration. If you use a prebuilt runtime library, a configuration file that matches the runtime library object file will automatically be used. See *Setting up your runtime environment*, page 119.

To override the default library configuration, use one of these methods:

-  Use a prebuilt configuration of your choice—to specify a runtime configuration explicitly:
-  Choose **Project>Options>General Options>Library Configuration>Library** and change the default setting.
-  Use the `--dlib_config` compiler option, see `--dlib_config`, page 269.
- The prebuilt libraries are based on the default configurations, see *Runtime library configurations*, page 125.

- 2 If you have built your own customized library, choose **Project>Options>Library Configuration>Library** and choose **Custom** to use your own configuration. For more information, see *Customizing and building your own runtime library*, page 123.

PREBUILT RUNTIME LIBRARIES

The prebuilt runtime libraries are configured for different combinations of these options:

- Library configuration—Normal or Full.

The linker will automatically include the correct library object file and library configuration file. To explicitly specify a library configuration, use the `--dlib_config` compiler option.

Library filename syntax

The names of the libraries are constructed from these elements:

<code>{library}</code>	d1 for the IAR DLIB runtime environment
<code>{cpu}</code>	for the ARM core
<code>{lib_config}</code>	Specifies the library configuration: n = Normal f = Full

You can find the library object files in the subdirectory `arm\lib\` and the library configuration files in the `arm\inc` subdirectory.

Groups of library files

The libraries are delivered in groups of library functions:

Library files for C library functions

These are the functions defined by Standard C, for example functions like `printf` and `scanf`. Note that this library does not include math functions.

The names of the library files are constructed in the following way:

`d1{architecture}_{cpu_mode}{byte_order}{lib-config}{rwp1}.a`

which more specifically means

`d1{4t|5E|6M|6Mx|7M|7Mx|7Sx}_{a|t}{l|b}{n|f}{s}.a`

Library files for C++ and Embedded C++ library functions

These are the functions defined by C++, compiled with support for either Standard C++ or Embedded C++.

The names of the library files are constructed in the following way:

```
dlpp{architecture}_{cpu_mode}{byte_order}
_{lib-config}{language}{rwp}.a
```

which more specifically means

```
dlpp{4t|5E|6M|6Mx|7M|7Mx|7Sx|4as|7as}_{a|t}{l|b}_{n|f}{c|e}{s}.a
```

Library files for math functions

These are the functions for floating-point arithmetic and functions with a floating-point type in its signature as defined by Standard C, for example functions like `sqrt`.

The names of the library files are constructed in the following way:

```
m{architecture}_{cpu_mode}{byte_order}{fp_implementation}.a
```

which more specifically means

```
m{4t|5E|6M|6Mx|7M|7Mx|7Sx}_{a|t}{l|b}{v|s}.a
```

Library files for runtime support functions

These are functions for system startup, initialization, non floating-point AEABI support routines, and some of the functions that are part of Standard C and C++.

The names of the library files are constructed in the following way:

```
rt{architecture}_{cpu_mode}{byte_order}.a
```

which more specifically means

```
rt{4t|5E|6M|6Mx|7M|7Mx|7Sx}_{a|t}{l|b}.a
```

Library files for debug support functions

These are functions for debug support for the semihosting interface. The names of the library files are constructed in the following way:

```
sh{debug_interface}_{byte_order}.a
```

or

```
sh{architecture}_{byte_order}.a
```

which more specifically means

```
sh{s|b|i}_{l|b}.a
```

```
or
sh{6Mx|7Mx|7Sx}_{l|b}.a
```

FORMATTERS FOR PRINTF

The `printf` function uses a formatter called `_Printf`. The full version is quite large, and provides facilities not required in many embedded applications. To reduce the memory consumption, three smaller, alternative versions are also provided. Note that the `wprintf` variants are not affected.

This table summarizes the capabilities of the different formatters:

Formatting capabilities	Tiny	Small/ SmallNoMb†	Large/ LargeNoMb†	Full/ FullNoMb†
Basic specifiers <code>c</code> , <code>d</code> , <code>i</code> , <code>o</code> , <code>p</code> , <code>s</code> , <code>u</code> , <code>X</code> , <code>x</code> , and <code>%</code>	Yes	Yes	Yes	Yes
Multibyte support	No	Yes/No	Yes/No	Yes/No
Floating-point specifiers <code>a</code> , and <code>A</code>	No	No	No	Yes
Floating-point specifiers <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , and <code>G</code>	No	No	Yes	Yes
Conversion specifier <code>n</code>	No	No	Yes	Yes
Format flag <code>+</code> , <code>-</code> , <code>#</code> , <code>0</code> , and space	No	Yes	Yes	Yes
Length modifiers <code>h</code> , <code>l</code> , <code>L</code> , <code>s</code> , <code>t</code> , and <code>Z</code>	No	Yes	Yes	Yes
Field width and precision, including <code>*</code>	No	Yes	Yes	Yes
<code>long long</code> support	No	No	Yes	Yes

Table 7: Formatters for `printf`

† NoMb means without multibytes.

Note: It is possible to optimize these functions even further, but that requires that you rebuild the library. See *Configuration symbols for `printf` and `scanf`*, page 147.

The compiler can automatically detect which formatting capabilities are needed in a direct call to `printf`, if the formatting string is a string literal. This information is passed to the linker, which combines the information from all modules to select a suitable formatter for the application. However, if the formatting string is a variable, or if the call is indirect through a function pointer, the compiler cannot perform the analysis, forcing the linker to select the Full formatter. In this case, you might want to override the automatically selected `printf` formatter.



To override the automatically selected `printf` formatter in the IDE:

Choose **Project>Options>General Options** to open the **Options** dialog box.

- 2 On the **Library Options** page, select the appropriate formatter.



To override the automatically selected printf formatter from the command line:

- I Use one of these ILINK command line options:

```
--redirect _Printf=_PrintfFull
--redirect _Printf=_PrintfFullNoMb
--redirect _Printf=_PrintfLarge
--redirect _Printf=_PrintfLargeNoMb
--redirect _Printf=_PrintfSmall
--redirect _Printf=_PrintfSmallNoMb
--redirect _Printf=_PrintfTiny
--redirect _Printf=_PrintfTinyNoMb
```

FORMATTERS FOR SCANF

In a similar way to the `printf` function, `scanf` uses a common formatter, called `_Scanf`. The full version is quite large, and provides facilities that are not required in many embedded applications. To reduce the memory consumption, two smaller, alternative versions are also provided. Note that the `wscanf` versions are not affected.

This table summarizes the capabilities of the different formatters:

Formatting capabilities	Small/ SmallNoMb†	Large/ LargeNoMb†	Full/ FullNoMb†
Basic specifiers <code>c</code> , <code>d</code> , <code>i</code> , <code>o</code> , <code>p</code> , <code>s</code> , <code>u</code> , <code>X</code> , <code>x</code> , and <code>%</code>	Yes	Yes	Yes
Multibyte support	Yes/No	Yes/No	Yes/No
Floating-point specifiers <code>a</code> , and <code>A</code>	No	No	Yes
Floating-point specifiers <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , and <code>G</code>	No	No	Yes
Conversion specifier <code>n</code>	No	No	Yes
Scan set <code>[</code> and <code>]</code>	No	Yes	Yes
Assignment suppressing <code>*</code>	No	Yes	Yes
<code>long long</code> support	No	No	Yes

Table 8: Formatters for `scanf`

† NoMb means without multibytes.

Note: It is possible to optimize these functions even further, but that requires that you rebuild the library. See *Configuration symbols for printf and scanf*, page 147.

The compiler can automatically detect which formatting capabilities are needed in a direct call to `scanf`, if the formatting string is a string literal. This information is passed to the linker, which combines the information from all modules to select a suitable formatter for the application. However, if the formatting string is a variable, or if the call is indirect through a function pointer, the compiler cannot perform the analysis, forcing

the linker to select the full formatter. In this case, you might want to override the automatically selected `scanf` formatter.



To manually specify the scanf formatter in the IDE:

- 1 Choose **Project>Options>General Options** to open the **Options** dialog box.
- 2 On the **Library Options** page, select the appropriate formatter.



To manually specify the scanf formatter from the command line:

- 1 Use one of these ILINK command line options:

```
--redirect _Scanf=_ScanfFull
--redirect _Scanf=_ScanfFullNoMb
--redirect _Scanf=_ScanfLarge
--redirect _Scanf=_ScanfLargeNoMb
--redirect _Scanf=_ScanfSmall
--redirect _Scanf=_ScanfSmallNoMb
```

THE C-SPY EMULATED I/O MECHANISM

- 1 The debugger will detect the presence of the function `__DebugBreak`, which will be part of the application if you linked it with the linker option for C-SPY emulated I/O.
- 2 In this case, the debugger will automatically set a breakpoint at the `__DebugBreak` function.
- 3 When your application calls a function in the DLIB low-level I/O interface, for example, `open`, the `__DebugBreak` function is called, which will cause the application to stop at the breakpoint and perform the necessary services.
- 4 The execution will then resume.

See also *Briefly about C-SPY emulated I/O*, page 117.

THE SEMIHOSTING MECHANISM

C-SPY emulated I/O is compatible with the semihosting interface provided by ARM Limited. When an application invokes a semihosting call, the execution stops at a debugger breakpoint. The debugger then handles the call, performs any necessary actions on the host computer and then resumes the execution.

There are three variants of semihosting mechanisms available:

- For Cortex-M, the interface uses `BKPT` instructions to perform semihosting calls
- For other ARM cores, `SVC` instructions are used for the semihosting calls
- *IAR breakpoint*, which is an IAR-specific alternative to semihosting that uses `SVC`.

To support semihosting via SVC, the debugger must set its semihosting breakpoint on the Supervisor Call vector to catch SVC calls. If your application uses SVC calls for other purposes than semihosting, the handling of this breakpoint will cause a severe performance penalty for each such call. IAR breakpoint is a way to get around this. By using a special function call instead of an SVC instruction to perform semihosting, the semihosting breakpoint can be set on that special function instead. This means that semihosting will not interfere with other uses of the Supervisor Call vector.

Note that IAR breakpoint is an IAR-specific extension of the semihosting standard. If you link your application with libraries built with toolchains from other vendors than IAR Systems and use IAR breakpoint, semihosting calls made from code in those libraries will not work.

MATH FUNCTIONS

Some C/C++ standard library math functions are available in different versions:

- The default versions
- Smaller versions (but less accurate)
- More accurate versions (but larger).

Smaller versions

The functions `cos`, `exp`, `log`, `log2`, `log10`, `__iar_Log` (a help function for `log`, `log2`, and `log10`), `pow`, `sin`, `tan`, and `__iar_Sin` (a help function for `sin` and `cos`) exist in additional, smaller versions in the library. They are about 20% smaller and about 20% faster than the default versions. The functions handle INF and NaN values. The drawbacks are that they almost always lose some precision and they do not have the same input range as the default versions.

The names of the functions are constructed like:

```
__iar_xxx_small<f|l>
```

where `f` is used for `float` variants, `l` is used for `long double` variants, and no suffix is used for `double` variants.



To specify individual math functions from the command line:

Redirect the default function names to these names when linking, using these options:

```
--redirect sin=__iar_sin_small
--redirect cos=__iar_cos_small
--redirect tan=__iar_tan_small
--redirect log=__iar_log_small
--redirect log2=__iar_log2_small
--redirect log10=__iar_log10_small
--redirect exp=__iar_exp_small
--redirect pow=__iar_pow_small
--redirect __iar_Sin=__iar_Sin_small
--redirect __iar_Log=__iar_Log_small

--redirect sinf=__iar_sin_smallf
--redirect cosf=__iar_cos_smallf
--redirect tanf=__iar_tan_smallf
--redirect logf=__iar_log_smallf
--redirect log2f=__iar_log2_smallf
--redirect log10f=__iar_log10_smallf
--redirect expf=__iar_exp_smallf
--redirect powf=__iar_pow_smallf
--redirect __iar_FSin=__iar_Sin_smallf
--redirect __iar_FLog=__iar_Log_smallf

--redirect sinl=__iar_sin_small1
--redirect cosl=__iar_cos_small1
--redirect tanl=__iar_tan_small1
--redirect logl=__iar_log_small1
--redirect log2l=__iar_log2_small1
--redirect log10l=__iar_log10_small1
--redirect expl=__iar_exp_small1
--redirect powl=__iar_pow_small1
--redirect __iar_LSin=__iar_Sin_small1
--redirect __iar_LLog=__iar_Log_small1
```

Note that if you want to redirect any of the functions `sin`, `cos`, or `__iar_Sin`, you must redirect all three functions.

Note that if you want to redirect any of the functions `log`, `log2`, `log10`, or `__iar_Log`, you must redirect all four functions.

More accurate versions

The functions `cos`, `pow`, `sin`, and `tan`, and the help functions `__iar_Sin` and `__iar_Pow` exist in versions in the library that are more exact and can handle larger

argument ranges. The drawback is that they are larger and slower than the default versions.

The names of the functions are constructed like:

```
__iar_xxx_accurate<f|l>
```

where *f* is used for float variants, *l* is used for long double variants, and no suffix is used for double variants.



To specify individual math functions from the command line:

- 1 Redirect the default function names to these names when linking, using these options:

```
--redirect sin=__iar_sin_accurate
--redirect cos=__iar_cos_accurate
--redirect tan=__iar_tan_accurate
--redirect pow=__iar_pow_accurate
--redirect __iar_Sin=__iar_Sin_accurate
--redirect __iar_Pow=__iar_Pow_accurate

--redirect sinf=__iar_sin_accuratef
--redirect cosf=__iar_cos_accuratef
--redirect tanf=__iar_tan_accuratef
--redirect powf=__iar_pow_accuratef
--redirect __iar_FSin=__iar_Sin_accuratef
--redirect __iar_FPow=__iar_Pow_accuratef

--redirect sinl=__iar_sin_accuratel
--redirect cosl=__iar_cos_accuratel
--redirect tanl=__iar_tan_accuratel
--redirect powl=__iar_pow_accuratel
--redirect __iar_LSin=__iar_Sin_accuratel
--redirect __iar_LPow=__iar_Pow_accuratel
```

Note that if you want to redirect any of the functions *sin*, *cos*, or *__iar_Sin*, you must redirect all three functions.

Note that if you want to redirect any of the functions *pow* or *__iar_Pow*, you must redirect both functions.

SYSTEM STARTUP AND TERMINATION

This section describes the runtime environment actions performed during startup and termination of your application.

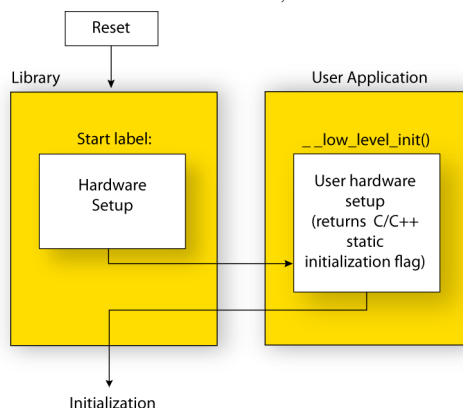
The code for handling startup and termination is located in the source files *cstartup.s*, *cmain.s*, *cexit.s*, and *low_level_init.c* located in the *arm\src\lib* directory.

For information about how to customize the system startup code, see *System initialization*, page 137.

System startup

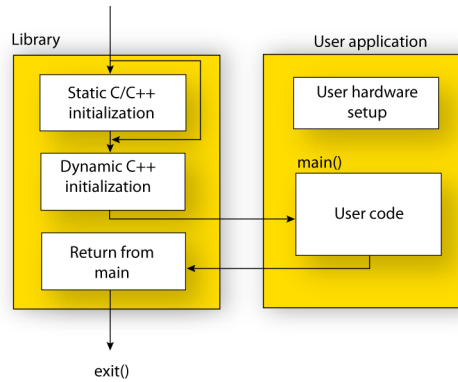
During system startup, an initialization sequence is executed before the `main` function is entered. This sequence performs initializations required for the target hardware and the C/C++ environment.

For the hardware initialization, it looks like this:



- When the CPU is reset it will start executing at the program entry label `__iar_program_start` in the system startup code.
- The stack pointer is initialized to the end of the `CSTACK` block
- For ARM7/9/11, Cortex-A, and Cortex-R devices, exception stack pointers are initialized to the end of each corresponding section
- The function `__low_level_init` is called if you defined it, giving the application a chance to perform early initializations.

For the C/C++ initialization, it looks like this:

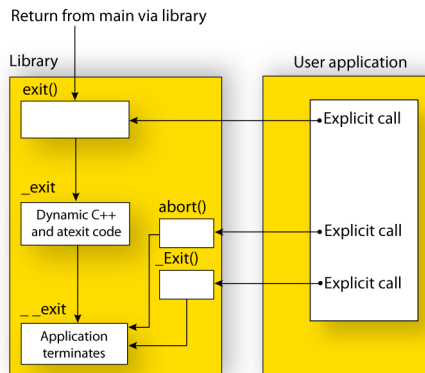


- Static and global variables are initialized. That is, zero-initialized variables are cleared and the values of other initialized variables are copied from ROM to RAM memory. This step is skipped if `__low_level_init` returns zero. For more information, see *Initialization at system startup*, page 90.
- Static C++ objects are constructed
- The `main` function is called, which starts the application.

For information about the initialization phase, see *Application execution—an overview*, page 58.

System termination

This illustration shows the different ways an embedded application can terminate in a controlled way:



An application can terminate normally in two different ways:

- Return from the `main` function
- Call the `exit` function.

Because the C standard states that the two methods should be equivalent, the system startup code calls the `exit` function if `main` returns. The parameter passed to the `exit` function is the return value of `main`.

The default `exit` function is written in C. It calls a small assembler function `_exit` that will:

- Call functions registered to be executed when the application ends. This includes C++ destructors for static and global variables, and functions registered with the standard function `atexit`. See also *Setting up the atexit limit*, page 106.
- Close all open files
- Call `__exit`
- When `__exit` is reached, stop the system.

An application can also exit by calling the `abort` or the `_Exit` function. The `abort` function just calls `__exit` to halt the system, and does not perform any type of cleanup. The `_Exit` function is equivalent to the `abort` function, except for the fact that `_Exit` takes an argument for passing exit status information.

If you want your application to do anything extra at exit, for example resetting the system (and if using `atexit` is not sufficient), you can write your own implementation of the `__exit(int)` function.

The library files that you can override with your own versions are located in the `arm\src\lib` directory. See *Overriding library modules*, page 122.

C-SPY debugging support for system termination

If you have enabled C-SPY emulated I/O during linking, the normal `__exit` function is replaced with a special one. C-SPY will then recognize when this function is called and can take appropriate actions to emulate program termination. For more information, see *Briefly about C-SPY emulated I/O*, page 117.

SYSTEM INITIALIZATION

It is likely that you need to adapt the system initialization. For example, your application might need to initialize memory-mapped special function registers (SFRs), or omit the default initialization of data sections performed by the system startup code.

You can do this by implementing your own version of the routine `__low_level_init`, which is called from the file `cmain.s` before the data sections are initialized. Modifying the file `cstartup.s` directly should be avoided.

For Cortex-M, the code for handling system startup is located in the source files `cstartup_M.s` and `low_level_init.c`, located in the `arm\src\lib` directory.

For other ARM devices, the code for handling system startup is located in the source files `cstartup.s` and `low_level_init.c`, located in the `arm\src\lib` directory.

Note: Normally, you do not need to customize either of the files `cmain.s` or `cexit.s`.

Note: Regardless of whether you implement your own version of `__low_level_init` or the file `cstartup.s`, you do not have to rebuild the library.

Customizing `__low_level_init`

A skeleton low-level initialization file is supplied with the product: `low_level_init.c`. Note that static initialized variables cannot be used within the file, because variable initialization has not been performed at this point.

The value returned by `__low_level_init` determines whether or not data sections should be initialized by the system startup code. If the function returns 0, the data sections will not be initialized.

Modifying the cstartup file

As noted earlier, you should not modify the file `cstartup.s` if implementing your own version of `__low_level_init` is enough for your needs. However, if you do need to modify the file `cstartup.s`, we recommend that you follow the general procedure for creating a modified copy of the file and adding it to your project, see *Overriding library modules*, page 122.

Note that you must make sure that the linker uses the start label used in your version of `cstartup.s`. For information about how to change the start label used by the linker, see *--entry*, page 309.

For Cortex-M, you must create a modified copy of `cstartup_M.s` or `cstartup_M.c` to use interrupts or other exception handlers.

THE DLIB LOW-LEVEL I/O INTERFACE

The runtime library uses a set of low-level functions—which are referred to as the *DLIB low-level I/O interface*—to communicate with the target system. Most of the low-level functions have no implementation.

For more information about this, see *Briefly about input and output (I/O)*, page 116.

This table lists the functions in the DLIB low-level I/O interface:

Function in DLIB low-level I/O interface	Defined in source file
abort	abort.c
__aeabi_assert	assert.c
clock	clock.c
__close	close.c
__exit	xxexit.c
getenv	getenv.c
__getzone	getzone.c
__lseek	lseek.c
__open	open.c
raise	raise.c
__read	read.c
remove	remove.c
rename	rename.c
signal	signal.c

Table 9: DLIB low-level I/O interface functions

Function in DLIB low-level I/O interface	Defined in source file
system	system.c
__time32, __time64	time.c
__write	write.c

Table 9: DLIB low-level I/O interface functions (Continued)

Note: You should not use the low-level functions prefixed with `__` directly in your application. Instead you should use the standard library functions that use these functions. For example, to write to `stdout`, you should use standard library functions like `printf` or `puts`, which in turn calls the low-level function `__write`. If you have forgot to implement a low-level function and your application will call that function via a standard library function, the linker will issues an error when you link in release build configuration.


Note: If you implement your own variants of the functions in this interface, your variants will be used even though you have enabled C-SPY emulated I/O, see *Briefly about C-SPY emulated I/O*, page 117.

abort

Source file	arm\src\lib\abort.c
Description	Standard C library function that aborts execution.
C-SPY debug action	Exits the application.
Default implementation	Calls <code>__exit(EXIT_FAILURE)</code> .
See also	<i>Briefly about retargeting</i> , page 118 <i>System termination</i> , page 136.

__aeabi_assert

Source file	arm\src\lib\assert.c
Description	Low-level function that handles a failed assert.
C-SPY debug action	Notifies the C-SPY debugger about the failed assert.

Default implementation	Failed asserts are reported by the function <code>__aeabi_assert</code> . By default, it prints an error message and calls <code>abort</code> . If this is not the behavior you require, you can implement your own version of the function. The assert macro is defined in the header file <code>assert.h</code> . To turn off assertions, define the symbol <code>NDEBUG</code> .  In the IDE, the symbol <code>NDEBUG</code> is by default defined in a Release project and <i>not</i> defined in a Debug project. If you build from the command line, you must explicitly define the symbol according to your needs. See <i>NDEBUG</i> , page 426.
See also	<i>Briefly about retargeting</i> , page 118.

clock

Source file	<code>arm\src\lib\clock.c</code>
Description	Standard C library function that accesses the processor time.
C-SPY debug action	Returns the clock on the host computer.
Default implementation	Returns <code>-1</code> to indicate that processor time is not available.
See also	<i>Briefly about retargeting</i> , page 118.

__close

Source file	<code>arm\src\lib\close.c</code>
Description	Low-level function that closes a file.
C-SPY debug action	Closes the associated host file on the host computer.
Default implementation	None.
See also	<i>Briefly about retargeting</i> , page 118.

__exit

Source file	<code>arm\src\lib__exit.c</code>
Description	Low-level function that halts execution.

C-SPY debug action	Notifies that the end of the application was reached.
Default implementation	Loops forever.
See also	<i>Briefly about retargeting</i> , page 118 <i>System termination</i> , page 136.

getenv

Source file	arm\src\lib\getenv.c arm\src\lib\environ.c
C-SPY debug action	Accesses the host environment.
Default implementation	<p>The <code>getenv</code> function in the library searches the string pointed to by the global variable <code>__environ</code>, for the key that was passed as argument. If the key is found, the value of it is returned, otherwise 0 (zero) is returned. By default, the string is empty.</p> <p>To create or edit keys in the string, you must create a sequence of null-terminated strings where each string has the format:</p> <pre>key=value\0</pre> <p>End the string with an extra null character (if you use a C string, this is added automatically). Assign the created sequence of strings to the <code>__environ</code> variable.</p> <p>For example:</p> <pre>const char MyEnv[] = "Key=Value\0Key2=Value2\0"; __environ = MyEnv;</pre> <p>If you need a more sophisticated environment variable handling, you should implement your own <code>getenv</code>, and possibly <code>putenv</code> function.</p> <p>Note: The <code>putenv</code> function is not required by the standard, and the library does not provide an implementation of it.</p>
See also	<i>Briefly about retargeting</i> , page 118.

__getzone

Source file	arm\src\lib\getzone.c
Description	Low-level function that returns the current time zone.

C-SPY debug action	Not applicable.
Default implementation	Returns " : ".
See also	<i>Briefly about retargeting</i> , page 118.

__lseek

Source file	<code>arm\src\lib\lseek.c</code>
Description	Low-level function for changing the location of the next access in an open file.
C-SPY debug action	Searches in the associated host file on the host computer.
Default implementation	None.
See also	<i>Briefly about retargeting</i> , page 118.

__open

Source file	<code>arm\src\lib\open.c</code>
Description	Low-level function that opens a file.
C-SPY debug action	Opens a file on the host computer.
Default implementation	None.
See also	<i>Briefly about retargeting</i> , page 118.

raise

Source file	<code>arm\src\lib\raise.c</code>
Description	Standard C library function that raises a signal.
C-SPY debug action	Not applicable.
Default implementation	Calls the signal handler for the raised signal, or terminates with call to <code>__exit(EXIT_FAILURE)</code> .

See also

Briefly about retargeting, page 118.

__read

Source file	arm\src\lib\read.c
Description	Low-level function that reads characters from <code>stdin</code> and from files.
C-SPY debug action	Directs <code>stdin</code> to the Terminal I/O window. All other files will read the associated host file.
Default implementation	None.

Example The code in this example uses memory-mapped I/O to read from a keyboard, whose port is assumed to be located at 0x1000:

```
#include <stddef.h>

__no_init volatile unsigned char kbIO @ 0x1000;

size_t __read(int handle,
              unsigned char *buf,
              size_t bufSize)
{
    size_t nChars = 0;

    /* Check for stdin
       (only necessary if FILE descriptors are enabled) */
    if (handle != 0)
    {
        return -1;
    }

    for (/*Empty*/; bufSize > 0; --bufSize)
    {
        unsigned char c = kbIO;
        if (c == 0)
            break;

        *buf++ = c;
        ++nChars;
    }

    return nChars;
}
```

For information about the handles associated with the streams, see *Retargeting—Adapting for your target system*, page 120.

For information about the @ operator, see *Controlling data and function placement in memory*, page 224.

See also *Briefly about retargeting*, page 118.

remove

Source file	arm\src\lib\remove.c
Description	Standard C library function that removes a file.
C-SPY debug action	Removes a file on the host computer.
Default implementation	Returns 0 to indicate success, but without removing a file.
See also	<i>Briefly about retargeting</i> , page 118.

rename

Source file	arm\src\lib\rename.c
Description	Standard C library function that renames a file.
C-SPY debug action	Renames a file on the host computer.
Default implementation	Returns -1 to indicate failure.
See also	<i>Briefly about retargeting</i> , page 118.

signal

Source file	arm\src\lib\signal.c
Description	Standard C library function that changes signal handlers.
C-SPY debug action	Not applicable.

Default implementation As specified by Standard C. You might want to modify this behavior if the environment supports some kind of asynchronous signals.

See also *Briefly about retargeting*, page 118.

system

Source file arm\src\lib\system.c

Description Standard C library function that executes commands.

C-SPY debug action Notifies the C-SPY debugger that `system` has been called and then returns `-1`.

Default implementation If you need to use the `system` function, you must implement it yourself. The `system` function available in the library returns `0` if a null pointer is passed to it to indicate that there is no command processor; otherwise it returns `-1` to indicate failure. If this is not the functionality that you require, you can implement your own version. This does not require that you rebuild the library.

See also *Briefly about retargeting*, page 118.

__time32, __time64

Source file arm\src\lib\time.c
arm\src\lib\time64.c

Description Low-level functions that return the current calendar time.

C-SPY debug action Returns the time on the host computer.

Default implementation Returns `-1` to indicate that calendar time is not available.

See also *Briefly about retargeting*, page 118.

__write

Source file arm\src\lib\write.c

Description Low-level function that writes to `stdout`, `stderr`, or a file.

C-SPY debug action	Directs <code>stdout</code> and <code>stderr</code> to the Terminal I/O window. All other files will write to the associated host file.
Default implementation	None.
Example	<p>The code in this example uses memory-mapped I/O to write to an LCD display, whose port is assumed to be located at address 0x1000:</p> <pre> #include <stddef.h> __no_init volatile unsigned char lcdIO @ 0x1000; size_t __write(int handle, const unsigned char *buf, size_t bufSize) { size_t nChars = 0; /* Check for the command to flush all handles */ if (handle == -1) { return 0; } /* Check for stdout and stderr (only necessary if FILE descriptors are enabled.) */ if (handle != 1 && handle != 2) { return -1; } for (/* Empty */; bufSize > 0; --bufSize) { lcdIO = *buf; ++buf; ++nChars; } return nChars; } </pre> <p>For information about the handles associated with the streams, see <i>Retargeting—Adapting for your target system</i>, page 120.</p>
See also	<i>Briefly about retargeting</i> , page 118.

CONFIGURATION SYMBOLS FOR PRINTF AND SCANF

If the provided formatters do not meet your requirements (*Formatters for printf*, page 128 and *Formatters for scanf*, page 129), you can customize the Full formatters. Note that this means you must rebuild the runtime library.

The default behavior of the `printf` and `scanf` formatters are defined by configuration symbols in the file `DLib_Defaults.h`.

These configuration symbols determine what capabilities the function `printf` should have:

Printf configuration symbols	Includes support for
_DLIB_PRINTF_MULTIBYTE	Multibyte characters
_DLIB_PRINTF_LONG_LONG	Long long (ll qualifier)
_DLIB_PRINTF_SPECIFIER_FLOAT	Floating-point numbers
_DLIB_PRINTF_SPECIFIER_A	Hexadecimal floating-point numbers
_DLIB_PRINTF_SPECIFIER_N	Output count (%n)
_DLIB_PRINTF_QUALIFIERS	Qualifiers h, l, L, v, t, and z
_DLIB_PRINTF_FLAGS	Flags -, +, #, and 0
_DLIB_PRINTF_WIDTH_AND_PRECISION	Width and precision
_DLIB_PRINTF_CHAR_BY_CHAR	Output char by char or buffered

Table 10: Descriptions of printf configuration symbols

When you build a library, these configurations determine what capabilities the function `scanf` should have:

Scanf configuration symbols	Includes support for
_DLIB_SCANF_MULTIBYTE	Multibyte characters
_DLIB_SCANF_LONG_LONG	Long long (ll qualifier)
_DLIB_SCANF_SPECIFIER_FLOAT	Floating-point numbers
_DLIB_SCANF_SPECIFIER_N	Output count (%n)
_DLIB_SCANF_QUALIFIERS	Qualifiers h, j, l, t, z, and L
_DLIB_SCANF_SCANSET	Scanset ([*])
_DLIB_SCANF_WIDTH	Width
_DLIB_SCANF_ASSIGNMENT_SUPPRESSING	Assignment suppressing ([*])

Table 11: Descriptions of scanf configuration symbols

To customize the formatting capabilities, you must;

- 1 Define the configuration symbols in the library configuration file according to your application requirements.
- 2 Rebuild the library, see *Customizing and building your own runtime library*, page 123.

CONFIGURATION SYMBOLS FOR FILE INPUT AND OUTPUT

File I/O is only supported by libraries with the Full library configuration, see *Runtime library configurations*, page 125, or in a customized library when the configuration symbol `__DLIB_FILE_DESCRIPTOR` is defined. If this symbol is not defined, functions taking a `FILE *` argument cannot be used.

To customize your library and rebuild it, see *Customizing and building your own runtime library*, page 123.

LOCALE

Locale is a part of the C language that allows language- and country-specific settings for several areas, such as currency symbols, date and time, and multibyte character encoding.

Depending on which library configuration you are using, you get different level of locale support. However, the more locale support, the larger your code will get. It is therefore necessary to consider what level of support your application needs. See *Runtime library configurations*, page 125.

The DLIB runtime library can be used in two main modes:

- With locale interface, which makes it possible to switch between different locales during runtime
- Without locale interface, where one selected locale is hardwired into the application.

Locale support in prebuilt libraries

The level of locale support in the prebuilt libraries depends on the library configuration.

- All prebuilt libraries support the C locale only
- All libraries with *full library configuration* have support for the locale interface. For prebuilt libraries with locale interface, it is by default only supported to switch multibyte character encoding at runtime.
- Libraries with *normal library configuration* do not have support for the locale interface.

If your application requires a different locale support, you must rebuild the library.

Customizing the locale support and rebuilding the library

If you decide to rebuild the library, you can choose between these locales:

- The C locale
- The POSIX locale
- A wide range of European locales.

Locale configuration symbols

The configuration symbol `_DLIB_FULL_LOCALE_SUPPORT`, which is defined in the library configuration file, determines whether a library has support for a locale interface or not. The locale configuration symbols `_LOCALE_USE_LANG_REGION` and `_ENCODING_USE_ENCODING` define all the supported locales and encodings:

```
#define _DLIB_FULL_LOCALE_SUPPORT 1
#define _LOCALE_USE_C           /* C locale */
#define _LOCALE_USE_EN_US      /* American English */
#define _LOCALE_USE_EN_GB      /* British English */
#define _LOCALE_USE_SV_SE      /* Swedish in Sweden */
```

See `DLib_Defaults.h` for a list of supported locale and encoding settings.

If you want to customize the locale support, you simply define the locale configuration symbols required by your application. For more information, see *Customizing and building your own runtime library*, page 123.

Note: If you use multibyte characters in your C or assembler source code, make sure that you select the correct locale symbol (the local host locale).

Building a library without support for locale interface

The locale interface is not included if the configuration symbol `_DLIB_FULL_LOCALE_SUPPORT` is set to 0 (zero). This means that a hardwired locale is used—by default the Standard C locale—but you can choose one of the supported locale configuration symbols. The `setlocale` function is not available and can therefore not be used for changing locales at runtime.

Building a library with support for locale interface

Support for the locale interface is obtained if the configuration symbol `_DLIB_FULL_LOCALE_SUPPORT` is set to 1. By default, the Standard C locale is used, but you can define as many configuration symbols as required. Because the `setlocale` function will be available in your application, it will be possible to switch locales at runtime.

Changing locales at runtime

The standard library function `setlocale` is used for selecting the appropriate portion of the application's locale when the application is running.

The `setlocale` function takes two arguments. The first one is a locale category that is constructed after the pattern `LC_CATEGORY`. The second argument is a string that describes the locale. It can either be a string previously returned by `setlocale`, or it can be a string constructed after the pattern:

lang_REGION

or

lang_REGION.encoding

The *lang* part specifies the language code, and the *REGION* part specifies a region qualifier, and *encoding* specifies the multibyte character encoding that should be used.

The *lang_REGION* part matches the `_LOCALE_USE_LANG_REGION` preprocessor symbols that can be specified in the library configuration file.

Example

This example sets the locale configuration symbols to Swedish to be used in Finland and UTF8 multibyte character encoding:

```
setlocale (LC_ALL, "sv_FI.Utf8");
```

Managing a multithreaded environment

In a multithreaded environment, the standard library must treat all library objects according to whether they are global or local to a thread. If an object is a true global object, any updates of its state must be guarded by a locking mechanism to make sure that only one thread can update it at any given time. If an object is local to a thread, the static variables containing the object state must reside in a variable area local to that thread. This area is commonly named *thread-local storage* (TLS).

The low-level implementations of locks and TLS are system-specific, and is not included in the DLIB runtime environment. If you are using an RTOS, check if it provides some or all of the required functions. Otherwise, you must provide your own.

MULTITHREAD SUPPORT IN THE DLIB RUNTIME ENVIRONMENT

The DLIB runtime environment uses two kinds of locks—*system locks* and *file stream locks*. The file stream locks are used as guards when the state of a file stream is updated,

and are only needed in the Full library configuration. The following objects are guarded with system locks:

- The heap (in other words when `malloc`, `new`, `free`, `delete`, `realloc`, or `calloc` is used).
- The file system (only available in the Full library configuration), but not the file streams themselves. The file system is updated when a stream is opened or closed, in other words when `fopen`, `fclose`, `fdopen`, `fflush`, or `freopen` is used.
- The signal system (in other words when `signal` is used).
- The temporary file system (in other words when `tmpnam` is used).
- Dynamically initialized function local objects with static storage duration.

These library objects use TLS:

Library objects using TLS	When these functions are used
Error functions	<code>errno</code> , <code>strerror</code>
Locale functions	<code>localeconv</code> , <code>setlocale</code>
Time functions	<code>asctime</code> , <code>localtime</code> , <code>gmtime</code> , <code>mktime</code>
Multibyte functions	<code>mbrlen</code> , <code>mbrtowc</code> , <code>mbsrtowc</code> , <code>mbtowc</code> , <code>wcrtomb</code> , <code>wcsrtomb</code> , <code>wctomb</code>
Rand functions	<code>rand</code> , <code>srand</code>
C++ exception engine	Not applicable
Miscellaneous functions	<code>atexit</code> , <code>strtok</code>

Table 12: Library objects using TLS

Note: If you are using `printf/scanf` (or any variants) with formatters, each individual formatter will be guarded, but the complete `printf/scanf` invocation will not be guarded.

If one of the C++ variants is used together with the DLIB runtime environment with multithread support, the compiler option `--guard_calls` must be used to make sure that function-static variables with dynamic initializers are not initialized simultaneously by several threads.

ENABLING MULTITHREAD SUPPORT

Prebuilt libraries with multithread support are provided in the product installation. See also, *Customizing and building your own runtime library*, page 123.

To configure multithread support for use with threaded applications:

- 1 To enable multithread support:



On the command line, use the linker option `--threaded_lib`.



In the IDE, choose **Project>Options>General Options>Library**

Configuration>Enable thread support in the library. The linker option `--threaded_lib` and the matching prebuilt library with thread support will automatically be used. If one of the C++ variants is used, the IDE will automatically use the compiler option `--guard_calls` to make sure that function-static variables with dynamic initializers are not initialized simultaneously by several threads.

- 2 To complement the built-in multithread support in the runtime library, you must also:
 - Implement code for the library's system locks interface, see *System locks interface*, page 152.
 - If file streams are used, implement code for the library's file stream locks interface or redirect the interface to the system locks interface (using the linker option `--redirect`), see *File streams locks interface*, page 153.
 - Implement code that handles thread creation, thread destruction, and TLS access methods for the library, see *TLS handling*, page 153.
 - Modify the linker configuration file accordingly.

See also *Lock usage*, page 153 for information about used file streams locks and system locks.

You can find the required declaration of functions and definitions of macros in the `DLib_Threads.h` file, which is included by `yvals.h`.

- 3 Build your project.

Note: If you are using a third-party RTOS, check their guidelines for how to enable multithread support with IAR Systems tools.

System locks interface

This interface must be fully implemented for system locks to work:

```
typedef void *__iar_Rmtx;                                /* Lock info object */

void __iar_system_Mtxinit(__iar_Rmtx *); /* Initialize a system
                                           lock */
void __iar_system_Mtxdst(__iar_Rmtx *); /* Destroy a system lock */
void __iar_system_Mtxlock(__iar_Rmtx *); /* Lock a system lock */
void __iar_system_Mtxunlock(__iar_Rmtx *); /* Unlock a system
                                           lock */
```


The lock and unlock implementation must survive nested calls.

File streams locks interface

This interface is only needed for the Full library configuration. If file streams are used, they can either be fully implemented or they can be redirected to the system locks interface. This interface must be implemented for file streams locks to work:

```
typedef void *__iar_Rmtx;                /* Lock info object */

void __iar_file_Mtxinit(__iar_Rmtx *); /* Initialize a file lock */
void __iar_file_Mtxdst(__iar_Rmtx *);  /* Destroy a file lock */
void __iar_file_Mtxlock(__iar_Rmtx *); /* Lock a file lock */
void __iar_file_Mtxunlock(__iar_Rmtx *); /* Unlock a file lock */
```

The lock and unlock implementation must survive nested calls.

Lock usage

The number of locks that the DLIB runtime environment assumes exist are:

- `_FOPEN_MAX`—the maximum number of file stream locks. These locks are only used in the Full library configuration, in other words only if both the macro symbols `_DLIB_FILE_DESCRIPTOR` and `_FILE_OP_LOCKS` are true.
- `_MAX_LOCK`—the maximum number of system locks.

Note that even if the application uses fewer locks, the runtime environment will initialize and destroy all the locks above.

For information about the initialization and destruction code, see `xsyslock.c`.

TLS handling

The DLIB runtime environment supports TLS memory areas for two types of threads: the *main thread* (the `main` function including the system startup and exit code) and *secondary threads*.

The main thread's TLS memory area:

- Is automatically created and initialized by your application's startup sequence
- Is automatically destructed by the application's destruct sequence
- Is located in the section `__DLIB_PERTHREAD`
- Exists also for non-threaded applications.

Each secondary thread's TLS memory area:

- Must be manually created and initialized
- Must be manually destructed

- Is located in a manually allocated memory area.

If you need the runtime library to support secondary threads, you must override the function:

```
void *__iar_dlib_perthread_access(void *symp);
```

The parameter is the address to the TLS variable to be accessed—in the main thread’s TLS area—and it should return the address to the symbol in the current TLS area.

Two interfaces can be used for creating and destroying secondary threads. You can use the following interface that allocates a memory area on the heap and initializes it. At deallocation, it destroys the objects in the area and then frees the memory.

```
void *__iar_dlib_perthread_allocate(void);  
void __iar_dlib_perthread_deallocate(void *);
```

Alternatively, if the application handles the TLS allocation, you can use this interface for initializing and destroying the objects in the memory area:

```
void __iar_dlib_perthread_initialize(void *);  
void __iar_dlib_perthread_destroy(void *);
```

These macros can be helpful when you implement an interface for creating and destroying secondary threads:

Macro	Description
__IAR_DLIB_PERTHREAD_SIZE	The size needed for the TLS memory area.
__IAR_DLIB_PERTHREAD_INIT_SIZE	The initializer size for the TLS memory area. You should initialize the rest of the TLS memory area, up to __IAR_DLIB_PERTHREAD_SIZE to zero.
__IAR_DLIB_PERTHREAD_SYMBOL_OFFSET(symbolptr)	The offset to the symbol in the TLS memory area.

Table 13: Macros for implementing TLS allocation

Note that the size needed for TLS variables depends on which resources in the DLIB runtime environment your application uses.

This is an example of how you can handle threads:

```
#include <yvals.h>

/* A thread's TLS pointer */
void _DLIB_TLS_MEMORY *TLSp;

/* Are we in a secondary thread? */
int InSecondaryThread = 0;

/* Allocate a thread-local TLS memory
   area and store a pointer to it in TLSp. */
void AllocateTLS()
{
    TLSp = __iar_dlib_perthread_allocate();
}

/* Deallocate the thread-local TLS memory area. */
void DeallocateTLS()
{
    __iar_dlib_perthread_deallocate(TLSp);
}

/* Access an object in the
   thread-local TLS memory area. */
void _DLIB_TLS_MEMORY *__iar_dlib_perthread_access(
    void _DLIB_TLS_MEMORY *sympb)
{
    char _DLIB_TLS_MEMORY *p = 0;
    if (InSecondaryThread)
        p = (char _DLIB_TLS_MEMORY *) TLSp;
    else
        p = (char _DLIB_TLS_MEMORY *)
            __segment_begin("__DLIB_PERTHREAD");

    p += __IAR_DLIB_PERTHREAD_SYMBOL_OFFSET(sympb);
    return (void _DLIB_TLS_MEMORY *) p;
}
```

The `TLSp` variable is unique for each thread, and must be exchanged by the RTOS or manually whenever a thread switch occurs.

TLS IN THE LINKER CONFIGURATION FILE

Normally, the linker automatically chooses how to initialize static data. If threads are used, the main thread's TLS memory area must be initialized by plain copying because

the initializers are used for each secondary thread's TLS memory area as well. This is controlled by the following statement in your linker configuration file:

```
initialize by copy with packing = none {section __DLIB_PERTHREAD };
```

Assembler language interface

- Mixing C and assembler
- Calling assembler routines from C
- Calling assembler routines from C++
- Calling convention
- Call frame information

Mixing C and assembler

The IAR C/C++ Compiler for ARM provides several ways to access low-level resources:

- Modules written entirely in assembler
- Intrinsic functions (the C alternative)
- Inline assembler.

It might be tempting to use simple inline assembler. However, you should carefully choose which method to use.

INTRINSIC FUNCTIONS

The compiler provides a few predefined functions that allow direct access to low-level processor operations without having to use the assembler language. These functions are known as intrinsic functions. They can be very useful in, for example, time-critical routines.

An intrinsic function looks like a normal function call, but it is really a built-in function that the compiler recognizes. The intrinsic functions compile into inline code, either as a single instruction, or as a short sequence of instructions.

For more information about the available intrinsic functions, see the chapter *Intrinsic functions*.

MIXING C AND ASSEMBLER MODULES

It is possible to write parts of your application in assembler and mix them with your C or C++ modules.

This causes some overhead in the form of function call and return instruction sequences, and the compiler will regard some registers as scratch registers. In many cases, the overhead of the extra instructions can be removed by the optimizer.

An important advantage is that you will have a well-defined interface between what the compiler produces and what you write in assembler. When using inline assembler, you will not have any guarantees that your inline assembler lines do not interfere with the compiler generated code.

When an application is written partly in assembler language and partly in C or C++, you are faced with several questions:

- How should the assembler code be written so that it can be called from C?
- Where does the assembler code find its parameters, and how is the return value passed back to the caller?
- How should assembler code call functions written in C?
- How are global C variables accessed from code written in assembler language?
- Why does not the debugger display the call stack when assembler code is being debugged?

The first question is discussed in the section *Calling assembler routines from C*, page 166. The following two are covered in the section *Calling convention*, page 169.

The answer to the final question is that the call stack can be displayed when you run assembler code in the debugger. However, the debugger requires information about the call frame, which must be supplied as annotations in the assembler source file. For more information, see *Call frame information*, page 176.

The recommended method for mixing C or C++ and assembler modules is described in *Calling assembler routines from C*, page 166, and *Calling assembler routines from C++*, page 168, respectively.

INLINE ASSEMBLER

Inline assembler can be used for inserting assembler instructions directly into a C or C++ function. Typically, this can be useful if you need to:

- Access hardware resources that are not accessible in C (in other words, when there is no definition for an SFR or there is no suitable intrinsic function available).
- Manually write a time-critical sequence of code that if written in C will not have the right timing.

- Manually write a speed-critical sequence of code that if written in C will be too slow.

An inline assembler statement is similar to a C function in that it can take input arguments (input operands), have return values (output operands), and read or write to C symbols (via the operands). An inline assembler statement can also declare *clobbered resources* (that is, values in registers and memory that have been overwritten).

Limitations

Most things you can do in normal assembler language are also possible with inline assembler, with the following differences:

- Alignment cannot be controlled; this means, for example, that `DC32` directives might be misaligned.
- The only accepted register synonyms are `SP` (for `R13`), `LR` (for `R14`), and `PC` (for `R15`).
- In general, assembler directives will cause errors or have no meaning. However, data definition directives will work as expected.
- Resources used (registers, memory, etc) that are also used by the C compiler must be declared as operands or clobbered resources.
- If you do not want to risk that the inline assembler statement to be optimized away by the compiler, you must declare it `volatile`.
- Accessing a C symbol or using a constant expression requires the use of operands.
- Dependencies between the expressions for the operands might result in an error.
- The pseudo-instruction `LDR Rd, =expr` is not available from inline assembler.

Risks with inline assembler

Without operands and clobbered resources, inline assembler statements have no interface with the surrounding C source code. This makes the inline assembler code fragile, and might also become a maintenance problem if you update the compiler in the future. There are also several limitations to using inline assembler without operands and clobbered resources:

- The compiler's various optimizations will disregard any effects of the inline statements, which will not be optimized at all.
- The inline assembler statement will be `volatile` and *clobbered memory* is not implied. This means that the compiler will not remove the assembler statement. It will simply be inserted at the given location in the program flow. The consequences or side-effects that the insertion might have on the surrounding code are not taken into consideration. If, for example, registers or memory locations are altered, they

might have to be restored within the sequence of inline assembler instructions for the rest of the code to work properly.



The following example (for ARM mode) demonstrates the risks of using the `asm` keyword without operands and clobbers:

```
int Add(int term1, int term2)
{
    asm("adds r0,r0,r1");
    return term1;
}
```

In this example:

- The function `Add` assumes that values are passed and returned in registers in a way that they might not always be, for example if the function is inlined.
- The `s` in the `adds` instruction implies that the condition flags are updated, which should be specified using the `cc` clobber operand. Otherwise, the compiler will assume that the condition flags are not modified.

Inline assembler without using operands or clobbered resources is therefore often best avoided.

Reference information for inline assembler

The `asm` and `__asm` keywords both insert inline assembler instructions. However, when you compile C source code, the `asm` keyword is not available when the option `--strict` is used. The `__asm` keyword is always available.

Syntax

The syntax of an inline assembler statement is (similar to the one used by GNU `gcc`):

```
asm [volatile]( string [assembler-interface])
```

string can contain one or more valid assembler instructions or data definition assembler directives, separated by `\n`.

For example:

```
asm("label:nop\n"
    "b label");
```

Note that you can define and use local labels in inline assembler instructions.

assembler-interface is:

```
: comma-separated list of output operands    /* optional */
: comma-separated list of input operands     /* optional */
: comma-separated list of clobbered resources /* optional */
```


Operands

An inline assembler statement can have one input and one output comma-separated list of operands. Each operand consists of an optional symbolic name in brackets, a quoted constraint, followed by a C expression in parentheses.

Syntax of operands

```
[ [ symbolic-name ] ] " [ modifiers ] constraint " ( expr )
```

For example:

```
int Add(int term1, int term2)
{
    int sum;

    asm("add %0,%1,%2"
        : "=r"(sum)
        : "r" (term1), "r" (term2));

    return sum;
}
```

In this example, the assembler instruction uses one output operand, `sum`, two input operands, `term1` and `term2`, and no clobbered resources.

It is possible to omit any list by leaving it empty. For example:

```
int matrix[M][N];

void MatrixPreloadRow(int row)
{
    asm volatile ("pld [%0]" : : "r" (&matrix[row][0]));
}
```

Operand constraints

Constraint	Description
<code>r</code>	Uses a general purpose register for the expression: R0–R12, R14 (for ARM and Thumb2) R0–R7 (for Thumb1)
<code>l</code>	R0–R7 (only valid for Thumb1)
<code>Rp</code>	Uses a pair of general purpose registers, for example R0, R1
<code>i</code>	An immediate integer operand with a constant value. Symbolic constants are allowed.
<code>j</code>	A 16-bit constant suitable for a MOVW instruction (valid for ARM and Thumb2).
<code>n</code>	An immediate operand, alias for <code>i</code> .

Table 14: Inline assembler operand constraints

Constraint	Description
I	An immediate constant valid for a data processing instruction (for ARM and Thumb2), or a constant in the range 0 to 255 (for Thumb1).
J	An immediate constant in the range -4095 to 4095 (for ARM and Thumb2), or a constant in the range -255 to -1 (for Thumb1).
K	An immediate constant that satisfies the I constraint if inverted (for ARM and Thumb2), or a constant that satisfies the I constraint multiplied by any power of 2 (for Thumb1).
L	An immediate constant that satisfies the I constraint if negated (for ARM and Thumb2), or a constant in the range -7 to 7 (for Thumb1).
M	An immediate constant that is a multiple of 4 in the range 0 to 1020 (only valid for Thumb1).
N	An immediate constant in the range 0 to 31 (only valid for Thumb1).
O	An immediate constant that is a multiple of 4 in the range -508 to 508 (only valid for Thumb1).
t	An S register.
w	A D register.
q	A Q register.
Dv	A 32-bit floating-point immediate constant for the VMOV.F32 instruction.
Dy	A 64-bit floating-point immediate constant for the VMOV.F64 instruction.
v2S ... v4Q	A vector of 2, 3, or 4 consecutive S, D, or Q registers. For example, v4Q is a vector of four Q registers. The vectors do not overlap, so the available v4Q register vectors are Q0-Q3, Q4-Q7, Q8-Q11, and Q12-Q15.

Table 14: Inline assembler operand constraints (Continued)

Constraint modifiers

Constraint modifiers can be used together with a constraint to modify its meaning. This table lists the supported constraint modifiers:

Modifier	Description
=	Write-only operand
+	Read-write operand
&	Early clobber output operand which is written to before the instruction has processed all the input operands.

Table 15: Supported constraint modifiers

Referring to operands

Assembler instructions refer to operands by prefixing their order number with %. The first operand has order number 0 and is referred to by %0.

If the operand has a symbolic name, you can refer to it using the syntax `%[operand.name]`. Symbolic operand names are in a separate namespace from C/C++ code and can be the same as a C/C++ variable names. Each operand name must however be unique in each assembler statement. For example:

```
int Add(int term1, int term2)
{
    int sum;

    asm("add %[Rd], %[Rn], %[Rm] "
        : [Rd] "=r" (sum)
        : [Rn] "r" (term1), [Rm] "r" (term2));

    return sum;
}
```

Operand modifiers

An operand modifier is a single letter between the % and the operand number, which is used for transforming the operand.

In the example below, the modifiers `L` and `H` are used for accessing the least and most significant 16 bits, respectively, of an immediate operand:

```
int Mov32()
{
    int a;
    asm("movw %0, %L1 \n"
        "movt %0, %H1 \n" : "=r" (a) : "i" (0x12345678UL));
    return a;
}
```

Some operand modifiers can be combined, in which case each letter will transform the result from the previous modifier. This table describes the transformation performed by each valid modifier:

Modifier	Description
L	The lowest-numbered register of a register pair, or the low 16 bits of an immediate constant.
H	The highest-numbered register of a register pair, or the high 16 bits of an immediate constant.
C	For an immediate operand, an integer or symbol address without a preceding # sign. Cannot be transformed by additional operand modifiers.

Table 16: Operand modifiers and transformations

Modifier	Description
B	For an immediate operand, the bitwise inverse of integer or symbol without a preceding # sign. Cannot be transformed by additional operand modifiers.
Q	The least significant register of a register pair.
R	The most significant register of a register pair.
M	For a register or a register pair, the register list suitable for <code>ldm</code> or <code>stm</code> . Cannot be transformed by additional operand modifiers.
a	Transforms a register <code>Rn</code> into a memory operand <code>[Rn, #0]</code> suitable for <code>pld</code> .
b	The low <code>S</code> register part of a <code>D</code> register.
p	The high <code>S</code> register part of a <code>D</code> register.
e	The low <code>D</code> register part of a <code>Q</code> register, or the low register in a vector of Neon registers.
f	The high <code>D</code> register part of a <code>Q</code> register, or the high register in a vector of Neon registers.
h	For a (vector of) <code>D</code> or <code>Q</code> registers, the corresponding list of <code>D</code> registers within curly braces. For example, <code>Q0</code> becomes <code>{D0, D1}</code> . Cannot be transformed by additional operand modifiers.
y	<code>S</code> register as indexed <code>D</code> register, for example <code>S7</code> becomes <code>D3[1]</code> . Cannot be transformed by additional operand modifiers.

Table 16: Operand modifiers and transformations

Input operands	<p>Input operands cannot have any modifiers, but they can have any valid C expression as long as the type of the expression fits the register.</p> <p>The C expression will be evaluated just before any of the assembler instructions in the inline assembler statement and assigned to the constraint, for example a register.</p>
Output operands	<p>Output operands must have <code>=</code> as a modifier and the C expression must be an l-value and specify a writable location. For example, <code>=r</code> for a write-only general purpose register. The constraint will be assigned to the evaluated C expression (as an l-value) immediately after the last assembler instruction in the inline assembler statement. Input operands are assumed to be consumed before output is produced and the compiler may use the same register for an input and output operand. To prohibit this, prefix the output constraint with <code>&</code> to make it an early clobber resource, for example <code>=&r</code>. This will ensure that the output operand will be allocated in a different register than the input operands.</p>
Input/output operands	<p>An operand that should be used both for input and output must be listed as an output operand and have the <code>+</code> modifier. The C expression must be an l-value and specify a</p>

writable location. The location will be read immediately before any assembler instructions and it will be written to right after the last assembler instruction.

This is an example of using a read-write operand:

```
int Double(int value)
{
    asm("add %0,%0,%0" : "+r"(value));

    return value;
}
```

In the example above, the input value for `value` will be placed in a general purpose register. After the assembler statement, the result from the `ADD` instruction will be placed in the same register.

Clobbered resources

An inline assembler statement can have a list of clobbered resources.

```
"resource1", "resource2", ...
```

Specify clobbered resources to inform the compiler about which resources the inline assembler statement destroys. Any value that resides in a clobbered resource and that is needed after the inline assembler statement will be reloaded.

Clobbered resources will not be used as input or output operands.

This is an example of how to use clobbered resources:

```
int Add(int term1, int term2)
{
    int sum;

    asm("adds %0,%1,%2"
        : "=r"(sum)
        : "r" (term1), "r" (term2)
        : "cc");

    return sum;
}
```

In this example the condition codes will be modified by the `ADDS` instruction. Therefore, `"cc"` must be listed in the clobber list.

This table lists valid clobbered resources:

Clobber	Description
R0-R12, R14 for ARM mode and Thumb2	General purpose registers
R0-R7, R12, R14 for Thumb1	

Table 17: List of valid clobbers

Clobber	Description
S0-S31, D0-D31, Q0-Q15	Floating-point registers
cc	The condition flags (N, Z, V, and C)
memory	To be used if the instructions modify any memory. This will avoid keeping memory values cached in registers across the inline assembler statement.

Table 17: List of valid clobbers (Continued)

AN EXAMPLE OF HOW TO USE CLOBBERED MEMORY

```
int StoreExclusive(unsigned long * location, unsigned long value)
{
    int failed;

    asm("strex %0,%2,[%1]"
        : "=&r"(failed)
        : "r"(location), "r"(value)
        : "memory");

    /* Note: 'strex' requires ARMv6 (ARM) or ARMv6T2 (THUMB) */

    return failed;
}
```

Calling assembler routines from C

An assembler routine that will be called from C must:

- Conform to the calling convention
- Have a PUBLIC entry-point label
- Be declared as external before any call, to allow type checking and optional promotion of parameters, as in these examples:

```
extern int foo(void);
or
extern int foo(int i, int j);
```

One way of fulfilling these requirements is to create skeleton code in C, compile it, and study the assembler list file.

CREATING SKELETON CODE

The recommended way to create an assembler language routine with the correct interface is to start with an assembler language source file created by the C compiler. Note that you must create skeleton code for each function prototype.

The following example shows how to create skeleton code to which you can easily add the functional body of the routine. The skeleton source code only needs to declare the variables required and perform simple accesses to them. In this example, the assembler routine takes an `int` and a `char`, and then returns an `int`:

```
extern int gInt;
extern char gChar;

int Func(int arg1, char arg2)
{
    int locInt = arg1;
    gInt = arg1;
    gChar = arg2;
    return locInt;
}

int main()
{
    int locInt = gInt;
    gInt = Func(locInt, gChar);
    return 0;
}
```

Note: In this example we use a low optimization level when compiling the code to show local and global variable access. If a higher level of optimization is used, the required references to local variables could be removed during the optimization. The actual function declaration is not changed by the optimization level.

COMPILING THE SKELETON CODE



In the IDE, specify list options on file level. Select the file in the workspace window. Then choose **Project>Options**. In the **C/C++ Compiler** category, select **Override inherited settings**. On the **List** page, deselect **Output list file**, and instead select the **Output assembler file** option and its suboption **Include source**. Also, be sure to specify a low level of optimization.



Use these options to compile the skeleton code:

```
icccarm skeleton.c -lA . -On -e
```

The `-lA` option creates an assembler language output file including C or C++ source lines as assembler comments. The `.` (period) specifies that the assembler file should be

named in the same way as the C or C++ module (`skeleton`), but with the filename extension `s`. The `-On` option means that no optimization will be used and `-e` enables language extensions. In addition, make sure to use relevant compiler options, usually the same as you use for other C or C++ source files in your project.

The result is the assembler source output file `skeleton.s`.

Note: The `-lA` option creates a list file containing call frame information (CFI) directives, which can be useful if you intend to study these directives and how they are used. If you only want to study the calling convention, you can exclude the CFI directives from the list file.



In the IDE, choose **Project>Options>C/C++ Compiler>List** and deselect the suboption **Include call frame information**.



On the command line, use the option `-lB` instead of `-lA`. Note that CFI information must be included in the source code to make the C-SPY **Call Stack** window work.

The output file

The output file contains the following important information:

- The calling convention
- The return values
- The global variables
- The function parameters
- How to create space on the stack (auto variables)
- Call frame information (CFI).

The CFI directives describe the call frame information needed by the **Call Stack** window in the debugger. For more information, see *Call frame information*, page 176.

Calling assembler routines from C++

The C calling convention does not apply to C++ functions. Most importantly, a function name is not sufficient to identify a C++ function. The scope and the type of the function are also required to guarantee type-safe linkage, and to resolve overloading.

Another difference is that non-static member functions get an extra, hidden argument, the `this` pointer.

However, when using C linkage, the calling convention conforms to the C calling convention. An assembler routine can therefore be called from C++ when declared in this manner:

```
extern "C"
{
    int MyRoutine(int);
}
```

The following example shows how to achieve the equivalent to a non-static member function, which means that the implicit `this` pointer must be made explicit. It is also possible to “wrap” the call to the assembler routine in a member function. Use an inline member function to remove the overhead of the extra call—this assumes that function inlining is enabled:

```
class MyClass;

extern "C"
{
    void DoIt(MyClass *ptr, int arg);
}

class MyClass
{
public:
    inline void DoIt(int arg)
    {
        ::DoIt(this, arg);
    }
};
```

Calling convention

A calling convention is the way a function in a program calls another function. The compiler handles this automatically, but, if a function is written in assembler language, you must know where and how its parameters can be found, how to return to the program location from where it was called, and how to return the resulting value.

It is also important to know which registers an assembler-level routine must preserve. If the program preserves too many registers, the program might be ineffective. If it preserves too few registers, the result would be an incorrect program.

This section describes the calling convention used by the compiler. These items are examined:

- Function declarations
- C and C++ linkage

- Preserved versus scratch registers
- Function entrance
- Function exit
- Return address handling

At the end of the section, some examples are shown to describe the calling convention in practice.

The calling convention used by the compiler adheres to the Procedure Call Standard for the ARM architecture, AAPCS, a part of AEABI; see *AEABI compliance*, page 217. AAPCS is not fully described here. For example, the use of floating-point coprocessor registers when using the VFP calling convention is not covered.

FUNCTION DECLARATIONS

In C, a function must be declared in order for the compiler to know how to call it. A declaration could look as follows:

```
int MyFunction(int first, char * second);
```

This means that the function takes two parameters: an integer and a pointer to a character. The function returns a value, an integer.

In the general case, this is the only knowledge that the compiler has about a function. Therefore, it must be able to deduce the calling convention from this information.

USING C LINKAGE IN C++ SOURCE CODE

In C++, a function can have either C or C++ linkage. To call assembler routines from C++, it is easiest if you make the C++ function have C linkage.

This is an example of a declaration of a function with C linkage:

```
extern "C"
{
    int F(int);
}
```

It is often practical to share header files between C and C++. This is an example of a declaration that declares a function with C linkage in both C and C++:

```
#ifdef __cplusplus
extern "C"
{
#endif

int F(int);

#ifdef __cplusplus
}
#endif
```

PRESERVED VERSUS SCRATCH REGISTERS

The general ARM CPU registers are divided into three separate sets, which are described in this section.

Scratch registers

Any function is permitted to destroy the contents of a scratch register. If a function needs the register value after a call to another function, it must store it during the call, for example on the stack.

Any of the registers R0 to R3, and R12, can be used as a scratch register by the function. Note that R12 is a scratch register also when calling between assembler functions only because of automatically inserted instructions for veneers.

Preserved registers

Preserved registers, on the other hand, are preserved across function calls. The called function can use the register for other purposes, but must save the value before using the register and restore it at the exit of the function.

The registers R4 through to R11 are preserved registers. They are preserved by the called function.

Special registers

For some registers, you must consider certain prerequisites:

- The stack pointer register, R13/SP, must at all times point to or below the last element on the stack. In the eventuality of an interrupt, everything below the point the stack pointer points to, can be destroyed. At function entry and exit, the stack pointer must be 8-byte aligned. In the function, the stack pointer must always be word aligned. At exit, SP must have the same value as it had at the entry.

- The register R15/PC is dedicated for the Program Counter.
- The link register, R14/LR, holds the return address at the entrance of the function.

FUNCTION ENTRANCE

Parameters can be passed to a function using one of these basic methods:

- In registers
- On the stack

It is much more efficient to use registers than to take a detour via memory, so the calling convention is designed to use registers as much as possible. Only a limited number of registers can be used for passing parameters; when no more registers are available, the remaining parameters are passed on the stack. These exceptions to the rules apply:

- Interrupt functions cannot take any parameters, except software interrupt functions that accept parameters and have return values
- Software interrupt functions cannot use the stack in the same way as ordinary functions. When an SVC instruction is executed, the processor switches to supervisor mode where the supervisor stack is used. Arguments can therefore not be passed on the stack if your application is not running in supervisor mode previous to the interrupt.

Hidden parameters

In addition to the parameters visible in a function declaration and definition, there can be hidden parameters:

- If the function returns a structure larger than 32 bits, the memory location where the structure is to be stored is passed as an extra parameter. Notice that it is always treated as the *first parameter*.
- If the function is a non-static C++ member function, then the `this` pointer is passed as the first parameter (but placed after the return structure pointer, if there is one). For more information, see *Calling assembler routines from C*, page 166.

Register parameters

The registers available for passing parameters are:

Parameters	Passed in registers
Scalar and floating-point values no larger than 32 bits, and single-precision (32-bits) floating-point values	Passed using the first free register: R0–R3

Table 18: Registers used for passing parameters

Parameters	Passed in registers
long long and double-precision (64-bit) values	Passed in first available register pair: R0 : R1, or R2 : R3

Table 18: Registers used for passing parameters


The assignment of registers to parameters is a straightforward process. Traversing the parameters from left to right, the first parameter is assigned to the available register or registers. Should there be no more available registers, the parameter is passed on the stack in reverse order.

When functions that have parameters smaller than 32 bits are called, the values are sign or zero extended to ensure that the unused bits have consistent values. Whether the values will be sign or zero extended depends on their type—signed or unsigned.

Stack parameters and layout

Stack parameters are stored in memory, starting at the location pointed to by the stack pointer. Below the stack pointer (towards low memory) there is free space that the called function can use. The first stack parameter is stored at the location pointed to by the stack pointer. The next one is stored at the next location on the stack that is divisible by four, etc. It is the responsibility of the caller to clean the stack after the called function has returned.

This figure illustrates how parameters are stored on the stack:



FUNCTION EXIT

A function can return a value to the function or program that called it, or it can have the return type `void`.

The return value of a function, if any, can be scalar (such as integers and pointers), floating-point, or a structure.

Registers used for returning values

The registers available for returning values are R0 and R0:R1.

Return values	Passed in registers
Scalar and structure return values no larger than 32 bits, and single-precision (32-bit) floating-point return values	R0
The memory address of a structure return value larger than 32 bits	R0
long long and double-precision (64-bit) return values	R0:R1

Table 19: Registers used for returning values

If the returned value is smaller than 32 bits, the value is sign- or zero-extended to 32 bits.

Stack layout at function exit

It is the responsibility of the caller to clean the stack after the called function has returned.

Return address handling

A function written in assembler language should, when finished, return to the caller, by jumping to the address pointed to by the register LR.

At function entry, non-scratch registers and the LR register can be pushed with one instruction. At function exit, all these registers can be popped with one instruction. The return address can be popped directly to PC.

The following example shows what this can look like:

```
name      call
section   .text:CODE
extern    func

push      {r4-r6,lr}    ; Preserve stack alignment 8
bl        func

; Do something here.

pop       {r4-r6,pc}    ; return

end
```

EXAMPLES

The following section shows a series of declaration examples and the corresponding calling conventions. The complexity of the examples increases toward the end.

Example 1

Assume this function declaration:

```
int add1(int);
```

This function takes one parameter in the register R0, and the return value is passed back to its caller in the register R0.

This assembler routine is compatible with the declaration; it will return a value that is one number higher than the value of its parameter:

```
name      return
section   .text:CODE
add       r0, r0, #1
bx        lr
end
```

Example 2

This example shows how structures are passed on the stack. Assume these declarations:

```
struct MyStruct
{
    short a;
    short b;
    short c;
    short d;
    short e;
};
```

```
int MyFunction(struct MyStruct x, int y);
```

The values of the structure members a, b, c, and d are passed in registers R0–R3. The last structure member e and the integer parameter y are passed on the stack. The calling function must reserve eight bytes on the top of the stack and copy the contents of the two stack parameters to that location. The return value is passed back to its caller in the register R0.

Example 3

The function below will return a structure of type `struct MyStruct`.

```
struct MyStruct
{
    int mA[20];
};
```

```
struct MyStruct MyFunction(int x);
```

It is the responsibility of the calling function to allocate a memory location for the return value and pass a pointer to it as a hidden first parameter. The pointer to the location where the return value should be stored is passed in `R0`. The parameter `x` is passed in `R1`.

Assume that the function instead was declared to return a pointer to the structure:

```
struct MyStruct *MyFunction(int x);
```

In this case, the return value is a scalar, so there is no hidden parameter. The parameter `x` is passed in `R0`, and the return value is returned in `R0`.

Call frame information

When you debug an application using C-SPY, you can view the *call stack*, that is, the chain of functions that called the current function. To make this possible, the compiler supplies debug information that describes the layout of the call frame, in particular information about where the return address is stored.

If you want the call stack to be available when debugging a routine written in assembler language, you must supply equivalent debug information in your assembler source using the assembler directive `CFI`. This directive is described in detail in the *IAR Assembler User Guide for ARM*.

CFI DIRECTIVES

The `CFI` directives provide C-SPY with information about the state of the calling function(s). Most important of this is the return address, and the value of the stack pointer at the entry of the function or assembler routine. Given this information, C-SPY can reconstruct the state for the calling function, and thereby unwind the stack.

A full description about the calling convention might require extensive call frame information. In many cases, a more limited approach will suffice.

When describing the call frame information, the following three components must be present:

- A *names block* describing the available resources to be tracked

- A *common block* corresponding to the calling convention
- A *data block* describing the changes that are performed on the call frame. This typically includes information about when the stack pointer is changed, and when permanent registers are stored or restored on the stack.

Resource	Description
CFA R13	The call frames of the stack
R0–R12	Processor general-purpose 32-bit registers
R13	Stack pointer, SP
R14	Link register, LR
D0–D31	Vector Floating Point (VFP) 64-bit coprocessor register
CPSR	Current program status register
SPSR	Saved program status register

Table 20: Call frame information resources defined in a names block

CREATING ASSEMBLER SOURCE WITH CFI SUPPORT

The recommended way to create an assembler language routine that handles call frame information correctly is to start with an assembler language source file created by the compiler.

- 1 Start with suitable C source code, for example:

```
int F(int);
int cfiExample(int i)
{
    return i + F(i);
}
```

- 2 Compile the C source code, and make sure to create a list file that contains call frame information—the CFI directives.



On the command line, use the option -lA.



In the IDE, choose **Project>Options>C/C++ Compiler>List** and make sure the suboption **Include call frame information** is selected.

For the source code in this example, the list file looks like this:

```

NAME Cfi

RTMODEL "__SystemLibrary", "DLib"

EXTERN F

PUBLIC cfiExample

CFI Names cfiNames0
CFI StackFrame CFA R13 DATA
CFI Resource R0:32, R1:32, R2:32, R3:32, R4:32, R5:32,
R6:32, R7:32
CFI Resource R8:32, R9:32, R10:32, R11:32, R12:32,
R13:32, R14:32
CFI EndNames cfiNames0

CFI Common cfiCommon0 Using cfiNames0
CFI CodeAlign 4
CFI DataAlign 4
CFI ReturnAddress R14 CODE
CFI CFA R13+0
CFI R0 Undefined
CFI R1 Undefined
CFI R2 Undefined
CFI R3 Undefined
CFI R4 SameValue
CFI R5 SameValue
CFI R6 SameValue
CFI R7 SameValue
CFI R8 SameValue
CFI R9 SameValue
CFI R10 SameValue
CFI R11 SameValue
CFI R12 Undefined
CFI R14 SameValue
CFI EndCommon cfiCommon0

SECTION `.text`:CODE:NOROOT(2)
CFI Block cfiBlock0 Using cfiCommon0
CFI Function cfiExample
ARM
cfiExample:
    PUSH        {R4,LR}
    CFI R14 Frame(CFA, -4)
    CFI R4 Frame(CFA, -8)

```

```

CFI CFA R13+8
MOVS    R4,R0
MOVS    R0,R4
BL      F
ADDS    R0,R0,R4
POP     {R4,PC}    ;; return
CFI EndBlock cfiBlock0

END

```

Note: The header file `Common.i` contains the macros `CFI_NAMES_BLOCK`, `CFI_COMMON_ARM`, and `CFI_COMMON_Thumb`, which declare a typical names block and a typical common block. These two macros declare several resources, both concrete and virtual.

Using C

- C language overview
- Extensions overview
- IAR C language extensions

C language overview

The IAR C/C++ Compiler for ARM supports the ISO/IEC 9899:1999 standard (including up to technical corrigendum No.3), also known as C99. In this guide, this standard is referred to as *Standard C* and is the default standard used in the compiler. This standard is stricter than C89.

In addition, the compiler also supports the ISO 9899:1990 standard (including all technical corrigenda and addenda), also known as C94, C90, C89, and ANSI C. In this guide, this standard is referred to as *C89*. Use the `--c89` compiler option to enable this standard.

The C99 standard is derived from C89, but adds features like these:

- The `inline` keyword advises the compiler that the function defined immediately after the keyword should be inlined
- Declarations and statements can be mixed within the same scope
- A declaration in the initialization expression of a `for` loop
- The `bool` data type
- The `long long` data type
- The `complex` floating-point type
- C++ style comments
- Compound literals
- Incomplete arrays at the end of structs
- Hexadecimal floating-point constants
- Designated initializers in structures and arrays
- The preprocessor operator `_Pragma()`
- Variadic macros, which are the preprocessor macro equivalents of `printf` style functions
- VLA (variable length arrays) must be explicitly enabled with the compiler option `--vla`

- Inline assembler using the `asm` or the `__asm` keyword, see .

Note: Even though it is a C99 feature, the IAR C/C++ Compiler for ARM does not support UCNs (universal character names).

Extensions overview

The compiler offers the features of Standard C and a wide set of extensions, ranging from features specifically tailored for efficient programming in the embedded industry to the relaxation of some minor standards issues.

This is an overview of the available extensions:

- IAR C language extensions

For information about available language extensions, see *IAR C language extensions*, page 183. For more information about the extended keywords, see the chapter *Extended keywords*. For information about C++, the two levels of support for the language, and C++ language extensions; see the chapter *Using C++*.

- Pragma directives

The `#pragma` directive is defined by Standard C and is a mechanism for using vendor-specific extensions in a controlled way to make sure that the source code is still portable.

The compiler provides a set of predefined pragma directives, which can be used for controlling the behavior of the compiler, for example how it allocates memory, whether it allows extended keywords, and whether it outputs warning messages. Most pragma directives are preprocessed, which means that macros are substituted in a pragma directive. The pragma directives are always enabled in the compiler. For several of them there is also a corresponding C/C++ language extension. For information about available pragma directives, see the chapter *Pragma directives*.

- Preprocessor extensions

The preprocessor of the compiler adheres to Standard C. The compiler also makes several preprocessor-related extensions available to you. For more information, see the chapter *The preprocessor*.

- Intrinsic functions

The intrinsic functions provide direct access to low-level processor operations and can be very useful in, for example, time-critical routines. The intrinsic functions compile into inline code, either as a single instruction or as a short sequence of instructions. For more information about using intrinsic functions, see *Mixing C and assembler*, page 157. For information about available functions, see the chapter *Intrinsic functions*.

- Library functions

The DLIB runtime environment provides the C and C++ library definitions in the C/C++ standard library that apply to embedded systems. For more information, see *DLIB runtime environment—implementation details*, page 429.

Note: Any use of these extensions, except for the pragma directives, makes your source code inconsistent with Standard C.

ENABLING LANGUAGE EXTENSIONS

You can choose different levels of language conformance by means of project options:

Command line	IDE*	Description
--strict	Strict	All IAR C language extensions are disabled; errors are issued for anything that is not part of Standard C.
None	Standard	All <i>extensions to Standard C</i> are enabled, but no <i>extensions for embedded systems programming</i> . For information about extensions, see <i>IAR C language extensions</i> , page 183.
-e	Standard with IAR extensions	All IAR C language extensions are enabled.

Table 21: Language extensions

* In the IDE, choose **Project>Options>C/C++ Compiler>Language>Language conformance** and select the appropriate option. Note that language extensions are enabled by default.

IAR C language extensions

The compiler provides a wide set of C language extensions. To help you to find the extensions required by your application, they are grouped like this in this section:

- *Extensions for embedded systems programming*—extensions specifically tailored for efficient embedded programming for the specific core you are using, typically to meet memory restrictions
- *Relaxations to Standard C*—that is, the relaxation of some minor Standard C issues and also some useful but minor syntax extensions, see *Relaxations to Standard C*, page 186.

EXTENSIONS FOR EMBEDDED SYSTEMS PROGRAMMING

The following language extensions are available both in the C and the C++ programming languages and they are well suited for embedded systems programming:

- Type attributes and object attributes
For information about the related concepts, the general syntax rules, and for reference information, see the chapter *Extended keywords*.
- Placement at an absolute address or in a named section
The @ operator or the directive #pragma location can be used for placing global and static variables at absolute addresses, or placing a variable or function in a named section. For more information about using these features, see *Controlling data and function placement in memory*, page 224, and *location*, page 370.
- Alignment control
Each data type has its own alignment; for more information, see *Alignment*, page 329. If you want to change the alignment, the __packed data type attribute, the #pragma pack, and the #pragma data_alignment directives are available. If you want to check the alignment of an object, use the __ALIGNOF__() operator.
The __ALIGNOF__ operator is used for accessing the alignment of an object. It takes one of two forms:
 - __ALIGNOF__ (type)
 - __ALIGNOF__ (expression)
 In the second form, the expression is not evaluated.
- Anonymous structs and unions
C++ includes a feature called anonymous unions. The compiler allows a similar feature for both structs and unions in the C programming language. For more information, see *Anonymous structs and unions*, page 223.
- Bitfields and non-standard types
In Standard C, a bitfield must be of the type int or unsigned int. Using IAR C language extensions, any integer type or enumeration can be used. The advantage is that the struct will sometimes be smaller. For more information, see *Bitfields*, page 332.
- static_assert()
The construction static_assert(const-expression, "message"); can be used in C/C++. The construction will be evaluated at compile time and if const-expression is false, a message will be issued including the message string.

- Parameters in variadic macros

Variadic macros are the preprocessor macro equivalents of `printf` style functions. The preprocessor accepts variadic macros with no arguments, which means if no parameter matches the `...` parameter, the comma is then deleted in the `" , ##__VA_ARGS__ "` macro definition. According to Standard C, the `...` parameter must be matched with at least one argument.

Dedicated section operators

The compiler supports getting the start address, end address, and size for a section with these built-in section operators:

<code>__section_begin</code>	Returns the address of the first byte of the named section or block.
<code>__section_end</code>	Returns the address of the first byte <i>after</i> the named section or block.
<code>__section_size</code>	Returns the size of the named section or block in bytes.

Note: The aliases `__segment_begin/__sfb`, `__segment_end/__sfe`, and `__segment_size/__sfs` can also be used.

The operators can be used on named sections or on named blocks defined in the linker configuration file.

These operators behave syntactically as if declared like:

```
void * __section_begin(char const * section)
void * __section_end(char const * section)
size_t __section_size(char const * section)
```

When you use the `@` operator or the `#pragma location` directive to place a data object or a function in a user-defined section, or when you use named blocks in the linker configuration file, the section operators can be used for getting the start and end address of the memory range where the sections or blocks were placed.

The named *section* must be a string literal and it must have been declared earlier with the `#pragma section` directive. The type of the `__section_begin` operator is a pointer to `void`. Note that you must enable language extensions to use these operators.

The operators are implemented in terms of *symbols* with dedicated names, and will appear in the linker map file under these names:

Operator	Symbol
<code>__section_begin(sec)</code>	<code>sec\$\$Base</code>
<code>__section_end(sec)</code>	<code>sec\$\$Limit</code>
<code>__section_size(sec)</code>	<code>sec\$\$Length</code>

Table 22: Section operators and their symbols

Note that the linker will not necessarily place sections with the same name consecutively when these operators are not used. Using one of these operators (or the equivalent symbols) will cause the linker to behave as if the sections were in a named block. This is to assure that the sections are placed consecutively, so that the operators can be assigned meaningful values. If this is in conflict with the section placement as specified in the linker configuration file, the linker will issue an error.

Example

In this example, the type of the `__section_begin` operator is `void *`.

```
#pragma section="MYSECTION"
...
section_start_address = __section_begin("MYSECTION");
```

See also *section*, page 376, and *location*, page 370.

RELAXATIONS TO STANDARD C

This section lists and briefly describes the relaxation of some Standard C issues and also some useful but minor syntax extensions:

- Arrays of incomplete types
An array can have an incomplete `struct`, `union`, or `enum` type as its element type. The types must be completed before the array is used (if it is), or by the end of the compilation unit (if it is not).
- Forward declaration of `enum` types
The extensions allow you to first declare the name of an `enum` and later resolve it by specifying the brace-enclosed list.
- Accepting missing semicolon at the end of a `struct` or `union` specifier
A warning—instead of an error—is issued if the semicolon at the end of a `struct` or `union` specifier is missing.
- Null and `void`
In operations on pointers, a pointer to `void` is always implicitly converted to another type if necessary, and a null pointer constant is always implicitly converted to a null

pointer of the right type if necessary. In Standard C, some operators allow this kind of behavior, while others do not allow it.

- Casting pointers to integers in static initializers

In an initializer, a pointer constant value can be cast to an integral type if the integral type is large enough to contain it. For more information about casting pointers, see *Casting*, page 338.

- Taking the address of a register variable

In Standard C, it is illegal to take the address of a variable specified as a register variable. The compiler allows this, but a warning is issued.

- `long float` means `double`

The type `long float` is accepted as a synonym for `double`.

- Repeated `typedef` declarations

Redeclarations of `typedef` that occur in the same scope are allowed, but a warning is issued.

- Mixing pointer types

Assignment and pointer difference is allowed between pointers to types that are interchangeable but not identical; for example, `unsigned char *` and `char *`. This includes pointers to integral types of the same size. A warning is issued.

Assignment of a string constant to a pointer to any kind of character is allowed, and no warning is issued.

- Non-top level `const`

Assignment of pointers is allowed in cases where the destination type has added type qualifiers that are not at the top level (for example, `int **` to `int const **`). Comparing and taking the difference of such pointers is also allowed.

- Non-lvalue arrays

A non-lvalue array expression is converted to a pointer to the first element of the array when it is used.

- Comments at the end of preprocessor directives

This extension, which makes it legal to place text after preprocessor directives, is enabled unless the strict Standard C mode is used. The purpose of this language extension is to support compilation of legacy code; we do not recommend that you write new code in this fashion.

- An extra comma at the end of `enum` lists

Placing an extra comma is allowed at the end of an `enum` list. In strict Standard C mode, a warning is issued.

- A label preceding a }

In Standard C, a label must be followed by at least one statement. Therefore, it is illegal to place the label at the end of a block. The compiler allows this, but issues a warning.

Note that this also applies to the labels of `switch` statements.

- Empty declarations

An empty declaration (a semicolon by itself) is allowed, but a remark is issued (provided that remarks are enabled).

- Single-value initialization

Standard C requires that all initializer expressions of static arrays, structs, and unions are enclosed in braces.

Single-value initializers are allowed to appear without braces, but a warning is issued. The compiler accepts this expression:

```
struct str
{
    int a;
} x = 10;
```

- Declarations in other scopes

External and static declarations in other scopes are visible. In the following example, the variable `y` can be used at the end of the function, even though it should only be visible in the body of the `if` statement. A warning is issued.

```
int test(int x)
{
    if (x)
    {
        extern int y;
        y = 1;
    }

    return y;
}
```

- Expanding function names into strings with the function as context

Use any of the symbols `__func__` or `__FUNCTION__` inside a function body to make the symbol expand into a string that contains the name of the current function.

Use the symbol `__PRETTY_FUNCTION__` to also include the parameter types and return type. The result might, for example, look like this if you use the `__PRETTY_FUNCTION__` symbol:

```
"void func(char) "
```

These symbols are useful for assertions and other trace utilities and they require that language extensions are enabled, see *-e*, page 269.

- Static functions in function and block scopes

Static functions may be declared in function and block scopes. Their declarations are moved to the file scope.

- Numbers scanned according to the syntax for numbers

Numbers are scanned according to the syntax for numbers rather than the `pp-number` syntax. Thus, `0x123e+1` is scanned as three tokens instead of one valid token. (If the `--strict` option is used, the `pp-number` syntax is used instead.)

Using C++

- Overview—EC++ and EEC++
- Overview—Standard C++
- Enabling support for C++ and variants
- C++ and EC++ feature descriptions
- EEC++ feature description
- EC++ and C++ language extensions

Overview—EC++ and EEC++

IAR Systems supports the C++ language. You can choose between these standards: Standard C++, the industry-standard Embedded C++, and Extended Embedded C++. *Using C++* describes what you need to consider when using the C++ language.

Embedded C++ is a proper subset of the ISO/IEC C++ standard which is intended for embedded systems programming. It was defined by an industry consortium, the Embedded C++ Technical Committee. Performance and portability are particularly important in embedded systems development, which was considered when defining the language. EC++ offers the same object-oriented benefits as C++, but without some features that can increase code size and execution time in ways that are hard to predict.

EMBEDDED C++

These C++ features are supported:

- Classes, which are user-defined types that incorporate both data structure and behavior; the essential feature of inheritance allows data structure and behavior to be shared among classes
- Polymorphism, which means that an operation can behave differently on different classes, is provided by virtual functions
- Overloading of operators and function names, which allows several operators or functions with the same name, provided that their argument lists are sufficiently different
- Type-safe memory management using the operators `new` and `delete`
- Inline functions, which are indicated as particularly suitable for inline expansion.

C++ features that are excluded are those that introduce overhead in execution time or code size that are beyond the control of the programmer. Also excluded are features added very late before Standard C++ was defined. Embedded C++ thus offers a subset of C++ which is efficient and fully supported by existing development tools.

Embedded C++ lacks these features of C++:

- Templates
- Multiple and virtual inheritance
- Exception handling
- Runtime type information
- New cast syntax (the operators `dynamic_cast`, `static_cast`, `reinterpret_cast`, and `const_cast`)
- Namespaces
- The `mutable` attribute.

The exclusion of these language features makes the runtime library significantly more efficient. The Embedded C++ library furthermore differs from the full C++ library in that:

- The standard template library (STL) is excluded
- Streams, strings, and complex numbers are supported without the use of templates
- Library features which relate to exception handling and runtime type information (the headers `except`, `stdexcept`, and `typeinfo`) are excluded.

Note: The library is not in the `std` namespace, because Embedded C++ does not support namespaces.

EXTENDED EMBEDDED C++

IAR Systems' Extended EC++ is a slightly larger subset of C++ which adds these features to the standard EC++:

- Full template support
- Multiple and virtual inheritance
- Namespace support
- The `mutable` attribute
- The cast operators `static_cast`, `const_cast`, and `reinterpret_cast`.

All these added features conform to the C++ standard.

To support Extended EC++, this product includes a version of the standard template library (STL), in other words, the C++ standard chapters utilities, containers, iterators, algorithms, and some numerics. This STL is tailored for use with the Extended EC++

language, which means no exceptions and no support for runtime type information (rtti). Moreover, the library is not in the `std` namespace.

Note: A module compiled with Extended EC++ enabled is fully link-compatible with a module compiled without Extended EC++ enabled.

Overview—Standard C++

The IAR C++ implementation fully complies with the ISO/IEC 14882:2003 C++ standard. In this guide, this standard is referred to as Standard C++.

The main reason for using Standard C++ instead of EC++ or EEC++ is when there is a need for either:

- Exception support
- Runtime type information (RTTI) support
- The standard C++ library (the EC++ library is a stripped version of the C++ library where streams and strings are not templates).

If code size is important and your application does not need any of these features, EC++ (or EEC++) is the better choice.

MODES FOR EXCEPTIONS AND RTTI SUPPORT

Both exceptions and runtime type information result in increased code size simply by being included in your application. You might want to disable either or both of these features to avoid this increase:

- Support for runtime type information constructs can be disabled by using the compiler option `--no_rtti`
- Support for exceptions can be disabled by using the compiler option `--no_exceptions`

Even if support is enabled while compiling, the linker can avoid including the extra code and tables in the final application. If no part of your application actually throws an exception, the code and tables supporting the use of exceptions are not included in the application code image. Also, if dynamic runtime type information constructs (`dynamic_cast/typeid`) are not used with polymorphic types, the objects needed to support them are not included in the application code image. To control this behavior, use the linker options `--no_exceptions`, `--force_exceptions`, and `--no_dynamic_rtti_elimination`.

Disabling exception support

When you use the compiler option `--no_exceptions`, the following will generate a compiler error:

- `throw` expressions
- `try-catch` statements
- Exception specifications on function definitions.

In addition, the extra code and tables needed to handle destruction of objects with auto storage duration when an exception is propagated through a function will not be generated when the compiler option `--no_exceptions` is used.

All functionality in system headers not directly involving exceptions is supported when the compiler option `--no_exceptions` is used.

The linker will produce an error if you try to link C++ modules compiled with exception support with modules compiled without exception support

For more information, see `--no_exceptions`, page 280.

Disabling RTTI support

When you use the compiler option `--no_rtti`, the following will generate a compiler error:

- The `typeid` operator
- The `dynamic_cast` operator.

Note: If `--no_rtti` is used but exception support is enabled, most RTTI support is still included in the compiler output object file because it is needed for exceptions to work.

For more information, see `--no_rtti`, page 283.

EXCEPTION HANDLING

Exception handling can be divided into three parts:

- Exception raise mechanisms—in C++ they are the `throw` and `rethrow` expressions.
- Exception catch mechanisms—in C++ they are the `try-catch` statements, the exception specifications for a function, and the implicit catch to prevent an exception leaking out from `main`.
- Information about currently active functions—if they have `try-catch` statements and the set of auto objects whose destructors need to be run if an exception is propagated through the function.

When an exception is raised, the function call stack is unwound, function by function, block by block. For each function or block, the destructors of auto objects that need destruction are run, and a check is made whether there is a catch handler for the exception. If there is, the execution will continue from that catch handler.

An application that mixes C++ code with assembler and C code, and that throws exceptions from one C++ function to another via assembler routines and C functions must use the linker option `--exception_tables` with the argument `unwind`.

The implementation of exceptions

Exceptions are implemented using a table method. For each function, the tables describe:

- How to unwind the function, that is, how to find its caller on the stack and restore registers that need restoring
- Which catch handlers that exist in the function
- Whether the function has an exception specification and which exceptions it allows to propagate
- The set of auto objects whose destructors must be run.

When an exception is raised, the runtime will proceed in two phases. The first phase will use the exception tables to search the stack for a function invocation containing a catch handler or exception specification that would cause stack unwinding to halt at that point. Once this point is found, the second phase is entered, doing the actual unwinding, and running the destructors of auto objects where that is needed.

The table method results in virtually no overhead in execution time or RAM usage when an exception is not actually thrown. It does incur a significant penalty in read-only memory usage for the tables and the extra code, and throwing and catching an exception is a relatively expensive operation.

The destruction of auto objects when the stack is being unwound as a result of an exception is implemented in code separated from the code that handles the normal operation of a function. This code, together with the code in catch handlers, is placed in a separate section (`.exc.text`) from the normal code (normally placed in `.text`). In some cases, for instance when there is fast and slow ROM memory, it can be advantageous to select on this difference when placing sections in the linker configuration file.

Enabling support for C++ and variants



In the compiler, the default language is C.

To compile files written in Standard C++, use the `--c++` compiler option. See `--c++`, page 264.

To compile files written in Embedded C++, use the `--ec++` compiler option. See `--ec++`, page 270.

To take advantage of *Extended* Embedded C++ features in your source code, use the `--eec++` compiler option. See `--eec++`, page 270.



To enable EC++, EEC++, or C++ in the IDE, choose **Project>Options>C/C++ Compiler>Language** and select the appropriate standard.

C++ and EC++ feature descriptions

When you write C++ source code for the IAR C/C++ Compiler for ARM, you must be aware of some benefits and some possible quirks when mixing C++ features—such as classes, and class members—with IAR language extensions, such as IAR-specific attributes.

USING IAR ATTRIBUTES WITH CLASSES

Static data members of C++ classes are treated the same way global variables are, and can have any applicable IAR type and object attribute.

Member functions are in general treated the same way free functions are, and can have any applicable IAR type and object attributes. Virtual member functions can only have attributes that are compatible with default function pointers, and constructors and destructors cannot have any such attributes.

The location operator `@` and the `#pragma location` directive can be used on static data members and with all member functions.

Example of using attributes with classes

```
class MyClass
{
public:
    // Locate a static variable in __memattr memory at address 60
    static __no_init int mI @ 60;

    // A static Thumb function
    static __thumb void F();

    // A Thumb function
    __thumb void G();

    // Interworking assumed
    virtual __arm void ArmH();

    // Locate a virtual function into SPECIAL
    virtual void M() const volatile @ "SPECIAL";
};
```

FUNCTION TYPES

A function type with `extern "C"` linkage is compatible with a function that has C++ linkage.

Example

```
extern "C"
{
    typedef void (*FpC)(void);    // A C function typedef
}

typedef void (*FpCpp)(void);    // A C++ function typedef

FpC F1;
FpCpp F2;
void MyF(FpC);

void MyG()
{
    MyF(F1);                    // Always works
    MyF(F2);                    // FpCpp is compatible with FpC
}
```

USING STATIC CLASS OBJECTS IN INTERRUPTS

If interrupt functions use static class objects that need to be constructed (using constructors) or destroyed (using destructors), your application will not work properly if the interrupt occurs before the objects are constructed, or, during or after the objects are destroyed.

To avoid this, make sure that these interrupts are not enabled until the static objects have been constructed, and are disabled when returning from `main` or calling `exit`. For information about system startup, see *System startup and termination*, page 133.

Function local static class objects are constructed the first time execution passes through their declaration, and are destroyed when returning from `main` or when calling `exit`.

USING NEW HANDLERS

To handle memory exhaustion, you can use the `set_new_handler` function.

New handlers in Standard C++

If you do not call `set_new_handler`, or call it with a NULL new handler, and `operator new` fails to allocate enough memory, `operator new` will throw `std::bad_alloc` if exceptions are enabled. If exceptions are not enabled, `operator new` will instead call `abort`.

If you call `set_new_handler` with a non-NULL new handler, the provided new handler will be called by `operator new` if the `operator new` fails to allocate enough memory. The new handler must then make more memory available and return, or abort execution in some manner. If exceptions are enabled, the new handler can also throw a `std::bad_alloc` exception. The `nothrow` variant of `operator new` will only return NULL in the presence of a new handler if exceptions are enabled and the new handler throws `std::bad_alloc`.

New handlers in Embedded C++

If you do not call `set_new_handler`, or call it with a NULL new handler, and `operator new` fails to allocate enough memory, it will call `abort`. The `nothrow` variant of the new operator will instead return NULL.

If you call `set_new_handler` with a non-NULL new handler, the provided new handler will be called by `operator new` if `operator new` fails to allocate memory. The new handler must then make more memory available and return, or abort execution in some manner. The `nothrow` variant of `operator new` will never return NULL in the presence of a new handler.

TEMPLATES

C++ and Extended EC++ supports templates according to the C++ standard, but not the `export` keyword. The implementation uses a two-phase lookup which means that the keyword `typename` must be inserted wherever needed. Furthermore, at each use of a template, the definitions of all possible templates must be visible. This means that the definitions of all templates must be in include files or in the actual source file.

DEBUG SUPPORT IN C-SPY

C-SPY® has built-in display support for the STL containers. The logical structure of containers is presented in the watch views in a comprehensive way that is easy to understand and follow.

Using C++, you can make C-SPY stop at a `throw` statement or if a raised exception does not have any corresponding `catch` statement.

For more information about this, see the *C-SPY® Debugging Guide for ARM*.

EEC++ feature description

This section describes features that distinguish Extended EC++ from EC++.

TEMPLATES

The compiler supports templates with the syntax and semantics as defined by Standard C++. However, note that the STL (standard template library) delivered with the product is tailored for Extended EC++, see *Extended Embedded C++*, page 192.

VARIANTS OF CAST OPERATORS

In Extended EC++ these additional variants of C++ cast operators can be used:

```
const_cast<to>(from)
static_cast<to>(from)
reinterpret_cast<to>(from)
```

MUTABLE

The `mutable` attribute is supported in Extended EC++. A `mutable` symbol can be changed even though the whole class object is `const`.

NAMESPACE

The namespace feature is only supported in *Extended EC++*. This means that you can use namespaces to partition your code. Note, however, that the library itself is not placed in the `std` namespace.

THE STD NAMESPACE

The `std` namespace is not used in either standard EC++ or in Extended EC++. If you have code that refers to symbols in the `std` namespace, simply define `std` as nothing; for example:

```
#define std
```

You must make sure that identifiers in your application do not interfere with identifiers in the runtime library.

EC++ and C++ language extensions

When you use the compiler in any C++ mode and enable IAR language extensions, the following C++ language extensions are available in the compiler:

- In a friend declaration of a class, the `class` keyword can be omitted, for example:

```
class B;
class A
{
    friend B;           //Possible when using IAR language
                       //extensions
    friend class B; //According to the standard
};
```

- Constants of a scalar type can be defined within classes, for example:

```
class A
{
    const int mSize = 10; //Possible when using IAR language
                       //extensions
    int mArr[mSize];
};
```

According to the standard, initialized static data members should be used instead.

- In the declaration of a class member, a qualified name can be used, for example:

```
struct A
{
    int A::F(); // Possible when using IAR language extensions
    int G();    // According to the standard
};
```

- It is permitted to use an implicit type conversion between a pointer to a function with C linkage (`extern "C"`) and a pointer to a function with C++ linkage (`extern "C++"`), for example:

```
extern "C" void F(); // Function with C linkage
void (*PF)()        // PF points to a function with C++ linkage
    = &F; // Implicit conversion of function pointer.
```


According to the standard, the pointer must be explicitly converted.

- If the second or third operands in a construction that contains the `?` operator are string literals or wide string literals (which in C++ are constants), the operands can be implicitly converted to `char *` or `wchar_t *`, for example:

```
bool X;

char *P1 = X ? "abc" : "def";           //Possible when using IAR
                                         //language extensions
char const *P2 = X ? "abc" : "def";    //According to the standard
```

- Default arguments can be specified for function parameters not only in the top-level function declaration, which is according to the standard, but also in `typedef` declarations, in pointer-to-function function declarations, and in pointer-to-member function declarations.
- In a function that contains a non-static local variable and a class that contains a non-evaluated expression (for example a `sizeof` expression), the expression can reference the non-static local variable. However, a warning is issued.
- An anonymous union can be introduced into a containing class by a `typedef` name. It is not necessary to first declare the union. For example:

```
typedef union
{
    int i,j;
} U; // U identifies a reusable anonymous union.

class A
{
public:
    U; // OK -- references to A::i and A::j are allowed.
};
```

In addition, this extension also permits *anonymous classes* and *anonymous structs*, as long as they have no C++ features (for example, no static data members or member functions, and no non-public members) and have no nested types other than other anonymous classes, structs, or unions. For example:

```
struct A
{
    struct
    {
        int i,j;
    }; // OK -- references to A::i and A::j are allowed.
};
```

- The friend class syntax allows nonclass types as well as class types expressed through a `typedef` without an elaborated type name. For example:

```
typedef struct S ST;

class C
{
public:
    friend S; // Okay (requires S to be in scope)
    friend ST; // Okay (same as "friend S;")
    // friend S const; // Error, cv-qualifiers cannot
                        // appear directly
};
```

Note: If you use any of these constructions without first enabling language extensions, errors are issued.

Application-related considerations

- Output format considerations
- Stack considerations
- Heap considerations
- Interaction between the tools and your application
- Checksum calculation for verifying image integrity
- Linker optimizations
- AEABI compliance
- CMSIS integration

Output format considerations

The linker produces an absolute executable image in the ELF/DWARF object file format.

You can use the IAR ELF Tool—`ielftool`— to convert an absolute ELF image to a format more suitable for loading directly to memory, or burning to a PROM or flash memory etc.

`ielftool` can produce these output formats:

- Plain binary
- Motorola S-records
- Intel hex.

Note: `ielftool` can also be used for other types of transformations, such as filling and calculating checksums in the absolute image.

The source code for `ielftool` is provided in the `arm/src` directory. For more information about `ielftool`, see *The IAR ELF Tool—ielftool*, page 480.

Stack considerations

To make your application use stack memory efficiently, there are some considerations to be made.

STACK SIZE CONSIDERATIONS

The required stack size depends heavily on the application’s behavior. If the given stack size is too large, RAM will be wasted. If the given stack size is too small, one of two things can happen, depending on where in memory you located your stack:

- Variable storage will be overwritten, leading to undefined behavior
- The stack will fall outside of the memory area, leading to an abnormal termination of your application.

Both alternatives are likely to result in application failure. Because the second alternative is easier to detect, you should consider placing your stack so that it grows toward the end of the memory.

For more information about the stack size, see *Setting up stack memory*, page 105, and *Saving stack space and RAM memory*, page 235.

STACK ALIGNMENT

The default `cstartup` code automatically initializes all stacks to an 8-byte aligned address.

For more information about aligning the stack, see *Calling convention*, page 169 and more specifically *Special registers*, page 171 and *Stack parameters and layout*, page 173.

EXCEPTION STACK

Cortex-M does not have individual exception stacks. By default, all exception stacks are placed in the `CSTACK` section.

The ARM7/9/11, Cortex-A, and Cortex-R devices support five exception modes which are entered when different exceptions occur. Each exception mode has its own stack to avoid corrupting the System/User mode stack.

The table shows proposed stack names for the various exception stacks, but any name can be used:

Processor mode	Proposed stack section name	Description
Supervisor	SVC_STACK	Operation system stack.

Table 23: Exception stacks for ARM7/9/11, Cortex-A, and Cortex-R

Processor mode	Proposed stack section name	Description
IRQ	IRQ_STACK	Stack for general-purpose (IRQ) interrupt handlers.
FIQ	FIQ_STACK	Stack for high-speed (FIQ) interrupt handlers.
Undefined	UND_STACK	Stack for undefined instruction interrupts. Supports software emulation of hardware coprocessors and instruction set extensions.
Abort	ABT_STACK	Stack for instruction fetch and data access memory abort interrupt handlers.

Table 23: Exception stacks for ARM7/9/11, Cortex-A, and Cortex-R

For each processor mode where a stack is needed, a separate stack pointer must be initialized in your startup code, and section placement should be done in the linker configuration file. The IRQ and FIQ stacks are the only exception stacks which are preconfigured in the supplied `cstartup.s` and `lnkarm.icf` files, but other exception stacks can easily be added.



To view any of these stacks in the Stack window available in the IDE, these preconfigured section names must be used instead of user-defined section names.

Heap considerations

The heap contains dynamic data allocated by use of the C function `malloc` (or a corresponding function) or the C++ operator `new`.

If your application uses dynamic memory allocation, you should be familiar with:

- Linker sections used for the heap
- Allocating the heap size, see *Setting up heap memory*, page 106.

ADVANCED VERSUS BASIC HEAP

The system library contains two separate heap memory handlers—the *basic* and the *advanced* heap handler. If your application appears to make explicit use of heap allocation, the linker automatically chooses the advanced heap; otherwise, the basic heap is used.

Note: If your product has a code size limitation, the basic heap is automatically chosen.

You can use a linker option to explicitly specify which handler you want to use:

- The basic heap (`--basic_heap`) is a very simple heap allocator, suitable for use in applications that do not use the heap very much. In particular, it can be used in applications that only allocate heap memory and never free it. The basic heap is not particularly speedy, and using it in applications that repeatedly free memory is quite likely to lead to unneeded fragmentation of the heap. The code for the basic heap is significantly smaller than that for the advanced heap. See *--basic_heap*, page 301.
- The advanced heap (`--advanced_heap`) provides efficient memory management for applications that use the heap extensively. In particular, applications that repeatedly allocate and free memory will likely get less overhead in both space and time. The code for the advanced heap is significantly larger than that for the basic heap. See *--advanced_heap*, page 301. For information about the definition, see *iar_dlmalloc.h*, page 434.

HEAP SIZE AND STANDARD I/O



If you excluded `FILE` descriptors from the DLIB runtime environment, as in the Normal configuration, there are no input and output buffers at all. Otherwise, as in the Full configuration, be aware that the size of the input and output buffers is set to 512 bytes in the `stdio` library header file. If the heap is too small, I/O will not be buffered, which is considerably slower than when I/O is buffered. If you execute the application using the simulator driver of the IAR C-SPY® Debugger, you are not likely to notice the speed penalty, but it is quite noticeable when the application runs on an ARM core. If you use the standard I/O library, you should set the heap size to a value which accommodates the needs of the standard I/O buffer.

Interaction between the tools and your application

The linking process and the application can interact symbolically in four ways:

- Creating a symbol by using the linker command line option `--define_symbol`. The linker will create a public absolute constant symbol that the application can use as a label, as a size, as setup for a debugger, etc.
- Creating an exported configuration symbol by using the command line option `--config_def` or the configuration directive `define symbol`, and exporting the symbol using the `export symbol` directive. ILINK will create a public absolute constant symbol that the application can use as a label, as a size, as setup for a debugger, etc.

One advantage of this symbol definition is that this symbol can also be used in expressions in the configuration file, for example to control the placement of sections into memory ranges.

- Using the compiler operators `__section_begin`, `__section_end`, or `__section_size`, or the assembler operators `SFB`, `SFE`, or `SIZEOF` on a named section or block. These operators provide access to the start address, end address, and size of a contiguous sequence of sections with the same name, or of a linker block specified in the linker configuration file.
- The command line option `--entry` informs the linker about the start label of the application. It is used by the linker as a root symbol and to inform the debugger where to start execution.

The following lines illustrate how to use `-D` to create a symbol. If you need to use this mechanism, add these options to your command line like this:

```
--define_symbol NrofElements=10
--config_def HEAP_SIZE=1024
```

The linker configuration file can look like this:

```
define memory Mem with size = 4G;
define region ROM = Mem:[from 0x00000 size 0x10000];
define region RAM = Mem:[from 0x20000 size 0x10000];

/* Export of symbol */
export symbol MY_HEAP_SIZE;

/* Setup a heap area with a size defined by an ILINK option */
define block MyHEAP with size = MY_HEAP_SIZE, alignment = 8 {};

place in RAM { block MyHEAP };
```

Add these lines to your application source code:

```
#include <stdlib.h>

/* Use symbol defined by ILINK option to dynamically allocate an
array of elements with specified size. The value takes the form
of a label.
*/
extern int NrofElements;

typedef char Elements;
Elements *GetElementArray()
{
    return malloc(sizeof(Elements) * (long) &NrofElements);
}
```

```

/* Use a symbol defined by ILINK option, a symbol that in the
 * configuration file was made available to the application.
 */
extern char MY_HEAP_SIZE;

/* Declare the section that contains the heap. */
#pragma section = "MYHEAP"

char *MyHeap()
{
    /* First get start of statically allocated section, */
    char *p = __section_begin("MYHEAP");

    /* ...then we zero it, using the imported size. */
    for (int i = 0; i < (int) &MY_HEAP_SIZE; ++i)
    {
        p[i] = 0;
    }
    return p;
}

```

Checksum calculation for verifying image integrity

This page contains information about checksum calculation:

- *Briefly about checksum calculation*, page 208
- *Calculating and verifying a checksum*, page 210
- *Troubleshooting checksum calculation*, page 216

For more information, see also *The IAR ELF Tool—ielftool*, page 480.

BRIEFLY ABOUT CHECKSUM CALCULATION

You can use a checksum to verify that the image is the same at runtime as when the image's original checksum was generated. In other words, to verify that the image has not been corrupted.

This works as follows:

- You need an initial checksum.
You can either use the IAR ELF Tool—*ielftool*—to generate an initial checksum or you might have a third-party checksum available.

- You must generate a second checksum during runtime.

You can either add specific code to your application source code for calculating a checksum during runtime or you can use some dedicated hardware on your device for calculating a checksum during runtime.

- You must add specific code to your application source code for comparing the two checksums and take an appropriate action if they differ.

If the two checksums have been calculated in the same way, and if there are no errors in the image, the checksums should be identical. If not, you should first suspect that the two checksums were not generated in the same way.

No matter which solutions you use for generating the two checksum, you must make sure that both checksums are calculated *in the exact same way*. If you use `ielftool` for the initial checksum and use a software-based calculation during runtime, you have full control of the generation for both checksums. However, if you are using a third-party checksum for the initial checksum or some hardware support for the checksum calculation during runtime, there might be additional requirements that you must consider.

For the two checksums, there are some choices that you must always consider and there are some choices to make only if there are additional requirements. Still, all of the details must be the same for both checksums.

Details always to consider:

- *Checksum range*

The memory range (or ranges) that you want to verify by means of checksums. Typically, you might want to calculate a checksum for all ROM memory. However, you might want to calculate a checksum only for specific ranges. Remember that:

- It is OK to have several ranges for one checksum.
- The checksum must be calculated from the lowest to the highest address for every memory range.
- Each memory range must be verified in the same order as defined (for example, `0x100-0x1FF,0x400-0x4FF` is not the same as `0x400-0x4FF,0x100-0x1FF`).
- If several checksums are used, you should place them in sections with unique names and use unique symbol names.
- A checksum should never be calculated on a memory range that contains a checksum or a software breakpoint.

- *Algorithm and size of checksum*

You should consider which algorithm is most suitable in your case. There are two basic choices, Sum (a simple arithmetic algorithm) or CRC (which is the most commonly used algorithm). For CRC there are different sizes to choose for the checksum, 2, 4, or 8 bytes where the predefined polynomials are wide enough to suit

the size, for more error detecting power. The predefined polynomials work well for most, but possibly not for all data sets. If not, you can specify your own polynomial. If you just want a decent error detecting mechanism, use the predefined CRC algorithm for your checksum size, typically CRC16 or CRC32.

Note that for an n -bit polynomial, the n :th bit is always considered to be set. For a 16-bit polynomial (for example, CRC16) this means that 0x11021 is the same as 0x1021.

For more information about selecting an appropriate polynomial for data sets with non-uniform distribution, see for example section 3.5.3 in *Tannenbaum, A.S., "Computer Networks," Prentice Hall 1981, ISBN: 0131646990*.

- *Fill*

Every byte in the checksum range must have a well-defined value before the checksum can be calculated. Typically, bytes with unknown values are *pad bytes* that have been added for alignment. This means that you must specify which fill pattern to be used during calculation, typically 0xFF or 0x00.

- *Initial value*

The checksum must always have an explicit initial value.

In addition to these mandatory details, there might be other details to consider. Typically, this might happen when you have a third-party checksum, you want the checksum be compliant with the Rocksoft™ checksum model, or when you use hardware support for generating a checksum during runtime. `ieftool` provides support for also controlling alignment, complement, bit order, byte order within words, and checksum unit size.

CALCULATING AND VERIFYING A CHECKSUM

In this example procedure, a checksum is calculated for ROM memory from 0x8002 up to 0x8FFF and the 2-byte calculated checksum is placed at 0x8000.

- I If you are using `ieftool` from the command line, you must first allocate a memory location for the calculated checksum.

Note: If you instead are using the IDE (and not the command line), the `__checksum`, `__checksum_begin`, and `__checksum_end` symbols, and the `.checksum` section are *automatically* allocated when you calculate the checksum, which means that you can skip this step.

You can allocate the memory location in two ways:


- By creating a global C/C++ or assembler constant symbol with a proper size, residing in a specific section (in this example `.checksum`)
- By using the linker option `--place_holder`.

For example, to allocate a 2-byte space for the symbol `__checksum` in the section `.checksum`, with alignment 4, specify:

```
--place_holder __checksum,2,.checksum,4
```

- 2 The `.checksum` section will only be included in your application if the section appears to be needed. If the checksum is not needed by the application itself, use the linker option `--keep=__checksum` (or the linker directive `keep`) to force the section to be included.

Alternatively, choose **Project>Options>Linker>Output** and specify `__checksum`:



- 3 To control the placement of the `.checksum` section, you must modify the linker configuration file. For example, it can look like this (note the handling of the block `CHECKSUM`):

```
define block CHECKSUM      { ro section .checksum };
place in ROM_region { ro, first block CHECKSUM };
```

Note: It is possible to skip this step, but in that case the `.checksum` section will automatically be placed with other read-only data.

- 4 When configuring `ielftool` to calculate a checksum, there are some basic choices to make:
 - Checksum algorithm
Choose which checksum algorithm you want to use. In this example, the CRC16 algorithm is used.
 - Memory range
Using the IDE, you can specify one memory range for which the checksum should be calculated. From the command line, you can specify any ranges.

- Fill pattern

Specify a fill pattern—typically 0xFF or 0x00—for bytes with unknown values. The fill pattern will be used in all checksum ranges.

For more information, see *Briefly about checksum calculation*, page 208.



To run `ielftool` from the IDE, choose **Project>Options>Linker>Checksum** and make your settings, for example:

In the simplest case, you can ignore (or leave with default settings) these options: **Complement**, **Bit order**, **Reverse byte order within word**, and **Checksum unit size**.



To run `ielftool` from the command line, specify the command, for example, like this:

```
ielftool --fill=0x00;0x8002-0x8FFF
--checksum=__checksum:2,crc16;0x8002-0x8FFF sourceFile.out
destinationFile.out
```

Note that `ielftool` needs an unstripped input ELF image. If you use the linker option `--strip`, remove it and use the `ielftool` option `--strip` instead.

The checksum will be created later on when you build your project and will be automatically placed in the specified symbol `__checksum` in the section `.checksum`.

5 You can specify several ranges instead of only one range.



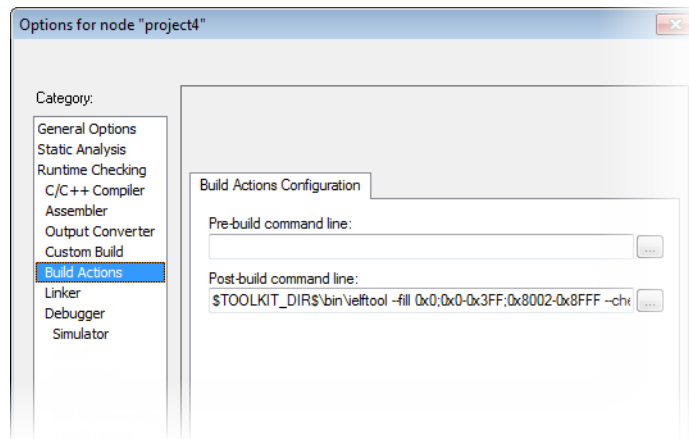
If you are using the IDE, perform these steps:

- Choose **Project>Options>Linker>Checksum** and make sure to deselect **Fill unused code memory**.

- Choose **Project>Options>Build Actions** and specify the ranges together with the rest of the required commands in the **Post-build command line** text field, for example like this:

```
$TOOLKIT_DIR$\bin\ielftool $PROJ_DIR$\debug\exe\output.out
$PROJ_DIR$\debug\exe\output.out
--fill 0x0;0x0-0x3FF;0x8002-0x8FFF
--checksum=__checksum:2,crc16;0x0-0x3FF;0x8002-0x8FFF
```

In your example, replace *output.out* with the name of your output file.



If you are using the command line, specify the ranges, for example like this:

```
ielftool output.out output.out --fill 0x0;0x0-0x3FF;0x8002-0x8FFF
--checksum=__checksum:2,crc16;0x0-0x3FF;0x8002-0x8FFF
```

In your example, replace *output.out* with the name of your output file.

- 6** Add a function for checksum calculation to your source code. Make sure that the function uses the same algorithm and settings as for the checksum calculated by `ielftool`. For example, a slow variant of the `crc16` algorithm but with small memory footprint (in contrast to the fast variant that uses more memory):

```
unsigned short SmallCrc16(uint16_t
    sum,
                                unsigned char *p,
                                unsigned int len)
{
    while (len--)
    {
        int i;
        unsigned char byte = *(p++);

        for (i = 0; i < 8; ++i)
        {
            unsigned long oSum = sum;
            sum <<= 1;
            if (byte & 0x80)
                sum |= 1;
            if (oSum & 0x8000)
                sum ^= 0x1021;
            byte <<= 1;
        }
    }
    return sum;
}
```

You can find the source code for this checksum algorithm in the `arm\src\linker` directory of your product installation.

- 7** Make sure that your application also contains a call to the function that calculates the checksum, compares the two checksums, and takes appropriate action if the checksum values do not match. This code gives an example of how the checksum can be calculated for your application and to be compared with the `ielftool` generated checksum:

```
/* The calculated checksum */

/* Linker generated symbols */
extern unsigned short const __checksum;
extern int __checksum_begin;
extern int __checksum_end;

void TestChecksum()
{
    unsigned short calc = 0;
    unsigned char zeros[2] = {0, 0};

    /* Run the checksum algorithm */
    calc = SmallCrc16(0,
                     (unsigned char *) &__checksum_begin,
                     ((unsigned char *) &__checksum_end -
                      ((unsigned char *) &__checksum_begin)+1));

    /* Fill the end of the byte sequence with zeros. */
    calc = SmallCrc16(calc, zeros, 2);

    /* Test the checksum */
    if (calc != __checksum)
    {
        printf("Incorrect checksum!\n");
        abort(); /* Failure */
    }

    /* Checksum is correct */
}
```

- 8** Build your application project and download it.

During the build, `ielftool` creates a checksum and places it in the specified symbol `__checksum` in the section `.checksum`.

- 9** Choose **Download and Debug** to start the C-SPY debugger.

During execution, the checksum calculated by `ielftool` and the checksum calculated by your application should be identical.

TROUBLESHOOTING CHECKSUM CALCULATION

If the two checksums do not match, there are several possible causes. These are some troubleshooting hints:

- If possible, start with a small example when trying to get the checksums to match.
- Verify that the exact same memory range or ranges are used in both checksum calculations.

To help you do this, `ielftool` lists the ranges for which the checksum is calculated on `stdout` about the exact addresses that were used and the order in which they were accessed.

- Make sure that all checksum symbols are excluded from all checksum calculations. Compare the checksum placement with the checksum range and make sure they do not overlap. You can find information in the **Build** message window after `ielftool` has generated a checksum.
- Verify that the checksum calculations use the same polynomial.
- Verify that the bits in the bytes are processed in the same order in both checksum calculations, from the least to the most significant bit or the other way around. You control this with the **Bit order** option (or from the command line, the `-m` parameter of the `--checksum` option).
- If you are using the small variant of CRC, check whether you need to feed additional bytes into the algorithm.

The number of zeros to add at the end of the byte sequence must match the size of the checksum, in other words, one zero for a 1-byte checksum, two zeros for a 2-byte checksum, four zeros for a 4-byte checksum, and eight zeros for an 8-byte checksum.

- Any breakpoints in flash memory change the content of the flash. This means that the checksum which is calculated by your application will no longer match the initial checksum calculated by `ielftool`. To make the two checksums match again, you must disable all your breakpoints in flash and any breakpoints set in flash by C-SPY internally. The stack plugin and the debugger option **Run to** both require C-SPY to set breakpoints. Read more about possible breakpoint consumers in the *C-SPY® Debugging Guide for ARM*.
- By default, a symbol that you have allocated in memory by using the linker option `--place_holder` is considered by C-SPY to be of the type `int`. If the size of the checksum is different than the size of an `int`, you can change the display format of the checksum symbol to match its size.

In the C-SPY **Watch** window, select the symbol and choose **Show As** from the context menu. Choose the display format that matches the size of the checksum symbol.

Linker optimizations

VIRTUAL FUNCTION ELIMINATION

Virtual Function Elimination (VFE) is a linker optimization that removes unneeded virtual functions and dynamic runtime type information.

In order for Virtual Function Elimination to work, all relevant modules must provide information about virtual function table layout, which virtual functions are called, and for which classes dynamic runtime type information is needed. If one or more modules do not provide this information, a warning is generated by the linker and Virtual Function Elimination is not performed.

If you know that modules that lack such information do not perform any virtual function calls and do not define any virtual function tables, you can use the `--vfe=forced` linker option to enable Virtual Function Elimination anyway.

Currently, tools from IAR Systems and from RealView provide the information needed for Virtual Function Elimination in a way that the linker can use.

Note that you can disable Virtual Function Elimination entirely by using the `--no_vfe` linker option. In this case, no warning will be issued for modules that lack VFE information.

For more information, see `--vfe`, page 326 and `--no_vfe`, page 320.

AEABI compliance

The IAR build tools for ARM support the Embedded Application Binary Interface for ARM, AEABI, defined by ARM Limited. This interface is based on the Intel IA64 ABI interface. The advantage of adhering to AEABI is that any such module can be linked with any other AEABI-compliant module, even modules produced by tools provided by other vendors.

The IAR build tools for ARM support the following parts of the AEABI:

AAPCS	Procedure Call Standard for the ARM architecture
CPPABI	C++ ABI for the ARM architecture (EC++ parts only)
AAELF	ELF for the ARM architecture
AADWARF	DWARF for the ARM architecture
RTABI	Runtime ABI for the ARM architecture
CLIBABI	C library ABI for the ARM architecture

The IAR build tools only support a *bare metal* platform, that is a ROM-based system that lacks an explicit operating system.

Note that:

- The AEABI is specified for C89 only
- The IAR build tools only support using the default and C locales
- The AEABI does not specify C++ library compatibility
- Neither the size of an `enum` or of `wchar_t` is constant in the AEABI.

If AEABI compliance is enabled, almost all optimizations performed in the system header files are turned off, and certain preprocessor constants become real constant variables instead.

LINKING AEABI-COMPLIANT MODULES USING THE IAR ILINK LINKER

When building an application using the IAR ILINK Linker, the following types of modules can be combined:

- Modules produced using IAR build tools, both AEABI-compliant modules as well as modules that are not AEABI-compliant
- AEABI-compliant modules produced using build tools from another vendor.

Note: To link a module produced by a compiler from another vendor, extra support libraries from that vendor might be required.

The IAR ILINK Linker automatically chooses the appropriate standard C/C++ libraries to use based on attributes from the object files. Imported object files might not have all these attributes. Therefore, you might need to help ILINK choose the standard library by verifying one or more of the following details:

- The used CPU by specifying the `--cpu` linker option
- If full I/O is needed; make sure to link with a Full library configuration in the standard library
- Explicitly specify runtime library file(s), possibly in combination with the `--no_library_search` linker option.

When linking, you should also consider virtual function elimination, see *Virtual function elimination*, page 217.

LINKING AEABI-COMPLIANT MODULES USING A THIRD-PARTY LINKER

If you have a module produced using the IAR C/C++ Compiler and you plan to link that module using a linker from a different vendor, that module must be AEABI-compliant, see *Enabling AEABI compliance in the compiler*, page 219.

In addition, if that module uses any of the IAR-specific compiler extensions, you must make sure that those features are also supported by the tools from the other vendor. Note specifically:

- Support for the following extensions must be verified: `#pragma pack`, `__no_init`, `__root`, and `__ramfunc`
- The following extensions are harmless to use: `#pragma location/@`, `__arm`, `__thumb`, `__swi`, `__irq`, `__fiq`, and `__nested`.

ENABLING AEABI COMPLIANCE IN THE COMPILER

You can enable AEABI compliance in the compiler by setting the `--aeabi` option. In this case, you must also use the `--guard_calls` option.



In the IDE, use the **Project>Options>C/C++ Compiler>Extra Options** page to specify the `--aeabi` and `--guard_calls` options.



On the command line, use the options `--aeabi` and `--guard_calls` to enable AEABI support in the compiler.

Alternatively, to enable support for AEABI for a specific system header file, you must define the preprocessor symbol `_AEABI_PORTABILITY_LEVEL` to non-zero prior to including a system header file, and make sure that the symbol `AEABI_PORTABLE` is set to non-zero after the inclusion of the header file:

```
#define _AEABI_PORTABILITY_LEVEL 1
#undef _AEABI_PORTABLE
#include <header.h>
#ifdef _AEABI_PORTABLE
    #error "header.h not AEABI compatible"
#endif
```

CMSIS integration

The `arm\CMSIS` subdirectory contains CMSIS (Cortex Microcontroller Software Interface Standard) and CMSIS DSP header and library files, and documentation. For more information about CMSIS, see <http://www.arm.com/cmsis>.

The special header file `inc\c\cmsis_iar.h` is provided as a CMSIS adaptation of the current version of the IAR C/C++ Compiler.

CMSIS DSP LIBRARY

IAR Embedded Workbench comes with prebuilt CMSIS DSP libraries in the `arm\CMSIS\Lib\IAR` directory. The names of the library files are constructed in this way:

```
iar_cortexM<0|3|4><1|b>[f]_math.a
```

where `<0|3|4>` selects the Cortex-M variant, `<1|b>` selects the byte order, and `[f]` indicates that the library is built for FPU (Cortex-M4 only).

The libraries for Cortex-M4 are applicable also to Cortex-M7.

CUSTOMIZING THE CMSIS DSP LIBRARY

The source code of the CMSIS DSP library is provided in the `arm\CMSIS\DSP_Lib\Source` directory. You can find an IAR Embedded Workbench project which is prepared for building a customized DSP library in the `arm\CMSIS\DSP_Lib\Source\IAR` directory.



BUILDING WITH CMSIS ON THE COMMAND LINE

This section contains examples of how to build your CMSIS-compatible application on the command line.

CMSIS only (that is without the DSP library)

```
iccarm -I $EW_DIR$\arm\CMSIS\Include
```

With the DSP library, for Cortex-M4, little-endian, and with FPU

```
iccarm --endian=little --cpu=Cortex-M4 --fpu=VFPv4_sp -I  
$EW_DIR$\arm\CMSIS\Include -D ARM_MATH_CM4
```

```
ilinkarm $EW_DIR$\arm\CMSIS\Lib\IAR\iar_cortexM31_math.a
```



BUILDING WITH CMSIS IN THE IDE

Choose **Project>Options>General Options>Library Configuration** to enable CMSIS support.

When enabled, CMSIS include paths and the DSP library will automatically be used. For more information, see the *IDE Project Management and Building Guide for ARM*.

Efficient coding for embedded applications

- Selecting data types
- Controlling data and function placement in memory
- Controlling compiler optimizations
- Facilitating good code generation

Selecting data types

For efficient treatment of data, you should consider the data types used and the most efficient placement of the variables.

USING EFFICIENT DATA TYPES

The data types you use should be considered carefully, because this can have a large impact on code size and code speed.

- Use `int` or `long` instead of `char` or `short` whenever possible, to avoid sign extension or zero extension. In particular, loop indexes should always be `int` or `long` to minimize code generation. Also, in Thumb mode, accesses through the stack pointer (`SP`) is restricted to 32-bit data types, which further emphasizes the benefits of using one of these data types.
- Use unsigned data types, unless your application really requires signed values.
- Be aware of the costs of using 64-bit data types, such as `double` and `long long`.
- Bitfields and packed structures generate large and slow code.
- Using floating-point types on a microprocessor without a math co-processor is very inefficient, both in terms of code size and execution speed.
- Declaring a pointer to `const` data tells the calling function that the data pointed to will not change, which opens for better optimizations.

For information about representation of supported data types, pointers, and structures types, see the chapter *Data representation*.

FLOATING-POINT TYPES

Using floating-point types on a microprocessor without a math coprocessor is very inefficient, both in terms of code size and execution speed. Thus, you should consider replacing code that uses floating-point operations with code that uses integers, because these are more efficient.

The compiler supports two floating-point formats—32 and 64 bits. The 32-bit floating-point type `float` is more efficient in terms of code size and execution speed. However, the 64-bit format `double` supports higher precision and larger numbers.

In the compiler, the floating-point type `float` always uses the 32-bit format, and the type `double` always uses the 64-bit format.

Unless the application requires the extra precision that 64-bit floating-point numbers give, we recommend using 32-bit floating-point numbers instead.

By default, a *floating-point constant* in the source code is treated as being of the type `double`. This can cause innocent-looking expressions to be evaluated in double precision. In the example below `a` is converted from a `float` to a `double`, the `double` constant `1.0` is added and the result is converted back to a `float`:

```
double Test(float a)
{
    return a + 1.0;
}
```

To treat a floating-point constant as a `float` rather than as a `double`, add the suffix `f` to it, for example:

```
double Test(float a)
{
    return a + 1.0f;
}
```

For more information about floating-point types, see *Basic data types—floating-point types*, page 336.

ALIGNMENT OF ELEMENTS IN A STRUCTURE

The ARM core requires that when accessing data in memory, the data must be aligned. Each element in a structure must be aligned according to its specified type requirements. This means that the compiler might need to insert *pad bytes* to keep the alignment correct.

There are situations when this can be a problem:

- There are external demands; for example, network communication protocols are usually specified in terms of data types with no padding in between

- You need to save data memory.

For information about alignment requirements, see *Alignment*, page 329.

Use the `#pragma pack` directive or the `__packed` data type attribute for a tighter layout of the structure. The drawback is that each access to an unaligned element in the structure will use more code.

Alternatively, write your own customized functions for *packing* and *unpacking* structures. This is a more portable way, which will not produce any more code apart from your functions. The drawback is the need for two views on the structure data—packed and unpacked.

For more information about the `#pragma pack` directive, see *pack*, page 373.

ANONYMOUS STRUCTS AND UNIONS

When a structure or union is declared without a name, it becomes anonymous. The effect is that its members will only be seen in the surrounding scope.

Anonymous structures are part of the C++ language; however, they are not part of the C standard. In the IAR C/C++ Compiler for ARM they can be used in C if language extensions are enabled.



In the IDE, language extensions are enabled by default.



Use the `-e` compiler option to enable language extensions. See *-e*, page 269, for additional information.

Example

In this example, the members in the anonymous union can be accessed, in function `F`, without explicitly specifying the union name:

```
struct S
{
    char mTag;
    union
    {
        long mL;
        float mF;
    };
} St;

void F(void)
{
    St.mL = 5;
}
```

The member names must be unique in the surrounding scope. Having an anonymous struct or union at file scope, as a global, external, or static variable is also allowed. This could for instance be used for declaring I/O registers, as in this example:

```
__no_init volatile
union
{
    unsigned char IOPORT;
    struct
    {
        unsigned char way: 1;
        unsigned char out: 1;
    };
} @ 0x1000;

/* The variables are used here. */
void Test(void)
{
    IOPORT = 0;
    way = 1;
    out = 1;
}
```

This declares an I/O register byte `IOPORT` at address `0x1000`. The I/O register has 2 bits declared, `Way` and `Out`. Note that both the inner structure and the outer union are anonymous.

Anonymous structures and unions are implemented in terms of objects named after the first field, with a prefix `_A_` to place the name in the implementation part of the namespace. In this example, the anonymous union will be implemented through an object named `_A_IOPORT`.

Controlling data and function placement in memory

The compiler provides different mechanisms for controlling placement of functions and data objects in memory. To use memory efficiently, you should be familiar with these mechanisms and know which one is best suited for different situations. You can use:

- The `@` operator and the `#pragma location` directive for absolute placement.

Using the `@` operator or the `#pragma location` directive, you can place individual global and static variables at absolute addresses. Note that it is not possible to use this notation for absolute placement of individual functions. For more information, see *Data placement at an absolute location*, page 225.

- The @ operator and the #pragma location directive for section placement.
Using the @ operator or the #pragma location directive, you can place individual functions, variables, and constants in named sections. The placement of these sections can then be controlled by linker directives. For more information, see *Data and function placement in sections*, page 226
- The @ operator and the #pragma location directive for register placement
Use the @ operator or the #pragma location directive to place individual global and static variables in registers. The variables must be declared __no_init. This is useful for individual data objects that must be located in a specific register.
- Using the --section option, you can set the default segment for functions, variables, and constants in a particular module. For more information, see *--section*, page 292.

DATA PLACEMENT AT AN ABSOLUTE LOCATION

The @ operator, alternatively the #pragma location directive, can be used for placing global and static variables at absolute addresses.

To place a variable at an absolute address, the argument to the @ operator and the #pragma location directive should be a literal number, representing the actual address. The absolute location must fulfill the alignment requirement for the variable that should be located.

Note: All declarations of __no_init variables placed at an absolute address are *tentative definitions*. Tentatively defined variables are only kept in the output from the compiler if they are needed in the module being compiled. Such variables will be defined in all modules in which they are used, which will work as long as they are defined in the same way. The recommendation is to place all such declarations in header files that are included in all modules that use the variables.

Other variables placed at an absolute address use the normal distinction between declaration and definition. For these variables, you must provide the definition in only one module, normally with an initializer. Other modules can refer to the variable by using an extern declaration, with or without an explicit address.

Examples

In this example, a __no_init declared variable is placed at an absolute address. This is useful for interfacing between multiple processes, applications, etc:

```
__no_init volatile char alpha @ 0xFF2000; /* OK */
```

The next example contains two const declared objects. The first one is not initialized, and the second one is initialized to a specific value. Both objects are placed in ROM. This is useful for configuration parameters, which are accessible from an external

interface. Note that in the second case, the compiler is not obliged to actually read from the variable, because the value is known.

```
#pragma location=0xFF2002
__no_init const int beta;           /* OK */

const int gamma @ 0xFF2004 = 3;     /* OK */
```

In the first case, the value is not initialized by the compiler; the value must be set by other means. The typical use is for configurations where the values are loaded to ROM separately, or for special function registers that are read-only.

```
__no_init int epsilon @ 0xFF2007; /* Error, misaligned. */
```

C++ considerations

In C++, module scoped `const` variables are static (module local), whereas in C they are global. This means that each module that declares a certain `const` variable will contain a separate variable with this name. If you link an application with several such modules all containing (via a header file), for instance, the declaration:

```
volatile const __no_init int x @ 0x100;           /* Bad in C++ */
```

the linker will report that more than one variable is located at address 0x100.

To avoid this problem and make the process the same in C and C++, you should declare these variables `extern`, for example:

```
/* The extern keyword makes x public. */
extern volatile const __no_init int x @ 0x100;
```

Note: C++ static member variables can be placed at an absolute address just like any other static variable.

DATA AND FUNCTION PLACEMENT IN SECTIONS

The following method can be used for placing data or functions in named sections other than default:

- The `@` operator, alternatively the `#pragma location` directive, can be used for placing individual variables or individual functions in named sections. The named section can either be a predefined section, or a user-defined section.
- The `--section` option can be used for placing variables and functions, which are parts of the whole compilation unit, in named sections.

C++ static member variables can be placed in named sections just like any other static variable.

If you use your own sections, in addition to the predefined sections, the sections must also be defined in the linker configuration file.

Note: Take care when explicitly placing a variable or function in a predefined section other than the one used by default. This is useful in some situations, but incorrect placement can result in anything from error messages during compilation and linking to a malfunctioning application. Carefully consider the circumstances; there might be strict requirements on the declaration and use of the function or variable.

The location of the sections can be controlled from the linker configuration file.

For more information about sections, see the chapter *Section reference*.

Examples of placing variables in named sections

In the following examples, a data object is placed in a user-defined section. Note that you must as always ensure that the section is placed in the appropriate memory area when linking.

```
__no_init int alpha @ "MY_NOINIT"; /* OK */

#pragma location="MY_CONSTANTS"
const int beta = 42;                /* OK */

const int gamma @ "MY_CONSTANTS" = 17; /* OK */
int theta @ "MY_ZEROS";              /* OK */
int phi @ "MY_INITED" = 4711;        /* OK */
```

The linker will normally arrange for the correct type of initialization for each variable. If you want to control or suppress automatic initialization, you can use the `initialize` and `do not initialize` directives in the linker configuration file.

Examples of placing functions in named sections

```
void f(void) @ "MY_FUNCTIONS";

void g(void) @ "MY_FUNCTIONS"
{
}

#pragma location="MY_FUNCTIONS"
void h(void);
```

DATA PLACEMENT IN REGISTERS

The `@` operator, alternatively the `#pragma location` directive, can be used for placing global and static variables in a register.

To place a variable in a register, the argument to the `@` operator and the `#pragma location` directive should be an identifier that corresponds to an ARM core register in

the range R4–R11 (R9 cannot be specified in combination with the `--rwp1` command line option).

A variable can be placed in a register only if it is declared as `__no_init`, has file scope, and its size is four bytes. A variable placed in a register does not have a memory address, so the address operator `&` cannot be used.

Within a module where a variable is placed in a register, the specified register will only be used for accessing that variable. The value of the variable is preserved across function calls to other modules because the registers R14–R11 are callee saved, and as such they are restored when execution returns. However, the value of a variable placed in a register is not always preserved as expected:

- In an exception handler or library callback routine (such as the comparator function passed to `qsort`) the value might not be preserved. The value will be preserved if the command line option `--lock_regs` is used for locking the register in all modules of the application, including library modules.
- In a fast interrupt handler, the value of a variable in R8–R11 is not preserved from outside the handler, because these registers are banked.
- The `longjmp` function and C++ exceptions might restore variables placed in registers to old values, unlike other variables with static storage duration which are not restored.

The linker does not prevent modules from placing different variables in the same register. Variables in different modules can be placed in the same register, and another module could use the register for other purposes.

Note: A variable placed in a register should be defined in an include file, to be included in every module that uses the variable. An unused definition in a module will cause the register to not be used in that module.

Controlling compiler optimizations

The compiler performs many transformations on your application to generate the best possible code. Examples of such transformations are storing values in registers instead of memory, removing superfluous code, reordering computations in a more efficient order, and replacing arithmetic operations by cheaper operations.

The linker should also be considered an integral part of the compilation system, because some optimizations are performed by the linker. For instance, all unused functions and variables are removed and not included in the final output.

SCOPE FOR PERFORMED OPTIMIZATIONS

You can decide whether optimizations should be performed on your whole application or on individual files. By default, the same types of optimizations are used for an entire project, but you should consider using different optimization settings for individual files. For example, put code that must execute very quickly into a separate file and compile it for minimal execution time, and the rest of the code for minimal code size. This will give a small program, which is still fast enough where it matters.

You can also exclude individual functions from the performed optimizations. The `#pragma optimize` directive allows you to either lower the optimization level, or specify another type of optimization to be performed. See *optimize*, page 372, for information about the pragma directive.

MULTI-FILE COMPILATION UNITS

In addition to applying different optimizations to different source files or even functions, you can also decide what a compilation unit consists of—one or several source code files.

By default, a compilation unit consists of one source file, but you can also use multi-file compilation to make several source files in a compilation unit. The advantage is that interprocedural optimizations such as inlining and cross jump have more source code to work on. Ideally, the whole application should be compiled as one compilation unit. However, for large applications this is not practical because of resource restrictions on the host computer. For more information, see *--mfc*, page 278.

Note: Only one object file is generated, and thus all symbols will be part of that object file.

If the whole application is compiled as one compilation unit, it is very useful to make the compiler also discard unused public functions and variables before the interprocedural optimizations are performed. Doing this limits the scope of the optimizations to functions and variables that are actually used. For more information, see *--discard_unused_publics*, page 268.

OPTIMIZATION LEVELS

The compiler supports different levels of optimizations. This table lists optimizations that are typically performed on each level:

Optimization level	Description
None (Best debug support)	Variables live through their entire scope
Low	Same as above but variables only live for as long as they are needed, not necessarily through their entire scope
Medium	Same as above, and: Live-dead analysis and optimization Dead code elimination Redundant label elimination Redundant branch elimination Code hoisting Peephole optimization Some register content analysis and optimization Common subexpression elimination Code motion Static clustering
High (Balanced)	Same as above, and: Instruction scheduling Cross jumping Advanced register content analysis and optimization Loop unrolling Function inlining Type-based alias analysis

Table 24: Compiler optimization levels

Note: Some of the performed optimizations can be individually enabled or disabled. For more information about these, see *Fine-tuning enabled transformations*, page 231.

A high level of optimization might result in increased compile time, and will most likely also make debugging more difficult, because it is less clear how the generated code relates to the source code. For example, at the low, medium, and high optimization levels, variables do not live through their entire scope, which means processor registers used for storing variables can be reused immediately after they were last used. Due to this, the C-SPY **Watch** window might not be able to display the value of the variable throughout its scope. At any time, if you experience difficulties when debugging your code, try lowering the optimization level.

SPEED VERSUS SIZE

At the high optimization level, the compiler balances between size and speed optimizations. However, it is possible to fine-tune the optimizations explicitly for either size or speed. They only differ in what thresholds that are used; speed will trade size for speed, whereas size will trade speed for size. Note that one optimization sometimes enables other optimizations to be performed, and an application might in some cases become smaller even when optimizing for speed rather than size.

If you use the optimization level High speed, the `--no_size_constraints` compiler option relaxes the normal restrictions for code size expansion and enables more aggressive optimizations.

You can choose an optimization goal for each module, or even individual functions, using command line options and pragma directives (see *-O*, page 287 and *optimize*, page 372). For a small embedded application, this makes it possible to achieve acceptable speed performance while minimizing the code size: Typically, only a few places in the application need to be fast, such as the most frequently executed inner loops, or the interrupt handlers.

Rather than compiling the whole application with High (Balanced) optimization, you can use High (Size) in general, but override this to get High (Speed) optimization only for those functions where the application needs to be fast.

Because of the unpredictable way in which different optimizations interact, where one optimization can enable other optimizations, sometimes a function becomes smaller when compiled with High (Speed) optimization than if High (Size) is used. Also, using multi-file compilation (see *--m/c*, page 278) can enable many optimizations to improve both speed and size performance. It is recommended that you experiment with different optimization settings so that you can pick the best ones for your project.

FINE-TUNING ENABLED TRANSFORMATIONS

At each optimization level you can disable some of the transformations individually. To disable a transformation, use either the appropriate option, for instance the command line option `--no_inline`, alternatively its equivalent in the IDE **Function inlining**, or the `#pragma optimize` directive. These transformations can be disabled individually:

- Common subexpression elimination
- Loop unrolling
- Function inlining
- Code motion
- Type-based alias analysis
- Static clustering

Common subexpression elimination

Redundant re-evaluation of common subexpressions is by default eliminated at optimization levels Medium and High. This optimization normally reduces both code size and execution time. However, the resulting code might be difficult to debug.

Note: This option has no effect at optimization levels None and Low.

For more information about the command line option, see `--no_cse`, page 280.

Loop unrolling

Loop unrolling means that the code body of a loop, whose number of iterations can be determined at compile time, is duplicated. Loop unrolling reduces the loop overhead by amortizing it over several iterations.

This optimization is most efficient for smaller loops, where the loop overhead can be a substantial part of the total loop body.

Loop unrolling, which can be performed at optimization level High, normally reduces execution time, but increases code size. The resulting code might also be difficult to debug.

The compiler heuristically decides which loops to unroll. Only relatively small loops where the loop overhead reduction is noticeable will be unrolled. Different heuristics are used when optimizing for speed, size, or when balancing between size and speed.

Note: This option has no effect at optimization levels None, Low, and Medium.

To disable loop unrolling, use the command line option `--no_unroll`, see `--no_unroll`, page 286.

Function inlining

Function inlining means that a function, whose definition is known at compile time, is integrated into the body of its caller to eliminate the overhead of the call. This optimization normally reduces execution time, but might increase the code size.

For more information, see *Inlining functions*, page 79.

Code motion

Evaluation of loop-invariant expressions and common subexpressions are moved to avoid redundant re-evaluation. This optimization, which is performed at optimization level Medium and above, normally reduces code size and execution time. The resulting code might however be difficult to debug.

Note: This option has no effect at optimization levels below Medium.

For more information about the command line option, see `--no_code_motion`, page 279.

Type-based alias analysis

When two or more pointers reference the same memory location, these pointers are said to be *aliases* for each other. The existence of aliases makes optimization more difficult because it is not necessarily known at compile time whether a particular value is being changed.

Type-based alias analysis optimization assumes that all accesses to an object are performed using its declared type or as a `char` type. This assumption lets the compiler detect whether pointers can reference the same memory location or not.

Type-based alias analysis is performed at optimization level High. For application code conforming to standard C or C++ application code, this optimization can reduce code size and execution time. However, non-standard C or C++ code might result in the compiler producing code that leads to unexpected behavior. Therefore, it is possible to turn this optimization off.

Note: This option has no effect at optimization levels None, Low, and Medium.

For more information about the command line option, see `--no_tbaa`, page 285.

Example

```
short F(short *p1, long *p2)
{
    *p2 = 0;
    *p1 = 1;
    return *p2;
}
```

With type-based alias analysis, it is assumed that a write access to the `short` pointed to by `p1` cannot affect the `long` value that `p2` points to. Thus, it is known at compile time that this function returns 0. However, in non-standard-conforming C or C++ code these pointers could overlap each other by being part of the same union. If you use explicit casts, you can also force pointers of different pointer types to point to the same memory location.

Static clustering

When static clustering is enabled, static and global variables that are defined within the same module are arranged so that variables that are accessed in the same function are stored close to each other. This makes it possible for the compiler to use the same base pointer for several accesses.

Note: This option has no effect at optimization levels None and Low.

For more information about the command line option, see `--no_clustering`, page 279.

Instruction scheduling

The compiler features an instruction scheduler to increase the performance of the generated code. To achieve that goal, the scheduler rearranges the instructions to minimize the number of pipeline stalls emanating from resource conflicts within the microprocessor.

For more information about the command line option, see *--no_scheduling*, page 283.

Facilitating good code generation

This section contains hints on how to help the compiler generate good code:

- *Writing optimization-friendly source code*, page 234
- *Saving stack space and RAM memory*, page 235
- *Function prototypes*, page 235
- *Integer types and bit negation*, page 236
- *Protecting simultaneously accessed variables*, page 236
- *Accessing special function registers*, page 237
- *Passing values between C and assembler objects*, page 239
- *Non-initialized variables*, page 239

WRITING OPTIMIZATION-FRIENDLY SOURCE CODE

The following is a list of programming techniques that will, when followed, enable the compiler to better optimize the application.

- Local variables—auto variables and parameters—are preferred over static or global variables. The reason is that the optimizer must assume, for example, that called functions can modify non-local variables. When the life spans for local variables end, the previously occupied memory can then be reused. Globally declared variables will occupy data memory during the whole program execution.
- Avoid taking the address of local variables using the `&` operator. This is inefficient for two main reasons. First, the variable must be placed in memory, and thus cannot be placed in a processor register. This results in larger and slower code. Second, the optimizer can no longer assume that the local variable is unaffected over function calls.
- Module-local variables—variables that are declared static—are preferred over global variables (non-static). Also avoid taking the address of frequently accessed static variables.
- The compiler is capable of inlining functions, see *Function inlining*, page 232. To maximize the effect of the inlining transformation, it is good practice to place the

definitions of small functions called from more than one module in the header file rather than in the implementation file. Alternatively, you can use multi-file compilation. For more information, see *Multi-file compilation units*, page 229.

- Avoid using inline assembler without operands and clobbered resources. Instead, use SFRs or intrinsic functions if available. Otherwise, use inline assembler *with* operands and clobbered resources or write a separate module in assembler language. For more information, see *Mixing C and assembler*, page 157.

SAVING STACK SPACE AND RAM MEMORY

The following is a list of programming techniques that save memory and stack space:

- If stack space is limited, avoid long call chains and recursive functions.
- Avoid using large non-scalar types, such as structures, as parameters or return type. To save stack space, you should instead pass them as pointers or, in C++, as references.

FUNCTION PROTOTYPES

It is possible to declare and define functions using one of two different styles:

- Prototyped
- Kernighan & Ritchie C (K&R C)

Both styles are valid C, however it is strongly recommended to use the prototyped style, and provide a prototype declaration for each public function in a header that is included both in the compilation unit defining the function and in all compilation units using it.

The compiler will not perform type checking on parameters passed to functions declared using K&R style. Using prototype declarations will also result in more efficient code in some cases, as there is no need for type promotion for these functions.

To make the compiler require that all function definitions use the prototyped style, and that all public functions have been declared before being defined, use the

Project>Options>C/C++ Compiler>Language 1>Require prototypes compiler option (`--require_prototypes`).

Prototyped style

In prototyped function declarations, the type for each parameter must be specified.

```
int Test(char, int); /* Declaration */

int Test(char ch, int i) /* Definition */
{
    return i + ch;
}
```

Kernighan & Ritchie style

In K&R style—pre-Standard C—it is not possible to declare a function prototyped. Instead, an empty parameter list is used in the function declaration. Also, the definition looks different.

For example:

```
int Test();      /* Declaration */

int Test(ch, i) /* Definition */
char ch;
int i;
{
    return i + ch;
}
```

INTEGER TYPES AND BIT NEGATION

In some situations, the rules for integer types and their conversion lead to possibly confusing behavior. Things to look out for are assignments or conditionals (test expressions) involving types with different size, and logical operations, especially bit negation. Here, *types* also includes types of constants.

In some cases there might be warnings (for example, for constant conditional or pointless comparison), in others just a different result than what is expected. Under certain circumstances the compiler might warn only at higher optimizations, for example, if the compiler relies on optimizations to identify some instances of constant conditionals. In this example an 8-bit character, a 32-bit integer, and two's complement is assumed:

```
void F1(unsigned char c1)
{
    if (c1 == ~0x80)
        ;
}
```

Here, the test is always false. On the right hand side, `0x80` is `0x00000080`, and `~0x00000080` becomes `0xFFFFF7F`. On the left hand side, `c1` is an 8-bit unsigned character, so it cannot be larger than 255. It also cannot be negative, which means that the integral promoted value can never have the topmost 24 bits set.

PROTECTING SIMULTANEOUSLY ACCESSED VARIABLES

Variables that are accessed asynchronously, for example by interrupt routines or by code executing in separate threads, must be properly marked and have adequate protection. The only exception to this is a variable that is always *read-only*.

To mark a variable properly, use the `volatile` keyword. This informs the compiler, among other things, that the variable can be changed from other threads. The compiler will then avoid optimizing on the variable (for example, keeping track of the variable in registers), will not delay writes to it, and be careful accessing the variable only the number of times given in the source code.

For more information about the `volatile` type qualifier and the rules for accessing volatile objects, see *Declaring objects volatile*, page 341.

ACCESSING SPECIAL FUNCTION REGISTERS

Specific header files for several ARM devices are included in the IAR product installation. The header files are named `iodevice.h` and define the processor-specific special function registers (SFRs).

Note: Each header file contains one section used by the compiler, and one section used by the assembler.

SFRs with bitfields are declared in the header file. This example is from `ioks32c5000a.h`:

```
__no_init volatile union
{
    unsigned short mwctl2;
    struct
    {
        unsigned short edr: 1;
        unsigned short edw: 1;
        unsigned short lee: 2;
        unsigned short lemd: 2;
        unsigned short lepl: 2;
    } mwctl2bit;
} @ 0x1000;

/* By including the appropriate include file in your code,
 * it is possible to access either the whole register or any
 * individual bit (or bitfields) from C code as follows.
 */

void Test()
{
    /* Whole register access */
    mwctl2 = 0x1234;

    /* Bitfield accesses */
    mwctl2bit.edw = 1;
    mwctl2bit.lepl = 3;
}
```

You can also use the header files as templates when you create new header files for other ARM devices.

PASSING VALUES BETWEEN C AND ASSEMBLER OBJECTS

The following example shows how you in your C source code can use inline assembler to set and get values from a special purpose register:

```
static unsigned long get_APSR( void )
{
    unsigned long value;
    asm volatile( "MRS %0, APSR" : "=r"(value) );
    return value;
}

static void set_APSR( unsigned long value)
{
    asm volatile( "MSR APSR, %0" :: "r"(value) );
}
```

The general purpose register is used for getting and setting the value of the special purpose register APSR. The same method can be used also for accessing other special purpose registers and specific instructions.

To read more about inline assembler, see *Inline assembler*, page 158.

NON-INITIALIZED VARIABLES

Normally, the runtime environment will initialize all global and static variables when the application is started.

The compiler supports the declaration of variables that will not be initialized, using the `__no_init` type modifier. They can be specified either as a keyword or using the `#pragma object_attribute` directive. The compiler places such variables in a separate section.

For `__no_init`, the `const` keyword implies that an object is read-only, rather than that the object is stored in read-only memory. It is not possible to give a `__no_init` object an initial value.


Variables declared using the `__no_init` keyword could, for example, be large input buffers or mapped to special RAM that keeps its content even when the application is turned off.

For more information, see `__no_init`, page 351. Note that to use this keyword, language extensions must be enabled; see `-e`, page 269. For more information, see also `object_attribute`, page 371.

Part 2. Reference information

This part of the *IAR C/C++ Development Guide for ARM* contains these chapters:

- External interface details
- Compiler options
- Linker options
- Data representation
- Extended keywords
- Pragma directives
- Intrinsic functions
- The preprocessor
- C/C++ standard library functions
- The linker configuration file
- Section reference
- The stack usage control file
- IAR utilities
- Implementation-defined behavior for Standard C
- Implementation-defined behavior for C89.





External interface details

- Invocation syntax
- Include file search procedure
- Compiler output
- LINK output
- Diagnostics

Invocation syntax

You can use the compiler and linker either from the IDE or from the command line. See the *IDE Project Management and Building Guide for ARM* for information about using the build tools from the IDE.

COMPILER INVOCATION SYNTAX

The invocation syntax for the compiler is:

```
icarm [options] [sourcefile] [options]
```

For example, when compiling the source file `prog.c`, use this command to generate an object file with debug information:

```
icarm prog.c --debug
```

The source file can be a C or C++ file, typically with the filename extension `c` or `cpp`, respectively. If no filename extension is specified, the file to be compiled must have the extension `c`.

Generally, the order of options on the command line, both relative to each other and to the source filename, is not significant. There is, however, one exception: when you use the `-I` option, the directories are searched in the same order as they are specified on the command line.

If you run the compiler from the command line without any arguments, the compiler version number and all available options including brief descriptions are directed to `stdout` and displayed on the screen.

ILINK INVOCATION SYNTAX

The invocation syntax for ILINK is:

```
ilinkarm [arguments]
```

Each argument is either a command-line option, an object file, or a library.

For example, when linking the object file `prog.o`, use this command:

```
ilinkarm prog.o --config configfile
```

If no filename extension is specified for the linker configuration file, the configuration file must have the extension `icf`.

Generally, the order of arguments on the command line is not significant. There is, however, one exception: when you supply several libraries, the libraries are searched in the same order that they are specified on the command line. The default libraries are always searched last.

The output executable image will be placed in a file named `a.out`, unless the `-o` option is used.

If you run ILINK from the command line without any arguments, the ILINK version number and all available options including brief descriptions are directed to `stdout` and displayed on the screen.

PASSING OPTIONS

There are three different ways of passing options to the compiler and to ILINK:

- Directly from the command line
Specify the options on the command line after the `iccarm` or `ilinkarm` commands; see *Invocation syntax*, page 243.
- Via environment variables
The compiler and linker automatically append the value of the environment variables to every command line; see *Environment variables*, page 245.
- Via a text file, using the `-f` option; see *-f*, page 273.

For general guidelines for the option syntax, an options summary, and a detailed description of each option, see the chapter *Compiler options*.

ENVIRONMENT VARIABLES

These environment variables can be used with the compiler:

Environment variable	Description
C_INCLUDE	Specifies directories to search for include files; for example: C_INCLUDE=c:\program files\iar systems\embedded workbench 7.n\arm\inc;c:\headers
QCCARM	Specifies command line options; for example: QCCARM=-lA asm.lst

Table 25: Compiler environment variables

This environment variable can be used with ILINK:

Environment variable	Description
ILINKARM_CMD_LINE	Specifies command line options; for example: ILINKARM_CMD_LINE=--config full.icf --silent

Table 26: ILINK environment variables

Include file search procedure

This is a detailed description of the compiler's `#include` file search procedure:

- If the name of the `#include` file is an absolute path specified in angle brackets or double quotes, that file is opened.
- If the compiler encounters the name of an `#include` file in angle brackets, such as:

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

it searches these directories for the file to include:
 - 1 The directories specified with the `-I` option, in the order that they were specified, see *-I*, page 275.
 - 2 The directories specified using the `C_INCLUDE` environment variable, if any; see *Environment variables*, page 245.
 - 3 The automatically set up library system include directories. See *--dlib_config*, page 269.
- If the compiler encounters the name of an `#include` file in double quotes, for example:

```
#include "vars.h"
```

it searches the directory of the source file in which the `#include` statement occurs, and then performs the same sequence as for angle-bracketed filenames.

If there are nested `#include` files, the compiler starts searching the directory of the file that was last included, iterating upwards for each included file, searching the source file directory last. For example:

```
src.c in directory dir\src
#include "src.h"
...
src.h in directory dir\include
#include "config.h"
...
```

When `dir\exe` is the current directory, use this command for compilation:

```
icarm ..\src\src.c -I..\include -I..\debugconfig
```

Then the following directories are searched in the order listed below for the file `config.h`, which in this example is located in the `dir\debugconfig` directory:

<code>dir\include</code>	Current file is <code>src.h</code> .
<code>dir\src</code>	File including current file (<code>src.c</code>).
<code>dir\include</code>	As specified with the first <code>-I</code> option.
<code>dir\debugconfig</code>	As specified with the second <code>-I</code> option.

Use angle brackets for standard header files, like `stdio.h`, and double quotes for files that are part of your application.

Note: Both `\` and `/` can be used as directory delimiters.

For more information, see *Overview of the preprocessor*, page 419.

Compiler output

The compiler can produce the following output:

- A linkable object file
The object files produced by the compiler use the industry-standard format ELF. By default, the object file has the filename extension `.o`.
- Optional list files
Various kinds of list files can be specified using the compiler option `-l`, see `-l`, page 275. By default, these files will have the filename extension `lst`.
- Optional preprocessor output files
A preprocessor output file is produced when you use the `--preprocess` option; by default, the file will have the filename extension `i`.

- Diagnostic messages

Diagnostic messages are directed to the standard error stream and displayed on the screen, and printed in an optional list file. For more information about diagnostic messages, see *Diagnostics*, page 248.

- Error return codes

These codes provide status information to the operating system which can be tested in a batch file, see *Error return codes*, page 247.

- Size information

Information about the generated amount of bytes for functions and data for each memory is directed to the standard output stream and displayed on the screen. Some of the bytes might be reported as *shared*.

Shared objects are functions or data objects that are shared between modules. If any of these occur in more than one module, only one copy is retained. For example, in some cases inline functions are not inlined, which means that they are marked as shared, because only one instance of each function will be included in the final application. This mechanism is sometimes also used for compiler-generated code or data not directly associated with a particular function or variable, and when only one instance is required in the final application.

ERROR RETURN CODES

The compiler and linker return status information to the operating system that can be tested in a batch file.

These command line error codes are supported:

Code	Description
0	Compilation or linking successful, but there might have been warnings.
1	Warnings were produced and the option <code>--warnings_affect_exit_code</code> was used.
2	Errors occurred.
3	Fatal errors occurred, making the tool abort.
4	Internal errors occurred, making the tool abort.

Table 27: Error return codes

ILINK output

ILINK can produce the following output:

- An absolute executable image

The final output produced by the IAR ILINK Linker is an absolute object file containing the executable image that can be put into an EPROM, downloaded to a hardware emulator, or executed on your PC using the IAR C-SPY Debugger Simulator. By default, the file has the filename extension `out`. The output format is always in ELF, which optionally includes debug information in the DWARF format.

- Optional logging information

During operation, ILINK logs its decisions on `stdout`, and optionally to a file. For example, if a library is searched, whether a required symbol is found in a library module, or whether a module will be part of the output. Timing information for each ILINK subsystem is also logged.

- Optional map files

A linker map file—containing summaries of linkage, runtime attributes, memory, and placement, as well as an entry list—can be generated by the ILINK option `--map`, see `--map`, page 315. By default, the map file has the filename extension `map`.

- Diagnostic messages

Diagnostic messages are directed to `stderr` and displayed on the screen, as well as printed in the optional map file. For more information about diagnostic messages, see *Diagnostics*, page 248.

- Error return codes

ILINK returns status information to the operating system which can be tested in a batch file, see *Error return codes*, page 247.

- Size information about used memory and amount of time

Information about the generated amount of bytes for functions and data for each memory is directed to `stdout` and displayed on the screen.

Diagnostics

This section describes the format of the diagnostic messages and explains how diagnostic messages are divided into different levels of severity.

MESSAGE FORMAT FOR THE COMPILER

All diagnostic messages are issued as complete, self-explanatory messages. A typical diagnostic message from the compiler is produced in the form:

```
filename,linenumber level[tag]: message
```


with these elements:

<i>filename</i>	The name of the source file in which the issue was encountered
<i>linenumber</i>	The line number at which the compiler detected the issue
<i>level</i>	The level of seriousness of the issue
<i>tag</i>	A unique tag that identifies the diagnostic message
<i>message</i>	An explanation, possibly several lines long

Diagnostic messages are displayed on the screen, as well as printed in the optional list file.

Use the option `--diagnostics_tables` to list all possible compiler diagnostic messages.

MESSAGE FORMAT FOR THE LINKER

All diagnostic messages are issued as complete, self-explanatory messages. A typical diagnostic message from ILINK is produced in the form:

```
level[tag]: message
```

with these elements:

<i>level</i>	The level of seriousness of the issue
<i>tag</i>	A unique tag that identifies the diagnostic message
<i>message</i>	An explanation, possibly several lines long

Diagnostic messages are displayed on the screen, as well as printed in the optional map file.

Use the option `--diagnostics_tables` to list all possible linker diagnostic messages.

SEVERITY LEVELS

The diagnostic messages are divided into different levels of severity:

Remark

A diagnostic message that is produced when the compiler or linker finds a construct that can possibly lead to erroneous behavior in the generated code. Remarks are by default not issued, but can be enabled, see `--remarks`, page 290.

Warning

A diagnostic message that is produced when the compiler or linker finds a potential problem which is of concern, but which does not prevent completion of the compilation or linking. Warnings can be disabled by use of the command line option

`--no_warnings`, see *--no_warnings*, page 287.

Error

A diagnostic message that is produced when the compiler or linker finds a serious error. An error will produce a non-zero exit code.

Fatal error

A diagnostic message that is produced when the compiler finds a condition that not only prevents code generation, but which makes further processing pointless. After the message is issued, compilation terminates. A fatal error will produce a non-zero exit code.

SETTING THE SEVERITY LEVEL

The diagnostic messages can be suppressed or the severity level can be changed for all diagnostics messages, except for fatal errors and some of the regular errors.

See the chapter *Compiler options*, for information about the compiler options that are available for setting severity levels.

For the compiler see also the chapter *Pragma directives*, for information about the pragma directives that are available for setting severity levels.

INTERNAL ERROR

An internal error is a diagnostic message that signals that there was a serious and unexpected failure due to a fault in the compiler or linker. It is produced using this form:

`Internal error: message`

where *message* is an explanatory message. If internal errors occur, they should be reported to your software distributor or IAR Systems Technical Support. Include enough information to reproduce the problem, typically:

- The product name
- The version number of the compiler or of ILINK, which can be seen in the header of the list or map files generated by the compiler or by ILINK, respectively
- Your license number
- The exact internal error message text
- The files involved of the application that generated the internal error

- A list of the options that were used when the internal error occurred.

Compiler options

- Options syntax
- Summary of compiler options
- Descriptions of compiler options

Options syntax

Compiler options are parameters you can specify to change the default behavior of the compiler. You can specify options from the command line—which is described in more detail in this section—and from within the IDE.



See the online help system for information about the compiler options available in the IDE and how to set them.

TYPES OF OPTIONS

There are two *types of names* for command line options, *short* names and *long* names. Some options have both.

- A short option name consists of one character, and it can have parameters. You specify it with a single dash, for example `-e`
- A long option name consists of one or several words joined by underscores, and it can have parameters. You specify it with double dashes, for example `--char_is_signed`.

For information about the different methods for passing options, see *Passing options*, page 244.

RULES FOR SPECIFYING PARAMETERS

There are some general syntax rules for specifying option parameters. First, the rules depending on whether the parameter is *optional* or *mandatory*, and whether the option has a short or a long name, are described. Then, the rules for specifying filenames and directories are listed. Finally, the remaining rules are listed.

Rules for optional parameters

For options with a short name and an optional parameter, any parameter should be specified without a preceding space, for example:

`-O` or `-Oh`

For options with a long name and an optional parameter, any parameter should be specified with a preceding equal sign (=), for example:

```
--misrac2004=n
```

Rules for mandatory parameters

For options with a short name and a mandatory parameter, the parameter can be specified either with or without a preceding space, for example:

```
-I..src or -I ..src\
```

For options with a long name and a mandatory parameter, the parameter can be specified either with a preceding equal sign (=) or with a preceding space, for example:

```
--diagnostics_tables=MyDiagnostics.lst
```

or

```
--diagnostics_tables MyDiagnostics.lst
```

Rules for options with both optional and mandatory parameters

For options taking both optional and mandatory parameters, the rules for specifying the parameters are:

- For short options, optional parameters are specified without a preceding space
- For long options, optional parameters are specified with a preceding equal sign (=)
- For short and long options, mandatory parameters are specified with a preceding space.

For example, a short option with an optional parameter followed by a mandatory parameter:

```
-lA MyList.lst
```

For example, a long option with an optional parameter followed by a mandatory parameter:

```
--preprocess=n PreprocOutput.lst
```

Rules for specifying a filename or directory as parameters

These rules apply for options taking a filename or directory as parameters:

- Options that take a filename as a parameter can optionally take a file path. The path can be relative or absolute. For example, to generate a listing to the file `List.lst` in the directory `..\listings\`:

```
iccarm prog.c -l ..\listings\List.lst
```

- For options that take a filename as the destination for output, the parameter can be specified as a path without a specified filename. The compiler stores the output in that directory, in a file with an extension according to the option. The filename will be the same as the name of the compiled source file, unless a different name was specified with the option `-o`, in which case that name is used. For example:

```
iccam prog.c -l ../listings\
```

The produced list file will have the default name `../listings\prog.lst`
- The *current directory* is specified with a period (`.`). For example:

```
iccam prog.c -l .
```
- `/` can be used instead of `\` as the directory delimiter.
- By specifying `-`, input files and output files can be redirected to the standard input and output stream, respectively. For example:

```
iccam prog.c -l -
```

Additional rules

These rules also apply:

- When an option takes a parameter, the parameter cannot start with a dash (`-`) followed by another character. Instead, you can prefix the parameter with two dashes; this example will create a list file called `-r`:

```
iccam prog.c -l ---r
```
- For options that accept multiple arguments of the same type, the arguments can be provided as a comma-separated list (without a space), for example:

```
--diag_warning=Be0001,Be0002
```

Alternatively, the option can be repeated for each argument, for example:

```
--diag_warning=Be0001
--diag_warning=Be0002
```

Summary of compiler options

This table summarizes the compiler command line options:

Command line option	Description
<code>--aapcs</code>	Specifies the calling convention
<code>--aeabi</code>	Enables AEABI-compliant code generation
<code>--align_sp_on_irq</code>	Generates code to align SP on entry to <code>__irq</code> functions
<code>--arm</code>	Sets the default function mode to ARM

Table 28: Compiler options summary

Command line option	Description
--c89	Specifies the C89 dialect
--char_is_signed	Treats <code>char</code> as signed
--char_is_unsigned	Treats <code>char</code> as unsigned
--cpu	Specifies a processor variant
--cpu_mode	Specifies the default CPU mode for functions
--c++	Specifies Standard C++
-D	Defines preprocessor symbols
--debug	Generates debug information
--dependencies	Lists file dependencies
--diag_error	Treats these as errors
--diag_remark	Treats these as remarks
--diag_suppress	Suppresses these diagnostics
--diag_warning	Treats these as warnings
--diagnostics_tables	Lists all diagnostic messages
--discard_unused_publics	Discards unused public symbols
--dlib_config	Uses the system include files for the DLIB library and determines which configuration of the library to use
-e	Enables language extensions
--ec++	Specifies Embedded C++
--eec++	Specifies Extended Embedded C++
--enable_hardware_workaround	Enables a specific hardware workaround
--enable_multibytes	Enables support for multibyte characters in source files
--enable_restrict	Enables the Standard C keyword <code>restrict</code>
--endian	Specifies the byte order of the generated code and data
--enum_is_int	Sets the minimum size on enumeration types
--error_limit	Specifies the allowed number of errors before compilation stops
-f	Extends the command line
--fpu	Selects the type of floating-point unit

Table 28: Compiler options summary (Continued)

Command line option	Description
<code>--generate_entries_without_bonds</code>	Generates extra functions for use from non-instrumented code. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>--guard_calls</code>	Enables guards for function static variable initialization
<code>--header_context</code>	Lists all referred source files and header files
<code>-I</code>	Specifies include file path
<code>--ignore_uninstrumented_pointers</code>	Disables checking of accesses via pointers from non-instrumented code. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>--interwork</code>	Generates interworking code
<code>-l</code>	Creates a list file
<code>--legacy</code>	Generates object code linkable with older tool chains
<code>--lock_regs</code>	Prevents the compiler from using specified registers
<code>--macro_positions_in_diagnostics</code>	Obtains positions inside macros in diagnostic messages
<code>--make_all_definitions_weak</code>	Turns all variable and function definitions into weak definitions.
<code>--mfc</code>	Enables multi-file compilation
<code>--misrac</code>	Enables error messages specific to MISRA-C:1998. This option is a synonym of <code>--misrac1998</code> and is only available for backwards compatibility.
<code>--misrac1998</code>	Enables error messages specific to MISRA-C:1998. See the <i>IAR Embedded Workbench® MISRA C:1998 Reference Guide</i> .
<code>--misrac2004</code>	Enables error messages specific to MISRA-C:2004. See the <i>IAR Embedded Workbench® MISRA C:2004 Reference Guide</i> .
<code>--misrac_verbose</code>	<i>IAR Embedded Workbench® MISRA C:1998 Reference Guide</i> or the <i>IAR Embedded Workbench® MISRA C:2004 Reference Guide</i> .

Table 28: Compiler options summary (Continued)

Command line option	Description
<code>--no_alignment_reduction</code>	Disables alignment reduction for simple Thumb functions
<code>--no_call_frame_info</code>	Disables output of call frame information
<code>--no_clustering</code>	Disables static clustering optimizations
<code>--no_code_motion</code>	Disables code motion optimization
<code>--no_const_align</code>	Disables the alignment optimization for constants.
<code>--no_cse</code>	Disables common subexpression elimination
<code>--no_exceptions</code>	Disables C++ exception support
<code>--no_fragments</code>	Disables section fragment handling
<code>--no_inline</code>	Disables function inlining
<code>--no_literal_pool</code>	Generates code that should run from a memory region where it is not allowed to read data, only to execute code
<code>--no_loop_align</code>	Disables the alignment of labels in loops
<code>--no_mem_idioms</code>	Makes the compiler not optimize certain memory access patterns
<code>--no_path_in_file_macros</code>	Removes the path from the return value of the symbols <code>__FILE__</code> and <code>__BASE_FILE__</code>
<code>--no_rtti</code>	Disables C++ RTTI support
<code>--no_rw_dynamic_init</code>	Disables runtime initialization of static C variables.
<code>--no_scheduling</code>	Disables the instruction scheduler
<code>--no_size_constraints</code>	Relaxes the normal restrictions for code size expansion when optimizing for speed.
<code>--no_static_destruction</code>	Disables destruction of C++ static variables at program exit
<code>--no_system_include</code>	Disables the automatic search for system include files
<code>--no_tbaa</code>	Disables type-based alias analysis
<code>--no_typedefs_in_diagnostics</code>	Disables the use of typedef names in diagnostics
<code>--no_unaligned_access</code>	Avoids unaligned accesses
<code>--no_unroll</code>	Disables loop unrolling
<code>--no_var_align</code>	Aligns variable objects based on the alignment of their type.
<code>--no_warnings</code>	Disables all warnings

Table 28: Compiler options summary (Continued)

Command line option	Description
<code>--no_wrap_diagnostics</code>	Disables wrapping of diagnostic messages
<code>-O</code>	Sets the optimization level
<code>-o</code>	Sets the object filename. Alias for <code>--output</code> .
<code>--only_stdout</code>	Uses standard output only
<code>--output</code>	Sets the object filename
<code>--predef_macros</code>	Lists the predefined symbols.
<code>--preinclude</code>	Includes an include file before reading the source file
<code>--preprocess</code>	Generates preprocessor output
<code>--public_equ</code>	Defines a global named assembler label
<code>-r</code>	Generates debug information. Alias for <code>--debug</code> .
<code>--relaxed_fp</code>	Relaxes the rules for optimizing floating-point expressions
<code>--remarks</code>	Enables remarks
<code>--require_prototypes</code>	Verifies that functions are declared before they are defined
<code>--ropi</code>	Generates code that uses PC-relative references to address code and read-only data.
<code>--runtime_checking</code>	Enables runtime error checking. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>--rwpi</code>	Generates code that uses an offset from the static base register to address-writable data.
<code>--section</code>	Changes a section name
<code>--separate_cluster_for_initialized_variables</code>	Separates initialized and non-initialized variables
<code>--silent</code>	Sets silent operation
<code>--strict</code>	Checks for strict compliance with Standard C/C++
<code>--system_include_dir</code>	Specifies the path for system include files
<code>--thumb</code>	Sets default function mode to Thumb
<code>--use_c++_inline</code>	Uses C++ inline semantics in C99
<code>--use_unix_directory_separators</code>	Uses <code>/</code> as directory separator in paths

Table 28: Compiler options summary (Continued)

Command line option	Description
--vectorize	Enables generation of NEON vector instructions
--vla	Enables C99 VLA support
--warn_about_c_style_casts	Makes the compiler warn when C-style casts are used in C++ source code
--warnings_affect_exit_code	Warnings affect exit code
--warnings_are_errors	Warnings are treated as errors

Table 28: Compiler options summary (Continued)


Descriptions of compiler options

The following section gives detailed reference information about each compiler option.



Note that if you use the options page **Extra Options** to specify specific command line options, the IDE does not perform an instant check for consistency problems like conflicting options, duplication of options, or use of irrelevant options.

--aapcs

Syntax	aapcs={std vfp}	
Parameters	std	Processor registers are used for floating-point parameters and return values in function calls according to standard AAPCS. <code>std</code> is the default when the software FPU is selected.
	vfp	VFP registers are used for floating-point parameters and return values. The generated code is not compatible with AEABI code. <code>vfp</code> is the default when a VFP unit is used.
Description	Use this option to specify the floating-point calling convention.	
		To set this option, use Project>Options>C/C++ Compiler>Extra Options .

--aeabi

Syntax	--aeabi
Description	Use this option to generate AEABI-compliant object code. Note that this option must be used together with the <code>--guard_calls</code> option.

See also

AEABI compliance, page 217 and *--guard_calls*, page 274.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--align_sp_on_irq

Syntax

`--align_sp_on_irq`

Description

Use this option to align the stack pointer (SP) on entry to `__irq` declared functions.

This is especially useful for nested interrupts, where the interrupted code uses the same SP as the interrupt handler. This means that the stack might only have 4-byte alignment, instead of the 8-byte alignment required by AEABI (and some instructions generated by the compiler for some cores).

See also

`__irq`, page 348.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--arm

Syntax

`--arm`

Description

Use this option to set default function mode to ARM. This setting must be the same for all files included in a program, unless they are interworking.

Note: This option has the same effect as the `--cpu_mode=arm` option.

See also

--interwork, page 275 and *__interwork*, page 348.



Project>Options>C/C++ Compiler>Code>Processor mode>Arm

--c89

Syntax

`--c89`

Description

Use this option to enable the C89 C dialect instead of Standard C.

Note: This option is mandatory when the MISRA C checking is enabled.

See also *C language overview*, page 181.



Project>Options>C/C++ Compiler>Language 1>C dialect>C89

--char_is_signed

Syntax `--char_is_signed`

Description By default, the compiler interprets the plain `char` type as unsigned. Use this option to make the compiler interpret the plain `char` type as signed instead. This can be useful when you, for example, want to maintain compatibility with another compiler.

Note: The runtime library is compiled without the `--char_is_signed` option and cannot be used with code that is compiled with this option.



Project>Options>C/C++ Compiler>Language 2>Plain ‘char’ is

--char_is_unsigned

Syntax `--char_is_unsigned`

Description Use this option to make the compiler interpret the plain `char` type as unsigned. This is the default interpretation of the plain `char` type.



Project>Options>C/C++ Compiler>Language 2>Plain ‘char’ is

--cpu

Syntax `--cpu=core`

Parameters `core` Specifies a specific processor variant

Description Use this option to select the processor variant for which the code is to be generated. The default is ARM7TDMI. The following cores and processor macrocells are recognized:

ARM7TDMI	ARM1020E	Cortex-A17
ARM7TDMI-S	ARM1022E	Cortex-M0

ARM710T	ARM1026EJ-S	Cortex-M0+
ARM720T	ARM1136J	Cortex-M1
ARM740T	ARM1136J-S	Cortex-Ms1*
ARM7EJ-S	ARM1136JF	Cortex-M3
ARM9TDMI	ARM1136JF-S	Cortex-M4
ARM920T	ARM1176J	Cortex-M4F
	(alias for ARM1176JZ)	
ARM922T	ARM1176J-S	Cortex-M7
	(alias for ARM1176JZ-S)	
ARM940T	ARM1176JF	Cortex-R4
	(alias for ARM1176JZF)	
ARM9E	ARM1176JF-S	Cortex-R4F
	(alias for ARM1176JZF-S)	
ARM9E-S	Cortex-A5	Cortex-R5
ARM926EJ-S	Cortex-A5F	Cortex-R5F
ARM966E-S	Cortex-A7	Cortex-R7
ARM968E-S	Cortex-A8	Cortex-R7F
ARM946E-S	Cortex-A9	XScale
ARM10E	Cortex-A15	XScale-IR7

* Cortex-M1 with Operating System extension.

See also

Processor variant, page 64



Project>Options>General Options>Target>Processor configuration

--cpu_mode

Syntax

--cpu_mode={ arm | a | thumb | t }

Parameters

arm, a (default) Selects the arm mode as the default mode for functions

`thumb, t` Selects the thumb mode as the default mode for functions

Description Use this option to select the default mode for functions. This setting must be the same for all files included in a program, unless they are interworking.

See also *--interwork*, page 275 and *__interwork*, page 348.



Project>Options>General Options>Target>Processor mode

--c++

Syntax `--c++`

Description By default, the language supported by the compiler is C. If you use Standard C++, you must use this option to set the language the compiler uses to C++.

See also *--ec++*, page 270, *--eec++*, page 270, and *Using C++*, page 191.



Project>Options>C/C++ Compiler>Language 1>C++

and

Project>Options>C/C++ Compiler>Language 1>C++ dialect>C++

-D

Syntax `-D symbol [=value]`

Parameters

symbol The name of the preprocessor symbol

value The value of the preprocessor symbol

Description Use this option to define a preprocessor symbol. If no value is specified, 1 is used. This option can be used one or more times on the command line.

The option `-D` has the same effect as a `#define` statement at the top of the source file:

`-Dsymbol`

is equivalent to:

`#define symbol 1`

To get the equivalence of:

```
#define FOO
```

specify the = sign but nothing after, for example:

```
-DFOO=
```



Project>Options>C/C++ Compiler>Preprocessor>Defined symbols

--debug, -r

Syntax

```
--debug  
-r
```

Description

Use the `--debug` or `-r` option to make the compiler include information in the object modules required by the IAR C-SPY® Debugger and other symbolic debuggers.

Note: Including debug information will make the object files larger than otherwise.



Project>Options>C/C++ Compiler>Output>Generate debug information

--dependencies

Syntax

```
--dependencies [= [i|m|n] [s]] {filename|directory|+}
```

Parameters

i (default)	Lists only the names of files
m	Lists in makefile style (multiple rules)
n	Lists in makefile style (one rule)
s	Suppresses system files
+	Gives the same output as <code>-o</code> , but with the filename extension <code>d</code>

See also *Rules for specifying a filename or directory as parameters*, page 254.

Description

Use this option to make the compiler list the names of all source and header files opened for input into a file with the default filename extension `i`.

Example

If `--dependencies` or `--dependencies=i` is used, the name of each opened input file, including the full path, if available, is output on a separate line. For example:

```
c:\iar\product\include\stdio.h
d:\myproject\include\foo.h
```

If `--dependencies=m` is used, the output is in makefile style. For each input file, one line containing a makefile dependency rule is produced. Each line consists of the name of the object file, a colon, a space, and the name of an input file. For example:

```
foo.o: c:\iar\product\include\stdio.h
foo.o: d:\myproject\include\foo.h
```

An example of using `--dependencies` with a popular make utility, such as `gmake` (GNU make):

- 1 Set up the rule for compiling files to be something like:

```
%.o : %.c
$(ICC) $(ICCFLAGS) $< --dependencies=m $*.d
```

That is, in addition to producing an object file, the command also produces a dependency file in makefile style (in this example, using the extension `.d`).

- 2 Include all the dependency files in the makefile using, for example:

```
-include $(sources:.c=.d)
```

Because of the dash (`-`) it works the first time, when the `.d` files do not yet exist.



This option is not available in the IDE.

--diag_error

Syntax

```
--diag_error=tag[, tag, ...]
```

Parameters

<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe117</code>
------------	---

Description

Use this option to reclassify certain diagnostic messages as errors. An error indicates a violation of the C or C++ language rules, of such severity that object code will not be generated. The exit code will be non-zero. This option may be used more than once on the command line.



Project>Options>C/C++ Compiler>Diagnostics>Treat these as errors

--diag_remark

Syntax	<code>--diag_remark=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe177</code>
Description	<p>Use this option to reclassify certain diagnostic messages as remarks. A remark is the least severe type of diagnostic message and indicates a source code construction that may cause strange behavior in the generated code. This option may be used more than once on the command line.</p> <p>Note: By default, remarks are not displayed; use the <code>--remarks</code> option to display them.</p>	



Project>Options>C/C++ Compiler>Diagnostics>Treat these as remarks

--diag_suppress

Syntax	<code>--diag_suppress=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe117</code>
Description	<p>Use this option to suppress certain diagnostic messages. These messages will not be displayed. This option may be used more than once on the command line.</p>	



Project>Options>C/C++ Compiler>Diagnostics>Suppress these diagnostics

--diag_warning

Syntax	<code>--diag_warning=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe826</code>
Description	<p>Use this option to reclassify certain diagnostic messages as warnings. A warning indicates an error or omission that is of concern, but which will not cause the compiler</p>	

to stop before compilation is completed. This option may be used more than once on the command line.



Project>Options>C/C++ Compiler>Diagnostics>Treat these as warnings

--diagnostics_tables

Syntax	<code>--diagnostics_tables {filename directory}</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	<p>Use this option to list all possible diagnostic messages in a named file. This can be convenient, for example, if you have used a pragma directive to suppress or change the severity level of any diagnostic messages, but forgot to document why.</p> <p>This option cannot be given together with other options.</p>



This option is not available in the IDE.

--discard_unused_publics

Syntax	<code>--discard_unused_publics</code>
Description	<p>Use this option to discard unused public functions and variables when compiling with the <code>--mfc</code> compiler option.</p> <p>Note: Do not use this option only on parts of the application, as necessary symbols might be removed from the generated output. Use the object attribute <code>__root</code> to keep symbols that are used from outside the compilation unit, for example interrupt handlers. If the symbol does not have the <code>__root</code> attribute and is defined in the library, the library definition will be used instead.</p>
See also	<code>--mfc</code> , page 278 and <i>Multi-file compilation units</i> , page 229.



Project>Options>C/C++ Compiler>Discard unused publics

--dlib_config

Syntax	<code>--dlib_config filename.h config</code>	
Parameters	<i>filename</i>	A DLIB configuration header file, see <i>Rules for specifying a filename or directory as parameters</i> , page 254.
	<i>config</i>	The default configuration file for the specified configuration will be used. Choose between:
		<code>none</code> , no configuration will be used
		<code>normal</code> , the normal library configuration will be used (default)
		<code>full</code> , the full library configuration will be used.
Description	Use this option to specify which library configuration to use, either by specifying an explicit file or by specifying a library configuration—in which case the default file for that library configuration will be used. Make sure that you specify a configuration that corresponds to the library you are using. If you do not specify this option, the default library configuration file will be used.	
	All prebuilt runtime libraries are delivered with corresponding configuration files. You can find the library object files and the library configuration files in the directory <code>arm\lib</code> . For examples and information about prebuilt runtime libraries, see <i>Prebuilt runtime libraries</i> , page 126.	
	If you build your own customized runtime library, you should also create a corresponding customized library configuration file, which must be specified to the compiler. For more information, see <i>Customizing and building your own runtime library</i> , page 123.	



To set related options, choose:

Project>Options>General Options>Library Configuration

-e

Syntax	<code>-e</code>
Description	In the command line version of the compiler, language extensions are disabled by default. If you use language extensions such as extended keywords and anonymous structs and unions in your source code, you must use this option to enable them.

Note: The `-e` option and the `--strict` option cannot be used at the same time.

See also

Enabling language extensions, page 183.



Project>Options>C/C++ Compiler>Language 1>Standard with IAR extensions

Note: By default, this option is selected in the IDE.

--ec++

Syntax

`--ec++`

Description

In the compiler, the default language is C. If you use Embedded C++, you must use this option to set the language the compiler uses to Embedded C++.



Project>Options>C/C++ Compiler>Language 1>C++

and

Project>Options>C/C++ Compiler>Language 1>C++ dialect>Embedded C++

--eec++

Syntax

`--eec++`

Description

In the compiler, the default language is C. If you take advantage of Extended Embedded C++ features like namespaces or the standard template library in your source code, you must use this option to set the language the compiler uses to Extended Embedded C++.

See also

Extended Embedded C++, page 192.



Project>Options>C/C++ Compiler>Language 1>C++


and

Project>Options>C/C++ Compiler>Language 1>C++ dialect>Extended Embedded C++


--enable_hardware_workaround

Syntax


`--enable_hardware_workaround=wait[,wait...]`

Parameters	<i>waid</i>	The ID number of the workaround to enable. For a list of available workarounds to enable, see the release notes.
Description	Use this option to make the compiler generate a workaround for a specific hardware problem.	
See also	The release notes for the compiler for a list of available parameters.	
		To set this option, use Project>Options>C/C++ Compiler>Extra Options .

--enable_multibytes

Syntax	<code>--enable_multibytes</code>	
Description	By default, multibyte characters cannot be used in C or C++ source code. Use this option to make multibyte characters in the source code be interpreted according to the host computer's default setting for multibyte support. Multibyte characters are allowed in C and C++ style comments, in string literals, and in character constants. They are transferred untouched to the generated code.	
		Project>Options>C/C++ Compiler>Language 2>Enable multibyte support

--enable_restrict

Syntax	<code>--enable_restrict</code>	
Description	Enables the Standard C keyword <code>restrict</code> . This option can be useful for improving analysis precision during optimization.	
		To set this option, use Project>Options>C/C++ Compiler>Extra options

--endian

Syntax	<code>--endian=</code>	
Parameters	<i>big, b</i>	Specifies big-endian as the default byte order

`little, 1 (default)` Specifies little-endian as the default byte order

Description Use this option to specify the byte order of the generated code and data. By default, the compiler generates code in little-endian byte order.

See also *Byte order*, page 330.



Project>Options>General Options>Target>Endian mode

--enum_is_int

Syntax `--enum_is_int`

Description Use this option to force the size of all enumeration types to be at least 4 bytes.

Note: This option will not consider the fact that an `enum` type can be larger than an integer type.

See also *The enum type*, page 331.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--error_limit

Syntax `--error_limit=n`

Parameters


`n` The number of errors before the compiler stops the compilation. *n* must be a positive integer; 0 indicates no limit.

Description Use the `--error_limit` option to specify the number of errors allowed before the compiler stops the compilation. By default, 100 errors are allowed.



This option is not available in the IDE.

-f

Syntax	<code>-f filename</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	<p>Use this option to make the compiler read command line options from the named file, with the default filename extension <code>.xcl</code>.</p> <p>In the command file, you format the items exactly as if they were on the command line itself, except that you may use multiple lines, because the newline character acts just as a space or tab character.</p> <p>Both C and C++ style comments are allowed in the file. Double quotes behave in the same way as in the Microsoft Windows command line environment.</p> <p> To set this option, use Project>Options>C/C++ Compiler>Extra Options.</p>

--fpu

Syntax	<code>--fpu={VFPv2 VFPv3 VFPv3_d16 VFPv4 VFPv4_sp VFPv4_d16 VFPv5_sp VFPv5_d16 VFP9-S none}</code>	
Parameters	VFPv2	For a system that implements a VFP unit that conforms to the architecture VFPv2.
	VFPv3	For a system that implements a VFP unit that conforms to the architecture VFPv3.
	VFPv3_d16	For a system that implements a VFP unit that conforms to the D16 variant of the architecture VFPv3.
	VFPv4	For a system that implements a VFP unit that conforms to the architecture VFPv4.
	VFPv4_d16	For a system that implements a VFP unit that conforms to the D16 variant of the architecture VFPv4.
	VFPv4_sp	For a system that implements a VFP unit that conforms to the single-precision variant of the architecture VFPv4.
	VFPv5_sp	For a system that implements a VFP unit that conforms to the single-precision variant of the architecture VFPv5.

	VFPv5_d16	For a system that implements a VFP unit that conforms to the D16 variant of the architecture VFPv5.
	VFP9-S	VFP9-S is an implementation of the VFPv2 architecture that can be used with the ARM9E family of CPU cores. Selecting the VFP9-S coprocessor is therefore identical to selecting the VFPv2 architecture.
	none (default)	The software floating-point library is used.
Description	Use this option to generate code that performs floating-point operations using a Vector Floating Point (VFP) coprocessor. By selecting a VFP coprocessor, you will override the use of the software floating-point library for all supported floating-point operations.	

See also *VFP and floating-point arithmetic*, page 64.



Project>Options>General Options>Target>FPU

--guard_calls

Syntax	--guard_calls
Description	Use this option to enable guards for function static variable initialization. This option should be used in a threaded C++ environment.
See also	<i>Managing a multithreaded environment</i> , page 150.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.


--header_context

Syntax	--header_context
Description	Occasionally, to find the cause of a problem it is necessary to know which header file that was included from which source line. Use this option to list, for each diagnostic message, not only the source position of the problem, but also the entire include stack at that point.




This option is not available in the IDE.

-I

Syntax	<code>-I path</code>	
Parameters	<code>path</code>	The search path for <code>#include</code> files
Description	Use this option to specify the search paths for <code>#include</code> files. This option can be used more than once on the command line.	
See also	<i>Include file search procedure</i> , page 245.	
	 Project>Options>C/C++ Compiler>Preprocessor>Additional include directories	

--interwork

Syntax	<code>--interwork</code>	
Description	Use this option to generate interworking code. In code compiled with this option, functions will by default be of the type <code>interwork</code> . It will be possible to mix files compiled as <code>arm</code> and <code>thumb</code> (using the <code>--cpu_mode</code> option) as long as they are all compiled with the <code>--interwork</code> option. Note: Source code compiled for an ARM architecture v5 or higher, or for AEABI compliance, is interworking by default.	
See also	<i>__interwork</i> , page 348.	
	 Project>Options>C/C++ Compiler>Code>Generate interwork code	

-l

Syntax	<code>-l [a A b B c C D] [N] [H] {filename directory}</code>	
Parameters	<code>a (default)</code>	Assembler list file
	<code>A</code>	Assembler list file with C or C++ source as comments

b	Basic assembler list file. This file has the same contents as a list file produced with -la, except that no extra compiler-generated information (runtime model attributes, call frame information, frame size information) is included *
B	Basic assembler list file. This file has the same contents as a list file produced with -LA, except that no extra compiler generated information (runtime model attributes, call frame information, frame size information) is included *
c	C or C++ list file
C (default)	C or C++ list file with assembler source as comments
D	C or C++ list file with assembler source as comments, but without instruction offsets and hexadecimal byte values
N	No diagnostics in file
H	Include source lines from header files in output. Without this option, only source lines from the primary source file are included

* This makes the list file less useful as input to the assembler, but more useful for reading by a human.

See also *Rules for specifying a filename or directory as parameters*, page 254.

Description Use this option to generate an assembler or C/C++ listing to a file. Note that this option can be used one or more times on the command line.



To set related options, choose:

Project>Options>C/C++ Compiler>List

--legacy

Syntax --legacy={RVCT3.0}

Parameters RVCT3.0 Generates object code linkable with the linker in RVCT3.0. Use this mode together with the --aeabi option to generate code that should be linked with the linker in RVCT3.0.

Description Use this option to generate code compatible with the specified toolchain.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--lock_regs

Syntax	<code>--lock_regs=register</code>	
Parameters	<i>registers</i>	A comma-separated list of register names and register intervals to be locked, in the range R4–R11.
Description	Use this option to prevent the compiler from generating code that uses the specified registers.	
Example	<pre>--lock_regs=R4 --lock_regs=R8–R11 --lock_regs=R4,R8–R11</pre>	



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--macro_positions_in_diagnostics

Syntax	<code>--macro_positions_in_diagnostics</code>	
Description	Use this option to obtain position references inside macros in diagnostic messages. This is useful for detecting incorrect source code constructs in macros.	



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--make_all_definitions_weak

Syntax	<code>--make_all_definitions_weak</code>	
Description	<p>Turns all variable and function definitions in the compilation unit into weak definitions.</p> <p>Note: Normally, it is better to use extended keywords or pragma directives to turn individual variable and function definitions into weak definitions.</p>	
See also	<code>__weak</code> , page 356.	



This option is not available in the IDE.

--mfc

Syntax	--mfc
Description	<p>Use this option to enable <i>multi-file compilation</i>. This means that the compiler compiles one or several source files specified on the command line as one unit, which enhances interprocedural optimizations.</p> <p>Note: The compiler will generate one object file per input source code file, where the first object file contains all relevant data and the other ones are empty. If you want only the first file to be produced, use the <code>-o</code> compiler option and specify a certain output file.</p>
Example	<code>icarm myfile1.c myfile2.c myfile3.c --mfc</code>
See also	<code>--discard_unused_publics</code> , page 268, <code>--output</code> , <code>-o</code> , page 288, and <i>Multi-file compilation units</i> , page 229.



Project>Options>C/C++ Compiler>Multi-file compilation

--no_alignment_reduction

Syntax	--no_alignment_reduction
Description	<p>Some simple Thumb/Thumb2 functions can be 2-byte aligned. Use this option to keep those functions 4-byte aligned.</p> <p>This option has no effect when compiling for ARM mode.</p>
	<p>To set this option, use Project>Options>C/C++ Compiler>Extra Options.</p>



--no_call_frame_info

Syntax	--no_call_frame_info
Description	<p>Normally, the compiler always generates call frame information in the output, to enable the debugger to display the call stack even in code from modules with no debug information. Use this option to disable the generation of call frame information.</p>

See also *Call frame information*, page 176.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--no_clustering

Syntax `--no_clustering`

Description Use this option to disable static clustering optimizations.

Note: This option has no effect at optimization levels below .

See also *Static clustering*, page 233.



Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Static clustering

--no_code_motion

Syntax `--no_code_motion`

Description Use this option to disable code motion optimizations.

Note: This option has no effect at optimization levels below Medium.

See also *Code motion*, page 232.



Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Code motion


--no_const_align

Syntax `--no_const_align`

Description By default, the compiler uses alignment 4 for objects with a size of 4 bytes or more. Use this option to make the compiler align `const` objects based on the alignment of their type.

For example, a string literal will get alignment 1, because it is an array with elements of the type `const char` which has alignment 1. Using this option might save ROM space, possibly at the expense of processing speed.

See also *Alignment*, page 329.

 To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.


--no_cse

Syntax `--no_cse`

Description Use this option to disable common subexpression elimination.

Note: This option has no effect at optimization levels below Medium.

See also *Common subexpression elimination*, page 232.

 **Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Common subexpression elimination**


--no_exceptions

Syntax `--no_exceptions`

Description Use this option to disable exception support in the C++ language. Exception statements like `throw` and `try-catch`, and exception specifications on function definitions will generate an error message. Exception specifications on function declarations are ignored. The option is only valid when used together with the `--c++` compiler option.

If exceptions are not used in your application, it is recommended to disable support for them by using this option, because exceptions cause a rather large increase in code size.

See also *Exception handling*, page 194.

 **Project>Options>C/C++ Compiler>Language 1>C++**
and
Project>Options>C/C++ Compiler>Language 1>C++ dialect>C++>With exceptions

--no_fragments

Syntax `--no_fragments`

Description Use this option to disable section fragment handling. Normally, the toolset uses IAR proprietary information for transferring section fragment information to the linker. The linker uses this information to remove unused code and data, and thus further minimize the size of the executable image. When you use this option, this information is not output in the object files.

See also *Keeping symbols and sections*, page 105.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**

--no_inline

Syntax `--no_inline`

Description Use this option to disable function inlining.

See also *Inlining functions*, page 79.



Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Function inlining

--no_literal_pool

Syntax `--no_literal_pool`

Description Use this option to generate code that should run from a memory region where it is not allowed to read data, only to execute code.

When this option is used, the compiler will construct addresses and large constants with the MOV32 pseudo instruction instead of using a literal pool: switch statements are no longer translated using tables, and constant data is placed in the `.rodata` section.

This option also affects the automatic library selection performed by the linker. An IAR-specific ELF attribute is used for determining whether libraries compiled with this option should be used.

This option is only allowed for ARMv6-M and ARMv7 cores, and can be combined with the options `--ropi` or `--rwpi` only for ARMv7 cores.

See also

--no_literal_pool, page 318.



Project>Options>C/C++ Compiler>Code>No data reads in code memory

--no_loop_align

Syntax

`--no_loop_align`

Description

Use this option to disable the 4-byte alignment of labels in loops. This option is only useful in Thumb2 mode.

In ARM/Thumb1 mode, this option is enabled but does not perform anything.

See also

Alignment, page 329



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--no_mem_idioms

Syntax

`--no_mem_idioms`

Description

Use this option to make the compiler not optimize code sequences that clear, set, or copy a memory region. These memory access patterns (idioms) can otherwise be aggressively optimized, in some cases using calls to the runtime library. In principle, the transformation can involve more than a library call.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--no_path_in_file_macros

Syntax

`--no_path_in_file_macros`

Description

Use this option to exclude the path from the return value of the predefined preprocessor symbols `__FILE__` and `__BASE_FILE__`.

See also

Description of predefined preprocessor symbols, page 420.



This option is not available in the IDE.

--no_rtti

Syntax `--no_rtti`

Description Use this option to disable the runtime type information (RTTI) support in the C++ language. RTTI statements like `dynamic_cast` and `typeid` will generate an error message. This option is only valid when used together with the `--c++` compiler option.

See also *Using C++*, page 191.



Project>Options>C/C++ Compiler>Language 1>C++

and

Project>Options>C/C++ Compiler>Language 1>C++ dialect>C++>With RTTI

--no_rw_dynamic_init

Syntax `--no_rw_dynamic_init`

Description Use this option to disable runtime initialization of static C variables.

C source code that is compiled with `--ropi` or `--rwpi` cannot have static pointer variables and constants initialized to addresses of objects that do not have a known address at link time. To solve this for writable static variables, the compiler generates code that performs the initialization at program startup (in the same way as dynamic initialization in C++).

See also `--ropi`, page 291 and `--rwpi`, page 292.



Project>Options>C/C++ Compiler>Code>No dynamic read/write/initialization

--no_scheduling

Syntax `--no_scheduling`

Description Use this option to disable the instruction scheduler.


Note: This option has no effect at optimization levels below High.

See also *Instruction scheduling*, page 234.




Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Instruction scheduling


--no_size_constraints

Syntax	<code>--no_size_constraints</code>
Description	<p>Use this option to relax the normal restrictions for code size expansion when optimizing for high speed.</p> <p>Note: This option has no effect unless used with <code>-Ohs</code>.</p>
See also	<p><i>Speed versus size</i>, page 231.</p> <p> Project>Options>C/C++ Compiler>Optimizations>Enable transformations>No size constraints</p>

--no_static_destruction

Syntax	<code>--no_static_destruction</code>
Description	<p>Normally, the compiler emits code to destroy C++ static variables that require destruction at program exit. Sometimes, such destruction is not needed.</p> <p>Use this option to suppress the emission of such code.</p>
See also	<p><i>Setting up the atexit limit</i>, page 106.</p> <p> To set this option, use Project>Options>C/C++ Compiler>Extra Options.</p>

--no_system_include

Syntax	<code>--no_system_include</code>
Description	<p>By default, the compiler automatically locates the system include files. Use this option to disable the automatic search for system include files. In this case, you might need to set up the search path by using the <code>-I</code> compiler option.</p>
See also	<p><i>--dlib_config</i>, page 269, and <i>--system_include_dir</i>, page 294.</p> <p> Project>Options>C/C++ Compiler>Preprocessor>Ignore standard include directories</p>

--no_tbaa

Syntax `--no_tbaa`

Description Use this option to disable type-based alias analysis.

Note: This option has no effect at optimization levels below High.

See also *Type-based alias analysis*, page 233.



Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Type-based alias analysis

--no_typedefs_in_diagnostics

Syntax `--no_typedefs_in_diagnostics`

Description Use this option to disable the use of typedef names in diagnostics. Normally, when a type is mentioned in a message from the compiler, most commonly in a diagnostic message of some kind, the typedef names that were used in the original declaration are used whenever they make the resulting text shorter.

Example

```
typedef int (*MyPtr)(char const *);
MyPtr p = "My text string";
```

will give an error message like this:

```
Error[Pe144]: a value of type "char *" cannot be used to
initialize an entity of type "MyPtr"
```

If the `--no_typedefs_in_diagnostics` option is used, the error message will be like this:

```
Error[Pe144]: a value of type "char *" cannot be used to
initialize an entity of type "int (*)(char const *)"
```



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--no_unaligned_access

Syntax `--no_unaligned_access`

Description Use this option to make the compiler avoid unaligned accesses. Data accesses are usually performed aligned for improved performance. However, some accesses, most

notably when reading from or writing to packed data structures, might be unaligned. When using this option, all such accesses will be performed using a smaller data size to avoid any unaligned accesses. This option is only useful for ARMv6 architectures and higher.

See also

Alignment, page 329.



To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--no_unroll

Syntax

--no_unroll

Description

Use this option to disable loop unrolling.

Note: This option has no effect at optimization levels below High.

See also

Loop unrolling, page 232.



Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Loop unrolling

--no_var_align

Syntax

--no_var_align

Description

By default, the compiler uses alignment 4 for variable objects with a size of 4 bytes or more. Use this option to make the compiler align variable objects based on the alignment of their type.

For example, a char array will get alignment 1, because its elements of the type char have alignment 1. Using this option might save RAM space, possibly at the expense of processing speed.


See also

Alignment, page 329 and *--no_const_align*, page 279.




To set this option, use **Project>Options>C/C++ Compiler>Extra Options**.

--no_warnings

Syntax	--no_warnings
Description	By default, the compiler issues warning messages. Use this option to disable all warning messages.
	 This option is not available in the IDE.

--no_wrap_diagnostics

Syntax	--no_wrap_diagnostics
Description	By default, long lines in diagnostic messages are broken into several lines to make the message easier to read. Use this option to disable line wrapping of diagnostic messages.
	 This option is not available in the IDE.

-O

Syntax	-O [n l m h hs hz]												
Parameters	<table><tr><td>n</td><td>None* (Best debug support)</td></tr><tr><td>l (default)</td><td>Low*</td></tr><tr><td>m</td><td>Medium</td></tr><tr><td>h</td><td>High, balanced</td></tr><tr><td>hs</td><td>High, favoring speed</td></tr><tr><td>hz</td><td>High, favoring size</td></tr></table>	n	None* (Best debug support)	l (default)	Low*	m	Medium	h	High, balanced	hs	High, favoring speed	hz	High, favoring size
n	None* (Best debug support)												
l (default)	Low*												
m	Medium												
h	High, balanced												
hs	High, favoring speed												
hz	High, favoring size												
	*The most important difference between None and Low is that at None, all non-static variables will live during their entire scope.												
Description	Use this option to set the optimization level to be used by the compiler when optimizing the code. If no optimization option is specified, the optimization level l is used by default. If only -O is used without any parameter, the optimization level High balanced is used. A low level of optimization makes it relatively easy to follow the program flow in the debugger, and, conversely, a high level of optimization makes it relatively hard.												

See also

Controlling compiler optimizations, page 228.



Project>Options>C/C++ Compiler>Optimizations

--only_stdout

Syntax

`--only_stdout`

Description

Use this option to make the compiler use the standard output stream (`stdout`) also for messages that are normally directed to the error output stream (`stderr`).



This option is not available in the IDE.

--output, -o

Syntax

`--output {filename|directory}`
`-o {filename|directory}`

Parameters

See Rules for specifying a filename or directory as parameters, page 254.

Description

By default, the object code output produced by the compiler is located in a file with the same name as the source file, but with the extension `o`. Use this option to explicitly specify a different output filename for the object code output.



This option is not available in the IDE.

--predef_macros

Syntax

`--predef_macros {filename|directory}`

Parameters

See Rules for specifying a filename or directory as parameters, page 254.

Description

Use this option to list the predefined symbols. When using this option, make sure to also use the same options as for the rest of your project.

If a filename is specified, the compiler stores the output in that file. If a directory is specified, the compiler stores the output in that directory, in a file with the `predef` filename extension.

Note that this option requires that you specify a source file on the command line.



This option is not available in the IDE.

--preinclude

Syntax

--preinclude *includefile*

Parameters

See *Rules for specifying a filename or directory as parameters*, page 254.

Description

Use this option to make the compiler read the specified include file before it starts to read the source file. This is useful if you want to change something in the source code for the entire application, for instance if you want to define a new symbol.



Project>Options>C/C++ Compiler>Preprocessor>Preinclude file

--preprocess

Syntax

--preprocess [= [c] [n] [l]] { *filename* | *directory* }

Parameters

c	Preserve comments
n	Preprocess only
l	Generate #line directives

See also *Rules for specifying a filename or directory as parameters*, page 254.

Description

Use this option to generate preprocessed output to a named file.



Project>Options>C/C++ Compiler>Preprocessor>Preprocessor output to file

--public_equ

Syntax

--public_equ *symbol* [= *value*]

Parameters

<i>symbol</i>	The name of the assembler symbol to be defined
<i>value</i>	An optional value of the defined assembler symbol

Description

This option is equivalent to defining a label in assembler language using the `EQU` directive and exporting it using the `PUBLIC` directive. This option can be used more than once on the command line.



This option is not available in the IDE.

--relaxed_fp

Syntax

`--relaxed_fp`

Description

Use this option to allow the compiler to relax the language rules and perform more aggressive optimization of floating-point expressions. This option improves performance for floating-point expressions that fulfill these conditions:

- The expression consists of both single- and double-precision values
- The double-precision values can be converted to single precision without loss of accuracy
- The result of the expression is converted to single precision.

Note that performing the calculation in single precision instead of double precision might cause a loss of accuracy.

Example

```
float F(float a, float b)
{
    return a + b * 3.0;
}
```

The C standard states that `3.0` in this example has the type `double` and therefore the whole expression should be evaluated in `double` precision. However, when the `--relaxed_fp` option is used, `3.0` will be converted to `float` and the whole expression can be evaluated in `float` precision.



To set related options, choose:

Project>Options>C/C++ Compiler>Language 2>Floating-point semantics

--remarks

Syntax

`--remarks`

Description

The least severe diagnostic messages are called remarks. A remark indicates a source code construct that may cause strange behavior in the generated code. By default, the

compiler does not generate remarks. Use this option to make the compiler generate remarks.

See also

Severity levels, page 249.



Project>Options>C/C++ Compiler>Diagnostics>Enable remarks

--require_prototypes

Syntax

--require_prototypes

Description

Use this option to force the compiler to verify that all functions have proper prototypes. Using this option means that code containing any of the following will generate an error:

- A function call of a function with no declaration, or with a Kernighan & Ritchie C declaration
- A function definition of a public function with no previous prototype declaration
- An indirect function call through a function pointer with a type that does not include a prototype.



Project>Options>C/C++ Compiler>Language 1>Require prototypes

--ropi

Syntax

--ropi

Description

Use this option to make the compiler generate code that uses PC-relative references to address code and read-only data.

When this option is used, these limitations apply:

- C++ constructions cannot be used
- The object attribute `__ramfunc` cannot be used
- Pointer constants cannot be initialized with the address of another constant, a string literal, or a function. However, writable variables can be initialized to constant addresses at runtime.

See also

--no_rw_dynamic_init, page 283, and *Description of predefined preprocessor symbols*, page 420.



Project>Options>C/C++ Compiler>Code>Code and read-only data (ropi)

--rwp*i*

Syntax

--rwp*i*

Description

Use this option to make the compiler generate code that uses the offset from the static base register (R9) to address-writable data.

When this option is used, these limitations apply:

- The object attribute `__ramfunc` cannot be used
- Pointer constants cannot be initialized with the address of a writable variable. However, static writable variables can be initialized to writable addresses at runtime.

See also

--no_rw_dynamic_init, page 283, and *Description of predefined preprocessor symbols*, page 420.



Project>Options>C/C++ Compiler>Code>Read/write data (rwpi)

--section

Syntax

--section OldName=NewName

Description

The compiler places functions and data objects into named sections which are referred to by the IAR ILINK Linker. Use this option to change the name of the section *OldName* to *NewName*.

This is useful if you want to place your code or data in different address ranges and you find the @ notation, alternatively the #pragma location directive, insufficient. Note that any changes to the section names require corresponding modifications in the linker configuration file.

Example

To place functions in the section MyText, use:

--section .text=MyText


See also

Controlling data and function placement in memory, page 224.




Project>Options>C/C++ Compiler>Output>Code section name


--separate_cluster_for_initialized_variables

Syntax	<code>--separate_cluster_for_initialized_variables</code>
Description	<p>Use this option to separate initialized and non-initialized variables when using variable clustering. This might reduce the number of bytes in the ROM area which are needed for data initialization, but it might lead to larger code.</p> <p>This option can be useful if you want to have your own data initialization routine, but want the IAR tools to arrange for the zero-initialized variables.</p>
See also	<p><i>Manual initialization</i>, page 107 and <i>initialize directive</i>, page 446.</p> <p> To set this option, use Project>Options>C/C++ Compiler>Extra Options.</p>

--silent

Syntax	<code>--silent</code>
Description	<p>By default, the compiler issues introductory messages and a final statistics report. Use this option to make the compiler operate without sending these messages to the standard output stream (normally the screen).</p> <p>This option does not affect the display of error and warning messages.</p> <p> This option is not available in the IDE.</p>

--strict

Syntax	<code>--strict</code>
Description	<p>By default, the compiler accepts a relaxed superset of Standard C and C++. Use this option to ensure that the source code of your application instead conforms to strict Standard C and C++.</p> <p>Note: The <code>-e</code> option and the <code>--strict</code> option cannot be used at the same time.</p>
See also	<p><i>Enabling language extensions</i>, page 183.</p> <p> Project>Options>C/C++ Compiler>Language 1>Language conformance>Strict</p>

--system_include_dir

Syntax	<code>--system_include_dir path</code>	
Parameters	<code>path</code>	The path to the system include files, see <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	By default, the compiler automatically locates the system include files. Use this option to explicitly specify a different path to the system include files. This might be useful if you have not installed IAR Embedded Workbench in the default location.	
See also	<code>--dlib_config</code> , page 269, and <code>--no_system_include</code> , page 284.	



This option is not available in the IDE.

--thumb

Syntax	<code>--thumb</code>	
Description	Use this option to set default function mode to Thumb. This setting must be the same for all files included in a program, unless they are interworking. Note: This option has the same effect as the <code>--cpu_mode=thumb</code> option.	
See also	<code>--interwork</code> , page 275 and <code>__interwork</code> , page 348.	



Project>Options>C/C++ Compiler>Code>Processor mode>Thumb


--use_c++_inline

Syntax	<code>--use_c++_inline</code>	
Description	Standard C uses slightly different semantics for the <code>inline</code> keyword than C++ does. Use this option if you want C++ semantics when you are using C.	
See also	<i>Inlining functions</i> , page 79	




Project>Options>C/C++ Compiler>Language 1>C dialect>C99>C++ inline semantics


--use_unix_directory_separators

Syntax	<code>--use_unix_directory_separators</code>
Description	<p>Use this option to make DWARF debug information use / (instead of \) as directory separators in file paths.</p> <p>This option can be useful if you have a debugger that requires directory separators in UNIX style.</p> <p> To set this option, use Project>Options>C/C++ Compiler>Extra Options.</p>


--vectorize

Syntax	<code>--vectorize</code>
Description	<p>Use this option to enable generation of NEON vector instructions.</p> <p>Loops will only be vectorized if the target processor has NEON capability and the optimization level is <code>-Ohs</code>.</p> <p> Project>Options>C/C++ Compiler>Optimizations>Enable transformations>Vectorize</p>


--vla

Syntax	<code>--vla</code>
Description	<p>Use this option to enable support for C99 variable length arrays. Such arrays are located on the heap. This option requires Standard C and cannot be used together with the <code>--c89</code> compiler option.</p> <p>Note: <code>--vla</code> should not be used together with the <code>longjmp</code> library function, as that can lead to memory leakages.</p>
See also	<p><i>C language overview</i>, page 181.</p> <p> Project>Options>C/C++ Compiler>Language 1>C dialect>Allow VLA</p>


--warn_about_c_style_casts

Syntax	<code>--warn_about_c_style_casts</code>
Description	Use this option to make the compiler warn when C-style casts are used in C++ source code.
	 This option is not available in the IDE.

--warnings_affect_exit_code

Syntax	<code>--warnings_affect_exit_code</code>
Description	By default, the exit code is not affected by warnings, because only errors produce a non-zero exit code. With this option, warnings will also generate a non-zero exit code.
	 This option is not available in the IDE.

--warnings_are_errors

Syntax	<code>--warnings_are_errors</code>
Description	<p>Use this option to make the compiler treat all warnings as errors. If the compiler encounters an error, no object code is generated. Warnings that have been changed into remarks are not treated as errors.</p> <p>Note: Any diagnostic messages that have been reclassified as warnings by the option <code>--diag_warning</code> or the <code>#pragma diag_warning</code> directive will also be treated as errors when <code>--warnings_are_errors</code> is used.</p>
See also	<code>--diag_warning</code> , page 267.
	 Project>Options>C/C++ Compiler>Diagnostics>Treat all warnings as errors

Linker options

- Summary of linker options
- Descriptions of linker options

For general syntax rules, see *Options syntax*, page 253.

Summary of linker options

This table summarizes the linker options:

Command line option	Description
--advanced_heap	Uses an advanced heap
--basic_heap	Uses a basic heap
--BE8	Uses the big-endian format BE8
--BE32	Uses the big-endian format BE32
--bounds_table_size	Specifies the size of the global bounds table. See the C-RUN documentation in the C-SPY® <i>Debugging Guide for ARM</i> .
--call_graph	Produces a call graph file in XML format
--config	Specifies the linker configuration file to be used by the linker
--config_def	Defines symbols for the configuration file
--config_search	Specifies more directories to search for linker configuration files
--cpp_init_routine	Specifies a user-defined C++ dynamic initialization routine
--cpu	Specifies a processor variant
--debug_heap	Uses the checked heap. See the C-RUN documentation in the C-SPY® <i>Debugging Guide for ARM</i> .
--default_to_complex_ranges	Makes <code>complex</code> ranges the default decompressor in <code>initialize</code> directives
--define_symbol	Defines symbols that can be used by the application

Table 29: Linker options summary

Command line option	Description
--dependencies	Lists file dependencies
--diag_error	Treats these message tags as errors
--diag_remark	Treats these message tags as remarks
--diag_suppress	Suppresses these diagnostic messages
--diag_warning	Treats these message tags as warnings
--diagnostics_tables	Lists all diagnostic messages
--do_segment_pad	Pads each ELF segment to n-byte alignment
--enable_hardware_workaround	Enables specified hardware workaround
--enable_stack_usage	Enables stack usage analysis
--entry	Treats the symbol as a root symbol and as the start of the application
--error_limit	Specifies the allowed number of errors before linking stops
--exception_tables	Generates exception tables for C code
--export_builtint_config	Produces an <code>icf</code> file for the default configuration
--extra_init	Specifies an extra initialization routine that will be called if it is defined.
-f	Extends the command line
--force_exceptions	Always includes exception runtime code
--force_output	Produces an output file even if errors occurred
--ignore_uninstrumented_pointers	Disables checking of accessing via pointers in memory for which no bounds have been set. See the C-RUN documentation in the C-SPY® <i>Debugging Guide for ARM</i> .
--image_input	Puts an image file in a section
--inline	Inlines small routines
--keep	Forces a symbol to be included in the application
-L	Specifies more directories to search for object and library files. Alias for <code>--search</code> .
--log	Enables log output for selected topics
--log_file	Directs the log to a file
--mangled_names_in_messages	Adds mangled names in messages
--map	Produces a map file

Table 29: Linker options summary (Continued)

Command line option	Description
<code>--merge_duplicate_sections</code>	Merges equivalent read-only sections
<code>--misrac</code>	Enables error messages specific to MISRA-C:1998. This option is a synonym to <code>--misrac1998</code> and is only available for backwards compatibility.
<code>--misrac1998</code>	Enables error messages specific to MISRA-C:1998. See the <i>IAR Embedded Workbench® MISRA C:1998 Reference Guide</i> .
<code>--misrac2004</code>	Enables error messages specific to MISRA-C:2004. See the <i>IAR Embedded Workbench® MISRA C:2004 Reference Guide</i> .
<code>--misrac_verbose</code>	Enables verbose logging of MISRA C checking. See the <i>IAR Embedded Workbench® MISRA C:1998 Reference Guide</i> and the <i>IAR Embedded Workbench® MISRA C:2004 Reference Guid</i> .
<code>--no_dynamic_rtti_elimination</code>	Includes dynamic runtime type information even when it is not needed.
<code>--no_entry</code>	Sets the entry point to zero
<code>--no_exceptions</code>	Generates an error if exceptions are used
<code>--no_fragments</code>	Disables section fragment handling
<code>--no_library_search</code>	Disables automatic runtime library search
<code>--no_literal_pool</code>	Generates code that should run from a memory region where it is not allowed to read data, only to execute code
<code>--no_locals</code>	Removes local symbols from the ELF executable image.
<code>--no_range_reservations</code>	Disables range reservations for absolute symbols
<code>--no_remove</code>	Disables removal of unused sections
<code>--no_veneers</code>	Disables generation of veneers
<code>--no_vfe</code>	Disables Virtual Function Elimination
<code>--no_warnings</code>	Disables generation of warnings
<code>--no_wrap_diagnostics</code>	Does not wrap long lines in diagnostic messages
<code>-o</code>	Sets the object filename. Alias for <code>--output</code> .
<code>--only_stdout</code>	Uses standard output only
<code>--output</code>	Sets the object filename
<code>--pi_veneers</code>	Generates position independent veneers.

Table 29: Linker options summary (Continued)

Command line option	Description
--place_holder	Reserve a place in ROM to be filled by some other tool, for example a checksum calculated by ielftool.
--redirect	Redirects a reference to a symbol to another symbol
--remarks	Enables remarks
--search	Specifies more directories to search for object and library files
--semihosting	Links with debug interface
--silent	Sets silent operation
--skip_dynamic_initialization	Suppresses dynamic initialization during system startup
--stack_usage_control	Specifies a stack usage control file
--strip	Removes debug information from the executable image
--threaded_lib	Configures the runtime library for use with threads
--use_full_std_template_names	Enables full names for standard C++ templates
--vfe	Controls Virtual Function Elimination
--warnings_affect_exit_code	Warnings affects exit code
--warnings_are_errors	Warnings are treated as errors
--whole_archive	Treats every object file in the archive as if it was specified on the command line.

Table 29: Linker options summary (Continued)


Descriptions of linker options

The following section gives detailed reference information about each linker option.




Note that if you use the options page **Extra Options** to specify specific command line options, the IDE does not perform an instant check for consistency problems like conflicting options, duplication of options, or use of irrelevant options.


--advanced_heap

Syntax	--advanced_heap
Description	Use this option to use an advanced heap.
See also	<i>Advanced versus basic heap</i> , page 205
	 To set this option, use Project>Options>Linker>Extra Options .

--basic_heap

Syntax	--basic_heap
Description	Use this option to use a basic heap.
See also	<i>Advanced versus basic heap</i> , page 205
	 To set this option, use Project>Options>Linker>Extra Options .

--BE8

Syntax	--BE8
Description	<p>Use this option to specify the Byte Invariant Addressing mode.</p> <p>This means that the linker reverses the byte order of the instructions, resulting in little-endian code and big-endian data. This is the default byte addressing mode for ARMv6 big-endian images. This is the only mode available for ARM v6M and ARM v7 with big-endian images.</p> <p>Byte Invariant Addressing mode is only available on ARM processors that support ARMv6, ARM v6M, and ARM v7.</p>
See also	<i>Byte order</i> , page 65, <i>Byte order</i> , page 330, <i>--BE32</i> , page 302, and <i>--endian</i> , page 271.
	 Project>Options>General Options>Target>Endian mode

--BE32

Syntax	--BE32
Description	Use this option to specify the legacy big-endian mode. This produces big-endian code and data. This is the only byte-addressing mode for all big-endian images prior to ARMv6. This mode is also available for ARM v6 with big-endian, but not for ARM v6M or ARM v7.
See also	<i>Byte order</i> , page 65, <i>Byte order</i> , page 330, --BE8, page 301, and --endian, page 271.



Project>Options>General Options>Target>Endian mode

--call_graph

Syntax	--call_graph { <i>filename</i> <i>directory</i> }
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	Use this option to produce a call graph file. If no filename extension is specified, the extension <code>cgx</code> is used. This option can only be used once on the command line. Using this option enables stack usage analysis in the linker.
See also	<i>Stack usage analysis</i> , page 92



Project>Options>Linker>Advanced>Call graph output (XML)

--config

Syntax	--config <i>filename</i>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	Use this option to specify the configuration file to be used by the linker (the default filename extension is <code>icf</code>). If no configuration file is specified, a default configuration is used. This option can only be used once on the command line.
See also	The chapter <i>The linker configuration file</i> .



Project>Options>Linker>Config>Linker configuration file

--config_def

Syntax	<code>--config_def symbol=constant_value</code>	
Parameters	<i>symbol</i>	The name of the symbol to be used in the configuration file.
	<i>constant_value</i>	The constant value of the configuration symbol.
Description	Use this option to define a constant configuration symbol to be used in the configuration file. This option has the same effect as the <code>define symbol</code> directive in the linker configuration file. This option can be used more than once on the command line.	
See also	<code>--define_symbol</code> , page 305 and <i>Interaction between ILINK and the application</i> , page 110.	



Project>Options>Linker>Config>Defined symbols for configuration file


--config_search

Syntax	<code>--config_search path</code>	
Parameters	<i>path</i>	A path to a directory where the linker should search for linker configuration include files.
Description	<p>Use this option to specify more directories to search for files when processing an <code>include</code> directive in a linker configuration file.</p> <p>By default, the linker searches for configuration include files only in the system configuration directory. To specify more than one search directory, use this option for each path.</p>	
See also	<code>include directive</code> , page 461.	




To set this option, use **Project>Options>Linker>Extra Options**.

--cpp_init_routine

Syntax	<code>--cpp_init_routine routine</code>	
Parameters	<i>routine</i>	A user-defined C++ dynamic initialization routine.
Description	<p>When using the IAR C/C++ compiler and the standard library, C++ dynamic initialization is handled automatically. In other cases you might need to use this option.</p> <p>If any sections with the section type <code>INIT_ARRAY</code> or <code>PREINIT_ARRAY</code> are included in your application, the C++ dynamic initialization routine is considered to be needed. By default, this routine is named <code>__iar_cstart_call_ctors</code> and is called by the startup code in the standard library. Use this option if you are not using the standard library and require another routine to handle these section types.</p> <p> To set this option, use Project>Options>Linker>Extra Options.</p>	

--cpu

Syntax	<code>--cpu=core</code>	
Parameters	<i>core</i>	Specifies a specific processor variant
Description	Use this option to select the processor variant to link your application for. The default is to use a processor or architecture compatible with the object file attributes.	
See also	<code>--cpu</code> , page 262	
		Project>Options>General Options>Target>Processor configuration

--default_to_complex_ranges

Syntax	<code>--default_to_complex_ranges</code>	
Description	Normally, if <code>initialize</code> directives in a linker configuration file do not specify <code>simple ranges</code> or <code>complex ranges</code> , the linker uses <code>simple ranges</code> if the associated section placement directives use single range regions.	

Use this option to make the linker always use `complex` ranges by default. This was the behavior of the linker before the introduction of `simple` ranges and `complex` ranges.

See also *initialize directive*, page 446.



Project>Options>Linker>Extra Options

--define_symbol

Syntax `--define_symbol symbol=constant_value`

Parameters

<i>symbol</i>	The name of the constant symbol that can be used by the application.
<i>constant_value</i>	The constant value of the symbol.

Description Use this option to define a constant symbol, that is a label, that can be used by your application. This option can be used more than once on the command line. Note that his option is different from the `define symbol` directive.

See also *--config_def*, page 303 and *Interaction between ILINK and the application*, page 110.



Project>Options>Linker>#define>Defined symbols

--dependencies

Syntax `--dependencies [= [i|m]] {filename|directory}`

Parameters

<i>i</i> (default)	Lists only the names of files
<i>m</i>	Lists in makefile style

See also *Rules for specifying a filename or directory as parameters*, page 254.

Description Use this option to make the linker list the names of the linker configuration, object, and library files opened for input into a file with the default filename extension `i`.

Example

If `--dependencies` or `--dependencies=i` is used, the name of each opened input file, including the full path, if available, is output on a separate line. For example:

```
c:\myproject\foo.o
d:\myproject\bar.o
```

If `--dependencies=m` is used, the output is in makefile style. For each input file, one line containing a makefile dependency rule is produced. Each line consists of the name of the output file, a colon, a space, and the name of an input file. For example:

```
a.out: c:\myproject\foo.o
a.out: d:\myproject\bar.o
```



This option is not available in the IDE.

--diag_error

Syntax

```
--diag_error=tag[, tag, ...]
```

Parameters

<i>tag</i>	The number of a diagnostic message, for example the message number Pe117
------------	--

Description

Use this option to reclassify certain diagnostic messages as errors. An error indicates a problem of such severity that an executable image will not be generated. The exit code will be non-zero. This option may be used more than once on the command line.



Project>Options>Linker>Diagnostics>Treat these as errors

--diag_remark

Syntax

```
--diag_remark=tag[, tag, ...]
```

Parameters

<i>tag</i>	The number of a diagnostic message, for example the message number Pe177
------------	--

Description

Use this option to reclassify certain diagnostic messages as remarks. A remark is the least severe type of diagnostic message and indicates a construction that may cause strange behavior in the executable image. This option may be used more than once on the command line.

Note: By default, remarks are not displayed; use the `--remarks` option to display them.



Project>Options>Linker>Diagnostics>Treat these as remarks

--diag_suppress

Syntax	<code>--diag_suppress=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe117</code>
Description	Use this option to suppress certain diagnostic messages. These messages will not be displayed. This option may be used more than once on the command line.	



Project>Options>Linker>Diagnostics>Suppress these diagnostics

--diag_warning


Syntax	<code>--diag_warning=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe826</code>
Description	Use this option to reclassify certain diagnostic messages as warnings. A warning indicates an error or omission that is of concern, but which will not cause the linker to stop before linking is completed. This option may be used more than once on the command line.	




Project>Options>Linker>Diagnostics>Treat these as warnings

--diagnostics_tables


Syntax	<code>--diagnostics_tables {filename directory}</code>	
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.	

Description	Use this option to list all possible diagnostic messages in a named file. This can be convenient, for example, if you have used a pragma directive to suppress or change the severity level of any diagnostic messages, but forgot to document why.
	This option cannot be given together with other options.
	This option is not available in the IDE.


--do_segment_pad

Syntax	--do_segment_pad
Description	Use this option to extend each ELF segment in the executable file with content, to make it an even multiple of bytes long (if possible). Some runtime library routines might access memory in units of bytes, and might, if the right data object is placed last in an ELF segment, access memory outside the strict bounds of the segment. If you are executing in an environment where this is a problem, you can use this option to extend the ELF segments appropriately so that this is not a problem.
	This option is not available in the IDE.


--enable_hardware_workaround

Syntax	--enable_hardware_workaround= <i>waid</i> [<i>waid</i> [...]]
Parameters	<div><div><i>waid</i></div><div>The ID number of the workaround that you want to enable. For a list of available workarounds, see the release notes available in the Information Center.</div></div>
Description	Use this option to make the linker generate a workaround for a specific hardware problem.
See also	The release notes for the linker for a list of available parameters.
	To set this option, use Project>Options>Linker>Extra Options .

--enable_stack_usage

Syntax	<code>--enable_stack_usage</code>	
Description	Use this option to enable stack usage analysis. If a linker map file is produced, a stack usage chapter is included in the map file. Note: If you use at least one of the <code>--stack_usage_control</code> or <code>--call_graph</code> options, stack usage analysis is automatically enabled.	
See also	<i>Stack usage analysis</i> , page 92	
		Project>Options>Linker>Advanced>Enable stack usage analysis

--entry

Syntax	<code>--entry <i>symbol</i></code>	
Parameters	<i>symbol</i>	The name of the symbol to be treated as a root symbol and start label
Description	Use this option to make a symbol be treated as a root symbol and the start label of the application. This is useful for loaders. If this option is not used, the default start symbol is <code>__iar_program_start</code> . A root symbol is kept whether or not it is referenced from the rest of the application, provided its module is included. A module in an object file is always included but a module part of a library is only included if needed. Note: The label referred to must be available in your application. You must also make sure that the reset vector refers to the new start label (for example <code>--redirect __iar_program_start=_myStartLabel</code>). 	
		Project>Options>Linker>Library>Override default program entry

--error_limit

Syntax	<code>--error_limit=<i>n</i></code>	
Parameters	<i>n</i>	The number of errors before the linker stops linking. <i>n</i> must be a positive integer; 0 indicates no limit.

Description Use the `--error_limit` option to specify the number of errors allowed before the linker stops the linking. By default, 100 errors are allowed.



This option is not available in the IDE.

--exception_tables

Syntax `--exception_tables={ncreate|unwind|cantunwind}`

Parameters

<code>ncreate</code> (default)	Does not generate entries. Uses the least amount of memory, but the result is undefined if an exception is propagated through a function without exception information.
<code>unwind</code>	Generates unwind entries that enable an exception to be correctly propagated through functions without exception information.
<code>cantunwind</code>	Generates no-unwind entries so that any attempt to propagate an exception through the function will result in a call to <code>terminate</code> .

Description Use this option to determine what the linker should do with functions that do not have exception information but which do have correct call frame information.

The compiler ensures that C functions get correct call frame information. For functions written in assembler language you need to use assembler directives to generate call frame information.

See also *Using C++*, page 191.



To set this option, use **Project>Options>Linker>Extra Options**.

--export_builtin_config

Syntax `--export_builtin_config filename`

Parameters See *Rules for specifying a filename or directory as parameters*, page 254.

Description Exports the configuration used by default to a file.



This option is not available in the IDE.

--extra_init

Syntax

```
--extra_init routine
```

Parameters

routine

A user-defined initialization routine.

Description

Use this option to make the linker add an entry for the specified routine at the end of the initialization table. The routine will be called during system startup, after other initialization routines have been called and before `main` is called. Note that the routine must preserve the value passed to it in register `R0`. No entry is added if the routine is not defined.



To set this option, use **Project>Options>Linker>Extra Options**.

-f

Syntax

```
-f filename
```

Parameters

See *Rules for specifying a filename or directory as parameters*, page 254.

Description

Use this option to make the linker read command line options from the named file, with the default filename extension `.xcl`.


In the command file, you format the items exactly as if they were on the command line itself, except that you may use multiple lines, because the newline character acts just as a space or tab character.

Both C and C++ style comments are allowed in the file. Double quotes behave in the same way as in the Microsoft Windows command line environment.




To set this option, use **Project>Options>Linker>Extra Options**.

--force_exceptions

Syntax	<code>--force_exceptions</code>
Description	<p>Use this option to make the linker include exception tables and exception code even when the linker heuristics indicate that exceptions are not used.</p> <p>The linker considers exceptions to be used if there is a <code>throw</code> expression that is not a <code>rethrow</code> in the included code. If there is no such <code>throw</code> expression in the rest of the code, the linker arranges for operator <code>new</code>, <code>dynamic_cast</code>, and <code>typeid</code> to call <code>abort</code> instead of throwing an exception on failure. If you need to catch exceptions from these constructs but your code contains no other throws, you might need to use this option.</p>
See also	<p><i>Using C++</i>, page 191.</p> <div> Project>Options>Linker>Optimizations>C++ Exceptions>Allow>Always include</div>

--force_output

Syntax	<code>--force_output</code>
Description	<p>Use this option to produce an output executable image regardless of any linking errors.</p> <div> To set this option, use Project>Options>Linker>Extra Options</div>

--image_input

Syntax	<code>--image_input filename [,symbol,[section[,alignment]]]</code>								
Parameters	<table><tr><td><i>filename</i></td><td>The pure binary file containing the raw image you want to link</td></tr><tr><td><i>symbol</i></td><td>The symbol which the binary data can be referenced with.</td></tr><tr><td><i>section</i></td><td>The section where the binary data will be placed; default is <code>.text</code>.</td></tr><tr><td><i>alignment</i></td><td>The alignment of the section; default is 1.</td></tr></table>	<i>filename</i>	The pure binary file containing the raw image you want to link	<i>symbol</i>	The symbol which the binary data can be referenced with.	<i>section</i>	The section where the binary data will be placed; default is <code>.text</code> .	<i>alignment</i>	The alignment of the section; default is 1.
<i>filename</i>	The pure binary file containing the raw image you want to link								
<i>symbol</i>	The symbol which the binary data can be referenced with.								
<i>section</i>	The section where the binary data will be placed; default is <code>.text</code> .								
<i>alignment</i>	The alignment of the section; default is 1.								

Description Use this option to link pure binary files in addition to the ordinary input files. The file's entire contents are placed in the section, which means it can only contain pure binary data.

Note: Just as for sections from object files, sections created by using the `--image_input` option are not included unless actually needed. You can either specify a symbol in the option and reference this symbol in your application (or by use of a `--keep` option), or you can specify a section name and use the `keep` directive in a linker configuration file to ensure that the section is included.

Example `--image_input bootstrap.abs,Bootstrap,CSTARTUPCODE,4`

The contents of the pure binary file `bootstrap.abs` are placed in the section `CSTARTUPCODE`. The section where the contents are placed is 4-byte aligned and will only be included if your application (or the command line option `--keep`) includes a reference to the symbol `Bootstrap`.

See also `--keep`, page 313.



Project>Options>Linker>Input>Raw binary image

--inline

Syntax `--inline`

Description Some routines are so small that they can fit in the space of the instruction that calls the routine. Use this option to make the linker replace the call of a routine with the body of the routine, where applicable.



Project>Options>Linker>Optimizations>Inline small routines

--keep

Syntax `--keep symbol`

Parameters

symbol The name of the symbol to be treated as a root symbol

Description Normally, the linker keeps a symbol only if it is needed by your application. Use this option to make a symbol always be included in the final application.



Project>Options>Linker>Input>Keep symbols

--log

Syntax	<code>--log <i>topic</i>[,<i>topic</i>,...]</code>	
Parameters	<i>topic</i> can be one of:	
	<code>call_graph</code>	Lists the call graph as seen by stack usage analysis.
	<code>initialization</code>	Lists copy batches and the compression selected for each batch.
	<code>libraries</code>	Lists all decisions taken by the automatic library selector. This might include extra symbols needed (<code>--keep</code>), redirections (<code>--redirect</code>), as well as which runtime libraries that were selected.
	<code>modules</code>	Lists each module that is selected for inclusion in the application, and which symbol that caused it to be included.
	<code>redirects</code>	Lists redirected symbols.
	<code>sections</code>	Lists each symbol and section fragment that is selected for inclusion in the application, and the dependence that caused it to be included.
	<code>unused_fragments</code>	Lists those section fragments that were not included in the application.
	<code>veeners</code>	Lists some veneer creation and usage statistics.
Description	Use this option to make the linker log information to <code>stdout</code> . The log information can be useful for understanding why an executable image became the way it is.	
See also	<code>--log_file</code> , page 315.	



Project>Options>Linker>List>Generate log

--log_file

Syntax	<code>--log_file filename</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	Use this option to direct the log output to the specified file.
See also	<code>--log</code> , page 314.



Project>Options>Linker>List>Generate log

--mangled_names_in_messages

Syntax	<code>--mangled_names_in_messages</code>
Description	Use this option to produce both mangled and unmangled names for C/C++ symbols in messages. Mangling is a technique used for mapping a complex C name or a C++ name (for example, for overloading) into a simple name. For example, <code>void h(int, char)</code> becomes <code>_Z1hic</code> .



This option is not available in the IDE.

--map

Syntax	<code>--map {filename directory}</code>
Description	<p>Use this option to produce a linker memory map file. The map file has the default filename extension <code>map</code>. The map file contains:</p> <ul style="list-style-type: none"> • Linking summary in the map file header which lists the version of the linker, the current date and time, and the command line that was used. • Runtime attribute summary which lists runtime attributes. • Placement summary which lists each section/block in address order, sorted by placement directives. • Initialization table layout which lists the data ranges, packing methods, and compression ratios. • Module summary which lists contributions from each module to the image, sorted by directory and library.

- Entry list which lists all public and some local symbols in alphabetical order, indicating which module they came from.
- Some of the bytes might be reported as *shared*.
Shared objects are functions or data objects that are shared between modules. If any of these occur in more than one module, only one copy is retained. For example, in some cases inline functions are not inlined, which means that they are marked as shared, because only one instance of each function will be included in the final application. This mechanism is sometimes also used for compiler-generated code or data not directly associated with a particular function or variable, and when only one instance is required in the final application.

This option can only be used once on the command line.



Project>Options>Linker>List>Generate linker map file

--merge_duplicate_sections

Syntax	<code>--merge_duplicate_sections</code>
Description	Use this option to keep only one copy of equivalent read-only sections. Note that this can cause different functions or constants to have the same address, so an application that depends on the addresses being different will not work correctly with this option enabled.



Project>Options>Linker>Optimizations>Merge duplicate sections

--no_dynamic_rtti_elimination

Syntax	<code>--no_dynamic_rtti_elimination</code>
Description	<p>Use this option to make the linker include dynamic (polymorphic) runtime type information (RTTI) data in the application image even when the linker heuristics indicate that it is not needed.</p> <p>The linker considers dynamic runtime type information to be needed if there is a <code>typeid</code> or <code>dynamic_cast</code> expression for a polymorphic type in the included code. By default, if the linker detects no such expression, RTTI data will not be included just to make dynamic RTTI requests work.</p>

Note: A `typeid` expression for a *non*-polymorphic type results in a direct reference to a particular RTTI object and will not cause the linker to include any potentially unneeded objects.

See also

Using C++, page 191.



To set this option, use **Project>Options>Linker>Extra Options**.

--no_entry

Syntax

`--no_entry`

Description

Use this option to set the entry point field to zero for produced ELF files.



This option is not available in the IDE.

--no_exceptions

Syntax

`--no_exceptions`

Description

Use this option to make the linker generate an error if there is a throw in the included code. This option is useful for making sure that your application does not use exceptions.

See also

Using C++, page 191.



To set related options, choose:

Project>Options>Linker>Optimizations>Allow C++ exceptions

--no_fragments

Syntax

`--no_fragments`

Description

Use this option to disable section fragment handling. Normally, the toolset uses IAR proprietary information for transferring section fragment information to the linker. The linker uses this information to remove unused code and data, and thus further minimize the size of the executable image. Use this option to disable the removal of fragments of sections, instead including or not including each section in its entirety, usually resulting in a larger application.

See also *Keeping symbols and sections*, page 105.



To set this option, use **Project>Options>Linker>Extra Options**

--no_library_search

Syntax `--no_library_search`

Description Use this option to disable the automatic runtime library search. This option turns off the automatic inclusion of the correct standard libraries. This is useful, for example, if the application needs a user-built standard library, etc.

Note that the option disables all steps of the automatic library selection, some of which might need to be reproduced if you are using the standard libraries. Use the `--log libraries` linker option together with automatic library selection enabled to determine which the steps are.



Project>Options>Linker>Library>Automatic runtime library selection

--no_literal_pool

Syntax `--no_literal_pool`

Description Use this option for code that should run from a memory region where it is not allowed to read data, only to execute code.

When this option is used, the linker will use the `MOV32` pseudo instruction in a mode-changing veneer, to avoid using the data bus to load the destination address. The option also means that libraries compiled with this option will be used.

The option `--no_literal_pool` is only allowed for ARMv7-M cores.

See also *--no_literal_pool*, page 281.



To set this option, use **Project>Options>Linker>Extra Options**.

--no_locals

Syntax `--no_locals`

Description Use this option to remove local symbols from the ELF executable image.

Note: This option does not remove any local symbols from the DWARF information in the executable image.



Project>Options>Linker>Output

--no_range_reservations

Syntax `--no_range_reservations`

Description Normally, the linker reserves any ranges used by absolute symbols with a non-zero size, excluding them from consideration for `place in` commands.

When this option is used, these reservations are disabled, and the linker is free to place sections in such a way as to overlap the extent of absolute symbols.



To set this option, use **Project>Options>Linker>Extra Options**.

--no_remove

Syntax `--no_remove`


Description When this option is used, unused sections are not removed. In other words, each module that is included in the executable image contains all its original sections.

See also *Keeping symbols and sections*, page 105.




To set this option, use **Project>Options>Linker>Extra Options**.


--no_veneers

Syntax	<code>--no_veneers</code>
Description	Use this option to disable the insertion of veneers even though the executable image needs it. In this case, the linker will generate a relocation error for each reference that needs a veneer.
See also	<i>Veneers</i> , page 111.
	 To set this option, use Project>Options>Linker>Extra Options .


--no_vfe

Syntax	<code>--no_vfe</code>
Description	Use this option to disable the Virtual Function Elimination optimization. All virtual functions in all classes with at least one instance will be kept, and Runtime Type Information data will be kept for all polymorphic classes. Also, no warning message will be issued for modules that lack VFE information.
See also	<code>--vfe</code> , page 326 and <i>Virtual function elimination</i> , page 217.
	 To set related options, choose: Project>Options>Linker>Optimizations>PerformC++ Virtual Function Elimination


--no_warnings

Syntax	<code>--no_warnings</code>
Description	By default, the linker issues warning messages. Use this option to disable all warning messages.
	 This option is not available in the IDE.


--no_wrap_diagnostics

Syntax	<code>--no_wrap_diagnostics</code>
Description	By default, long lines in diagnostic messages are broken into several lines to make the message easier to read. Use this option to disable line wrapping of diagnostic messages.
	 This option is not available in the IDE.

--only_stdout

Syntax	<code>--only_stdout</code>
Description	Use this option to make the linker use the standard output stream (<code>stdout</code>) also for messages that are normally directed to the error output stream (<code>stderr</code>).
	 This option is not available in the IDE.

--output, -o

Syntax	<code>--output {filename directory}</code> <code>-o {filename directory}</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	By default, the object executable image produced by the linker is located in a file with the name <code>a.out</code> . Use this option to explicitly specify a different output filename, which by default will have the filename extension <code>out</code> .
	 Project>Options>Linker>Output>Output file

--pi_veneers

Syntax	<code>--pi_veneers</code>
Description	Use this option to make the linker generate position-independent veneers. Note that this type of veneer is larger and slower than normal veneers.
See also	<i>Veneers</i> , page 111.



To set this option, use **Project>Options>Linker>Extra Options**.

--place_holder

Syntax	<code>--place_holder symbol[,size[,section[,alignment]]]</code>	
Parameters	<i>symbol</i>	The name of the symbol to create
	<i>size</i>	Size in ROM; by default 4 bytes
	<i>section</i>	Section name to use; by default <code>.text</code>
	<i>alignment</i>	Alignment of section; by default 1
Description	Use this option to reserve a place in ROM to be filled by some other tool, for example a checksum calculated by <code>ie1ftool</code> . Each use of this linker option results in a section with the specified name, size, and alignment. The symbol can be used by your application to refer to the section.	
	Note: Like any other section, sections created by the <code>--place_holder</code> option will only be included in your application if the section appears to be needed. The <code>--keep</code> linker option, or the <code>keep</code> linker directive can be used for forcing such section to be included.	
See also	<i>IAR utilities</i> , page 477.	



To set this option, use **Project>Options>Linker>Extra Options**

--redirect

Syntax	<code>--redirect <i>from_symbol</i>=<i>to_symbol</i></code>		
Parameters	<i>from_symbol</i>	The name of the source symbol	
	<i>to_symbol</i>	The name of the destination symbol	
Description	Use this option to change references to an external symbol so that they refer to another symbol.		
	Note: Redirection will normally not affect references within a module.		



To set this option, use **Project>Options>Linker>Extra Options**

--remarks

Syntax	<code>--remarks</code>
Description	The least severe diagnostic messages are called remarks. A remark indicates a source code construct that may cause strange behavior in the generated code. By default, the linker does not generate remarks. Use this option to make the linker generate remarks.
See also	<i>Severity levels</i> , page 249.



Project>Options>Linker>Diagnostics>Enable remarks

--search, -L


Syntax	<code>--search path</code> <code>-L path</code>
Parameters	<i>path</i> A path to a directory where the linker should search for object and library files.
Description	Use this option to specify more directories for the linker to search for object and library files in. By default, the linker searches for object and library files only in the working directory. Each use of this option on the command line adds another search directory.
See also	<i>The linking process</i> , page 56.




This option is not available in the IDE.

--semihosting


Syntax	<code>--semihosting[=iar_breakpoint]</code>
--------	---

Parameters	<code>iar_breakpoint</code>	The IAR-specific mechanism can be used when debugging applications that use SWI/SVC extensively.
Description	Use this option to include the debug interface—breakpoint mechanism—in the output image. If no parameter is specified, the SWI/SVC mechanism is included for ARM7/9/11, and the BKPT mechanism is included for Cortex-M.	
See also	<i>The semihosting mechanism</i> , page 130.	
		Project>Options>General Options>Library Configuration>Semihosted

--silent

Syntax	<code>--silent</code>
Description	<p>By default, the linker issues introductory messages and a final statistics report. Use this option to make the linker operate without sending these messages to the standard output stream (normally the screen).</p> <p>This option does not affect the display of error and warning messages.</p> <p> This option is not available in the IDE.</p>

--skip_dynamic_initialization

Syntax	<code>--skip_dynamic_initialization</code>
Description	<p>When using the IAR C/C++ compiler and the standard library, C++ dynamic initialization is handled automatically.</p> <p>Use this option to suppress dynamic initialization to be performed during system startup. Typically, this can be useful if you need to set up, for example, heap management for an RTOS before the initialization takes place.</p> <p>In this case you must add a call to the library function <code>__iar_dynamic_initialization</code> in your application source code. Initialization will then take place at the time of the call to this function.</p> <p> To set this option, use Project>Options>Linker>Extra Options.</p>

--stack_usage_control

Syntax	<code>--stack_usage_control=filename</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	Use this option to specify a stack usage control file. This file controls stack usage analysis, or provides more stack usage information for modules or functions. You can use this option multiple times to specify multiple stack usage control files. If no filename extension is specified, the extension <code>suc</code> is used. Using this option enables stack usage analysis in the linker.
See also	<i>Stack usage analysis</i> , page 92



Project>Options>Linker>Advanced>Control file

--strip

Syntax	<code>--strip</code>
Description	By default, the linker retains the debug information from the input object files in the output executable image. Use this option to remove that information.



To set related options, choose:

Project>Options>Linker>Output>Include debug information in output


--threaded_lib

Syntax	<code>--threaded_lib</code>
Description	Use this option to automatically configure the runtime library for use with threads.




Project>Options>General Options>Library Configuration>Enable thread support in library


--use_full_std_template_names

Syntax	<code>--use_full_std_template_names</code>	
Description	In the unmangled names of C++ entities, the linker by default uses shorter names for some classes. For example, "std::string" instead of "std::basic_string<char, std::char_traits<char>, std::allocator<char>>". Use this option to make the linker instead use the full, unabbreviated names.	
		This option is not available in the IDE.

--vfe

Syntax	<code>--vfe=[forced]</code>	
Parameters	<code>forced</code>	Performs Virtual Function Elimination even if one or more modules lack the needed virtual function elimination information.
Description	By default, Virtual Function Elimination is always performed but requires that all object files contain the necessary virtual function elimination information. Use <code>--vfe=forced</code> to perform Virtual Function Elimination even if one or more modules do not have the necessary information. Forcing the use of Virtual Function Elimination can be unsafe if some of the modules that lack the needed information perform virtual function calls or use dynamic Runtime Type Information.	
See also	<code>--no_vfe</code> , page 320 and <i>Virtual function elimination</i> , page 217.	
		To set related options, choose: Project>Options>Linker>Optimizations>Perform C++ Virtual Function Elimination

--warnings_affect_exit_code

Syntax	<code>--warnings_affect_exit_code</code>
Description	By default, the exit code is not affected by warnings, because only errors produce a non-zero exit code. With this option, warnings will also generate a non-zero exit code.
	 This option is not available in the IDE.

--warnings_are_errors

Syntax	<code>--warnings_are_errors</code>
Description	Use this option to make the linker treat all warnings as errors. If the linker encounters an error, no executable image is generated. Warnings that have been changed into remarks are not treated as errors. Note: Any diagnostic messages that have been reclassified as warnings by the option <code>--diag_warning</code> will also be treated as errors when <code>--warnings_are_errors</code> is used.
See also	<code>--diag_warning</code> , page 267 and <code>--diag_warning</code> , page 307.



Project>Options>Linker>Diagnostics>Treat all warnings as errors

--whole_archive

Syntax	<code>--whole_archive filename</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
Description	Use this option to make the linker treat every object file in the archive as if it was specified on the command line. This is useful when an archive contains root content that is always included from an object file (filename extension <code>.o</code>), but only included from an archive if some entry from the module is referred to.
Example	If <code>archive.a</code> contains the object files <code>file1.o</code> , <code>file2.o</code> , and <code>file3.o</code> , using <code>--whole_archive archive.a</code> is equivalent to specifying <code>file1.o file2.o file3.o</code> .
See also	<i>Keeping modules</i> , page 105



To set this option, use **Project>Options>Linker>Extra Options**

Data representation

- Alignment
- Byte order
- Basic data types—integer types
- Basic data types—floating-point types
- Pointer types
- Structure types
- Type qualifiers
- Data types in C++

See the chapter *Efficient coding for embedded applications* for information about which data types provide the most efficient code for your application.

Alignment

Every C data object has an alignment that controls how the object can be stored in memory. Should an object have an alignment of, for example, 4, it must be stored on an address that is divisible by 4.

The reason for the concept of alignment is that some processors have hardware limitations for how the memory can be accessed.

Assume that a processor can read 4 bytes of memory using one instruction, but only when the memory read is placed on an address divisible by 4. Then, 4-byte objects, such as `long` integers, will have alignment 4.

Another processor might only be able to read 2 bytes at a time; in that environment, the alignment for a 4-byte `long` integer might be 2.

A structure type will have the same alignment as the structure member with the most strict alignment. To decrease the alignment requirements on the structure and its members, use `#pragma pack` or the `__packed` data type attribute.

All data types must have a size that is a multiple of their alignment. Otherwise, only the first element of an array would be guaranteed to be placed in accordance with the alignment requirements. This means that the compiler might add pad bytes at the end of the structure. For more information about pad bytes, see *Packed structure types*, page 339.

Note that with the `#pragma data_alignment` directive you can increase the alignment demands on specific variables.

ALIGNMENT ON THE ARM CORE

The alignment of a data object controls how it can be stored in memory. The reason for using alignment is that the ARM core can access 4-byte objects more efficiently when the object is stored at an address divisible by 4.

Objects with alignment 4 must be stored at an address divisible by 4, while objects with alignment 2 must be stored at addresses divisible by 2.

The compiler ensures this by assigning an alignment to every data type, ensuring that the ARM core will be able to read the data.

For related information, see `--align_sp_on_irq`, page 261 and `--no_const_align`, page 279.

Byte order

In the little-endian byte order, which is default, the *least* significant byte is stored at the lowest address in memory. The *most* significant byte is stored at the highest address.

In the big-endian byte order, the *most* significant byte is stored at the lowest address in memory. The *least* significant byte is stored at the highest address. If you use the big-endian byte order, it might be necessary to use the `#pragma bitfields=reversed` directive to be compatible with code for other compilers and I/O register definitions of some devices, see *Bitfields*, page 332.

Note: There are two variants of the big-endian mode, BE8 and BE32, which you specify at link time. In BE8 data is big-endian and code is little-endian. In BE32 both data and code are big-endian. In architectures before v6, the BE32 endian mode is used, and after v6 the BE8 mode is used. In the v6 (ARM11) architecture, both big-endian modes are supported.

Basic data types—integer types

The compiler supports both all Standard C basic data types and some additional types.

INTEGER TYPES—AN OVERVIEW

This table gives the size and range of each integer data type:

Data type	Size	Range	Alignment
<code>bool</code>	8 bits	0 to 1	1
<code>char</code>	8 bits	0 to 255	1
<code>signed char</code>	8 bits	-128 to 127	1
<code>unsigned char</code>	8 bits	0 to 255	1
<code>signed short</code>	16 bits	-32768 to 32767	2
<code>unsigned short</code>	16 bits	0 to 65535	2
<code>signed int</code>	32 bits	-2 ³¹ to 2 ³¹ -1	4
<code>unsigned int</code>	32 bits	0 to 2 ³² -1	4
<code>signed long</code>	32 bits	-2 ³¹ to 2 ³¹ -1	4
<code>unsigned long</code>	32 bits	0 to 2 ³² -1	4
<code>signed long long</code>	64 bits	-2 ⁶³ to 2 ⁶³ -1	8
<code>unsigned long long</code>	64 bits	0 to 2 ⁶⁴ -1	8

Table 30: Integer types

Signed variables are represented using the two’s complement form.

BOOL

The `bool` data type is supported by default in the C++ language. If you have enabled language extensions, the `bool` type can also be used in C source code if you include the file `stdbool.h`. This will also enable the boolean values `false` and `true`.

THE ENUM TYPE

The compiler will use the smallest type required to hold `enum` constants, preferring signed rather than unsigned.

When IAR Systems language extensions are enabled, and in C++, the `enum` constants and types can also be of the type `long`, `unsigned long`, `long long`, or `unsigned long long`.

To make the compiler use a larger type than it would automatically use, define an `enum` constant with a large enough value. For example:

```
/* Disables usage of the char type for enum */
enum Cards{Spade1, Spade2,
           DontUseChar=257};
```

For related information, see `--enum_is_int`, page 272.

THE CHAR TYPE

The `char` type is by default unsigned in the compiler, but the `--char_is_signed` compiler option allows you to make it signed. Note, however, that the library is compiled with the `char` type as unsigned.

THE WCHAR_T TYPE

The `wchar_t` data type is an integer type whose range of values can represent distinct codes for all members of the largest extended character set specified among the supported locals.

The `wchar_t` data type is supported by default in the C++ language. To use the `wchar_t` type also in C source code, you must include the file `stdint.h` from the runtime library.

BITFIELDS

In Standard C, `int`, `signed int`, and `unsigned int` can be used as the base type for integer bitfields. In standard C++, and in C when language extensions are enabled in the compiler, any integer or enumeration type can be used as the base type. It is implementation defined whether a plain integer type (`char`, `short`, `int`, etc) results in a signed or unsigned bitfield.

In the IAR C/C++ Compiler for ARM, plain integer types are treated as unsigned.

Bitfields in expressions are treated as `int` if `int` can represent all values of the bitfield. Otherwise, they are treated as the bitfield base type.

Each bitfield is placed in the next suitably aligned container of its base type that has enough available bits to accommodate the bitfield. Within each container, the bitfield is placed in the first available byte or bytes, taking the byte order into account. Note that containers can overlap if needed, as long as they are suitably aligned for their type.

In addition, the compiler supports an alternative bitfield allocation strategy (disjoint types), where bitfield containers of different types are not allowed to overlap. Using this allocation strategy, each bitfield is placed in a new container if its type is different from that of the previous bitfield, or if the bitfield does not fit in the same container as the previous bitfield. Within each container, the bitfield is placed from the least significant

bit to the most significant bit (disjoint types) or from the most significant bit to the least significant bit (reverse disjoint types). This allocation strategy will never use less space than the default allocation strategy (joined types), and can use significantly more space when mixing bitfield types.

Use the `#pragma bitfields` directive to choose which bitfield allocation strategy to use, see *bitfields*, page 361.

Assume this example:

```
struct BitfieldExample
{
    uint32_t a : 12;
    uint16_t b : 3;
    uint16_t c : 7;
    uint8_t d;
};
```

The example in the joined types bitfield allocation strategy

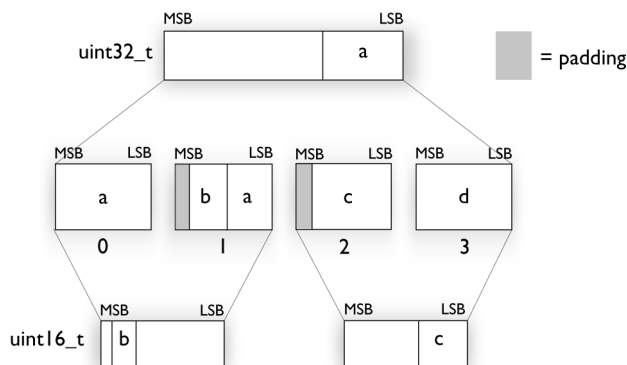
To place the first bitfield, `a`, the compiler allocates a 32-bit container at offset 0 and puts `a` into the first and second bytes of the container.

For the second bitfield, `b`, a 16-bit container is needed and because there are still four bits free at offset 0, the bitfield is placed there.

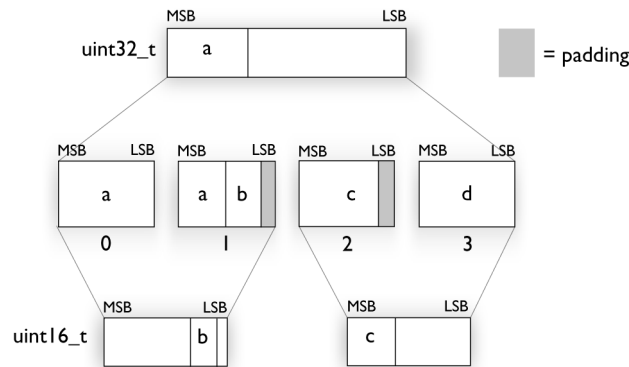
For the third bitfield, `c`, as there is now only one bit left in the first 16-bit container, a new container is allocated at offset 2, and `c` is placed in the first byte of this container.

The fourth member, `d`, can be placed in the next available full byte, which is the byte at offset 3.

In little-endian mode, each bitfield is allocated starting from the least significant free bit of its container to ensure that it is placed into bytes from left to right.



In big-endian mode, each bitfield is allocated starting from the most significant free bit of its container to ensure that it is placed into bytes from left to right.



The example in the disjoint types bitfield allocation strategy

To place the first bitfield, `a`, the compiler allocates a 32-bit container at offset 0 and puts `a` into the least significant 12 bits of the container.

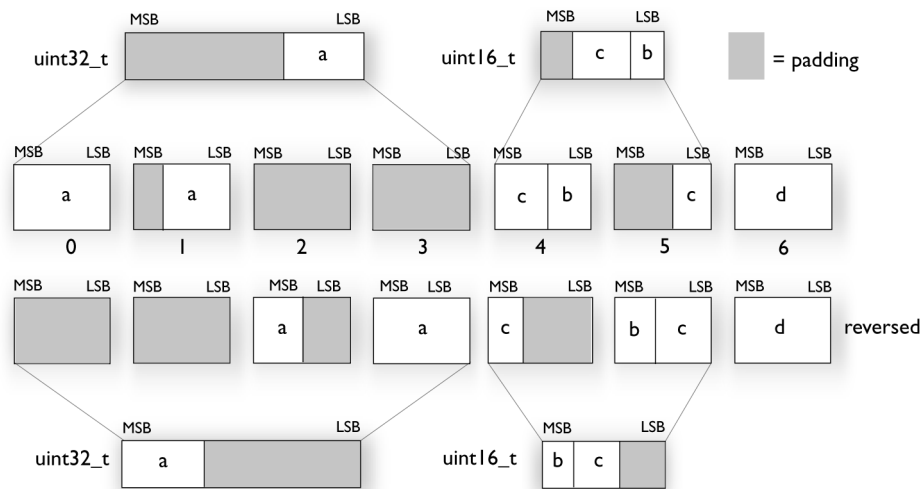
To place the second bitfield, `b`, a new container is allocated at offset 4, because the type of the bitfield is not the same as that of the previous one. `b` is placed into the least significant three bits of this container.

The third bitfield, `c`, has the same type as `b` and fits into the same container.


The fourth member, `d`, is allocated into the byte at offset 6. `d` cannot be placed into the same container as `b` and `c` because it is not a bitfield, it is not of the same type, and it would not fit.

When using reverse order (reverse disjoint types), each bitfield is instead placed starting from the most significant bit of its container.

This is the layout of `bitfield_example` in little-endian mode:



This is the layout of `bitfield_example` in big-endian mode:



Basic data types—floating-point types

In the IAR C/C++ Compiler for ARM, floating-point values are represented in standard IEEE 754 format. The sizes for the different floating-point types are:

Type	Size	Range (+/-)	Decimals	Exponent	Mantissa	Alignment
float	32 bits	±1.18E-38 to ±3.40E+38	7	8 bits	23 bits	4
double	64 bits	±2.23E-308 to ±1.79E+308	15	11 bits	52 bits	8
long double	64 bits	±2.23E-308 to ±1.79E+308	15	11 bits	52 bits	8

Table 31: Floating-point types

For Cortex-M0 and Cortex-M1, the compiler does not support subnormal numbers. All operations that should produce subnormal numbers will instead generate zero. For information about the representation of subnormal numbers for other cores, see *Representation of special floating-point numbers*, page 337.


FLOATING-POINT ENVIRONMENT

Exception flags for floating-point values are supported for devices with a VFP unit, and they are defined in the `fenv.h` file. For devices without a VFP unit, the functions defined in the `fenv.h` file exist but have no functionality.

The `feraiseexcept` function does not raise an `inexact` floating-point exception when called with `FE_OVERFLOW` or `FE_UNDERFLOW`.

32-BIT FLOATING-POINT FORMAT

The representation of a 32-bit floating-point number as an integer is:



The exponent is 8 bits, and the mantissa is 23 bits.

The value of the number is:

$$(-1)^S * 2^{(Exponent-127)} * 1.Mantissa$$


The range of the number is at least:

$$\pm 1.18E-38 \text{ to } \pm 3.39E+38$$

The precision of the float operators (+, -, *, and /) is approximately 7 decimal digits.

64-BIT FLOATING-POINT FORMAT

The representation of a 64-bit floating-point number as an integer is:



The exponent is 11 bits, and the mantissa is 52 bits.

The value of the number is:

$$(-1)^S * 2^{(\text{Exponent}-1023)} * 1.\text{Mantissa}$$

The range of the number is at least:

$$\pm 2.23\text{E}-308 \text{ to } \pm 1.79\text{E}+308$$

The precision of the float operators (+, -, *, and /) is approximately 15 decimal digits.

REPRESENTATION OF SPECIAL FLOATING-POINT NUMBERS

This list describes the representation of special floating-point numbers:

- Zero is represented by zero mantissa and exponent. The sign bit signifies positive or negative zero.
- Infinity is represented by setting the exponent to the highest value and the mantissa to zero. The sign bit signifies positive or negative infinity.
- Not a number (NaN) is represented by setting the exponent to the highest positive value and the most significant bit in the mantissa to 1. The value of the sign bit is ignored.
- Subnormal numbers are used for representing values smaller than what can be represented by normal values. The drawback is that the precision will decrease with smaller values. The exponent is set to 0 to signify that the number is subnormal, even though the number is treated as if the exponent was 1. Unlike normal numbers, subnormal numbers do not have an implicit 1 as the most significant bit (the MSB) of the mantissa. The value of a subnormal number is:

$$(-1)^S * 2^{(1-\text{BIAS})} * 0.\text{Mantissa}$$

where BIAS is 127 and 1023 for 32-bit and 64-bit floating-point values, respectively.

Pointer types

The compiler has two basic types of pointers: function pointers and data pointers.

FUNCTION POINTERS

The size of function pointers is always 32 bits, and the range is 0x0-0xFFFFFFFF.

When function pointer types are declared, attributes are inserted before the `*` sign, for example:

```
typedef void (__thumb __interwork * IntHandler) (void);
```

This can be rewritten using `#pragma` directives:

```
#pragma type_attribute=__thumb __interwork
typedef void IntHandler_function(void);
typedef IntHandler_function *IntHandler;
```

DATA POINTERS

There is one data pointer available. Its size is 32 bits and the range is `0x0–0xFFFFFFFF`.

CASTING

Casts between pointers have these characteristics:

- Casting a *value* of an integer type to a pointer of a smaller type is performed by truncation
- Casting a pointer type to a smaller integer type is performed by truncation
- Casting a *pointer type* to a larger integer type is performed by zero extension
- Casting a *data pointer* to a function pointer and vice versa is illegal
- Casting a *function pointer* to an integer type gives an undefined result
- Casting a value of an unsigned integer type to a pointer of a larger type is performed by zero extension

size_t

`size_t` is the unsigned integer type of the result of the `sizeof` operator. In the IAR C/C++ Compiler for ARM, the type used for `size_t` is unsigned `int`.

ptrdiff_t

`ptrdiff_t` is the signed integer type of the result of subtracting two pointers. In the IAR C/C++ Compiler for ARM, the type used for `ptrdiff_t` is the signed integer variant of the `size_t` type.

intptr_t

`intptr_t` is a signed integer type large enough to contain a `void *`. In the IAR C/C++ Compiler for ARM, the type used for `intptr_t` is signed long `int`.

uintptr_t

`uintptr_t` is equivalent to `intptr_t`, with the exception that it is unsigned.

Structure types

The members of a `struct` are stored sequentially in the order in which they are declared: the first member has the lowest memory address.

ALIGNMENT OF STRUCTURE TYPES


The `struct` and `union` types have the same alignment as the member with the highest alignment requirement. Note that this alignment requirement also applies to a member that is a structure. To allow arrays of aligned structure objects, the size of a `struct` is adjusted to an even multiple of the alignment.

GENERAL LAYOUT

Members of a `struct` are always allocated in the order specified in the declaration. Each member is placed in the `struct` according to the specified alignment (offsets).

```
struct First
{
    char c;
    short s;
} s;
```

This diagram shows the layout in memory:



The alignment of the structure is 2 bytes, and a pad byte must be inserted to give `short s` the correct alignment.

PACKED STRUCTURE TYPES

The `packed` keyword is used for relaxing the alignment requirements of the members of a structure. This changes the layout of the structure. The members are placed in the same order as when declared, but there might be less pad space between members.

Note that accessing an object that is not correctly aligned requires code that is both larger and slower. If such structure members are accessed many times, it is usually better to construct the correct values in a `struct` that is not packed, and access this `struct` instead.

Special care is also needed when creating and using pointers to misaligned members. For direct access to misaligned members in a packed `struct`, the compiler will emit the correct (but slower and larger) code when needed. However, when a misaligned member is accessed through a pointer to the member, the normal (smaller and faster) code is


used. In the general case, this will not work, because the normal code might depend on the alignment being correct.

This example declares a packed structure:

```
#pragma pack(1)
struct S
{
    char c;
    short s;
};

#pragma pack()
```


The structure `S` has this memory layout:



The next example declares a new non-packed structure, `S2`, that contains the structure `s` declared in the previous example:

```
struct S2
{
    struct S s;
    long l;
};
```

The structure `S2` has this memory layout



The structure `S` will use the memory layout, size, and alignment described in the previous example. The alignment of the member `l` is 4, which means that alignment of the structure `S2` will become 4.

For more information, see *Alignment of elements in a structure*, page 222.

Type qualifiers

According to the C standard, `volatile` and `const` are type qualifiers.

DECLARING OBJECTS VOLATILE

By declaring an object `volatile`, the compiler is informed that the value of the object can change beyond the compiler's control. The compiler must also assume that any accesses can have side effects—thus all accesses to the `volatile` object must be preserved.

There are three main reasons for declaring an object `volatile`:

- Shared access; the object is shared between several tasks in a multitasking environment
- Trigger access; as for a memory-mapped SFR where the fact that an access occurs has an effect
- Modified access; where the contents of the object can change in ways not known to the compiler.

Definition of access to volatile objects

The C standard defines an abstract machine, which governs the behavior of accesses to `volatile` declared objects. In general and in accordance to the abstract machine:

- The compiler considers each read and write access to an object declared `volatile` as an access
- The unit for the access is either the entire object or, for accesses to an element in a composite object—such as an array, struct, class, or union—the element. For example:

```
char volatile a;
a = 5;    /* A write access */
a += 6;   /* First a read then a write access */
```
- An access to a bitfield is treated as an access to the underlying type
- Adding a `const` qualifier to a `volatile` object will make write accesses to the object impossible. However, the object will be placed in RAM as specified by the C standard.

However, these rules are not detailed enough to handle the hardware-related requirements. The rules specific to the IAR C/C++ Compiler for ARM are described below.

Rules for accesses

In the IAR C/C++ Compiler for ARM, accesses to `volatile` declared objects are subject to these rules:

- All accesses are preserved
- All accesses are complete, that is, the whole object is accessed

- All accesses are performed in the same order as given in the abstract machine
- All accesses are atomic, that is, they cannot be interrupted.

The compiler adheres to these rules for accesses to all 8-, 16-, and 32-bit scalar types, except for accesses to unaligned 16- and 32-bit fields in packed structures.

For all combinations of object types not listed, only the rule that states that all accesses are preserved applies.

DECLARING OBJECTS VOLATILE AND CONST

If you declare a `volatile` object `const`, it will be write-protected but it will still be stored in RAM memory as the C standard specifies.

To store the object in read-only memory instead, but still make it possible to access it as a `const volatile` object, define the variable like this:

```
const volatile int x @ "FLASH";
```

The compiler will generate the read/write section `FLASH`. That section should be placed in ROM and is used for manually initializing the variables when the application starts up.

Thereafter, the initializers can be reflashed with other values at any time.

DECLARING OBJECTS CONST

The `const` type qualifier is used for indicating that a data object, accessed directly or via a pointer, is non-writable. A pointer to `const` declared data can point to both constant and non-constant objects. It is good programming practice to use `const` declared pointers whenever possible because this improves the compiler's possibilities to optimize the generated code and reduces the risk of application failure due to erroneously modified data.

Static and global objects declared `const` are allocated in ROM.

In C++, objects that require runtime initialization cannot be placed in ROM.

Data types in C++

In C++, all plain C data types are represented in the same way as described earlier in this chapter. However, if any Embedded C++ features are used for a type, no assumptions can be made concerning the data representation. This means, for example, that it is not supported to write assembler code that accesses class members.

Extended keywords

- General syntax rules for extended keywords
- Summary of extended keywords
- Descriptions of extended keywords

General syntax rules for extended keywords

The compiler provides a set of attributes that can be used on functions or data objects to support specific features of the ARM core. There are two types of attributes—*type attributes* and *object attributes*:

- Type attributes affect the *external functionality* of the data object or function
- Object attributes affect the *internal functionality* of the data object or function.

The syntax for the keywords differs slightly depending on whether it is a type attribute or an object attribute, and whether it is applied to a data object or a function.

For more information about each attribute, see *Descriptions of extended keywords*, page 347.

Note: The extended keywords are only available when language extensions are enabled in the compiler.



In the IDE, language extensions are enabled by default.



Use the `-e` compiler option to enable language extensions. See `-e`, page 269.

TYPE ATTRIBUTES

Type attributes define how a function is called, or how a data object is accessed. This means that if you use a type attribute, it must be specified both when a function or data object is defined and when it is declared.

You can either place the type attributes explicitly in your declarations, or use the pragma directive `#pragma type_attribute`.

Type attributes can be further divided into *memory type attributes* and *general type attributes*. Memory type attributes are referred to as simply *memory attributes* in the rest of the documentation.

General type attributes

Available *function type attributes* (affect how the function should be called):

```
__arm, __fiq, __interwork, __irq, __swi, __task, __thumb
```

Available *data type attributes*:

```
__big_endian, __little_endian __packed
```

You can specify as many type attributes as required for each level of pointer indirection.

Syntax for type attributes used on data objects

Type attributes use almost the same syntax rules as the type qualifiers `const` and `volatile`. For example:

```
__little_endian int i;
int __little_endian j;
```

Both `i` and `j` will be accessed with little-endian byte order.

Unlike `const` and `volatile`, when a type attribute is used before the type specifier in a derived type, the type attribute applies to the object, or typedef itself, except in structure member declarations.

Using a type definition can sometimes make the code clearer:

```
typedef __packed int packed_int;
packed_int *q1;
```

`packed_int` is a typedef for packed integers. The variable `q1` can point to such integers.

You can also use the `#pragma type_attributes` directive to specify type attributes for a declaration. The type attributes specified in the pragma directive are applied to the data object or typedef being declared.

```
#pragma type_attribute=__packed
int * q2;
```

The variable `q2` is packed.

Syntax for type attributes used on functions

The syntax for using type attributes on functions differs slightly from the syntax of type attributes on data objects. For functions, the attribute must be placed either in front of the return type, or in parentheses, for example:

```
__irq __arm void my_handler(void);

or

void (__irq __arm my_handler)(void);
```


This declaration of `my_handler` is equivalent with the previous one:

```
#pragma type_attribute=__irq __arm
void my_handler(void);
```

OBJECT ATTRIBUTES

Normally, object attributes affect the internal functionality of functions and data objects, but not directly how the function is called or how the data is accessed. This means that an object attribute does not normally need to be present in the declaration of an object.

These object attributes are available:

- Object attributes that can be used for variables:


```
__absolute, __no_alloc, __no_alloc16, __no_alloc_str,
__no_alloc_str16, __no_init, __ro_placement
```
- Object attributes that can be used for functions and variables:


```
location, @, __root, __weak
```
- Object attributes that can be used for functions:


```
__intrinsic, __nested, __noreturn, __ramfunc, __stackless
```

You can specify as many object attributes as required for a specific function or data object.

For more information about `location` and `@`, see *Controlling data and function placement in memory*, page 224.

Syntax for object attributes

The object attribute must be placed in front of the type. For example, to place `myarray` in memory that is not initialized at startup:

```
__no_init int myarray[10];
```

The `#pragma object_attribute` directive can also be used. This declaration is equivalent to the previous one:

```
#pragma object_attribute=__no_init
int myarray[10];
```

Note: Object attributes cannot be used in combination with the `typedef` keyword.

Summary of extended keywords

This table summarizes the extended keywords:

Extended keyword	Description
<code>__absolute</code>	Makes references to the object use absolute addressing
<code>__arm</code>	Makes a function execute in ARM mode
<code>__big_endian</code>	Declares a variable to use the big-endian byte order
<code>__fiq</code>	Declares a fast interrupt function
<code>__interwork</code>	Declares a function to be callable from both ARM and Thumb mode
<code>__intrinsic</code>	Reserved for compiler internal use only
<code>__irq</code>	Declares an interrupt function
<code>__little_endian</code>	Declares a variable to use the little-endian byte order
<code>__nested</code>	Allows an <code>__irq</code> declared interrupt function to be nested, that is, interruptible by the same type of interrupt
<code>__no_alloc,</code> <code>__no_alloc16</code>	Makes a constant available in the execution file
<code>__no_alloc_str,</code> <code>__no_alloc_str16</code>	Makes a string literal available in the execution file
<code>__no_init</code>	Places a data object in non-volatile memory
<code>__noreturn</code>	Informs the compiler that the function will not return
<code>__packed</code>	Decreases data type alignment to 1
<code>__pcrel</code>	Used internally by the compiler for constant data when the <code>--ropi</code> compiler option is used
<code>__ramfunc</code>	Makes a function execute in RAM
<code>__root</code>	Ensures that a function or variable is included in the object code even if unused
<code>__ro_placement</code>	Places <code>const volatile</code> data in read-only memory.
<code>__sbrel</code>	Used internally by the compiler for constant data when the <code>--rwpi</code> compiler option is used
<code>__stackless</code>	Makes a function callable without a working stack
<code>__swi</code>	Declares a software interrupt function
<code>__task</code>	Relaxes the rules for preserving registers
<code>__thumb</code>	Makes a function execute in Thumb mode
<code>__weak</code>	Declares a symbol to be externally weakly linked

Table 32: Extended keywords summary

Descriptions of extended keywords

This section gives detailed information about each extended keyword.

__absolute

Syntax

See *Syntax for object attributes*, page 345.

Description

The `__absolute` keyword makes references to the object use absolute addressing.

The following limitations apply:

- Only available when the `--ropi` or `--rwpi` compiler option is used
- Can only be used on external declarations.

Example

```
extern __absolute char otherBuffer[100];
```

__arm

Syntax

See *Syntax for type attributes used on functions*, page 344.

Description

The `__arm` keyword makes a function execute in ARM mode. An `__arm` declared function can, unless it is also declared `__interwork`, only be called from functions that also execute in ARM mode.

A function declared `__arm` cannot be declared `__thumb`.

Note: Non-interwork ARM functions cannot be called from Thumb mode.

Example

```
__arm int func1(void);
```

See also

`__interwork`, page 348.

__big_endian

Syntax

See *Syntax for type attributes used on data objects*, page 344.

Description

The `__big_endian` keyword is used for accessing a variable that is stored in the big-endian byte order regardless of what byte order the rest of the application uses. The `__big_endian` keyword is available when you compile for ARMv6 or higher.

Note that this keyword cannot be used on pointers. Also, this attribute cannot be used on arrays.

Example `__big_endian long my_variable;`

See also `__little_endian`, page 349.

__fiq

Syntax See *Syntax for type attributes used on functions*, page 344.

Description The `__fiq` keyword declares a fast interrupt function. All interrupt functions must be compiled in ARM mode. A function declared `__fiq` does not accept parameters and does not have a return value. This keyword is not available when you compile for Cortex-M devices.

Example `__fiq __arm void interrupt_function(void);`

__interwork

Syntax See *Syntax for type attributes used on functions*, page 344.

Description A function declared `__interwork` can be called from functions executing in either ARM or Thumb mode.

Note: By default, functions are interwork when the `--interwork` compiler option is used, and when the `--cpu` option is used and it specifies a core where interwork is default.

Example `typedef void (__thumb __interwork *IntHandler)(void);`

__intrinsic

Description The `__intrinsic` keyword is reserved for compiler internal use only.

__irq

Syntax See *Syntax for type attributes used on functions*, page 344.

Description The `__irq` keyword declares an interrupt function. All interrupt functions must be compiled in ARM mode. A function declared `__irq` does not accept parameters and does not have a return value. This keyword is not available when you compile for Cortex-M devices.

Example `__irq __arm void interrupt_function(void);`

See also `--align_sp_on_irq`, page 261

__little_endian

Syntax See *Syntax for type attributes used on data objects*, page 344.

Description The `__little_endian` keyword is used for accessing a variable that is stored in the little-endian byte order regardless of what byte order the rest of the application uses. The `__little_endian` keyword is available when you compile for ARMv6 or higher.

Note that this keyword cannot be used on pointers. Also, this attribute cannot be used on arrays.

Example `__little_endian long my_variable;`

See also `__big_endian`, page 347.

__nested

Syntax See *Syntax for object attributes*, page 345.

Description The `__nested` keyword modifies the enter and exit code of an interrupt function to allow for nested interrupts. This allows interrupts to be enabled, which means new interrupts can be served inside an interrupt function, without overwriting the SPSR and return address in R14. Nested interrupts are only supported for `__irq` declared functions.

Note: The `__nested` keyword requires the processor mode to be in either User or System mode.

Example `__irq __nested __arm void interrupt_handler(void);`

See also *Nested interrupts*, page 77 and `--align_sp_on_irq`, page 261.

__no_alloc, __no_alloc16

Syntax	See <i>Syntax for object attributes</i> , page 345.
Description	<p>Use the <code>__no_alloc</code> or <code>__no_alloc16</code> object attribute on a constant to make the constant available in the executable file without occupying any space in the linked application.</p> <p>You cannot access the contents of such a constant from your application. You can take its address, which is an integer offset to the section of the constant. The type of the offset is <code>unsigned long</code> when <code>__no_alloc</code> is used, and <code>unsigned short</code> when <code>__no_alloc16</code> is used.</p>
Example	<code>__no_alloc const struct MyData my_data @ "XXX" = {...};</code>
See also	<code>__no_alloc_str</code> , <code>__no_alloc_str16</code> , page 350.

__no_alloc_str, __no_alloc_str16

Syntax	<code>__no_alloc_str(string_literal @ section)</code> and <code>__no_alloc_str16(string_literal @ section)</code> where				
	<table> <tr> <td><i>string_literal</i></td><td>The string literal that you want to make available in the executable file.</td></tr> <tr> <td><i>section</i></td><td>The name of the section to place the string literal in.</td></tr> </table>	<i>string_literal</i>	The string literal that you want to make available in the executable file.	<i>section</i>	The name of the section to place the string literal in.
<i>string_literal</i>	The string literal that you want to make available in the executable file.				
<i>section</i>	The name of the section to place the string literal in.				
Description	<p>Use the <code>__no_alloc_str</code> or <code>__no_alloc_str16</code> operators to make string literals available in the executable file without occupying any space in the linked application.</p> <p>The value of the expression is the offset of the string literal in the section. For <code>__no_alloc_str</code>, the type of the offset is <code>unsigned long</code>. For <code>__no_alloc_str16</code>, the type of the offset is <code>unsigned short</code>.</p>				

Example

```
#define MYSEG "YYY"
#define X(str) __no_alloc_str(str @ MYSEG)

extern void dbg_printf(unsigned long fmt, ...)

#define DBGPRINTF(fmt, ...) dbg_printf(X(fmt), __VA_ARGS__)

void
foo(int i, double d)
{
    DBGPRINTF("The value of i is: %d, the value of d is: %f",i,d);
}
```

Depending on your debugger and the runtime support, this could produce trace output on the host computer. Note that there is no such runtime support in C-SPY, unless you use an external plugin module.

See also

`__no_alloc`, `__no_alloc16`, page 350.

__no_init**Syntax**

See *Syntax for object attributes*, page 345.

Description

Use the `__no_init` keyword to place a data object in non-volatile memory. This means that the initialization of the variable, for example at system startup, is suppressed.

Example

```
__no_init int myarray[10];
```

See also

Non-initialized variables, page 239 and *do not initialize directive*, page 449.

__noreturn**Syntax**

See *Syntax for object attributes*, page 345.

Description

The `__noreturn` keyword can be used on a function to inform the compiler that the function will not return. If you use this keyword on such functions, the compiler can optimize more efficiently. Examples of functions that do not return are `abort` and `exit`.

Note: At optimization levels medium or high, the `__noreturn` keyword might cause incorrect call stack debug information at any point where it can be determined that the current function cannot return.

Example

```
__noreturn void terminate(void);
```

__packed

Syntax

See *Syntax for type attributes used on data objects*, page 344. An exception is when the keyword is used for modifying the structure type in a `struct` or `union` declarations, see below.

Description

Use the `__packed` keyword to specify a data alignment of 1 for a data type. `__packed` can be used in two ways:

- When used before the `struct` or `union` keyword in a structure definition, the maximum alignment of each member in the structure is set to 1, eliminating the need for gaps between the members.
You can also use the `__packed` keyword with structure declarations, but it is illegal to refer to a structure type defined without the `__packed` keyword using a structure declaration with the `__packed` keyword.
- When used in any other position, it follows the syntax rules for type attributes, and affects a type in its entirety. A type with the `__packed` type attribute is the same as the type attribute without the `__packed` type attribute, except that it has a data alignment of 1. Types that already have an alignment of 1 are not affected by the `__packed` type attribute.

A normal pointer can be implicitly converted to a pointer to `__packed`, but the reverse conversion requires a cast.

Note: Accessing data types at other alignments than their natural alignment can result in code that is significantly larger and slower.

Use either `__packed` or `#pragma pack` to relax the alignment restrictions for a type and the objects defined using that type. Mixing `__packed` and `#pragma pack` might lead to unexpected behavior.

Example

```
/* No pad bytes in X: */
__packed struct X { char ch; int i; };
/* __packed is optional here: */
struct X * xp;

/* NOTE: no __packed: */
struct Y { char ch; int i; };
/* ERROR: Y not defined with __packed: */
__packed struct Y * yp ;

/* Member 'i' has alignment 1: */
struct Z { char ch; __packed int i; };
```



```
void Foo(struct X * xp)
{
    /* Error:"int *" -> "int __packed *" not allowed: */
    int * p1 = xp->l;
    /* OK: */
    int __packed * p2 = &xp->i;
    /* OK, char not affected */
    char * p3 = &xp->ch;
}
```

See also *pack*, page 373.

__ramfunc

Syntax

See *Syntax for type attributes used on functions*, page 344.

Description

The `__ramfunc` keyword makes a function execute in RAM. Two code sections will be created: one for the RAM execution (`.textrw`), and one for the ROM initialization (`.textrw_init`).

If a function declared `__ramfunc` tries to access ROM, the compiler will issue a warning. This behavior is intended to simplify the creation of *upgrade* routines, for instance, rewriting parts of flash memory. If this is not why you have declared the function `__ramfunc`, you can safely ignore or disable these warnings.

Functions declared `__ramfunc` are by default stored in the section named `.textrw`.

Example

```
__ramfunc int FlashPage(char * data, char * page);
```

See also

The *C-SPY® Debugging Guide for ARM* to read more about `__ramfunc` declared functions in relation to breakpoints.

__ro_placement

Syntax

See *Syntax for object attributes*, page 345.

Description

The `__ro_placement` attribute specifies that a data object should be placed in read-only memory. There are two cases where you might want to use this object attribute:

- Data objects declared `const volatile` are by default placed in read-write memory. Use the `__ro_placement` object attribute to place the data object in read-only memory instead.

- In C++, a data object declared `const` and that needs dynamic initialization is placed in read-write memory and initialized at system startup. If you use the `__ro_placement` object attribute, the compiler will give an error message if the data object needs dynamic initialization.

You can only use the `__ro_placement` object attribute on `const` objects.

In some cases (primarily involving simple constructors), the compiler will be able to optimize C++ dynamic initialization of a data object into static initialization. In that case no error message will be issued for the object.

Example

```
__ro_placement const volatile int x = 10;
```

__root

Syntax

See *Syntax for object attributes*, page 345.

Description

A function or variable with the `__root` attribute is kept whether or not it is referenced from the rest of the application, provided its module is included. Program modules are always included and library modules are only included if needed.

Example

```
__root int myarray[10];
```

See also

For more information about root symbols and how they are kept, see *Keeping symbols and sections*, page 105.

__stackless

Syntax

See *Syntax for object attributes*, page 345.

Description

The `__stackless` keyword declares a function that can be called without a working stack.



A function declared `__stackless` violates the calling convention in such a way that it is not possible to return from it. However, the compiler cannot reliably detect if the function returns and will not issue an error if it does.

Example

```
__stackless void start_application(void);
```

__swi

Syntax

See *Syntax for type attributes used on functions*, page 344.

Description

The `__swi` declares a software interrupt function. It inserts an `SVC` (formerly `SWI`) instruction and the specified software interrupt number to make a proper function call. A function declared `__swi` accepts arguments and returns values. The `__swi` keyword makes the compiler generate the correct return sequence for a specific software interrupt function. Software interrupt functions follow the same calling convention regarding parameters and return values as an ordinary function, except for the stack usage.

The `__swi` keyword also expects a software interrupt number which is specified with the `#pragma swi_number=number` directive. The `swi_number` is used as an argument to the generated assembler `SVC` instruction, and can be used by the `SVC` interrupt handler, for example `SWI_Handler`, to select one software interrupt function in a system containing several such functions. Note that the software interrupt number should only be specified in the function declaration—typically, in a header file that you include in the source code file that calls the interrupt function—not in the function definition.

Note: All interrupt functions must be compiled in ARM mode, except for Cortex-M. Use either the `__arm` keyword or the `#pragma type_attribute=__arm` directive to alter the default behavior if needed.

Example

To declare your software interrupt function, typically in a header file, write for example like this:

```
#pragma swi_number=0x23
__swi int swi0x23_function(int a, int b);
...
```

To call the function:

```
...
int x = swi0x23_function(1, 2); /* Will be replaced by SVC 0x23,
                                hence the linker will never
                                try to locate the
                                swi0x23_function */
...
```

Somewhere in your application source code, you define your software interrupt function:

```
...
__swi __arm int the_actual_swi0x23_function(int a, int b)
{
    ...
    return 42;
}
```

See also

Software interrupts, page 78 and *Calling convention*, page 169.

__task

Syntax	See <i>Syntax for type attributes used on functions</i> , page 344.
Description	<p>This keyword allows functions to relax the rules for preserving registers. Typically, the keyword is used on the start function for a task in an RTOS.</p> <p>By default, functions save the contents of used preserved registers on the stack upon entry, and restore them at exit. Functions that are declared <code>__task</code> do not save all registers, and therefore require less stack space.</p> <p>Because a function declared <code>__task</code> can corrupt registers that are needed by the calling function, you should only use <code>__task</code> on functions that do not return or call such a function from assembler code.</p> <p>The function <code>main</code> can be declared <code>__task</code>, unless it is explicitly called from the application. In real-time applications with more than one task, the root function of each task can be declared <code>__task</code>.</p>
Example	<pre>__task void my_handler(void);</pre>

__thumb

Syntax	See <i>Syntax for type attributes used on functions</i> , page 344.
Description	<p>The <code>__thumb</code> keyword makes a function execute in Thumb mode. Unless the function is also declared <code>__interwork</code>, the function declared <code>__thumb</code> can only be called from functions that also execute in Thumb mode.</p> <p>A function declared <code>__thumb</code> cannot be declared <code>__arm</code>.</p> <p>Note: Non-interwork Thumb functions cannot be called from ARM mode.</p>
Example	<pre>__thumb int func2(void);</pre>
See also	<code>__interwork</code> , page 348.

__weak

Syntax	See <i>Syntax for object attributes</i> , page 345.
Description	Using the <code>__weak</code> object attribute on an external declaration of a symbol makes all references to that symbol in the module weak.

Using the `__weak` object attribute on a public definition of a symbol makes that definition a weak definition.

The linker will not include a module from a library solely to satisfy weak references to a symbol, nor will the lack of a definition for a weak reference result in an error. If no definition is included, the address of the object will be zero.

When linking, a symbol can have any number of weak definitions, and at most one non-weak definition. If the symbol is needed, and there is a non-weak definition, this definition will be used. If there is no non-weak definition, one of the weak definitions will be used.

Example

```
extern __weak int foo; /* A weak reference. */

__weak void bar(void) /* A weak definition. */
{
    /* Increment foo if it was included. */
    if (&foo != 0)
        ++foo;
}
```


Pragma directives

- Summary of pragma directives
- Descriptions of pragma directives

Summary of pragma directives

The `#pragma` directive is defined by Standard C and is a mechanism for using vendor-specific extensions in a controlled way to make sure that the source code is still portable.

The pragma directives control the behavior of the compiler, for example how it allocates memory for variables and functions, whether it allows extended keywords, and whether it outputs warning messages.

The pragma directives are always enabled in the compiler.

This table lists the pragma directives of the compiler that can be used either with the `#pragma` preprocessor directive or the `_Pragma()` preprocessor operator:

Pragma directive	Description
<code>bitfields</code>	Controls the order of bitfield members.
<code>calls</code>	Lists possible called functions for indirect calls.
<code>call_graph_root</code>	Specifies that the function is a call graph root.
<code>data_alignment</code>	Gives a variable a higher (more strict) alignment.
<code>default_function_attributes</code>	Sets default type and object attributes for declarations and definitions of functions.
<code>default_no_bounds</code>	Applies <code>#pragma no_bounds</code> to a whole set of functions. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>default_variable_attributes</code>	Sets default type and object attributes for declarations and definitions of variables.
<code>define_with_bounds</code>	Instruments a function to track pointer bounds. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .

Table 33: Pragma directives summary

Pragma directive	Description
<code>define_without_bounds</code>	Defines the version of a function that does not have extra bounds information. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>diag_default</code>	Changes the severity level of diagnostic messages.
<code>diag_error</code>	Changes the severity level of diagnostic messages.
<code>diag_remark</code>	Changes the severity level of diagnostic messages.
<code>diag_suppress</code>	Suppresses diagnostic messages.
<code>diag_warning</code>	Changes the severity level of diagnostic messages.
<code>disable_check</code>	Specifies that the immediately following function does not check accesses against bounds. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>error</code>	Signals an error while parsing.
<code>generate_entry_without_bounds</code>	Enables generation of an extra entry without bounds for the immediately following function. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>include_alias</code>	Specifies an alias for an include file.
<code>inline</code>	Controls inlining of a function.
<code>language</code>	Controls the IAR Systems language extensions.
<code>location</code>	Specifies the absolute address of a variable, places a variable in a register, or places groups of functions or variables in named sections.
<code>message</code>	Prints a message.
<code>no_bounds</code>	Specifies that the immediately following function is not instrumented for bounds checking. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>object_attribute</code>	Adds object attributes to the declaration or definition of a variable or function.
<code>optimize</code>	Specifies the type and level of an optimization.
<code>pack</code>	Specifies the alignment of structures and union members.
<code>__printf_args</code>	Verifies that a function with a printf-style format string is called with the correct arguments.

Table 33: Pragma directives summary (Continued)

Pragma directive	Description
public_equ	Defines a public assembler label and gives it a value.
required	Ensures that a symbol that is needed by another symbol is included in the linked output.
rtmodel	Adds a runtime model attribute to the module.
__scanf_args	Verifies that a function with a scanf-style format string is called with the correct arguments.
section	Declares a section name to be used by intrinsic functions.
segment	This directive is an alias for <code>#pragma section</code> .
STDC CX_LIMITED_RANGE	Specifies whether the compiler can use normal complex mathematical formulas or not.
STDC FENV_ACCESS	Specifies whether your source code accesses the floating-point environment or not.
STDC FP_CONTRACT	Specifies whether the compiler is allowed to contract floating-point expressions or not.
swi_number	Sets the interrupt number of a software interrupt.
unroll	Unrolls loops.
vectorize	Enables or disables generation of NEON vector instructions for a loop.
weak	Makes a definition a weak definition, or creates a weak alias for a function or a variable.
type_attribute	Adds type attributes to a declaration or to definitions.

Table 33: Pragma directives summary (Continued)

Note: For portability reasons, see also *Recognized pragma directives (6.10.6)*, page 522.

Descriptions of pragma directives

This section gives detailed information about each pragma directive.

bitfields

Syntax

```
#pragma bitfields=disjoint_types|joined_types|
                    reversed_disjoint_types|reversed|default}
```

Parameters	<code>disjoint_types</code>	Bitfield members are placed from the least significant bit to the most significant bit in the container type. Storage containers of bitfields with different base types will not overlap.
	<code>joined_types</code>	Bitfield members are placed depending on the byte order. Storage containers of bitfields will overlap other structure members. For more information, see <i>Bitfields</i> , page 332.
	<code>reversed_disjoint_types</code>	Bitfield members are placed from the most significant bit to the least significant bit in the container type. Storage containers of bitfields with different base types will not overlap.
	<code>reversed</code>	This is an alias for <code>reversed_disjoint_types</code> .
	<code>default</code>	Restores the default layout of bitfield members. The default behavior for the compiler is <code>joined_types</code> .

Description Use this pragma directive to control the layout of bitfield members.

Example

```
#pragma bitfields=disjoint_types
/* Structure that uses disjoint bitfield types. */
struct S
{
    unsigned char  error : 1;
    unsigned char  size  : 4;
    unsigned short code  : 10;
};
#pragma bitfields=default /* Restores to default setting. */
```

See also *Bitfields*, page 332.

calls

Syntax `#pragma calls=function[, function...]`

Parameters

<code>function</code>	Any declared function
-----------------------	-----------------------

Description Use this pragma directive to list the functions that can be indirectly called in the following statement. This information can be used for stack usage analysis in the linker.

Note: For an accurate result, you must list all possible called functions.

Example

```
void Fun1(), Fun2();

void Caller(void (*fp)(void))
{
    #pragma calls = Fun1, Fun2
    (*fp)();
}
```

See also

Stack usage analysis, page 92

call_graph_root

Syntax

```
#pragma call_graph_root[=category]
```

Parameters

category A string that identifies an optional call graph root category

Description

Use this pragma directive to specify that, for stack usage analysis purposes, the immediately following function is a call graph root. You can also specify an optional category. The compiler will usually automatically assign a call graph root category to interrupt and task functions. If you use the `#pragma call_graph_root` directive on such a function you will override the default category. You can specify any string as a category.

Example

```
#pragma call_graph_root="interrupt"
```

See also

Stack usage analysis, page 92

data_alignment

Syntax

```
#pragma data_alignment=expression
```

Parameters

expression A constant which must be a power of two (1, 2, 4, etc.).

Description

Use this pragma directive to give a variable a higher (more strict) alignment of the start address than it would otherwise have. This directive can be used on variables with static and automatic storage duration.

When you use this directive on variables with automatic storage duration, there is an upper limit on the allowed alignment for each function, determined by the calling convention used.

Note: Normally, the size of a variable is a multiple of its alignment. The `data_alignment` directive only affects the alignment of the variable's start address, and not its size, and can thus be used for creating situations where the size is not a multiple of the alignment.

default_function_attributes

Syntax	<pre>#pragma default_function_attributes=[attribute...]</pre> <p>where <i>attribute</i> can be:</p> <pre>type_attribute object_attribute @ section_name</pre>	
Parameters	<pre>type_attribute</pre> <p>See <i>Type attributes</i>, page 343.</p>	<pre>object_attribute</pre> <p>See <i>Object attributes</i>, page 345.</p>
	<pre>@ section_name</pre> <p>See <i>Data and function placement in sections</i>, page 226.</p>	
Description	<p>Use this pragma directive to set default section placement, type attributes, and object attributes for function declarations and definitions. The default settings are only used for declarations and definitions that do not specify type or object attributes or location in some other way.</p> <p>Specifying a <code>default_function_attributes</code> pragma directive with no attributes, restores the initial state where no such defaults have been applied to function declarations and definitions.</p>	
Example	<pre>/* Place following functions in section MYSEC */ #pragma default_function_attributes = @ "MYSEC" int fun1(int x) { return x + 1; } int fun2(int x) { return x - 1; } /* Stop placing functions into MYSEC */ #pragma default_function_attributes =</pre> <p>has the same effect as:</p> <pre>int fun1(int x) @ "MYSEC" { return x + 1; } int fun2(int x) @ "MYSEC" { return x - 1; }</pre>	

See also *location*, page 370
object_attribute, page 371
type_attribute, page 378

default_variable_attributes

Syntax	<pre>#pragma default_variable_attributes=[attribute...]</pre> <p>where <i>attribute</i> can be:</p> <pre>type_attribute object_attribute @ section_name</pre>	
Parameters	<pre>type_attribute</pre> <p>See <i>Type attributes</i>, page 343.</p> <pre>object_attributes</pre> <p>See <i>Object attributes</i>, page 345.</p> <pre>@ section_name</pre> <p>See <i>Data and function placement in sections</i>, page 226.</p>	
Description	<p>Use this pragma directive to set default section placement, type attributes, and object attributes for declarations and definitions of variables with static storage duration. The default settings are only used for declarations and definitions that do not specify type or object attributes or location in some other way.</p> <p>Specifying a <code>default_variable_attributes</code> pragma with no attributes restores the initial state of no such defaults being applied to variables with static storage duration.</p>	
Example	<pre>/* Place following variables in section MYSEC */ #pragma default_variable_attributes = @ "MYSEC" int var1 = 42; int var2 = 17; /* Stop placing variables into MYSEC */ #pragma default_variable_attributes =</pre> <p>has the same effect as:</p> <pre>int var1 @ "MYSEC" = 42; int var2 @ "MYSEC" = 17;</pre>	
See also	<p><i>location</i>, page 370 <i>object_attribute</i>, page 371 <i>type_attribute</i>, page 378</p>	

diag_default

Syntax	#pragma diag_default=tag[, tag, ...]	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number Pe177.
Description	Use this pragma directive to change the severity level back to the default, or to the severity level defined on the command line by any of the options --diag_error, --diag_remark, --diag_suppress, or --diag_warnings, for the diagnostic messages specified with the tags.	
See also	<i>Diagnostics</i> , page 248.	

diag_error

Syntax	#pragma diag_error=tag[, tag, ...]	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number Pe177.
Description	Use this pragma directive to change the severity level to <code>error</code> for the specified diagnostics.	
See also	<i>Diagnostics</i> , page 248.	

diag_remark

Syntax	#pragma diag_remark=tag[, tag, ...]	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number Pe177.
Description	Use this pragma directive to change the severity level to <code>remark</code> for the specified diagnostic messages.	
See also	<i>Diagnostics</i> , page 248.	

diag_suppress

Syntax	<code>#pragma diag_suppress=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe117</code> .
Description	Use this pragma directive to suppress the specified diagnostic messages.	
See also	<i>Diagnostics</i> , page 248.	

diag_warning

Syntax	<code>#pragma diag_warning=tag[, tag, ...]</code>	
Parameters	<i>tag</i>	The number of a diagnostic message, for example the message number <code>Pe826</code> .
Description	Use this pragma directive to change the severity level to <code>warning</code> for the specified diagnostic messages.	
See also	<i>Diagnostics</i> , page 248.	

error

Syntax	<code>#pragma error message</code>	
Parameters	<i>message</i>	A string that represents the error message.
Description	Use this pragma directive to cause an error message when it is parsed. This mechanism is different from the preprocessor directive <code>#error</code> , because the <code>#pragma error</code> directive can be included in a preprocessor macro using the <code>_Pragma</code> form of the directive and only causes an error if the macro is used.	

Example	<pre>#if FOO_AVAILABLE #define FOO ... #else #define FOO _Pragma("error\"Foo is not available\"") #endif</pre> <p>If <code>FOO_AVAILABLE</code> is zero, an error will be signaled if the <code>FOO</code> macro is used in actual source code.</p>
---------	---

include_alias

Syntax	<pre>#pragma include_alias ("orig_header" , "subst_header") #pragma include_alias (<orig_header> , <subst_header>)</pre>				
Parameters	<table><tr><td><i>orig_header</i></td><td>The name of a header file for which you want to create an alias.</td></tr><tr><td><i>subst_header</i></td><td>The alias for the original header file.</td></tr></table>	<i>orig_header</i>	The name of a header file for which you want to create an alias.	<i>subst_header</i>	The alias for the original header file.
<i>orig_header</i>	The name of a header file for which you want to create an alias.				
<i>subst_header</i>	The alias for the original header file.				
Description	<p>Use this pragma directive to provide an alias for a header file. This is useful for substituting one header file with another, and for specifying an absolute path to a relative file.</p> <p>This pragma directive must appear before the corresponding <code>#include</code> directives and <code>subst_header</code> must match its corresponding <code>#include</code> directive exactly.</p>				
Example	<pre>#pragma include_alias (<stdio.h> , <C:\MyHeaders\stdio.h>) #include <stdio.h></pre> <p>This example will substitute the relative file <code>stdio.h</code> with a counterpart located according to the specified path.</p>				
See also	<i>Include file search procedure</i> , page 245.				

inline

Syntax	<pre>#pragma inline[=forced =never]</pre>				
Parameters	<table><tr><td>No parameter</td><td>Has the same effect as the <code>inline</code> keyword.</td></tr><tr><td><code>forced</code></td><td>Disables the compiler's heuristics and forces inlining.</td></tr></table>	No parameter	Has the same effect as the <code>inline</code> keyword.	<code>forced</code>	Disables the compiler's heuristics and forces inlining.
No parameter	Has the same effect as the <code>inline</code> keyword.				
<code>forced</code>	Disables the compiler's heuristics and forces inlining.				

	<code>never</code>	Disables the compiler's heuristics and makes sure that the function will not be inlined.
Description	<p>Use <code>#pragma inline</code> to advise the compiler that the function defined immediately after the directive should be inlined according to C++ inline semantics.</p> <p>Specifying <code>#pragma inline=forced</code> will always inline the defined function. If the compiler fails to inline the function for some reason, for example due to recursion, a warning message is emitted.</p> <p>Inlining is normally performed only on the High optimization level. Specifying <code>#pragma inline=forced</code> will enable inlining of the function also on the Medium optimization level.</p>	
See also	<i>Inlining functions</i> , page 79.	

language

Syntax	#pragma language={extended default save restore}	
Parameters	extended	Enables the IAR Systems language extensions from the first use of the pragma directive and onward.
	default	From the first use of the pragma directive and onward, restores the settings for the IAR Systems language extensions to whatever that was specified by compiler options.
	save restore	Saves and restores, respectively, the IAR Systems language extensions setting around a piece of source code. Each use of save must be followed by a matching restore in the same file without any intervening #include directive.
Description	Use this pragma directive to control the use of language extensions.	
Example	At the top of a file that needs to be compiled with IAR Systems extensions enabled: <pre>#pragma language=extended /* The rest of the file. */</pre>	

Around a particular part of the source code that needs to be compiled with IAR Systems extensions enabled, but where the state before the sequence cannot be assumed to be the same as that specified by the compiler options in use:

```
#pragma language=save
#pragma language=extended
/* Part of source code. */
#pragma language=restore
```

See also `-e`, page 269 and `--strict`, page 293.

location

Syntax `#pragma location={address|register|NAME}`

Parameters

<i>address</i>	The absolute address of the global or static variable or function for which you want an absolute location.
<i>register</i>	An identifier that corresponds to one of the ARM core registers R4–R11.
<i>NAME</i>	A user-defined section name; cannot be a section name predefined for use by the compiler and linker.

Description

Use this pragma directive to specify:

- The location—the absolute address—of the global or static variable whose declaration follows the pragma directive. The variables must be declared `__no_init`.
- An identifier specifying a register. The variable defined after the pragma directive is placed in the register. The variable must be declared as `__no_init` and have file scope.

A string specifying a section for placing either a variable or function whose declaration follows the pragma directive. Do not place variables that would normally be in different sections (for example, variables declared as `__no_init` and variables declared as `const`) in the same named section.

Example

```
#pragma location=0xFFFF0400
__no_init volatile char PORT1; /* PORT1 is located at address
                                0xFFFF0400 */

#pragma location=R8
__no_init int TASK; /* TASK is placed in R8 */

#pragma location="FLASH"
char PORT2; /* PORT2 is located in section FLASH */

/* A better way is to use a corresponding mechanism */
#define FLASH _Pragma("location=\"FLASH\"")
/* ... */
FLASH int i; /* i is placed in the FLASH section */
```

See also

Controlling data and function placement in memory, page 224 and *Declare and place your own sections*, page 104.

message**Syntax**

```
#pragma message(message)
```

Parameters

message The message that you want to direct to the standard output stream.

Description

Use this pragma directive to make the compiler print a message to the standard output stream when the file is compiled.

Example

```
#ifdef TESTING
#pragma message("Testing")
#endif
```

object_attribute**Syntax**

```
#pragma object_attribute=object_attribute[ object_attribute...]
```

Parameters

For information about object attributes that can be used with this pragma directive, see *Object attributes*, page 345.

Description

Use this pragma directive to add one or more IAR-specific object attributes to the declaration or definition of a variable or function. Object attributes affect the actual

variable or function and not its type. When you define a variable or function, the union of the object attributes from all declarations including the definition, is used.

Example

```
#pragma object_attribute=__no_init
char bar;
```

is equivalent to:

```
__no_init char bar;
```

See also

General syntax rules for extended keywords, page 343.

optimize

Syntax

```
#pragma optimize=[goal] [level] [no_optimization...]
```

Parameters

goal	Choose between: size, optimizes for size balanced, optimizes balanced between speed and size speed, optimizes for speed. no_size_constraints, optimizes for speed, but relaxes the normal restrictions for code size expansion.
level	Specifies the level of optimization; choose between none, low, medium, or high.
no_optimization	Disables one or several optimizations; choose between: no_code_motion, disables code motion no_cse, disables common subexpression elimination no_inline, disables function inlining no_tbaa, disables type-based alias analysis no_unroll, disables loop unrolling no_vectorize, disables generation of NEON vector instructions no_scheduling, disables instruction scheduling. vectorize, enables generation of NEON vector instructions

Description	<p>Use this pragma directive to decrease the optimization level, or to turn off some specific optimizations. This pragma directive only affects the function that follows immediately after the directive.</p> <p>The parameters <code>size</code>, <code>balanced</code>, <code>speed</code>, and <code>no_size_constraints</code> only have effect on the <code>high</code> optimization level and only one of them can be used as it is not possible to optimize for speed and size at the same time. It is also not possible to use preprocessor macros embedded in this pragma directive. Any such macro will not be expanded by the preprocessor.</p> <p>Note: If you use the <code>#pragma optimize</code> directive to specify an optimization level that is higher than the optimization level you specify using a compiler option, the pragma directive is ignored.</p>
Example	<pre>#pragma optimize=speed int SmallAndUsedOften() { /* Do something here. */ } #pragma optimize=size int BigAndSeldomUsed() { /* Do something here. */ }</pre>
See also	<i>Fine-tuning enabled transformations</i> , page 231.

pack

Syntax	<pre>#pragma pack(<i>n</i>) #pragma pack() #pragma pack({push pop}[<i>, name</i>] [<i>, n</i>])</pre>										
Parameters	<table> <tr> <td><i>n</i></td><td>Sets an optional structure alignment; one of: 1, 2, 4, 8, or 16</td></tr> <tr> <td>Empty list</td><td>Restores the structure alignment to default</td></tr> <tr> <td><i>push</i></td><td>Sets a temporary structure alignment</td></tr> <tr> <td><i>pop</i></td><td>Restores the structure alignment from a temporarily pushed alignment</td></tr> <tr> <td><i>name</i></td><td>An optional pushed or popped alignment label</td></tr> </table>	<i>n</i>	Sets an optional structure alignment; one of: 1, 2, 4, 8, or 16	Empty list	Restores the structure alignment to default	<i>push</i>	Sets a temporary structure alignment	<i>pop</i>	Restores the structure alignment from a temporarily pushed alignment	<i>name</i>	An optional pushed or popped alignment label
<i>n</i>	Sets an optional structure alignment; one of: 1, 2, 4, 8, or 16										
Empty list	Restores the structure alignment to default										
<i>push</i>	Sets a temporary structure alignment										
<i>pop</i>	Restores the structure alignment from a temporarily pushed alignment										
<i>name</i>	An optional pushed or popped alignment label										

Description	<p>Use this pragma directive to specify the maximum alignment of <code>struct</code> and <code>union</code> members.</p> <p>The <code>#pragma pack</code> directive affects declarations of structures following the pragma directive to the next <code>#pragma pack</code> or the end of the compilation unit.</p> <p>Note: This can result in significantly larger and slower code when accessing members of the structure.</p> <p>Use either <code>__packed</code> or <code>#pragma pack</code> to relax the alignment restrictions for a type and the objects defined using that type. Mixing <code>__packed</code> and <code>#pragma pack</code> might lead to unexpected behavior.</p>
See also	<i>Structure types</i> , page 339.

__printf_args

Syntax	<code>#pragma __printf_args</code>
Description	Use this pragma directive on a function with a <code>printf</code> -style format string. For any call to that function, the compiler verifies that the argument to each conversion specifier (for example <code>%d</code>) is syntactically correct.
Example	<pre>#pragma __printf_args int printf(char const *,...); void PrintNumbers(unsigned short x) { printf("%d", x); /* Compiler checks that x is an integer */ }</pre>

public_equ

Syntax	<code>#pragma public_equ= "symbol", value</code>				
Parameters	<table><tr><td><i>symbol</i></td><td>The name of the assembler symbol to be defined (string).</td></tr><tr><td><i>value</i></td><td>The value of the defined assembler symbol (integer constant expression).</td></tr></table>	<i>symbol</i>	The name of the assembler symbol to be defined (string).	<i>value</i>	The value of the defined assembler symbol (integer constant expression).
<i>symbol</i>	The name of the assembler symbol to be defined (string).				
<i>value</i>	The value of the defined assembler symbol (integer constant expression).				
Description	Use this pragma directive to define a public assembler label and give it a value.				
Example	<code>#pragma public_equ= "MY_SYMBOL", 0x123456</code>				

See also `--public_equ`, page 289.

required

Syntax	<code>#pragma required=symbol</code>
Parameters	<p><i>symbol</i> Any statically linked function or variable.</p>
Description	<p>Use this pragma directive to ensure that a symbol which is needed by a second symbol is included in the linked output. The directive must be placed immediately before the second symbol.</p> <p>Use the directive if the requirement for a symbol is not otherwise visible in the application, for example if a variable is only referenced indirectly through the section it resides in.</p>
Example	<pre>const char copyright[] = "Copyright by me"; #pragma required=copyright int main() { /* Do something here. */ }</pre> <p>Even if the copyright string is not used by the application, it will still be included by the linker and available in the output.</p>

rtmodel

Syntax	<code>#pragma rtmodel="key", "value"</code>
Parameters	<p><i>"key"</i> A text string that specifies the runtime model attribute.</p> <p><i>"value"</i> A text string that specifies the value of the runtime model attribute. Using the special value <i>*</i> is equivalent to not defining the attribute at all.</p>
Description	<p>Use this pragma directive to add a runtime model attribute to a module, which can be used by the linker to check consistency between modules.</p> <p>This pragma directive is useful for enforcing consistency between modules. All modules that are linked together and define the same runtime attribute key must have the same</p>

value for the corresponding key, or the special value *. It can, however, be useful to state explicitly that the module can handle any runtime model.

A module can have several runtime model definitions.

Note: The predefined compiler runtime model attributes start with a double underscore. To avoid confusion, this style must not be used in the user-defined attributes.

Example

```
#pragma rtmodel="I2C", "ENABLED"
```

The linker will generate an error if a module that contains this definition is linked with a module that does not have the corresponding runtime model attributes defined.

See also

Checking module consistency, page 113.

__scanf_args

Syntax

```
#pragma __scanf_args
```

Description

Use this pragma directive on a function with a scanf-style format string. For any call to that function, the compiler verifies that the argument to each conversion specifier (for example %d) is syntactically correct.

Example

```
#pragma __scanf_args
int scanf(char const *,...);

int GetNumber()
{
    int nr;
    scanf("%d", &nr); /* Compiler checks that
                       the argument is a
                       pointer to an integer */

    return nr;
}
```

section

Syntax

```
#pragma section="NAME"

alias

#pragma segment="NAME"
```


Parameters	<i>NAME</i>	The name of the section.
Description	<p>Use this pragma directive to define a section name that can be used by the section operators <code>__section_begin</code>, <code>__section_end</code>, and <code>__section_size</code>. All section declarations for a specific section must have the same memory type attribute and alignment.</p> <p>Note: To place variables or functions in a specific section, use the <code>#pragma location</code> directive or the <code>@</code> operator.</p>	
Example	<code>#pragma section="MYSECTION"</code>	
See also	<i>Dedicated section operators</i> , page 185 and the chapter <i>Linking your application</i> .	

STDC CX_LIMITED_RANGE

Syntax	<code>#pragma STDC CX_LIMITED_RANGE {ON OFF DEFAULT}</code>	
Parameters	ON	Normal complex mathematic formulas can be used.
	OFF	Normal complex mathematic formulas cannot be used.
	DEFAULT	Sets the default behavior, that is OFF.
Description	<p>Use this pragma directive to specify that the compiler can use the normal complex mathematic formulas for <code>*</code> (multiplication), <code>/</code> (division), and <code>abs</code>.</p> <p>Note: This directive is required by Standard C. The directive is recognized but has no effect in the compiler.</p>	

STDC FENV_ACCESS

Syntax	<code>#pragma STDC FENV_ACCESS {ON OFF DEFAULT}</code>	
Parameters	ON	Source code accesses the floating-point environment. Note that this argument is not supported by the compiler.
	OFF	Source code does not access the floating-point environment.
	DEFAULT	Sets the default behavior, that is OFF.

Description Use this pragma directive to specify whether your source code accesses the floating-point environment or not.

Note: This directive is required by Standard C.

STDC FP_CONTRACT

Syntax #pragma STDC FP_CONTRACT {ON|OFF|DEFAULT}

Parameters

ON	The compiler is allowed to contract floating-point expressions.
OFF	The compiler is not allowed to contract floating-point expressions. Note that this argument is not supported by the compiler.
DEFAULT	Sets the default behavior, that is ON.

Description Use this pragma directive to specify whether the compiler is allowed to contract floating-point expressions or not. This directive is required by Standard C.

Example #pragma STDC FP_CONTRACT=ON

swi_number

Syntax #pragma swi_number=number

Parameters

number	The software interrupt number
--------	-------------------------------

Description Use this pragma directive together with the __swi extended keyword. It is used as an argument to the generated SVC assembler instruction, and is used for selecting one software interrupt function in a system containing several such functions.

Example #pragma swi_number=17

See also *Software interrupts*, page 78.

type_attribute

Syntax #pragma type_attribute=type_attr[type_attr...]

Parameters	For information about type attributes that can be used with this pragma directive, see <i>Type attributes</i> , page 343.
Description	<p>Use this pragma directive to specify IAR-specific <i>type attributes</i>, which are not part of Standard C. Note however, that a given type attribute might not be applicable to all kind of objects.</p> <p>This directive affects the declaration of the identifier, the next variable, or the next function that follows immediately after the pragma directive.</p>
Example	<p>In this example, thumb-mode code is generated for the function <code>foo</code>:</p> <pre>#pragma type_attribute=__thumb void foo(void) { }</pre> <p>This declaration, which uses extended keywords, is equivalent:</p> <pre>__thumb void foo(void) { }</pre>
See also	The chapter <i>Extended keywords</i> .

unroll

Syntax	<code>#pragma unroll=<i>n</i></code>
Parameters	<p><i>n</i></p> <p>The number of loop bodies in the unrolled loop, a constant integer. <code>#pragma unroll = 1</code> will prevent the unrolling of a loop.</p>
Description	<p>Use this pragma directive to specify that the loop following immediately after the directive should be unrolled and that the unrolled loop should have <i>n</i> copies of the loop body. The pragma directive can only be placed immediately before a <code>for</code>, <code>do</code>, or <code>while</code> loop, whose number of iterations can be determined at compile time.</p> <p>Normally, unrolling is most effective for relatively small loops. However, in some cases, unrolling larger loops can be beneficial if it exposes opportunities for further optimizations between the unrolled loop iterations, for example common subexpression elimination or dead code elimination.</p>

The `#pragma unroll` directive can be used to force a loop to be unrolled if the unrolling heuristics are not aggressive enough. The pragma directive can also be used to reduce the aggressiveness of the unrolling heuristics.

Example

```
#pragma unroll=4
for (i = 0; i < 64; ++i)
{
    foo(i * k; (i + 1) * k);
}
```

See also

Loop unrolling, page 232

vectorize

Syntax

```
#pragma vectorize [= never]
```

Parameters	No parameter	Enables generation of NEON vector instructions.
	<code>never</code>	Disables generation of NEON vector instructions.

Description

Use this pragma directive to enable or disable generation of NEON vector instructions for the loop that follows immediately after the pragma directive. This pragma directive can only be placed immediately before a `for`, `do`, or `while` loop. If the optimization level is lower than High, the pragma directive has no effect.

Example

```
#pragma vectorize
for (i = 0; i < 1024; ++i)
{
    a[i] = b[i] * c[i];
}
```

weak

Syntax

```
#pragma weak symbol1 [= symbol2]
```

Parameters	<i>symbol1</i>	A function or variable with external linkage.
	<i>symbol2</i>	A defined function or variable.

Description	<p>This pragma directive can be used in one of two ways:</p> <ul style="list-style-type: none">● To make the definition of a function or variable with external linkage a weak definition. The <code>__weak</code> attribute can also be used for this purpose.● To create a weak alias for another function or variable. You can make more than one alias for the same function or variable.
Example	<p>To make the definition of <code>foo</code> a weak definition, write:</p> <pre>#pragma weak foo</pre> <p>To make <code>NMI_Handler</code> a weak alias for <code>Default_Handler</code>, write:</p> <pre>#pragma weak NMI_Handler=Default_Handler</pre> <p>If <code>NMI_Handler</code> is not defined elsewhere in the program, all references to <code>NMI_Handler</code> will refer to <code>Default_Handler</code>.</p>
See also	<p><code>__weak</code>, page 356.</p>

Intrinsic functions

- Summary of intrinsic functions
- Descriptions of intrinsic functions

Summary of intrinsic functions

To use intrinsic functions in an application, include the header file `intrinsics.h`.

To use Neon intrinsic functions in an application, include the header file `arm_neon.h`. For more information, see *Intrinsic functions for Neon instructions*, page 389.

Note that the intrinsic function names start with double underscores, for example:

`__disable_interrupt`

This table summarizes the intrinsic functions:

Intrinsic function	Description
<code>__as_get_base</code>	Creates a pointer of the same type as the parameter, representing the base of the area pointed to by the parameter. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>__as_get_bounds</code>	Creates a pointer of the same type as the parameter, representing the upper bound of the area pointed to by the parameter. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>__as_make_bounds</code>	Creates a pointer with bounds information. See the C-RUN documentation in the <i>C-SPY® Debugging Guide for ARM</i> .
<code>__CLREX</code>	Inserts a CLREX instruction
<code>__CLZ</code>	Inserts a CLZ instruction
<code>__disable_fiq</code>	Disables fast interrupt requests (fiq)
<code>__disable_interrupt</code>	Disables interrupts
<code>__disable_irq</code>	Disables interrupt requests (irq)
<code>__DMB</code>	Inserts a DMB instruction
<code>__DSB</code>	Inserts a DSB instruction
<code>__enable_fiq</code>	Enables fast interrupt requests (fiq)
<code>__enable_interrupt</code>	Enables interrupts

Table 34: Intrinsic functions summary

Intrinsic function	Description
<code>__enable_irq</code>	Enables interrupt requests (irq)
<code>__get_BASEPRI</code>	Returns the value of the Cortex-M3/Cortex-M4/Cortex-M7 <code>BASEPRI</code> register
<code>__get_CONTROL</code>	Returns the value of the Cortex-M <code>CONTROL</code> register
<code>__get_CPSR</code>	Returns the value of the ARM <code>CPSR</code> (Current Program Status Register)
<code>__get_FAULTMASK</code>	Returns the value of the Cortex-M3/Cortex-M4/Cortex-M7 <code>FAULTMASK</code> register
<code>__get_FPSCR</code>	Returns the value of <code>FPSCR</code>
<code>__get_interrupt_state</code>	Returns the interrupt state
<code>__get_IPSR</code>	Returns the value of the <code>IPSR</code> register
<code>__get_LR</code>	Returns the value of the link register
<code>__get_MSP</code>	Returns the value of the <code>MSP</code> register
<code>__get_PRIMASK</code>	Returns the value of the Cortex-M <code>PRIMASK</code> register
<code>__get_PSP</code>	Returns the value of the <code>PSP</code> register
<code>__get_PSR</code>	Returns the value of the <code>PSR</code> register
<code>__get_SB</code>	Returns the value of the static base register
<code>__get_SP</code>	Returns the value of the stack pointer register
<code>__ISB</code>	Inserts an <code>ISB</code> instruction
<code>__LDC</code>	Inserts the coprocessor load instruction <code>LDC</code>
<code>__LDCL</code>	Inserts the coprocessor load instruction <code>LDCL</code>
<code>__LDC2</code>	Inserts the coprocessor load instruction <code>LDC2</code>
<code>__LDC2L</code>	Inserts the coprocessor load instruction <code>LDC2L</code>
<code>__LDC_noidx</code>	Inserts the coprocessor load instruction <code>LDC</code>
<code>__LDCL_noidx</code>	Inserts the coprocessor load instruction <code>LDCL</code>
<code>__LDC2_noidx</code>	Inserts the coprocessor load instruction <code>LDC2</code>
<code>__LDC2L_noidx</code>	Inserts the coprocessor load instruction <code>LDC2L</code>
<code>__LDREX</code>	Inserts an <code>LDREX</code> instruction
<code>__LDREXB</code>	Inserts an <code>LDREXB</code> instruction
<code>__LDREXD</code>	Inserts an <code>LDREXD</code> instruction
<code>__LDREXH</code>	Inserts an <code>LDREXH</code> instruction
<code>__MCR</code>	Inserts the coprocessor write instruction <code>MCR</code>
<code>__MCR2</code>	Inserts the coprocessor write instruction <code>MCR2</code>

Table 34: Intrinsic functions summary (Continued)

Intrinsic function	Description
<code>__MRC</code>	Inserts the coprocessor read instruction <code>MRC</code>
<code>__MRC2</code>	Inserts the coprocessor read instruction <code>MRC2</code>
<code>__no_operation</code>	Inserts a <code>NOP</code> instruction
<code>__PKHBT</code>	Inserts a <code>PKHBT</code> instruction
<code>__PKHTB</code>	Inserts a <code>PKHTB</code> instruction
<code>__PLD</code>	Inserts the preload data instruction <code>PLD</code>
<code>__PLDW</code>	Inserts the preload data instruction <code>PLDW</code>
<code>__PLI</code>	Inserts a <code>PLI</code> instruction
<code>__QADD</code>	Inserts a <code>QADD</code> instruction
<code>__QADD8</code>	Inserts a <code>QADD8</code> instruction
<code>__QADD16</code>	Inserts a <code>QADD16</code> instruction
<code>__QASX</code>	Inserts a <code>QASX</code> instruction
<code>__QCFflag</code>	Returns the value of the cumulative saturation flag of the <code>FPSCR</code> register
<code>__QDADD</code>	Inserts a <code>QDADD</code> instruction
<code>__QDOUBLE</code>	Inserts a <code>QDOUBLE</code> instruction
<code>__QDSUB</code>	Inserts a <code>QDSUB</code> instruction
<code>__QFlag</code>	Returns the <code>Q</code> flag that indicates if overflow/saturation has occurred
<code>__QSAX</code>	Inserts a <code>QSAX</code> instruction
<code>__QSUB</code>	Inserts a <code>QSUB</code> instruction
<code>__QSUB8</code>	Inserts a <code>QSUB8</code> instruction
<code>__QSUB16</code>	Inserts a <code>QSUB16</code> instruction
<code>__RBIT</code>	Inserts an <code>RBIT</code> instruction
<code>__reset_Q_flag</code>	Clears the <code>Q</code> flag that indicates if overflow/saturation has occurred
<code>__reset_QC_flag</code>	Clears the value of the cumulative saturation flag <code>QC</code> of the <code>FPSCR</code> register
<code>__REV</code>	Inserts an <code>REV</code> instruction
<code>__REV16</code>	Inserts an <code>REV16</code> instruction
<code>__REVSH</code>	Inserts an <code>REVSH</code> instruction
<code>__SADD8</code>	Inserts an <code>SADD8</code> instruction
<code>__SADD16</code>	Inserts an <code>SADD16</code> instruction

Table 34: Intrinsic functions summary (Continued)

Intrinsic function	Description
__SASX	Inserts an SASX instruction
__SEL	Inserts an SEL instruction
__set_BASEPRI	Sets the value of the Cortex-M3/Cortex-M4/Cortex-M7 BASEPRI register
__set_CONTROL	Sets the value of the Cortex-M CONTROL register
__set_CPSR	Sets the value of the ARM CPSR (Current Program Status Register)
__set_FAULTMASK	Sets the value of the Cortex-M3/Cortex-M4/Cortex-M7 FAULTMASK register
__set_FPSCR	Sets the value of the FPSCR register
__set_interrupt_state	Restores the interrupt state
__set_LR	Assigns a new address to the link register
__set_MSP	Sets the value of the MSP register
__set_PRIMASK	Sets the value of the Cortex-M PRIMASK register
__set_PSP	Sets the value of the PSP register
__set_SB	Assigns a new address to the static base register
__set_SP	Assigns a new address to the stack pointer register
__SEV	Inserts an SEV instruction
__SHADD8	Inserts an SHADD8 instruction
__SHADD16	Inserts an SHADD16 instruction
__SHASX	Inserts an SHASX instruction
__SHSAX	Inserts an SHSAX instruction
__SHSUB8	Inserts an SHSUB8 instruction
__SHSUB16	Inserts an SHSUB16 instruction
__SMLABB	Inserts an SMLABB instruction
__SMLABT	Inserts an SMLABT instruction
__SMLAD	Inserts an SMLAD instruction
__SMLADX	Inserts an SMLADX instruction
__SMLALBB	Inserts an SMLALBB instruction
__SMLALBT	Inserts an SMLALBT instruction
__SMLALD	Inserts an SMLALD instruction
__SMLALDX	Inserts an SMLALDX instruction
__SMLALTB	Inserts an SMLALTB instruction

Table 34: Intrinsic functions summary (Continued)

Intrinsic function	Description
__SMLALTT	Inserts an SMLALTT instruction
__SMLATB	Inserts an SMLATB instruction
__SMLATT	Inserts an SMLATT instruction
__SMLAWB	Inserts an SMLAWB instruction
__SMLAWT	Inserts an SMLAWT instruction
__SMLSD	Inserts an SMLSD instruction
__SMLSDX	Inserts an SMLSDX instruction
__SMLSXD	Inserts an SMLSXD instruction
__SMLSXD	Inserts an SMLSXD instruction
__SMMLA	Inserts an SMMLA instruction
__SMMLAR	Inserts an SMMLAR instruction
__SMMLS	Inserts an SMMLS instruction
__SMMLSR	Inserts an SMMLSR instruction
__SMMUL	Inserts an SMMUL instruction
__SMMULR	Inserts an SMMULR instruction
__SMUAD	Inserts an SMUAD instruction
__SMUADX	Inserts an SMUADX instruction
__SMUL	Inserts a signed 16-bit multiplication
__SMULBB	Inserts an SMULBB instruction
__SMULBT	Inserts an SMULBT instruction
__SMULTB	Inserts an SMULTB instruction
__SMULTT	Inserts an SMULTT instruction
__SMULWB	Inserts an SMULWB instruction
__SMULWT	Inserts an SMULWT instruction
__SMUSD	Inserts an SMUSD instruction
__SMUSDX	Inserts an SMUSDX instruction
__SSAT	Inserts an SSAT instruction
__SSAT16	Inserts an SSAT16 instruction
__SSAX	Inserts an SSAX instruction
__SSUB8	Inserts an SSUB8 instruction
__SSUB16	Inserts an SSUB16 instruction
__STC	Inserts the coprocessor store instruction STC

Table 34: Intrinsic functions summary (Continued)

Intrinsic function	Description
__STCL	Inserts the coprocessor store instruction STCL
__STC2	Inserts the coprocessor store instruction STC2
__STC2L	Inserts the coprocessor store instruction STC2L
__STC_noidx	Inserts the coprocessor store instruction STC
__STCL_noidx	Inserts the coprocessor store instruction STCL
__STC2_noidx	Inserts the coprocessor store instruction STC2
__STC2L_noidx	Inserts the coprocessor store instruction STC2L
__STREX	Inserts a STREX instruction
__STREXB	Inserts a STREXB instruction
__STREXD	Inserts a STREXD instruction
__STREXH	Inserts a STREXH instruction
__SWP	Inserts an SWP instruction
__SWPB	Inserts an SWPB instruction
__SXTAB	Inserts an SXTAB instruction
__SXTAB16	Inserts an SXTAB16 instruction
__SXTAH	Inserts an SXTAH instruction
__SXTB16	Inserts an SXTB16 instruction
__UADD8	Inserts a UADD8 instruction
__UADD16	Inserts a UADD16 instruction
__UASX	Inserts a UASX instruction
__UHADD8	Inserts a UHADD8 instruction
__UHADD16	Inserts a UHADD16 instruction
__UHASX	Inserts a UHASX instruction
__UHSAX	Inserts a UHSAX instruction
__UHSUB8	Inserts a UHSUB8 instruction
__UHSUB16	Inserts a UHSUB16 instruction
__UMAAL	Inserts a UMAAL instruction
__UQADD8	Inserts a UQADD8 instruction
__UQADD16	Inserts a UQADD16 instruction
__UQASX	Inserts a UQASX instruction
__UQSAX	Inserts a UQSAX instruction
__UQSUB8	Inserts a UQSUB8 instruction

Table 34: Intrinsic functions summary (Continued)

Intrinsic function	Description
__UQSUB16	Inserts a UQSUB16 instruction
__USAD8	Inserts a USAD8 instruction
__USADA8	Inserts a USAD8A8 instruction
__USAT	Inserts a USAT instruction
__USAT16	Inserts a USAT16 instruction
__USAX	Inserts a USAX instruction
__USUB8	Inserts a USUB8 instruction
__USUB16	Inserts a USUB16 instruction
__UXTAB	Inserts a UXTAB instruction
__UXTAB16	Inserts a UXTAB16 instruction
__UXTAH	Inserts a UXTAH instruction
__UXTB16	Inserts a UXTB16 instruction
__WFE	Inserts a WFE instruction
__WFI	Inserts a WFI instruction
__YIELD	Inserts a YIELD instruction

Table 34: Intrinsic functions summary (Continued)

INTRINSIC FUNCTIONS FOR NEON INSTRUCTIONS

The Neon co-processor implements the Advanced SIMD instruction set extension, as defined by the ARM architecture. To use Neon intrinsic functions in an application, include the header file `arm_neon.h`. The functions use vector types that are named according to this pattern:

`<type><size>x<number_of_lanes>_t`

where:

- *type* is `int`, `unsigned int`, `float`, or `poly`
- *size* is 8, 16, 32, or 64
- *number_of_lanes* is 1, 2, 4, 8, or 16.

The total bit width of a vector type is *size* times *number_of_lanes*, and should fit in a D register (64 bits) or a Q register (128 bits).

For example:

```
__intrinsic float32x2_t vsub_f32(float32x2_t, float32x2_t);
```

The intrinsic function `vsub_f32` inserts a `VSUB.F32` instruction that operates on two 64-bit vectors (D registers), each with two elements (lanes) of 32-bit floating-point type.

Some functions use an array of vector types. As an example, the definition of an array type with four elements of type `float32x2_t` is:

```
typedef struct
{
    float32x2_t val[4];
}
float32x2x4_t;
```

Descriptions of intrinsic functions

This section gives reference information about each intrinsic function.

`__CLREX`

Syntax	<code>void __CLREX(void);</code>
Description	Inserts a <code>CLREX</code> instruction. This intrinsic function requires architecture ARMv6K or ARMv7 for ARM mode, and AVRv7 for Thumb mode.

`__CLZ`

Syntax	<code>unsigned char __CLZ(unsigned long);</code>
Description	Inserts a <code>CLZ</code> instruction. This intrinsic function requires an ARMv5 architecture or higher for ARM mode, and ARMv6T2 or higher for Thumb mode.

`__disable_fiq`

Syntax	<code>void __disable_fiq(void);</code>
Description	Disables fast interrupt requests (fiq). This intrinsic function can only be used in privileged mode and is not available for Cortex-M devices.

__disable_interrupt

Syntax	<code>void __disable_interrupt(void);</code>
Description	Disables interrupts. For Cortex-M devices, it raises the execution priority level to 0 by setting the priority mask bit, <code>PRIMASK</code> . For other devices, it disables interrupt requests (irq) and fast interrupt requests (fiq). This intrinsic function can only be used in privileged mode.

__disable_irq

Syntax	<code>void __disable_irq(void);</code>
Description	Disables interrupt requests (irq). This intrinsic function can only be used in privileged mode and is not available for Cortex-M devices.

__DMB

Syntax	<code>void __DMB(void);</code>
Description	Inserts a <code>DMB</code> instruction. This intrinsic function requires an ARMv6M architecture, or an ARMv7 architecture or higher.

__DSB

Syntax	<code>void __DSB(void);</code>
Description	Inserts a <code>DSB</code> instruction. This intrinsic function requires an ARMv6M architecture, or an ARMv7 architecture or higher.

__enable_fiq

Syntax	<code>void __enable_fiq(void);</code>
Description	Enables fast interrupt requests (fiq). This intrinsic function can only be used in privileged mode, and it is not available for Cortex-M devices.

__enable_interrupt

Syntax	<code>void __enable_interrupt(void);</code>
Description	<p>Enables interrupts. For Cortex-M devices, it resets the execution priority level to default by clearing the priority mask bit, <code>PRIMASK</code>. For other devices, it enables interrupt requests (irq) and fast interrupt requests (fiq).</p> <p>This intrinsic function can only be used in privileged mode.</p>

__enable_irq

Syntax	<code>void __enable_irq(void);</code>
Description	<p>Enables interrupt requests (irq).</p> <p>This intrinsic function can only be used in privileged mode, and it is not available for Cortex-M devices.</p>

__get_BASEPRI

Syntax	<code>unsigned long __get_BASEPRI(void);</code>
Description	<p>Returns the value of the <code>BASEPRI</code> register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M3, Cortex-M4, or Cortex-M7 device.</p>

__get_CONTROL

Syntax	<code>unsigned long __get_CONTROL(void);</code>
Description	<p>Returns the value of the <code>CONTROL</code> register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M device.</p>

__get_CPSR

Syntax	<code>unsigned long __get_CPSR(void);</code>
Description	<p>Returns the value of the ARM <code>CPSR</code> (Current Program Status Register). This intrinsic function can only be used in privileged mode, is not available for Cortex-M devices, and it requires ARM mode.</p>

__get_FAULTMASK

Syntax	<code>unsigned long __get_FAULTMASK(void);</code>
Description	Returns the value of the <code>FAULTMASK</code> register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M3, Cortex-M4, or Cortex-M7 device.

__get_FPSCR

Syntax	<code>unsigned long __get_FPSCR(void);</code>
Description	Returns the value of <code>FPSCR</code> (floating-point status and control register). This intrinsic function is only available for devices with a VFP coprocessor.

__get_interrupt_state

Syntax	<code>__istate_t __get_interrupt_state(void);</code>
Description	Returns the global interrupt state. The return value can be used as an argument to the <code>__set_interrupt_state</code> intrinsic function, which will restore the interrupt state. This intrinsic function can only be used in privileged mode, and cannot be used when using the <code>--aeabi</code> compiler option.

Example	<pre>#include "intrinsics.h" void CriticalFn() { __istate_t s = __get_interrupt_state(); __disable_interrupt(); /* Do something here. */ __set_interrupt_state(s); }</pre>
---------	---

The advantage of using this sequence of code compared to using `__disable_interrupt` and `__enable_interrupt` is that the code in this example will not enable any interrupts disabled before the call of `__get_interrupt_state`.

__get_IPSR

Syntax	<code>unsigned long __get_IPSR(void);</code>
Description	Returns the value of the <code>IPSR</code> register (Interrupt Program Status Register). This intrinsic function can only be used in privileged mode, and is only available for Cortex-M devices.

__get_LR

Syntax	<code>unsigned long __get_LR(void);</code>
Description	Returns the value of the link register (<code>R14</code>).

__get_MSP

Syntax	<code>unsigned long __get_MSP(void);</code>
Description	Returns the value of the <code>MSP</code> register (Main Stack Pointer). This intrinsic function can only be used in privileged mode, and is only available for Cortex-M devices.

__get_PRIMASK

Syntax	<code>unsigned long __get_PRIMASK(void);</code>
Description	Returns the value of the <code>PRIMASK</code> register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M device.

__get_PSP

Syntax	<code>unsigned long __get_PSP(void);</code>
Description	Returns the value of the <code>PSP</code> register (Process Stack Pointer). This intrinsic function can only be used in privileged mode, and is only available for Cortex-M devices.

__get_PSR

Syntax `unsigned long __get_PSR(void);`

Description Returns the value of the `PSR` register (combined Program Status Register). This intrinsic function can only be used in privileged mode, and is only available for Cortex-M devices.

__get_SB

Syntax `unsigned long __get_SB(void);`

Description Returns the value of the static base register (`R9`).

__get_SP

Syntax `unsigned long __get_SP(void);`

Description Returns the value of the stack pointer register (`R13`).

__ISB

Syntax `void __ISB(void);`

Description Inserts an `ISB` instruction. This intrinsic function requires an ARMv6M architecture, or an ARMv7 architecture or higher.

__LDC**__LDCL****__LDC2****__LDC2L**

Syntax `void __LDCxxx(__ul coproc, __ul CRn, __ul const *src);`

Parameters	<div><div><i>coproc</i></div><div><i>CRn</i></div><div><i>src</i></div></div> <div><div>The coprocessor number 0 . . 15.</div><div>The coprocessor register to load.</div><div>A pointer to the data to load.</div></div>
Description	<div><p>Inserts the coprocessor load instruction <code>LDC</code>—or one of its variants—which means that a value will be loaded into a coprocessor register. The parameters <i>coproc</i> and <i>CRn</i> will be encoded in the instruction and must therefore be constants.</p><p>The intrinsic functions <code>__LDC</code> and <code>__LDCL</code> require architecture ARMv4 or higher for ARM mode, and ARMv6T2 or higher for Thumb mode.</p><p>The intrinsic functions <code>__LDC2</code> and <code>__LDC2L</code> require architecture ARMv5 or higher for ARM mode, and ARMv6T2 or higher for Thumb mode.</p></div>
<div><div>__LDC_noidx</div><div>__LDCL_noidx</div><div>__LDC2_noidx</div><div>__LDC2L_noidx</div></div>	
Syntax	<div><div><pre>void __LDCxxx_noidx(__ul coproc, __ul CRn, __ul const *src, __ul option);</pre></div></div>
Parameters	<div><div><i>coproc</i></div><div><i>CRn</i></div><div><i>src</i></div><div><i>option</i></div></div> <div><div>The coprocessor number 0 . . 15.</div><div>The coprocessor register to load.</div><div>A pointer to the data to load.</div><div>Additional coprocessor option 0 . . 255.</div></div>
Description	<div><p>Inserts the coprocessor load instruction <code>LDC</code>, or one of its variants. A value will be loaded into a coprocessor register. The parameters <i>coproc</i>, <i>CRn</i>, and <i>option</i> will be encoded in the instruction and must therefore be constants.</p><p>The intrinsic functions <code>__LDC_noidx</code> and <code>__LDCL_noidx</code> require architecture ARMv4 or higher for ARM mode, and ARMv6T2 or higher for Thumb mode.</p></div>

The intrinsic functions `__LDC2_noidx` and `__LDC2L_noidx` require architecture ARMv5 or higher for ARM mode, and ARMv6T2 or higher for Thumb mode.

`__LDREX`

`__LDREXB`

`__LDREXD`

`__LDREXH`

Syntax

```
unsigned long __LDREX(unsigned long *);
unsigned char __LDREXB(unsigned char *);
unsigned long long __LDREXD(unsigned long long *);
unsigned short __LDREXH(unsigned short *);
```

Description

Inserts the specified instruction.

The `__LDREX` intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv6T2 or higher for Thumb mode.

The `__LDREXB` and the `__LDREXH` intrinsic functions require architecture ARMv6K or ARMv7 for ARM mode, and ARMv7 for Thumb mode.

The `__LDREX` intrinsic function requires architecture ARMv6K or ARMv7 for ARM mode, and ARMv7 but not ARMv7-M for Thumb mode.

`__MCR`

`__MCR2`

Syntax

```
void __MCR(__ul coproc, __ul opcode_1, __ul src, __ul CRn, __ul CRm, __ul opcode_2);
void __MCR2(__ul coproc, __ul opcode_1, __ul src, __ul CRn, __ul CRm, __ul opcode_2);
```

Parameters

<i>coproc</i>	The coprocessor number 0..15.
<i>opcode_1</i>	Coprocessor-specific operation code.

<i>src</i>	The value to be written to the coprocessor.
<i>CRn</i>	The coprocessor register to write to.
<i>CRm</i>	Additional coprocessor register; set to zero if not used.
<i>opcode_2</i>	Additional coprocessor-specific operation code; set to zero if not used.

Description

Inserts a coprocessor write instruction (MCR or MCR2). A value will be written to a coprocessor register. The parameters *coproc*, *opcode_1*, *CRn*, *CRm*, and *opcode_2* will be encoded in the instruction and must therefore be constants.

The intrinsic function `__MCR` requires either ARM mode, or an ARMv6T2 or higher for Thumb mode.

The intrinsic function `__MCR2` requires an ARMv5T architecture or higher for ARM mode, or ARMv6T2 or higher for Thumb mode.

__MRC

__MRC2

Syntax

```
unsigned long __MRC(__ul coproc, __ul opcode_1, __ul CRn, __ul CRm, __ul opcode_2);
unsigned long __MRC2(__ul coproc, __ul opcode_1, __ul CRn, __ul CRm, __ul opcode_2);
```

Parameters

<i>coproc</i>	The coprocessor number 0..15.
<i>opcode_1</i>	Coprocessor-specific operation code.
<i>CRn</i>	The coprocessor register to write to.
<i>CRm</i>	Additional coprocessor register; set to zero if not used.
<i>opcode_2</i>	Additional coprocessor-specific operation code; set to zero if not used.

Description

Inserts a coprocessor read instruction (MRC or MRC2). Returns the value of the specified coprocessor register. The parameters *coproc*, *opcode_1*, *CRn*, *CRm*, and *opcode_2* will be encoded in the instruction and must therefore be constants.

The intrinsic function `__MRC` requires either ARM mode, or an ARMv6T2 or higher for Thumb mode.

The intrinsic function `__MRC2` requires an ARMv5T architecture or higher for ARM mode, or ARMv6T2 or higher for Thumb mode.

__no_operation

Syntax	<code>void __no_operation(void);</code>
Description	Inserts a NOP instruction.

__PKHBT

Syntax	<code>unsigned long __PKHBT(unsigned long x, unsigned long y, unsigned long count);</code>	
Parameters	<i>x</i>	First operand.
	<i>y</i>	Second operand, optionally shifted left.
	<i>count</i>	Shift count 0–31, where 0 means no shift.
Description	Inserts a <code>PKHBT</code> instruction, with an optionally shifted operand (LSL) for count in the range 1–31. This intrinsic function requires an ARM v6 architecture or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.	

__PKHTB

Syntax	<code>unsigned long __PKHTB(unsigned long x, unsigned long y, unsigned long count);</code>	
Parameters	<i>x</i>	First operand.
	<i>y</i>	Second operand, optionally shifted right (arithmetic shift).
	<i>count</i>	Shift count 0–32, where 0 means no shift.
Description	Inserts a <code>PKHTB</code> instruction, with an optionally shifted operand (ASR) for count in the range 1–32.	

This intrinsic function requires an ARM v6 architecture or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARM v7E-M for Thumb mode.

__PLD

__PLDW

Syntax

```
void __PLD(void const *);  
void __PLDW(void const *);
```

Description

Inserts a preload data instruction (PLD or PLDW).

The intrinsic function `__PLD` requires an ARMv7 architecture. `__PLDW` requires an ARMv7 architecture with MP extensions (for example Cortex-A5).

__PLI

Syntax

```
void __PLI(void const *);
```

Description

Inserts a PLI instruction.

This intrinsic function requires an ARM v7 architecture.

__QADD

__QDADD

__QDSUB

__QSUB

Syntax

```
signed long __Qxxx(signed long, signed long);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv5E or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__QADD8**__QADD16****__QASX****__QSAX****__QSUB8****__QSUB16**

Syntax

```
unsigned long __Qxxx(unsigned long, unsigned long);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__QCFlag

Syntax

```
unsigned long __QCFlag(void);
```

Description

Returns the value of the cumulative saturation flag QC of the FPSCR register (Floating-point Status and Control Register). This intrinsic function is only available for devices with Neon (Advanced SIMD).

__QDOUBLE

Syntax

```
signed long __QDOUBLE(signed long);
```

Description

Inserts an instruction `QADD Rd, Rs, Rs` for a source register `Rs`, and a destination register `Rd`.

This intrinsic function requires architecture ARMv5E or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__QFlag

Syntax

```
int __QFlag(void);
```

Description

Returns the `Q` flag that indicates if overflow/saturation has occurred.

This intrinsic function requires architecture ARMv5E or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__RBIT

Syntax

```
unsigned long __RBIT(unsigned long);
```

Description

Inserts an `RBIT` instruction, which reverses the bit order in a 32-bit register.

This intrinsic function requires architecture ARMv6T2 or higher.

__reset_Q_flag

Syntax

```
void __reset_Q_flag(void);
```

Description

Clears the `Q` flag that indicates if overflow/saturation has occurred.

This intrinsic function requires an ARM v5E architecture or higher for ARM mode, and ARM v7A, ARM v7R, or ARM v7E-M for Thumb mode.

__reset_QC_flag

Syntax

```
void __reset_QC_flag(void);
```

Description

Clears the value of the cumulative saturation flag `QC` of the `FPSCR` register (Floating-point Status and Control Register). This intrinsic function is only available for devices with Neon (Advanced SIMD).

__REV**__REV16****__REVSH**

Syntax

```
unsigned long __REV(unsigned long);
unsigned long __REV16(unsigned long);
signed long __REVSH(short);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher.

__SADD8**__SADD16****__SASX****__SSAX****__SSUB8****__SSUB16**

Syntax

```
unsigned long __Sxxx(unsigned long, unsigned long);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SEL

Syntax

```
unsigned long __SEL(unsigned long, unsigned long);
```

Description

Inserts an SEL instruction.

This intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__set_BASEPRI

Syntax

```
void __set_BASEPRI(unsigned long);
```

Description

Sets the value of the `BASEPRI` register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M3, Cortex-M4, or Cortex-M7 device.

__set_CONTROL

Syntax

```
void __set_CONTROL(unsigned long);
```

Description

Sets the value of the `CONTROL` register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M device.

__set_CPSR

Syntax

```
void __set_CPSR(unsigned long);
```

Description

Sets the value of the ARM `CPSR` (Current Program Status Register). Only the control field is changed (bits 0-7). This intrinsic function can only be used in privileged mode, is not available for Cortex-M devices, and it requires ARM mode.

__set_FAULTMASK

Syntax

```
void __set_FAULTMASK(unsigned long);
```

Description

Sets the value of the `FAULTMASK` register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M3, Cortex-M4, or Cortex-M7 device.

__set_FPSCR

Syntax

```
void __set_FPSCR(unsigned long);
```

Description

Sets the value of `FPSCR` (floating-point status and control register)

This intrinsic function is only available for devices with a VFP coprocessor.

__set_interrupt_state

Syntax	<code>void __set_interrupt_state(__istate_t);</code>
Description	Restores the interrupt state to a value previously returned by the <code>__get_interrupt_state</code> function. For information about the <code>__istate_t</code> type, see <i>__get_interrupt_state</i> , page 393.

__set_LR

Syntax	<code>void __set_LR(unsigned long);</code>
Description	Assigns a new address to the link register (R14).

__set_MSP

Syntax	<code>void __set_MSP(unsigned long);</code>
Description	Sets the value of the <code>MSP</code> register (Main Stack Pointer). This intrinsic function can only be used in privileged mode, and is only available for Cortex-M devices.

__set_PRIMASK

Syntax	<code>void __set_PRIMASK(unsigned long);</code>
Description	Sets the value of the <code>PRIMASK</code> register. This intrinsic function can only be used in privileged mode and it requires a Cortex-M device.

__set_PSP

Syntax	<code>void __set_PSP(unsigned long);</code>
Description	Sets the value of the <code>PSP</code> register (Process Stack Pointer). This intrinsic function can only be used in privileged mode, and is only available for Cortex-M devices.

__set_SB

Syntax	<code>void __set_SB(unsigned long);</code>
Description	Assigns a new address to the static base register (R9).

__set_SP

Syntax	<code>void __set_SP(unsigned long);</code>
Description	Assigns a new address to the stack pointer register (R13).

__SEV

Syntax	<code>void __SEV(void);</code>
Description	Inserts an <code>SEV</code> instruction. This intrinsic function requires architecture ARMv7 for ARM mode, and ARMv6-M or ARMv7 for Thumb mode.

__SHADD8

__SHADD16

__SHASX

__SHSAX

__SHSUB8

__SHSUB16

Syntax	<code>unsigned long __SHxxx(unsigned long, unsigned long);</code>
Description	Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SMLABB

__SMLABT

__SMLATB

__SMLATT

__SMLAWB

__SMLAWT

Syntax	<code>unsigned long __SMLAxxx(unsigned long, unsigned long, unsigned long);</code>
Description	<p>Inserts the specified instruction.</p> <p>These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.</p>

__SMLAD

__SMLADX

__SMLSD

__SMLSDX

Syntax	<code>unsigned long __SMLxxx(unsigned long, unsigned long, unsigned long);</code>
Description	Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SMLALBB

__SMLALBT

__SMLALTB

__SMLALTT

Syntax	<code>unsigned long long __SMLALxxx(unsigned long, unsigned long, unsigned long long);</code>
--------	---

Description	<p>Inserts the specified instruction.</p> <p>These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.</p>
-------------	---

__SMLALD

__SMLALDX

__SMLS LD

__SMLS LD X

Syntax	<code>unsigned long long __SMLxxx(unsigned long, unsigned long, unsigned long long);</code>
--------	---

Description	<p>Inserts the specified instruction.</p> <p>These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.</p>
-------------	---

__SMMLA

__SMMLAR

__SMMLS

__SMMLSR

Syntax

```
unsigned long __SMMLxxx(unsigned long, unsigned long, unsigned long);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SMMUL

__SMMULR

Syntax

```
unsigned long __SMMULxxx(unsigned long, unsigned long);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SMUAD

__SMUADX

__SMUSD

__SMUSDX

Syntax

```
unsigned long __SMUxxx(unsigned long, unsigned long);
```

Description

Inserts the specified instruction.

__SMUL

```
signed long __SMUL(signed short, signed short);
```

This intrinsic function requires architecture ARMv5-E or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SMULBB

__SMULBT

__SMULTB

__SMULTT

__SMULWB

__SMULWT

```
unsigned long __SMULxxx(unsigned long, unsigned long);
```

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__SSAT

```
unsigned long __SSAT(unsigned long, unsigned long);
```

The compiler will incorporate a shift instruction into the operand when possible. For example, `__SSAT(x << 3, 11)` compiles to `SSAT Rd, #11, Rn, LSL #3`, where the

value of *x* has been placed in register *Rn* and the return value of `__SSAT` will be placed in register *Rd*.

This intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7-M for Thumb mode.

__SSAT16

Syntax

```
unsigned long __SSAT16(unsigned long, unsigned long);
```

Description

Inserts an `SSAT16` instruction.

This intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARM v7E-M for Thumb mode.

__STC

__STCL

__STC2

__STC2L

Syntax

```
void __STCxxx(__ul coproc, __ul CRn, __ul const *dst);
```

Parameters

<i>coproc</i>	The coprocessor number 0..15.
<i>CRn</i>	The coprocessor register to load.
<i>dst</i>	A pointer to the destination.

Description

Inserts the coprocessor store instruction `STC`—or one of its variants—which means that the value of the specified coprocessor register will be written to a memory location. The parameters *coproc* and *CRn* will be encoded in the instruction and must therefore be constants.

The intrinsic functions `__STC` and `__STCL` require architecture ARMv4 or higher for ARM mode, and ARM v6T2 or higher for Thumb mode.

The intrinsic functions `__STC2` and `__STC2L` require architecture ARMv5 or higher for ARM mode, and ARMv6-T2 or higher for Thumb mode.

__STC_noidx

__STCL_noidx

__STC2_noidx

__STC2L_noidx

Syntax	<pre>void __STCxxx_noidx(__ul coproc, __ul CRn, __ul const *dst, __ul option);</pre>	
Parameters	<i>coproc</i>	The coprocessor number 0..15.
	<i>CRn</i>	The coprocessor register to load.
	<i>dst</i>	A pointer to the destination.
	<i>option</i>	Additional coprocessor option 0..255.
Description	<p>Inserts the coprocessor store instruction <i>STC</i>—or one of its variants—which means that the value of the specified coprocessor register will be written to a memory location. The parameters <i>coproc</i>, <i>CRn</i>, and <i>option</i> will be encoded in the instruction and must therefore be constants.</p> <p>The intrinsic functions <i>__STC_noidx</i> and <i>__STCL_noidx</i> require architecture ARMv4 or higher for ARM mode, and ARMv6-T2 or higher for Thumb mode.</p> <p>The intrinsic functions <i>__STC2_noidx</i> and <i>__STC2L_noidx</i> require architecture ARMv5 or higher for ARM mode, and ARMv6-T2 or higher for Thumb mode.</p>	

__STREX**__STREXB****__STREXD****__STREXH**

Syntax

```
unsigned long __STREX(unsigned long, unsigned long *);
unsigned long __STREXB(unsigned char, unsigned char *);
unsigned long __STREXD(unsigned long long, unsigned long long*);
unsigned long __STREXH(unsigned short, unsigned short *);
```

Description

Inserts the specified instruction.

The `__STREX` intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv6-T2 or higher for Thumb mode.

The `__STREXB` and the `__STREXH` intrinsic functions require architecture ARMv6K or ARMv7 for ARM mode, and ARMv7 for Thumb mode.

The `__STREXD` intrinsic function requires architecture ARMv6K or ARMv7 for ARM mode, and ARMv7 except for ARMv7-M for Thumb mode.

__SWP**__SWPB**

Syntax

```
unsigned long __SWP(unsigned long, unsigned long *);
char __SWPB(unsigned char, unsigned char *);
```

Description

Inserts the specified instruction.

These intrinsic functions require ARM mode.

SXTBI6

Description	Inserts the specified instruction.
-------------	------------------------------------

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

USUBI6

Description	Inserts the specified instruction.
-------------	------------------------------------

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__UHADD8**__UHADD16****__UHASX****__UHSAX****__UHSUB8****__UHSUB16**

Syntax

`unsigned long __UHxxx(unsigned long, unsigned long);`

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__UMAAL

Syntax

`unsigned long long __UMAAL(unsigned long, unsigned long,
unsigned long, unsigned long);`

Description

Inserts an `UMAAL` instruction.

This intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

__UQSUBI6

```
unsigned long __UQxxx(unsigned long, unsigned long);
```

Inserts the specified instruction.

__USADA8

```
unsigned long __USADxxx(unsigned long, unsigned long);
```

Inserts the specified instruction.

```
unsigned long __USAT(unsigned long, unsigned long);
```

Inserts a USAT instruction.

416

value of x has been placed in register R_n and the return value of `__USAT` will be placed in register R_d .

This intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7-M for Thumb mode.

`__USAT16`

Syntax

```
unsigned long __USAT16(unsigned long, unsigned long);
```

Description

Inserts a `USAT16` instruction.

This intrinsic function requires architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

`__UXTAB`

`__UXTAB16`

`__UXTAH`

`__UXTB16`

Syntax

```
unsigned long __UXTxxx(unsigned long, unsigned long);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv6 or higher for ARM mode, and ARMv7-A, ARMv7-R, or ARMv7E-M for Thumb mode.

`__WFE`

`__WFI`

`__YIELD`

Syntax

```
void long __xxx(void);
```

Description

Inserts the specified instruction.

These intrinsic functions require architecture ARMv7 for ARM mode, and ARMv6-M, or ARMv7 for Thumb mode.

The preprocessor

- Overview of the preprocessor
- Description of predefined preprocessor symbols
- Descriptions of miscellaneous preprocessor extensions

Overview of the preprocessor

The preprocessor of the IAR C/C++ Compiler for ARM adheres to Standard C. The compiler also makes these preprocessor-related features available to you:

- Predefined preprocessor symbols
These symbols allow you to inspect the compile-time environment, for example the time and date of compilation. For more information, see *Description of predefined preprocessor symbols*, page 420.
- User-defined preprocessor symbols defined using a compiler option
In addition to defining your own preprocessor symbols using the `#define` directive, you can also use the option `-D`, see *-D*, page 264.
- Preprocessor extensions
There are several preprocessor extensions, for example many pragma directives; for more information, see the chapter *Pragma directives*. For information about the corresponding `_Pragma` operator and the other extensions related to the preprocessor, see *Descriptions of miscellaneous preprocessor extensions*, page 426.
- Preprocessor output
Use the option `--preprocess` to direct preprocessor output to a named file, see *--preprocess*, page 289.

To specify a path for an include file, use forward slashes:

```
#include "mydirectory/myfile"
```

In source code, use forward slashes:

```
file = fopen("mydirectory/myfile", "rt");
```

Note that backslashes can also be used. In this case, use one in include file paths and two in source code strings.

Description of predefined preprocessor symbols

This section lists and describes the preprocessor symbols.

__AAPCS__

Description An integer that is set based on the `--aapcs` option. The symbol is set to 1 if the AAPCS base standard is the selected calling convention (`--aapcs=std`). The symbol is undefined for other calling conventions.

See also `--aapcs`, page 260.

__AAPCS_VFP__

Description An integer that is set based on the `--aapcs` option. The symbol is set to 1 if the VFP variant of AAPCS is the selected calling convention (`--aapcs=vfp`). The symbol is undefined for other calling conventions.

See also `--aapcs`, page 260.

__ARM_ADVANCED_SIMD__

Description An integer that is set based on the `--cpu` option. The symbol is set to 1 if the selected processor architecture has the Advanced SIMD architecture extension. The symbol is undefined for other cores.

See also `--cpu`, page 262.

__ARM_MEDIA__

Description An integer that is set based on the `--cpu` option. The symbol is set to 1 if the selected processor architecture has the ARMv6 SIMD extensions for multimedia. The symbol is undefined for other cores.

See also `--cpu`, page 262.

__ARM_PROFILE_M__

Description An integer that is set based on the `--cpu` option. The symbol is set to 1 if the selected processor architecture is a profile M core. The symbol is undefined for other cores.

See also `--cpu`, page 262.

__ARMVFP__

Description An integer that reflects the `--fpu` option and is defined to `__ARMVFPV2__`, `__ARMVFPV3__`, or `__ARMVFPV4__`. These symbolic names can be used when testing the `__ARMVFP__` symbol. If VFP code generation is disabled (default), the symbol will be undefined.

See also `--fpu`, page 273.

__ARMVFP_D16__

Description An integer that is set based on the `--fpu` option. The symbol is set to 1 if the selected FPU is a VFPv3 or VFPv4 unit with only 16 D registers. Otherwise, the symbol is undefined.

See also `--fpu`, page 273.

__ARMVFP_FPI6__

Description An integer that is set based on the `--fpu` option. The symbol is set to 1 if the selected FPU only supports 16-bit floating-point numbers. Otherwise, the symbol is undefined.

See also `--fpu`, page 273.

__ARMVFP_SP__

Description An integer that is set based on the `--fpu` option. The symbol is set to 1 if the selected FPU only supports 32-bit single-precision. Otherwise, the symbol is undefined.

See also `--fpu`, page 273.

__BASE_FILE__

Description A string that identifies the name of the base source file (that is, not the header file), being compiled.

See also `__FILE__`, page 423, and `--no_path_in_file_macros`, page 282.

__BUILD_NUMBER__

Description A unique integer that identifies the build number of the compiler currently in use.

__CORE__

Description An integer that identifies the chip core in use. The value reflects the setting of the `--cpu` option and is defined to `__ARM4TM__`, `__ARM5__`, `__ARM5E__`, `__ARM6__`, `__ARM6M__`, `__ARM6SM__`, `__ARM7M__`, `__ARM7EM__`, `__ARM7A__`, or `__ARM7R__`. These symbolic names can be used when testing the `__CORE__` symbol.

__COUNTER__

Description A macro that expands to a new integer each time it is expanded, starting at zero (0) and counting up.

__cplusplus

Description An integer which is defined when the compiler runs in any of the C++ modes, otherwise it is undefined. When defined, its value is 199711L. This symbol can be used with `#ifdef` to detect whether the compiler accepts C++ code. It is particularly useful when creating header files that are to be shared by C and C++ code.

This symbol is required by Standard C.

__CPU_MODE__

Description An integer that reflects the selected CPU mode and is defined to 1 for Thumb and 2 for ARM.

__DATE__

Description

A string that identifies the date of compilation, which is returned in the form "Mmm dd yyyy", for example "Oct 30 2014"

This symbol is required by Standard C.

__embedded_cplusplus

Description

An integer which is defined to 1 when the compiler runs in Embedded C++ or Extended Embedded C++ mode, otherwise the symbol is undefined. This symbol can be used with `#ifdef` to detect whether the compiler accepts C++ code. It is particularly useful when creating header files that are to be shared by C and C++ code.

This symbol is required by Standard C.

__FILE__

Description

A string that identifies the name of the file being compiled, which can be both the base source file and any included header file.

This symbol is required by Standard C.

See also

`__BASE_FILE__`, page 422, and `--no_path_in_file_macros`, page 282.

__func__

Description

A predefined string identifier that is initialized with the name of the function in which the symbol is used. This is useful for assertions and other trace utilities. The symbol requires that language extensions are enabled.

This symbol is required by Standard C.

See also

`-e`, page 269 and `__PRETTY_FUNCTION__`, page 424.

__FUNCTION__

Description

A predefined string identifier that is initialized with the name of the function in which the symbol is used. This is useful for assertions and other trace utilities. The symbol requires that language extensions are enabled.

See also

`-e`, page 269 and `__PRETTY_FUNCTION__`, page 424.

__IAR_SYSTEMS_ICC__

Description	An integer that identifies the IAR compiler platform. The current value is 8. Note that the number could be higher in a future version of the product. This symbol can be tested with <code>#ifdef</code> to detect whether the code was compiled by a compiler from IAR Systems.
-------------	---

__ICCARM__

Description	An integer that is set to 1 when the code is compiled with the IAR C/C++ Compiler for ARM.
-------------	--

__LINE__

Description	<p>An integer that identifies the current source line number of the file being compiled, which can be both the base source file and any included header file.</p> <p>This symbol is required by Standard C.</p>
-------------	---

__LITTLE_ENDIAN__

Description	<p>An integer that reflects the setting of the <code>--endian</code> option and is defined to 1 when the byte order is little-endian. The symbol is defined to 0 when the byte order is big-endian.</p> <p>An integer that reflects the byte order and is defined to 0 (big-endian).</p>
-------------	--

__PRETTY_FUNCTION__

Description	A predefined string identifier that is initialized with the function name, including parameter types and return type, of the function in which the symbol is used, for example <code>"void func(char) "</code> . This symbol is useful for assertions and other trace utilities. The symbol requires that language extensions are enabled.
-------------	--

See also	<code>-e</code> , page 269 and <code>__func__</code> , page 423.
----------	--

__ROPI__

Description	An integer that is defined when the <code>--ropi</code> compiler option is used.
-------------	--

See also	<code>--ropi</code> , page 291.
----------	---------------------------------

__RWPI__

Description An integer that is defined when the `--rwp`i compiler option is used.

See also `--rwp`i, page 292.

__STDC__

Description An integer that is set to 1, which means the compiler adheres to Standard C. This symbol can be tested with `#ifdef` to detect whether the compiler in use adheres to Standard C.*
This symbol is required by Standard C.

__STDC_VERSION__

Description An integer that identifies the version of the C standard in use. The symbol expands to 199901L, unless the `--c89` compiler option is used in which case the symbol expands to 199409L. This symbol does not apply in EC++ mode.

This symbol is required by Standard C.

__TIME__

Description A string that identifies the time of compilation in the form "hh:mm:ss".

This symbol is required by Standard C.

__TIMESTAMP__

Description A string constant that identifies the date and time of the last modification of the current source file. The format of the string is the same as that used by the `asctime` standard function (in other words, "Tue Sep 16 13:03:52 2014").

__VER__

Description An integer that identifies the version number of the IAR compiler in use. For example, version 5.11.3 is returned as 5011003.

Descriptions of miscellaneous preprocessor extensions

This section gives reference information about the preprocessor extensions that are available in addition to the predefined symbols, pragma directives, and Standard C directives.

NDEBUG

Description

This preprocessor symbol determines whether any assert macros you have written in your application shall be included or not in the built application.

If this symbol is not defined, all assert macros are evaluated. If the symbol is defined, all assert macros are excluded from the compilation. In other words, if the symbol is:

- **defined**, the assert code will *not* be included
- **not defined**, the assert code will be included

This means that if you write any assert code and build your application, you should define this symbol to exclude the assert code from the final application.

Note that the assert macro is defined in the `assert.h` standard include file.

In the IDE, the `NDEBUG` symbol is automatically defined if you build your application in the Release build configuration.

See also

`__aeabi_assert`, page 139.

#warning message

Syntax

`#warning message`

where *message* can be any string.

Description

Use this preprocessor directive to produce messages. Typically, this is useful for assertions and other trace utilities, similar to the way the Standard C `#error` directive is used. This directive is not recognized when the `--strict` compiler option is used.

C/C++ standard library functions

- C/C++ standard library overview
- DLIB runtime environment—implementation details

For detailed reference information about the library functions, see the online help system.

C/C++ standard library overview

The IAR DLIB Runtime Environment is a complete implementation of the C/C++ standard library, compliant with Standard C and C++. This library also supports floating-point numbers in IEEE 754 format and it can be configured to include different levels of support for locale, file descriptors, multibyte characters, etc.

For more information about customization, see the chapter *The DLIB runtime environment*.

For detailed information about the library functions, see the online documentation supplied with the product. There is also keyword reference information for the DLIB library functions. To obtain reference information for a function, select the function name in the editor window and press F1.

For more information about library functions, see the chapter *Implementation-defined behavior for Standard C* in this guide.

HEADER FILES

Your application program gains access to library definitions through header files, which it incorporates using the `#include` directive. The definitions are divided into several different header files, each covering a particular functional area, letting you include just those that are required.

It is essential to include the appropriate header file before making any reference to its definitions. Failure to do so can cause the call to fail during execution, or generate error or warning messages at compile time or link time.

LIBRARY OBJECT FILES

Most of the library definitions can be used without modification, that is, directly from the library object files that are supplied with the product. For information about how to set up a runtime library, see *Setting up the runtime environment*, page 119. The linker will include only those routines that are required—directly or indirectly—by your application.

See also *Overriding library modules*, page 122 for information about how you can override library modules with your own versions.

ALTERNATIVE MORE ACCURATE LIBRARY FUNCTIONS

The default implementation of `cos`, `sin`, `tan`, and `pow` is designed to be fast and small. As an alternative, there are versions designed to provide better accuracy. They are named `__iar_xxx_accuratef` for float variants of the functions and `__iar_xxx_accuratel` for long double variants of the functions, and where `xxx` is `cos`, `sin`, etc.

To use any of these more accurate versions, use the `--redirect` linker option.

REENTRANCY

A function that can be simultaneously invoked in the main application and in any number of interrupts is reentrant. A library function that uses statically allocated data is therefore not reentrant.

Most parts of the DLIB runtime environment are reentrant, but the following functions and parts are not reentrant because they need static data:

- Heap functions—`malloc`, `free`, `realloc`, `calloc`, and the C++ operators `new` and `delete`
- Locale functions—`localeconv`, `setlocale`
- Multibyte functions—`mbrlen`, `mbrtowc`, `mbsrtowc`, `mbtowc`, `wcrtomb`, `wcsrtomb`, `wctomb`
- Rand functions—`rand`, `srand`
- Time functions—`asctime`, `localtime`, `gmtime`, `mktime`
- The miscellaneous functions `atexit`, `strerror`, `strtok`
- Functions that use files or the heap in some way. This includes `scanf`, `sscanf`, `getchar`, and `putchar`. In addition, if you are using the options `--enable_multibyte` and `--dlib_config=Full`, the `printf` and `sprintf` functions (or any variants) can also use the heap.

Functions that can set `errno` are not reentrant, because an `errno` value resulting from one of these functions can be destroyed by a subsequent use of the function before it is read. This applies to math and string conversion functions, among others.

Remedies for this are:

- Do not use non-reentrant functions in interrupt service routines
- Guard calls to a non-reentrant function by a mutex, or a secure region, etc.

THE LONGJMP FUNCTION



A `longjmp` is in effect a jump to a previously defined `setjmp`. Any variable length arrays or C++ objects residing on the stack during stack unwinding will not be destroyed. This can lead to resource leaks or incorrect application behavior.

DLIB runtime environment—implementation details

The DLIB runtime environment provides most of the important C and C++ standard library definitions that apply to embedded systems. These are of the following types:

- Adherence to a free-standing implementation of Standard C. The library supports most of the hosted functionality, but you must implement some of its base functionality. For additional information, see the chapter *Implementation-defined behavior for Standard C* in this guide.
- Standard C library definitions, for user programs.
- C++ library definitions, for user programs.
- `CSTARTUP`, the module containing the start-up code, see the chapter *The DLIB runtime environment* in this guide.
- Runtime support libraries; for example low-level floating-point routines.
- Intrinsic functions, allowing low-level use of ARM features. See the chapter *Intrinsic functions* for more information.

In addition, the DLIB runtime environment includes some added C functionality, see *Added C functionality*, page 433.

C HEADER FILES

This section lists the C header files specific to the DLIB runtime environment. Header files may additionally contain target-specific definitions; these are documented in the chapter *Using C*.

This table lists the C header files:

Header file	Usage
assert.h	Enforcing assertions when functions execute
complex.h	Computing common complex mathematical functions
ctype.h	Classifying characters
errno.h	Testing error codes reported by library functions
fenv.h	Floating-point exception flags
float.h	Testing floating-point type properties
inttypes.h	Defining formatters for all types defined in stdint.h
iso646.h	Using Amendment 1—iso646.h standard header
limits.h	Testing integer type properties
locale.h	Adapting to different cultural conventions
math.h	Computing common mathematical functions
setjmp.h	Executing non-local goto statements
signal.h	Controlling various exceptional conditions
stdarg.h	Accessing a varying number of arguments
stdbool.h	Adds support for the bool data type in C.
stddef.h	Defining several useful types and macros
stdint.h	Providing integer characteristics
stdio.h	Performing input and output
stdlib.h	Performing a variety of operations
string.h	Manipulating several kinds of strings
tgmath.h	Type-generic mathematical functions
time.h	Converting between various time and date formats
uchar.h	Unicode functionality (IAR extension to Standard C)
wchar.h	Support for wide characters
wctype.h	Classifying wide characters

Table 35: Traditional Standard C header files—DLIB

C++ HEADER FILES

This section lists the C++ header files:

- The C++ library header files
The header files that constitute the Standard C++ and the Embedded C++ library.
- The C++ standard template library (STL) header files

The header files that constitute STL for the Standard C++ and the Extended Embedded C++ library.

- The C++ C header files

The C++ header files that provide the resources from the C library.

The C++ library header files

This table lists the header files that can be used in C++ as well as in Embedded C++:

Header file	Usage
<code>complex</code>	Defining a class that supports complex arithmetic
<code>exception</code>	Defining several functions that control exception handling; only usable in C++
<code>fstream</code>	Defining several I/O stream classes that manipulate external files
<code>iomanip</code>	Declaring several I/O stream manipulators that take an argument
<code>ios</code>	Defining the class that serves as the base for many I/O streams classes
<code>iosfwd</code>	Declaring several I/O stream classes before they are necessarily defined
<code>iostream</code>	Declaring the I/O stream objects that manipulate the standard streams
<code>istream</code>	Defining the class that performs extractions
<code>limits</code>	Defining numerical values; only usable in C++
<code>locale</code>	Adapting to different cultural conventions; only usable in C++
<code>new</code>	Declaring several functions that allocate and free storage
<code>ostream</code>	Defining the class that performs insertions
<code>sstream</code>	Defining several I/O stream classes that manipulate string containers
<code>stdexcept</code>	Defining several classes useful for reporting exceptions; only usable in C++
<code>streambuf</code>	Defining classes that buffer I/O stream operations
<code>string</code>	Defining a class that implements a string container
<code>strstream</code>	Defining several I/O stream classes that manipulate in-memory character sequences
<code>typeinfo</code>	Defining type information support; only usable in C++

Table 36: C++ header files

The C++ standard template library (STL) header files

The following table lists the standard template library (STL) header files that can be used in C++ as well as in Extended Embedded C++:

Header file	Description
algorithm	Defines several common operations on sequences
bitset	Defining a container with fixed-sized sequences of bits; only usable in C++
deque	A deque sequence container
functional	Defines several function objects
hash_map	A map associative container, based on a hash algorithm
hash_set	A set associative container, based on a hash algorithm
iterator	Defines common iterators, and operations on iterators
list	A doubly-linked list sequence container
map	A map associative container
memory	Defines facilities for managing memory
numeric	Performs generalized numeric operations on sequences
queue	A queue sequence container
set	A set associative container
slist	A singly-linked list sequence container
stack	A stack sequence container
utility	Defines several utility components
valarray	Defining varying length array container; only usable in C++
vector	A vector sequence container

Table 37: Standard template library header files

Using Standard C libraries in C++

The C++ library works in conjunction with some of the header files from the Standard C library, sometimes with small alterations. The header files come in two forms—new and traditional—for example, `cassert` and `assert.h`.

This table shows the new header files:

Header file	Usage
cassert	Enforcing assertions when functions execute
complex	Computing common complex mathematical functions

Table 38: New Standard C header files—DLIB

Header file	Usage
<code>cctype</code>	Classifying characters
<code>cerrno</code>	Testing error codes reported by library functions
<code>cfenv</code>	Floating-point exception flags
<code>cfloat</code>	Testing floating-point type properties
<code>cinttypes</code>	Defining formatters for all types defined in <code>stdint.h</code>
<code>ciso646</code>	Using Amendment 1— <code>iso646.h</code> standard header
<code>climits</code>	Testing integer type properties
<code>clocale</code>	Adapting to different cultural conventions
<code>cmath</code>	Computing common mathematical functions
<code>csetjmp</code>	Executing non-local goto statements
<code>csignal</code>	Controlling various exceptional conditions
<code>cstdarg</code>	Accessing a varying number of arguments
<code>cstdbool</code>	Adds support for the <code>bool</code> data type in C.
<code>cstddef</code>	Defining several useful types and macros
<code>stdint</code>	Providing integer characteristics
<code>stdio</code>	Performing input and output
<code>stdlib</code>	Performing a variety of operations
<code>cstring</code>	Manipulating several kinds of strings
<code>ctgmth</code>	Type-generic mathematical functions
<code>ctime</code>	Converting between various time and date formats
<code>wchar</code>	Support for wide characters
<code>wctype</code>	Classifying wide characters

Table 38: New Standard C header files—DLIB (Continued)

LIBRARY FUNCTIONS AS INTRINSIC FUNCTIONS

Certain C library functions will under some circumstances be handled as intrinsic functions and will generate inline code instead of an ordinary function call, for example `memcpy`, `memset`, and `strcat`.

ADDED C FUNCTIONALITY

The DLIB runtime environment includes some added C functionality.

The following include files provide these features:

- `fenv.h`

- `iar_dlmalloc.h`
- `stdio.h`
- `stdlib.h`
- `string.h`
- `time.h`

fenv.h

In `fenv.h`, trap handling support for floating-point numbers is defined with the functions `fegettrapenable` and `fegettrapdisable`.

iar_dlmalloc.h

The `iar_dlmalloc.h` header file contains support for the advanced (`dlmalloc`) heap handler. For more information, see *Heap considerations*, page 205.

stdio.h

These functions provide additional I/O functionality:

<code>fdopen</code>	Opens a file based on a low-level file descriptor.
<code>fileno</code>	Gets the low-level file descriptor from the file descriptor (<code>FILE*</code>).
<code>__gets</code>	Corresponds to <code>fgets</code> on <code>stdin</code> .
<code>getw</code>	Gets a <code>wchar_t</code> character from <code>stdin</code> .
<code>putw</code>	Puts a <code>wchar_t</code> character to <code>stdout</code> .
<code>__ungetchar</code>	Corresponds to <code>ungetc</code> on <code>stdout</code> .
<code>__write_array</code>	Corresponds to <code>fwrite</code> on <code>stdout</code> .

string.h

These are the additional functions defined in `string.h`:

<code>strdup</code>	Duplicates a string on the heap.
<code>strcasecmp</code>	Compares strings case-insensitive.
<code>strncasecmp</code>	Compares strings case-insensitive and bounded.
<code>strlen</code>	Bounded string length.

time.h

There are two interfaces for using `time_t` and the associated functions `time`, `ctime`, `difftime`, `gmtime`, `localtime`, and `mktime`:

- The 32-bit interface supports years from 1900 up to 2035 and uses a 32-bit integer for `time_t`. The type and function have names like `__time32_t`, `__time32`, etc. This variant is mainly available for backwards compatibility.
- The 64-bit interface supports years from -9999 up to 9999 and uses a signed `long long` for `time_t`. The type and function have names like `__time64_t`, `__time64`, etc.

In both interfaces, `time_t` starts at the year 1970.

The interfaces are defined in the system header file `time.h`.

An application can use either interface, and even mix them by explicitly using the 32- or 64-bit variants. By default, the library and the header redirect `time_t`, `time` etc. to the 32-bit variants. However, to explicitly redirect them to their 64-bit variants, define `_DLIB_TIME_USES_64` in front of the inclusion of `time.h` or `ctime`.

See also, `__time32`, `__time64`, page 145.

`clock_t` is represented by a 32-bit integer type.

SYMBOLS USED INTERNALLY BY THE LIBRARY

The following symbols are used by the library, which means that they are visible in library source files, etc:

`__assignment_by_bitwise_copy_allowed`

This symbol determines properties for class objects.

`__code`, `__data`

These symbols are used as memory attributes internally by the compiler, and they might have to be used as arguments in certain templates.

`__constrange()`

Determines the allowed range for a parameter to an intrinsic function and that the parameter must be of type `const`.

`__construction_by_bitwise_copy_allowed`

This symbol determines properties for class objects.

`__has_constructor`, `__has_destructor`

These symbols determine properties for class objects and they function like the `sizeof` operator. The symbols are true when a class, base class, or member (recursively) has a user-defined constructor or destructor, respectively.

`__memory_of`

Determines the class memory. A class memory determines which memory a class object can reside in. This symbol can only occur in class definitions as a class memory.

Note: The symbols are reserved and should only be used by the library.

Use the compiler option `--predef_macros` to determine the value for any predefined symbols.

The linker configuration file

- Overview
- Defining memories and regions
- Regions
- Section handling
- Section selection
- Using symbols, expressions, and numbers
- Structural configuration

Before you read this chapter you must be familiar with the concept of *sections*, see *Modules and sections*, page 84.

Overview

To link and locate an application in memory according to your requirements, ILINK needs information about how to handle sections and how to place them into the available memory regions. In other words, ILINK needs a *configuration*, passed to it by means of the *linker configuration file*.

This file consists of a sequence of directives and typically, provides facilities for:

- Defining available addressable memories
 - giving the linker information about the maximum size of possible addresses and defining the available physical memory, as well as dealing with memories that can be addressed in different ways.
- Defining the regions of the available memories that are populated with ROM or RAM
 - giving the start and end address for each region.
- Section groups

dealing with how to group sections into blocks and overlays depending on the section requirements.

- Defining how to handle initialization of the application
giving information about which sections that are to be initialized, and how that initialization should be made.
- Memory allocation
defining where—in what memory region—each set of sections should be placed.
- Using symbols, expressions, and numbers
expressing addresses and sizes, etc, in the other configuration directives. The symbols can also be used in the application itself.
- Structural configuration
meaning that you can include or exclude directives depending on a condition, and to split the configuration file into several different files.

Comments can be written either as C comments (`/* . . . */`) or as C++ comments (`// . . .`).

Defining memories and regions

ILINK needs information about the available memory spaces, or more specifically it needs information about:

- The maximum size of possible addressable memories
The `define memory` directive defines a *memory space* with a given size, which is the maximum possible amount of addressable memory, not necessarily physically available. See *define memory directive*, page 439.
- Available physical memory
The `define region` directive defines a region in the available memories in which specific sections of application code and sections of application data can be placed. See *define region directive*, page 439.

A region consists of one or several memory ranges. A range is a continuous sequence of bytes in a memory and several ranges can be expressed by using region expressions. See *Region expression*, page 441.

This section gives detailed information about each linker directive specific to defining memories and regions.

define memory directive

Syntax	<pre>define memory [name] with size = size_expr [,unit-size];</pre> <p>where <i>unit-size</i> is one of:</p> <pre>unitbitsize = bitsize_expr unitbytesize = bytesize_expr</pre> <p>and where <i>expr</i> is an expression, see <i>expressions</i>, page 459.</p>							
Parameters	<table><tr><td><i>size_expr</i></td><td>Specifies how many <i>units</i> the memory space contains; always counted from address zero.</td></tr><tr><td><i>bitsize_expr</i></td><td>Specifies how many bits each unit contains.</td></tr><tr><td><i>bytesize_expr</i></td><td>Specifies how many bytes each unit contains. Each byte contains 8 bits.</td></tr></table>	<i>size_expr</i>	Specifies how many <i>units</i> the memory space contains; always counted from address zero.	<i>bitsize_expr</i>	Specifies how many bits each unit contains.	<i>bytesize_expr</i>	Specifies how many bytes each unit contains. Each byte contains 8 bits.	
<i>size_expr</i>	Specifies how many <i>units</i> the memory space contains; always counted from address zero.							
<i>bitsize_expr</i>	Specifies how many bits each unit contains.							
<i>bytesize_expr</i>	Specifies how many bytes each unit contains. Each byte contains 8 bits.							
Description	<p>The <code>define memory</code> directive defines a <i>memory space</i> with a given size, which is the maximum possible amount of addressable memory, not necessarily physically available. This sets the limits for the possible addresses to be used in the linker configuration file. For many microcontrollers, one memory space is sufficient. However, some microcontrollers require two or more. For example, a Harvard architecture usually requires two different memory spaces, one for code and one for data. If only one memory is defined, the memory name is optional. If no <i>unit-size</i> is given, the unit contains 8 bits.</p>							
Example	<pre>/* Declare the memory space Mem of four Gigabytes */ define memory Mem with size = 4G;</pre>							

define region directive

Syntax	<pre>define region <i>name</i> = <i>region-expr</i>;</pre> <p>where <i>region-expr</i> is a region expression, see also <i>Regions</i>, page 440.</p>		
Parameters	<table><tr><td><i>name</i></td><td>The name of the region.</td></tr></table>	<i>name</i>	The name of the region.
<i>name</i>	The name of the region.		
Description	The <code>define region</code> directive defines a region in which specific sections of code and sections of data can be placed. A region consists of one or several memory ranges, where each memory range consists of a continuous sequence of bytes in a specific memory.		

Several ranges can be combined by using region expressions. Note that those ranges do not need to be consecutive or even in the same memory.

Example

```
/* Define the 0x10000-byte code region ROM located at address
   0x10000 in memory Mem */
define region ROM = Mem:[from 0x10000 size 0x10000];
```

Regions

A *region* is a set of non-overlapping memory ranges. A *region expression* is built up out of *region literals* and set operations (union, intersection, and difference) on regions.

Region literal

Syntax

```
[ memory-name: ] [from expr { to expr | size expr }
                  [ repeat expr [ displacement expr ] ]]
```

where *expr* is an expression, see *expressions*, page 459.

Parameters	
<i>memory-name</i>	The name of the memory space in which the region literal will be located. If there is only one memory, the name is optional.
from <i>expr</i>	<i>expr</i> is the start address of the memory range (inclusive).
to <i>expr</i>	<i>expr</i> is the end address of the memory range (inclusive).
size <i>expr</i>	<i>expr</i> is the size of the memory range.
repeat <i>expr</i>	<i>expr</i> defines several ranges in the same memory for the region literal.
displacement <i>expr</i>	<i>expr</i> is the displacement from the previous range start in the repeat sequence. Default displacement is the same value as the range size.

Description

A region literal consists of one memory range. When you define a range, the memory it resides in, a start address, and a size must be specified. The range size can be stated explicitly by specifying a size, or implicitly by specifying the final address of the range. The final address is included in the range and a zero-sized range will only contain an address. A range can span over the address zero and such a range can even be expressed by unsigned values, because it is known where the memory wraps.

The `repeat` parameter will create a region literal that contains several ranges, one for each repeat. This is useful for *banked* or *far* regions.

Example

```
/* The 5-byte size range spans over the address zero */
Mem:[from -2 to 2]

/* The 512-byte size range spans over zero, in a 64-Kbyte memory
*/
Mem:[from 0xFF00 to 0xFF]

/* Defining several ranges in the same memory, a repeating
   literal */
Mem:[from 0 size 0x100 repeat 3 displacement 0x1000]

/* Resulting in a region containing:
   Mem:[from 0 size 0x100]
   Mem:[from 0x1000 size 0x100]
   Mem:[from 0x2000 size 0x100]
*/
```

See also

define region directive, page 439, and *Region expression*, page 441.

Region expression

Syntax

```
region-operand
| region-expr | region-operand
| region-expr - region-operand
| region-expr & region-operand
```

where *region-operand* is one of:

```
( region-expr )
region-name
region-literal
empty-region
```

where *region-name* is a region, see *define region directive*, page 439

where *region-literal* is a region literal, see *Region literal*, page 440

and where *empty-region* is an empty region, see *Empty region*, page 442.

Description

Normally, a region consists of one memory range, which means a *region literal* is sufficient to express it. When a region contains several ranges, possibly in different memories, it is instead necessary to use a *region expression* to express it. Region expressions are actually set expressions on sets of memory ranges.

To create region expressions, three operators are available: union (`|`), intersection (`&`), and difference (`-`). These operators work as in *set theory*. For example, if you have the sets A and B, then the result of the operators would be:

- `A | B`: all elements in either set A or set B
- `A & B`: all elements in both set A and B
- `A - B`: all elements in set A but not in B.

Example

```
/* Resulting in a range starting at 1000 and ending at 2FFF, in
   memory Mem */
Mem:[from 0x1000 to 0x1FFF] | Mem:[from 0x1500 to 0x2FFF]

/* Resulting in a range starting at 1500 and ending at 1FFF, in
   memory Mem */
Mem:[from 0x1000 to 0x1FFF] & Mem:[from 0x1500 to 0x2FFF]

/* Resulting in a range starting at 1000 and ending at 14FF, in
   memory Mem */
Mem:[from 0x1000 to 0x1FFF] - Mem:[from 0x1500 to 0x2FFF]

/* Resulting in two ranges. The first starting at 1000 and ending
   at 1FFF, the second starting at 2501 and ending at 2FFF.
   Both located in memory Mem */
Mem:[from 0x1000 to 0x2FFF] - Mem:[from 0x2000 to 0x24FF]
```

Empty region

Syntax

```
[ ]
```

Description

The empty region does not contain any memory ranges. If the empty region is used in a placement directive that actually is used for placing one or more sections, ILINK will issue an error.

Example

```
define region Code = Mem:[from 0 size 0x10000];
if (Banked) {
    define region Bank = Mem:[from 0x8000 size 0x1000];
} else {
    define region Bank = [];
}
define region NonBanked = Code - Bank;

/* Depending on the Banked symbol, the NonBanked region is either
   one range with 0x10000 bytes, or two ranges with 0x8000 and
   0x7000 bytes, respectively. */
```

See also

Region expression, page 441.

Section handling

Section handling describes how ILINK should handle the sections of the execution image, which means:

- Placing sections in regions

The `place at` and `place in` directives place sets of sections with similar attributes into previously defined regions. See *place at directive*, page 450 and *place in directive*, page 451.

- Making sets of sections with special requirements

The `block` directive makes it possible to create empty sections with specific sizes and alignments, sequentially sorted sections of different types, etc.

The `overlay` directive makes it possible to create an area of memory that can contain several overlay images. See *define block directive*, page 444, and *define overlay directive*, page 445.

- Initializing the application

The directives `initialize` and `do not initialize` control how the application should be started. With these directives, the application can initialize global symbols at startup, and copy pieces of code. The initializers can be stored in several ways, for example they can be compressed. See *initialize directive*, page 446 and *do not initialize directive*, page 449.

- Keeping removed sections

The `keep` directive retains sections even though they are not referred to by the rest of the application, which means it is equivalent to the *root* concept in the assembler and compiler. See *keep directive*, page 449.

- Additional more specialized directives:

`use init table directive`

This section gives detailed information about each linker directive specific to section handling.

define block directive

Syntax

```
define [movable] block name
  [ with param, param... ]
{
  extended-selectors
}
[except
{
  section_selectors
}];
```

where *param* can be one of:

```
size = expr
maximum size = expr
alignment = expr
fixed order
static base [basename]
```

and where the rest of the directive selects sections to include in the block, see *Section selection*, page 452.

Parameters

<i>name</i>	The name of the block to be defined.
<i>size</i>	Customizes the size of the block. By default, the size of a block is the sum of its parts dependent of its contents.
maximum <i>size</i>	Specifies an upper limit for the size of the block. An error is generated if the sections in the block do not fit.
alignment	Specifies a minimum alignment for the block. If any section in the block has a higher alignment than the minimum alignment, the block will have that alignment.
fixed order	Places sections in fixed order; if not specified, the order of the sections will be arbitrary.
static base [<i>basename</i>]	Specifies that the static base with the name <i>basename</i> will be placed at the start of the block or in the middle of the block, as appropriate for the particular static base. The startup code must ensure that the register that holds the static base is initialized to the correct value. If there is only one static base, the name can be omitted.

Description

The `block` directive defines a contiguous area of memory that contains a possibly empty set of sections or other blocks. Blocks with no content are useful for allocating

space for stacks or heaps. Blocks with content are usually used to group together sections that must to be consecutive.

You can access the start, end, and size of a block from an application by using the `__section_begin`, `__section_end`, or `__section_size` operators. If there is no block with the specified name, but there are sections with that name, a block will be created by the linker, containing all such sections.

movable blocks are for use with read-only and read-write position independence. Making blocks movable enables the linker to validate the application’s use of addresses. Movable blocks are located in exactly the same way as other blocks, but the linker will check that the appropriate relocations are used when referring to symbols in movable blocks.

Example

```
/* Create a 0x1000-byte block for the heap */
define block HEAP with size = 0x1000, alignment = 8 { };
```

See also

Interaction between the tools and your application, page 206. See *define overlay directive*, page 445 for an Accessing example.

define overlay directive

Syntax

```
define overlay name [ with param, param... ]
{
    extended-selectors;
}
[except
{
    section_selectors
}];
```

For information about extended selectors and except clauses, see *Section selection*, page 452.

Parameters	<div> <div><code>name</code></div> <div>The name of the overlay.</div> </div> <div> <div><code>size</code></div> <div>Customizes the size of the overlay. By default, the size of a overlay is the sum of its parts dependent of its contents.</div> </div> <div> <div><code>maximum size</code></div> <div>Specifies an upper limit for the size of the overlay. An error is generated if the sections in the overlay do not fit.</div> </div> <div> <div><code>alignment</code></div> <div>Specifies a minimum alignment for the overlay. If any section in the overlay has a higher alignment than the minimum alignment, the overlay will have that alignment.</div> </div>
------------	---

<code>fixed order</code>	Places sections in fixed order; if not specified, the order of the sections will be arbitrary.
--------------------------	--

Description	<p>The <code>overlay</code> directive defines a named set of sections. In contrast to the <code>block</code> directive, the <code>overlay</code> directive can define the same name several times. Each definition will then be grouped in memory at the same place as all other definitions of the same name. This creates an <i>overlaid</i> memory area, which can be useful for an application that has several independent sub-applications.</p> <p>Place each sub-application image in ROM and reserve a RAM overlay area that can hold all sub-applications. To execute a sub-application, first copy it from ROM to the RAM overlay. Note that ILINK does not help you with managing interdependent overlay definitions, apart from generating a diagnostic message for any reference from one overlay to another overlay.</p> <p>The size of an overlay will be the same size as the largest definition of that overlay name and the alignment requirements will be the same as for the definition with the highest alignment requirements.</p> <p>Note: Sections that were overlaid must be split into a RAM and a ROM part and you must take care of all the copying needed.</p>
-------------	--

See also	<i>Manual initialization</i> , page 107.
----------	--

initialize directive

Syntax	<pre>initialize { by copy manually } [with param, param...]</pre>
--------	---

where *param* can be one of:

`packing = algorithm`
`simple ranges`
`complex ranges`

Parameters	<table><tr><td><code>by copy</code></td><td>Splits the section into sections for initializers and initialized data, and handles the initialization at application startup automatically.</td></tr><tr><td><code>manually</code></td><td>Splits the section into sections for initializers and initialized data. The initialization at application startup is not handled automatically.</td></tr></table>	<code>by copy</code>	Splits the section into sections for initializers and initialized data, and handles the initialization at application startup automatically.	<code>manually</code>	Splits the section into sections for initializers and initialized data. The initialization at application startup is not handled automatically.
<code>by copy</code>	Splits the section into sections for initializers and initialized data, and handles the initialization at application startup automatically.				
<code>manually</code>	Splits the section into sections for initializers and initialized data. The initialization at application startup is not handled automatically.				

<i>algorithm</i>	<p>Specifies how to handle the initializers. Choose between:</p> <ul style="list-style-type: none"> <i>none</i> - Disables compression of the selected section contents. This is the default method for initialize manually. <i>zeros</i> - Compresses consecutive bytes with the value zero. <i>packbits</i> - Compresses with the PackBits algorithm. This method generates good results for data with many identical consecutive bytes. <i>lz77</i> - Compresses with the Lempel-Ziv-77 algorithm. This method handles a larger variety of inputs well, but has a slightly larger decompressor. <i>auto</i> - ILINK estimates the resulting size using each packing method (except for <i>auto</i>), and then chooses the packing method that produces the smallest estimated size. Note that the size of the decompressor is also included. This is the default method for <i>initialize by copy</i>. <i>smallest</i> - This is a synonym for <i>auto</i>.
Description	<p>The <i>initialize</i> directive splits each selected section into one section that holds initializer data and another section that holds the space for the initialized data. The section that holds the space for the initialized data retains the original section name, and the section that holds initializer data gets the name suffix <i>_init</i>. You can choose whether the initialization at startup should be handled automatically (<i>initialize by copy</i>) or whether you should handle it yourself (<i>initialize manually</i>).</p> <p>When you use the packing method <i>auto</i> (default for <i>initialize by copy</i>), ILINK will automatically choose an appropriate packing algorithm for the initializers. To override this, specify a different packing method. The <i>--log initialization</i> option shows how ILINK decided which packing algorithm to use.</p> <p>When initializers are compressed, a decompressor is automatically added to the image.</p> <p>Each decompressor has two variants: one that can only handle a single source and destination range at a time, and one that can handle more complex cases. By default, the linker chooses a decompressor variant based on whether the associated section placement directives specify a single- or multi-range memory region. In general, this is the desired behavior, but you can use the <i>with complex ranges</i> or the <i>with simple ranges</i> modifier on an <i>initialize</i> directive to specify which decompressor variant to use. You can also use the command line option <i>--default_to_complex_ranges</i> to make <i>initialize</i> directives by default use complex ranges. The simple ranges decompressors are normally hundreds of bytes smaller than the complex ranges variants.</p>

When initializers are compressed, the exact size of the compressed initializers is unknown until the exact content of the uncompressed data is known. If this data contains other addresses, and some of these addresses are dependent on the size of the compressed initializers, the linker fails with error Lp017. To avoid this, place compressed initializers last, or in a memory region together with sections whose addresses do not need to be known.

Due to an internal dependence, generation of compressed initializers can also fail (with error LP021) if the address of the initialized area depends on the size of its initializers. To avoid this, place the initializers and the initialized area in different parts of the memory (for example, the initializers are placed in ROM and the initialized area in RAM).

Unless `initialize manually` is used, ILINK will arrange for initialization to occur during system startup by including an initialization table. Startup code calls an initialization routine that reads this table and performs the necessary initializations.

Zero-initialized sections are not affected by the `initialize` directive.

The `initialize` directive is normally used for initialized variables, but can be used for copying any sections, for example copying executable code from slow ROM to fast RAM, or for overlays. For another example, see *define overlay directive*, page 445.

Sections that are needed for initialization are not affected by the `initialize by copy` directive. This includes the `__low_level_init` function and anything it references.

Anything reachable from the program entry label is considered *needed for initialization* unless reached via a section fragment with a label starting with `__iar_init$$done`. The `--log sections` option, in addition to logging the marking of section fragments to be included in the application, also logs the process of determining which sections are needed for initialization.

Example

```
/* Copy all read-write sections automatically from ROM to RAM at
   program start */
initialize by copy { rw };
place in RAM { rw };
place in ROM { ro };
```

See also

Initialization at system startup, page 90, and *do not initialize directive*, page 449.

do not initialize directive

Syntax

```
do not initialize
{
    section-selectors
}
[except
{
    section-selectors
}];
```

For information about extended selectors and except clauses, see *Section selection*, page 452.

Description

Use the `do not initialize` directive to specify the sections that you do not want to be automatically zero-initialized by the system startup code. The directive can only be used on `zeroinit` sections.

Typically, this is useful if you want to handle zero-initialization in some other way for all or some `zeroinit` sections.

This can also be useful if you want to suppress zero-initialization of variables entirely. Normally, this is handled automatically for variables specified as `__no_init` in the source, but if you link with object files produced by older tools from IAR Systems or other tool vendors, you might need to suppress zero-initialization specifically for some sections.

Example

```
/* Do not initialize read-write sections whose name ends with
   _noinit at program start */
do not initialize { rw section .*_noinit };
place in RAM { rw section .*_noinit };
```

See also

Initialization at system startup, page 90, and *initialize directive*, page 446.

keep directive

Syntax

```
keep
{
    section-selectors
}
[except
{
    section-selectors
}];
```

For information about extended selectors and except clauses, see *Section selection*, page 452.

Description	The <code>keep</code> directive specifies that all selected sections should be kept in the executable image, even if there are no references to the sections.
Example	<pre>keep { section .keep* } except {section .keep};</pre>

place at directive

Syntax	<pre>["name":] place at { address [memory:] <i>expr</i> start of <i>region_expr</i> end of <i>region_expr</i> } { <i>extended-selectors</i> } [except { <i>section-selectors</i> }];</pre> <p>For information about extended selectors and except clauses, see <i>Section selection</i>, page 452.</p>
--------	--

Parameters	<table><tr><td><i>memory: expr</i></td><td>A specific address in a specific memory. The address must be available in the supplied memory defined by the <code>define memory</code> directive. The memory specifier is optional if there is only one memory.</td></tr><tr><td><i>start of region_expr</i></td><td>A region expression that results in a single-internal region. The start of the interval is used.</td></tr><tr><td><i>end of region_expr</i></td><td>A region expression that results in a single-internal region. The end of the interval is used.</td></tr></table>	<i>memory: expr</i>	A specific address in a specific memory. The address must be available in the supplied memory defined by the <code>define memory</code> directive. The memory specifier is optional if there is only one memory.	<i>start of region_expr</i>	A region expression that results in a single-internal region. The start of the interval is used.	<i>end of region_expr</i>	A region expression that results in a single-internal region. The end of the interval is used.
<i>memory: expr</i>	A specific address in a specific memory. The address must be available in the supplied memory defined by the <code>define memory</code> directive. The memory specifier is optional if there is only one memory.						
<i>start of region_expr</i>	A region expression that results in a single-internal region. The start of the interval is used.						
<i>end of region_expr</i>	A region expression that results in a single-internal region. The end of the interval is used.						

Description	<p>The <code>place at</code> directive places sections and blocks either at a specific address or, at the beginning or the end of a region. The same address cannot be used for two different <code>place at</code> directives. It is also not possible to use an empty region in a <code>place at</code> directive. If placed in a region, the sections and blocks will be placed before any other sections or blocks placed in the same region with a <code>place in</code> directive.</p> <p>The <i>name</i>, if specified, is used in the map file and in some log messages.</p>
-------------	--

Example

```
/* Place the read-only section .startup at the beginning of the
   code_region */
"START": place at start of ROM { readonly section .startup };
```

See also *place in directive*, page 451.

place in directive

Syntax

```
[ "name": ]
place in region-expr
{
    extended-selectors
}
[except{
    section-selectors
}];
```

where *region-expr* is a region expression, see also *Regions*, page 440.

and where the rest of the directive selects sections to include in the block. See *Section selection*, page 452.

Description

The `place in` directive places sections and blocks in a specific region. The sections and blocks will be placed in the region in an arbitrary order.

To specify a specific order, use the `block` directive. The region can have several ranges.

The *name*, if specified, is used in the map file and in some log messages.

Example

```
/* Place the read-only sections in the code_region */
"ROM": place in ROM { readonly };
```

See also *place at directive*, page 450.

use init table directive

Syntax

```
use init table name for
{
    extended-selectors
}
[except
{
    section-selectors
}];
```

For information about extended selectors and except clauses, see *Section selection*, page 452.

Parameters	<div><div><i>name</i></div><div>The name of the <code>init</code> table.</div></div>
Description	<div><p>Normally, all initialization entries are generated into a single initialization table (called <code>Table</code>). Use this directive to cause some of the entries to be put into a separate table. You can then use this initialization table at another time, or under different circumstances, than the normal initialization table.</p><p>Initialization entries for all variables not mentioned in a <code>use init table</code> directive are put into the normal initialization table. By having multiple <code>use init table</code> directives you can have multiple initialization tables.</p><p>The start, end, and size of the <code>init</code> table can be accessed in the application program by using <code>__section_begin</code>, <code>__section_end</code>, or <code>__section_size</code> of "<code>Region\$\$name</code>", respectively, or via the symbols <code>Region\$\$name\$\$Base</code>, <code>Region\$\$name\$\$Limit</code>, and <code>Region\$\$name\$\$Length</code>.</p></div>
Example	<div><pre>use init table Core2 for { section *.core2}; /* __section_begin("Region\$\$Core2") can be used to get the start of the Core2 init table. */</pre></div>

Section selection

The purpose of *section selection* is to specify—by means of *section selectors* and *except clauses*—the sections that an `ILINK` directive should be applied to. All sections that match one or more of the section selectors will be selected, and none of the sections selectors in the except clause, if any. Each section selector can match sections on section attributes, section name, and object or library name.

Some directives provide functionality that requires more detailed selection capabilities, for example directives that can be applied on both sections and blocks. In this case, the *extended-selectors* are used.

This section gives detailed information about each linker directive specific to section selection.

section-selectors

Syntax	<code>[section-selector [, section-selector...]]</code>
--------	--

```

section-selector is:

[ section-attribute ][ section-type ]
[ symbol symbol-name ][ section section-name ]
[ object {module|filename} ]

section-attribute is:

ro [ code|data ] | rw [ code|data ] | zi

section-type is:

[ preinit_array|init_array ]

```

Parameters

<i>section-attribute</i>	<p>Only sections with the specified attribute will be selected. <i>section-attribute</i> can consist of:</p> <p>ro readonly, for read-only sections. rw readwrite, for read/write sections. zi zeroinit, for zero-initialized sections. These sections have no content and should possibly be initialized with zeros during system startup. code, for sections that contain code. data, for sections that contain data.</p>
<i>section-type</i>	<p>Only sections with that ELF section type will be selected. <i>section-type</i> can be:</p> <p>preinit_array, sections of the ELF section type SHT_PREINIT_ARRAY. init_array, sections of the ELF section type SHT_INIT_ARRAY.</p>
<i>symbol symbol-name</i>	<p>Only sections that define at least one public symbol that matches the symbol name pattern will be selected. <i>symbol-name</i> is the symbol name pattern. Two wildcards are allowed:</p> <p>? matches any single character. * matches zero or more characters.</p>
<i>section section-name</i>	<p>Only sections whose names match the <i>section-name</i> will be selected. Two wildcards are allowed:</p> <p>? matches any single character * matches zero or more characters.</p>

`object module-spec` Only sections that originate from library modules or object files that matches `module-spec` will be selected.
`module-spec` can be in one of two forms:

`module`, a name in the form `objectname(libraryname)`. Sections from object modules where both the object name and the library name match their respective patterns are selected. An empty library name pattern selects only sections from object files.

`filename`, the name of an object file, or an object in a library.
Two wildcards are allowed:

? matches any single character

* matches zero or more characters.

Description

A section selector selects all sections that match the section attribute, section type, symbol name, section name, and the name of the module. Up to four of the five conditions can be omitted.

It is also possible to use only { } without any section selectors, which can be useful when defining blocks.

Note that a section selector with narrower scope has higher priority than a more generic section selector. If more than one section selector matches for the same purpose, one of them must be more specific. A section selector is more specific than another one if in priority order:

- It specifies a symbol name with no wildcards and the other one does not.
- It specifies a section name or object name with no wildcards and the other one does not
- It specifies a section type and the other one does not
- There could be sections that match the other selector that also match this one, and the reverse is not true.

Selector 1	Selector 2	More specific
ro	ro code	Selector 2
symbol mysym	section foo	Selector 1
ro code section f*	ro section f*	Selector 1
section foo*	section f*	Selector 1

Table 39: Examples of section selector specifications

Selector 1	Selector 2	More specific
section *x	section f*	Neither
init_array	section f*	Selector 1
section .intvec	ro section .int*	Selector 1
section .intvec	object foo.o	Neither

Table 39: Examples of section selector specifications (Continued)

Example

```
{ rw } /* Selects all read-write sections */

{ section .mydata* } /* Selects only .mydata* sections */
/* Selects .mydata* sections available in the object special.o */
{ section .mydata* object special.o }
```

Assuming a section in an object named `foo.o` in a library named `lib.a`, any of these selectors will select that section:

```
object foo.o(lib.a)
object f*(lib*)
object foo.o
object lib.a
```

See also

initialize directive, page 446, *do not initialize directive*, page 449, and *keep directive*, page 449.

extended-selectors

Syntax

```
{ [ extended-selector ] [ , extended-selector... ] }
```

where *extended-selector* is:

```
[ first | last | midway ]
    { section-selector |
      block name [ inline-block-def ] |
      overlay name }
```

where *inline-block-def* is:

```
[ block-params ] extended-selectors
```

Parameters

first	Places the selected sections, block, or overlay first in the containing placement directive, block, or overlay.
last	Places the selected sections, block or overlay last in the containing placement directive, block, or overlay.

	<code>midway</code>	Places the selected sections, block, or overlay so that they are no further than half the maximum size of the containing block away from either edge of the block. Note that this parameter can only be used inside a block that has a maximum size.
	<code>name</code>	The name of the block or overlay.
Description		<p>Use <code>extended-selectors</code> to select content for inclusion in a placement directive, block, or overlay. In addition to using section selection patterns, you can also explicitly specify blocks or overlays for inclusion.</p> <p>Using the <code>first</code> or <code>last</code> keyword, you can specify one pattern, block, or overlay that is to be placed first or last in the containing placement directive, block, or overlay. If you need more precise control of the placement order you can instead use a block with fixed order.</p> <p>Blocks can be defined separately, using the <code>define block</code> directive, or <code>inline</code>, as part of an <code>extended-selector</code>.</p> <p>The <code>midway</code> parameter is primarily useful together with a static base that can have both negative and positive offsets.</p>
Example		<pre>define block First { ro section .f* }; /* Define a block holding any read-only section*/ matching ".f*" */ define block Table { first block First, ro section .b }; /* Define a block where the block First comes before the sections matching ".b*". */</pre> <p>You can also define the block <code>First</code> inline, instead of in a separate <code>define block</code> directive:</p> <pre>define block Table { first block First { ro section .f* }, ro section .b* };</pre>
See also		<i>define block directive</i> , page 444, <i>define overlay directive</i> , page 445, and <i>place at directive</i> , page 450.

Using symbols, expressions, and numbers

In the linker configuration file, you can also:

- Define and export symbols

The `define symbol` directive defines a symbol with a specified value that can be used in expressions in the configuration file. The symbol can also be exported to be used by the application or the debugger. See *define symbol directive*, page 457, and *export directive*, page 458.

- Use expressions and numbers

In the linker configuration file, expressions and numbers are used for specifying addresses, sizes, etc. See *expressions*, page 459.

This section gives detailed information about each linker directive specific to defining symbols, expressions and numbers.

check that directive

Syntax	<code>check that <i>expression</i>;</code>							
Parameters	<i>expression</i>	A boolean expression.						
Description	<p>You can use the <code>check that</code> directive to compare the results of stack usage analysis against the sizes of blocks and regions. If the expression evaluates to zero, an error is emitted.</p> <p>Three extra operators are available for use only in <code>check that</code> expressions:</p> <table><tr><td><code>maxstack(<i>category</i>)</code></td><td>The stack depth of the deepest call chain for any call graph root function in the category.</td></tr><tr><td><code>totalstack(<i>category</i>)</code></td><td>The sum of the stack depths of the deepest call chains for each call graph root function in the category.</td></tr><tr><td><code>size(<i>block</i>)</code></td><td>The size of the block.</td></tr></table>		<code>maxstack(<i>category</i>)</code>	The stack depth of the deepest call chain for any call graph root function in the category.	<code>totalstack(<i>category</i>)</code>	The sum of the stack depths of the deepest call chains for each call graph root function in the category.	<code>size(<i>block</i>)</code>	The size of the block.
<code>maxstack(<i>category</i>)</code>	The stack depth of the deepest call chain for any call graph root function in the category.							
<code>totalstack(<i>category</i>)</code>	The sum of the stack depths of the deepest call chains for each call graph root function in the category.							
<code>size(<i>block</i>)</code>	The size of the block.							
Example	<pre>check that maxstack("Program entry") + totalstack("interrupt") + 1K <= size(block CSTACK);</pre>							
See also	<i>Stack usage analysis</i> , page 92.							

define symbol directive

Syntax	<code>define [<i>exported</i>] symbol <i>name</i> = <i>expr</i>;</code>
--------	---

Parameters	<div><div>exported</div><div>Exports the symbol to be usable by the executable image.</div></div> <div><div>name</div><div>The name of the symbol.</div></div> <div><div>expr</div><div>The symbol value.</div></div>
Description	<p>The <code>define symbol</code> directive defines a symbol with a specified value. The symbol can then be used in expressions in the configuration file. The symbols defined in this way work exactly like the symbols defined with the option <code>--config_def</code> outside of the configuration file.</p> <p>The <code>define exported symbol</code> variant of this directive is a shortcut for using the directive <code>define symbol</code> in combination with the <code>export symbol</code> directive. On the command line this would require both a <code>--config_def</code> option and a <code>--define_symbol</code> option to achieve the same effect.</p> <p>Note:</p> <ul style="list-style-type: none">• A symbol cannot be redefined• Symbols that are either prefixed by <code>_X</code>, where <i>X</i> is a capital letter, or that contain <code>__</code> (double underscore) are reserved for toolset vendors.
Example	<pre>/* Define the symbol my_symbol with the value 4 */ define symbol my_symbol = 4;</pre>
See also	<i>export directive</i> , page 458 and <i>Interaction between ILINK and the application</i> , page 110.

export directive

Syntax	<code>export symbol name;</code>
Parameters	<div><div>name</div><div>The name of the symbol.</div></div>
Description	<p>The <code>export</code> directive defines a symbol to be exported, so that it can be used both from the executable image and from a global label. The application, or the debugger, can then refer to it for setup purposes etc.</p>
Example	<pre>/* Define the symbol my_symbol to be exported */ export symbol my_symbol;</pre>

expressions

Syntax

An expression is built up of the following constituents:

```
expression binop expression
unop expression
expression ? expression : expression
(expression)
number
symbol
func-operator
```

where *binop* is one of these binary operators:

`+, -, *, /, %, <<, >>, <, >, ==, !=, &, ^, |, &&, ||`

where *unop* is one of this unary operators:

`+, -, !, ~`

where *number* is a number, see *numbers*, page 460

where *symbol* is a defined symbol, see *define symbol directive*, page 457 and *--config_def*, page 303

and where *func-operator* is one of these function-like operators:

<code>minimum(expr, expr)</code>	Returns the smallest of the two parameters.
<code>maximum(expr, expr)</code>	Returns the largest of the two parameters.
<code>isempty(r)</code>	Returns True if the region is empty, otherwise False.
<code>isdefinedsymbol(expr-symbol)</code>	Returns True if the expression symbol is defined, otherwise False.
<code>start(r)</code>	Returns the lowest address in the region.
<code>end(r)</code>	Returns the highest address in the region.
<code>size(r)</code>	Returns the size of the complete region.

where *expr* is an expression, and *r* is a region expression, see *Region expression*, page 441.

Description

In the linker configuration file, an expression is a 65-bit value with the range -2^{64} to 2^{64} . The expression syntax closely follows C syntax with some minor exceptions. There are no assignments, casts, pre- or post-operations, and no address operations (`*`, `&`, `[]`, `->`, and `.`). Some operations that extract a value from a region expression, etc, use

a syntax resembling that of a function call. A boolean expression returns 0 (False) or 1 (True).

numbers

Syntax

`nr [nr-suffix]`

where `nr` is either a decimal number or a hexadecimal number (0x... or 0X...).

and where `nr-suffix` is one of:

```
K          /* Kilo = (1 << 10) 1024 */
M          /* Mega = (1 << 20) 1048576 */
G          /* Giga = (1 << 30) 1073741824 */
T          /* Tera = (1 << 40) 1099511627776 */
P          /* Peta = (1 << 50) 1125899906842624 */
```

Description

A number can be expressed either by normal C means or by suffixing it with a set of useful suffixes, which provides a compact way of specifying numbers.

Example

1024 is the same as 0x400, which is the same as 1K.

Structural configuration

The structural directives provide means for creating structure within the configuration, such as:

- Conditional inclusion
An `if` directive includes or excludes other directives depending on a condition, which makes it possible to have directives for several different memory configurations in the same file. See *if directive*, page 461.
- Dividing the linker configuration file into several different files
The `include` directive makes it possible to divide the configuration file into several logically distinct files. See *include directive*, page 461.
- Signaling an error for unsupported cases

This section gives detailed information about each linker directive specific to structural configuration.

error directive

Syntax

`error string`

Parameters	<i>string</i>	The error message.
Description	An <code>error</code> directive can be used for signaling an error if the directive occurs in the active part of a conditional directive.	
Example	<code>error "Unsupported configuration"</code>	

if directive

Syntax	<pre> if (<i>expr</i>) { <i>directives</i> [} else if (<i>expr</i>) { <i>directives</i>] [} else { <i>directives</i>] } </pre> <p>where <i>expr</i> is an expression, see <i>expressions</i>, page 459.</p>	
Parameters	<i>directives</i>	Any ILINK directive.
Description	<p>An <code>if</code> directive includes or excludes other directives depending on a condition, which makes it possible to have directives for several different memory configurations, for example both a banked and non-banked memory configuration, in the same file.</p> <p>The directives inside an <code>if</code> part, <code>else if</code> part, or an <code>else</code> part are syntax checked and processed regardless of whether the conditional expression was true or false, but only the directives in the part where the conditional expression was true, or the <code>else</code> part if none of the conditions were true, will have any effect outside the <code>if</code> directive. The <code>if</code> directives can be nested.</p>	
Example	See <i>Empty region</i> , page 442.	

include directive

Syntax	<code>include "<i>filename</i>";</code>	
Parameters	<i>filename</i>	A path where both / and \ can be used as the directory delimiter.

Description	<p>The <code>include</code> directive makes it possible to divide the configuration file into several logically distinct parts, each in a separate file. For instance, there might be parts that you need to change often and parts that you seldom edit.</p> <p>Normally, the linker searches for configuration include files in the system configuration directory. You can use the <code>--config_search</code> linker option to add more directories to search.</p>
See also	<i>--config_search</i> , page 303

Section reference

- Summary of sections
- Descriptions of sections and blocks

For more information about sections, see the chapter *Modules and sections*, page 84.

Summary of sections

This table lists the ELF sections and blocks that are used by the IAR build tools:

Section	Description
<code>.bss</code>	Holds zero-initialized static and global variables.
<code>CSTACK</code>	Holds the stack used by C or C++ programs.
<code>.data</code>	Holds static and global initialized variables.
<code>.data_init</code>	Holds initial values for <code>.data</code> sections when the linker directive <code>initialize</code> is used.
<code>__DLIB_PERTHREAD</code>	Holds variables that contain static states for DLIB modules.
<code>.exc.text</code>	Holds exception-related code.
<code>HEAP</code>	Holds the heap used for dynamically allocated data.
<code>.iar.dynexit</code>	Holds the <code>atexit</code> table.
<code>.init_array</code>	Holds a table of dynamic initialization functions.
<code>.intvec</code>	Holds the reset vector table
<code>IRQ_STACK</code>	Holds the stack for interrupt requests, IRQ, and exceptions.
<code>.noinit</code>	Holds <code>__no_init</code> static and global variables.
<code>.preinit_array</code>	Holds a table of dynamic initialization functions.
<code>.prepreinit_array</code>	Holds a table of dynamic initialization functions.
<code>.rodata</code>	Holds constant data.
<code>.text</code>	Holds the program code.
<code>.textrw</code>	Holds <code>__ramfunc</code> declared program code.
<code>.textrw_init</code>	Holds initializers for the <code>.textrw</code> declared section.

Table 40: Section summary

In addition to the ELF sections used for your application, the tools use a number of other ELF sections for a variety of purposes:

- Sections starting with `.debug` generally contain debug information in the DWARF format
- Sections starting with `.iar.debug` contain supplemental debug information in an IAR format
- The section `.comment` contains the tools and command lines used for building the file
- Sections starting with `.rel` or `.rela` contain ELF relocation information
- The section `.symtab` contains the symbol table for a file
- The section `.strtab` contains the names of the symbol in the symbol table
- The section `.shstrtab` contains the names of the sections.

Descriptions of sections and blocks

This section gives reference information about each section, where the:

- *Description* describes what type of content the section is holding and, where required, how the section is treated by the linker
- *Memory placement* describes memory placement restrictions.

For information about how to allocate sections in memory by modifying the linker configuration file, see *Placing code and data—the linker configuration file*, page 87.

.bss

Description	Holds zero-initialized static and global variables.
Memory placement	This section can be placed anywhere in memory.

CSTACK

Description	Block that holds the internal data stack.
Memory placement	This block can be placed anywhere in memory.
See also	<i>Setting up stack memory</i> , page 105.

.data

Description	Holds static and global initialized variables. In object files, this includes the initial values. When the linker directive <code>initialize</code> is used, a corresponding <code>.data_init</code> section is created for each <code>.data</code> section, holding the possibly compressed initial values.
Memory placement	This section can be placed anywhere in memory.

.data_init

Description	Holds the possibly compressed initial values for <code>.data</code> sections. This section is created by the linker if the <code>initialize</code> linker directive is used.
Memory placement	This section can be placed anywhere in memory.

__DLIB_PERTHREAD

Description	Holds thread-local static and global initialized variables used by the main thread. This section is placed automatically. If you change the placement, you must not change its initialization. The initialization of this section must be controlled with the <code>initialize</code> directive.
Memory placement	This section can be placed anywhere in memory.
See also	<i>Managing a multithreaded environment</i> , page 150.

.exc.text

Description	Holds code that is only executed when your application handles an exception.
Memory placement	In the same memory as <code>.text</code> .
See also	<i>Exception handling</i> , page 194.

HEAP

Description	Holds the heap used for dynamically allocated data in memory, in other words data allocated by <code>malloc</code> and <code>free</code> , and in C++, <code>new</code> and <code>delete</code> .
-------------	---

Memory placement	This section can be placed anywhere in memory.
See also	<i>Setting up heap memory</i> , page 106.

.iar.dynexit

Description	Holds the table of calls to be made at exit.
Memory placement	This section can be placed anywhere in memory.
See also	<i>Setting up the atexit limit</i> , page 106.

.init_array

Description	Holds pointers to routines to call for initializing one or more C++ objects with static storage duration.
Memory placement	This section can be placed anywhere in memory.

.intvec

Description	Holds the reset vector table and exception vectors which contain branch instructions to <i>cstartup</i> , interrupt service routines etc.
Memory placement	Must be placed at address range 0x00 to 0x3F.

IRQ_STACK

Description	Holds the stack which is used when servicing IRQ exceptions. Other stacks may be added as needed for servicing other exception types: FIQ, SVC, ABT, and UND. The <i>cstartup.s</i> file must be modified to initialize the exception stack pointers used. Note: This section is not used when compiling for Cortex-M.
Memory placement	This section can be placed anywhere in memory.
See also	<i>Exception stack</i> , page 204

.noinit

Description	Holds static and global <code>__no_init</code> variables.
Memory placement	This section can be placed anywhere in memory.

.preinit_array

Description	Like <code>.init_array</code> , but is used by the library to make some C++ initializations happen before the others.
Memory placement	This section can be placed anywhere in memory.
See also	<code>.init_array</code> , page 466.

.prepreinit_array

Description	Like <code>.init_array</code> , but is used when C static initialization is rewritten as dynamic initialization. Performed before all C++ dynamic initialization.
Memory placement	This section can be placed anywhere in memory.
See also	<code>.init_array</code> , page 466.

.rodata

Description	Holds constant data. This can include constant variables, string and aggregate literals, etc.
Memory placement	This section can be placed anywhere in memory.

.text

Description	Holds program code, including the code for system initialization.
Memory placement	This section can be placed anywhere in memory.

.textrw

Description	Holds __ramfunc declared program code.
Memory placement	This section can be placed anywhere in memory.
See also	__ramfunc, page 353.

.textrw_init

Description	Holds initializers for the .textrw declared sections.
Memory placement	This section can be placed anywhere in memory.
See also	__ramfunc, page 353.

The stack usage control file

- Overview
- Stack usage control directives
- Syntactic components

Before you read this chapter, see *Stack usage analysis*, page 92.

Overview

A stack usage control file consists of a sequence of directives that control stack usage analysis. You can use C ("*/...**") and C++ ("*//...*") comments in these files.

The default filename extension for stack usage control files is `.suc`.

C++ NAMES

When you specify the name of a C++ function in a stack usage control file, you must use the name exactly as used by the linker. Both the number and names of parameters, as well as the names of types must match. However, most non-significant white-space differences are accepted. In particular, you must enclose the name in quote marks because all C++ function names include non-identifier characters.

You can also use wildcards in function names. "*##*" matches any sequence of characters, and "*##*" matches a single character. This makes it possible to write function names that will match any instantiation of a template function.

Examples:

```
"operator new(unsigned int) "  
"ostream::flush()" // EC++  
"std::ostream::flush()" // C++  
"operator <<(ostream &, char const *)" "  
"void _Sort<##>(*, ##, *)" "
```

Stack usage control directives

This section gives detailed reference information about each stack usage control directive.

call graph root directive

Syntax	<code>call graph root [category] : func-spec [, func-spec...] ;</code>	
Parameters	<i>category</i>	See <i>category</i> , page 473
	<i>func-spec</i>	See <i>func-spec</i> , page 473
Description	<p>Specifies that the listed functions are call graph roots. You can optionally specify a call graph root category. Call graph roots are listed under their category in the <i>Stack Usage</i> chapter in the linker map file.</p> <p>The linker will normally issue a warning for functions needed in the application that are not call graph roots and which do not appear to be called.</p>	
Example	<code>call graph root [task]: MyFunc10, MyFunc11;</code>	
See also	<i>call_graph_root</i> , page 363.	

exclude directive

Syntax	<code>exclude func-spec [, func-spec...] ;</code>	
Parameters	<i>func-spec</i>	See <i>func-spec</i> , page 473
Description	Excludes the specified functions, and call trees originating with them, from stack usage calculations.	
Example	<code>exclude MyFunc5, MyFunc6;</code>	

function directive

Syntax	<code>[override] function [category] func-spec : stack-size [, call-info...] ;</code>	
Parameters	<i>category</i>	See <i>category</i> , page 473
	<i>func-spec</i>	See <i>func-spec</i> , page 473
	<i>call-info</i>	See <i>call-info</i> , page 474

	<i>stack-size</i>	See <i>stack-size</i> , page 474
Description	<p>Specifies what the maximum stack usage is in a function and which other functions that are called from that function.</p> <p>Normally, an error is issued if there already is stack usage information for the function, but if you start with <code>override</code>, the error will be suppressed and the information supplied in the directive will be used instead of the previous information.</p>	
Example	<pre>function MyFunc1: 32, calls MyFunc2, calls MyFunc3, MyFunc4: 16; function [interrupt] MyInterruptHandler: 44;</pre>	

max recursion depth directive

Syntax	<code>max recursion depth <i>func-spec</i> : <i>size</i>;</code>	
Parameters	<i>func-spec</i>	See <i>func-spec</i> , page 473
	<i>size</i>	See <i>size</i> , page 475
Description	<p>Specifies the maximum number of iterations through any of the cycles in the recursion nest of which the function is a member.</p> <p>A recursion nest is a set of cycles in the call graph where each cycle shares at least one node with another cycle in the nest.</p> <p>Stack usage analysis will base its result on the max recursion depth multiplied by the stack usage of the deepest cycle in the nest. If the nest is not entered on a point along one of the deepest cycles, no stack usage result will be calculated for such calls.</p>	
Example	<code>max recursion depth MyFunc12: 10;</code>	

no calls from directive

Syntax	<code>no calls from <i>module-spec</i> to <i>func-spec</i> [<i>,</i> <i>func-spec</i>... <i>]</i>;</code>	
Parameters	<i>func-spec</i>	See <i>func-spec</i> , page 473

	<i>module-spec</i>	See <i>module-spec</i> , page 473
Description	<p>When you provide stack usage information for some functions in a module without stack usage information, the linker warns about functions that are referenced from the module but not listed as called. This is primarily to help avoid problems with C runtime routines, calls to which are generated by the compiler, beyond user control.</p> <p>If there actually is no call to some of these functions, use the <code>no calls from</code> directive to selectively suppress the warning for the specified functions. You can also disable the warning entirely (<code>--diag_suppress</code> or Project>Options>Linker>Diagnostics>Suppress these diagnostics).</p>	
Example	<pre>no calls from [file.o] to MyFunc13, MyFunc14;</pre>	

possible calls directive

Syntax	<pre>possible calls <i>calling-func</i> : <i>called-func</i> [, <i>called-func</i>...];</pre>	
Parameters	<i>calling-func</i>	See <i>func-spec</i> , page 473
	<i>called-func</i>	See <i>func-spec</i> , page 473
Description	<p>Specifies an exhaustive list of possible destinations for all indirect calls in one function. Use this for functions which are known to perform indirect calls and where you know exactly which functions that might be called in this particular application. Consider using the <code>#pragma calls</code> directive if the information about which functions that might be called is available when compiling.</p>	
Example	<pre>possible calls MyFunc7: MyFunc8, MyFunc9;</pre>	
	<p>When the function does not perform any calls, the list is empty:</p> <pre>possible calls MyFunc8: ;</pre>	
See also	<i>calls</i> , page 362.	

Syntactic components

This section describes the syntactical components that can be used by the stack usage control directives.

category

Syntax	[<i>name</i>]
Description	A call graph root category. You can use any name you like. Categories are not case-sensitive.
Example	category examples: [interrupt] [task]

func-spec

Syntax	[?] <i>name</i> [<i>module-spec</i>]
Description	Specifies the name of a symbol, and for module-local symbols, the name of the module it is defined in. Normally, if <i>func-spec</i> does not match a symbol in the program, a warning is emitted. Prefixing with ? suppresses this warning.
Example	<i>func-spec</i> examples: xFun MyFun [file.o] ?"fun1(int) "

module-spec

Syntax	[<i>name</i> [(<i>name</i>)]]
Description	<p>Specifies the name of a module, and optionally, in parentheses, the name of the library it belongs to. To distinguish between modules with the same name, you can specify:</p> <ul style="list-style-type: none"> • The complete path of the file ("D:\C1\test\file.o") • As many path elements as are needed at the end of the path ("test\file.o") • Some path elements at the start of the path, followed by "...", followed by some path elements at the end ("D:\...\file.o"). <p>Note that when using multi-file compilation (<code>--mfc</code>), multiple files are compiled into a single module, named after the first file.</p>

Example	<i>module-spec</i> examples: [file.o] [file.o(lib.a)] ["D:\C1\test\file.o"]
---------	--

name

Description	<p>A name can be either an identifier or a quoted string.</p> <p>The first character of an identifier must be either a letter or one of the characters "_", "\$", or ".". The rest of the characters can also be digits.</p> <p>A quoted string starts and ends with " and can contain any character. Two consecutive " characters can be used inside a quoted string to represent a single ".</p>
-------------	--

Example	<i>name</i> examples: MyFun file.o "file-1.o"
---------	--

call-info

Syntax	<i>calls func-spec [, func-spec...] [: stack-size]</i>
Description	Specifies one or more called functions, and optionally, the stack size at the calls.
Example	<i>call-info</i> examples: calls MyFunc1 : stack 16 calls MyFunc2, MyFunc3, MyFunc4

stack-size

Syntax	<i>[stack] size</i>
Description	Specifies the size of a stack frame.
Example	<i>stack-size</i> examples: 24 stack 28

size

Description	A decimal integer, or 0x followed by a hexadecimal integer. Either alternative can optionally be followed by a suffix indicating a power of two (K=2 ¹⁰ , M=2 ²⁰ , G=2 ³⁰ , T=2 ⁴⁰ , P=2 ⁵⁰).
Example	<div>size examples: 24 0x18 2048 2K</div>

IAR utilities

- The IAR Archive Tool—`iarchive`—creates and manipulates a library (an archive) of several ELF object files
- The IAR ELF Tool—`ielftool`—performs various transformations on an ELF executable image (such as fill, checksum, format conversions, etc)
- The IAR ELF Dumper—`ielfdump`—creates a text representation of the contents of an ELF relocatable or executable image
- The IAR ELF Object Tool—`iobjmanip`—is used for performing low-level manipulation of ELF object files
- The IAR Absolute Symbol Exporter—`ismexport`—exports absolute symbols from a ROM image file, so that they can be used when you link an add-on application.

The IAR Archive Tool—`iarchive`

The IAR Archive Tool, `iarchive`, can create a library (an archive) file from several ELF object files. You can also use `iarchive` to manipulate ELF libraries.

A library file contains several relocatable ELF object modules, each of which can be independently used by a linker. In contrast with object modules specified directly to the linker, each module in a library is only included if it is needed.

For information about how to build a library in the IDE, see the *IDE Project Management and Building Guide for ARM*.

INVOCATION SYNTAX

The invocation syntax for the archive builder is:

```
iarchive parameters
```

Parameters

The parameters are:

Parameter	Description
<i>command</i>	Command line options that define an operation to be performed. Such an option must be specified before the name of the library file.
<i>libraryfile</i>	The library file to be operated on.
<i>objectfile1 ... objectfileN</i>	The object file(s) that the specified command operates on.
<i>options</i>	Command line options that define actions to be performed. These options can be placed anywhere on the command line.

Table 41: iarchive parameters

Examples

This example creates a library file called `mylibrary.a` from the source object files `module1.o`, `module2.o`, and `module3.o`:

```
iarchive mylibrary.a module1.o module2.o module3.o.
```

This example lists the contents of `mylibrary.a`:

```
iarchive --toc mylibrary.a
```

This example replaces `module3.o` in the library with the content in the `module3.o` file and appends `module4.o` to `mylibrary.a`:

```
iarchive --replace mylibrary.a module3.o module4.o
```

SUMMARY OF IARCHIVE COMMANDS

This table summarizes the `iarchive` commands:

Command line option	Description
<code>--create</code>	Creates a library that contains the listed object files.
<code>--delete, -d</code>	Deletes the listed object files from the library.
<code>--extract, -x</code>	Extracts the listed object files from the library.
<code>--replace, -r</code>	Replaces or appends the listed object files to the library.
<code>--symbols</code>	Lists all symbols defined by files in the library.
<code>--toc, -t</code>	Lists all files in the library.

Table 42: iarchive commands summary

For more information, see *Descriptions of options*, page 492.

SUMMARY OF IARCHIVE OPTIONS

This table summarizes the `iarchive` options:

Command line option	Description
<code>-f</code>	Extends the command line.
<code>--output, -o</code>	Specifies the library file.
<code>--verbose, -V</code>	Reports all performed operations.
<code>--version</code>	Sends tool output to the console and then exits.

Table 43: *iarchive* options summary

For more information, see *Descriptions of options*, page 492.

DIAGNOSTIC MESSAGES

This section lists the messages produced by `iarchive`:

La001: could not open file *filename*

`iarchive` failed to open an object file.

La002: illegal path *pathname*

The path *pathname* is not a valid path.

La006: too many parameters to *cmd* command

A list of object modules was specified as parameters to a command that only accepts a single library file.

La007: too few parameters to *cmd* command

A command that takes a list of object modules was issued without the expected modules.

La008: *lib* is not a library file

The library file did not pass a basic syntax check. Most likely the file is not the intended library file.

La009: *lib* has no symbol table

The library file does not contain the expected symbol information. The reason might be that the file is not the intended library file, or that it does not contain any ELF object modules.

La010: no library parameter given

The tool could not identify which library file to operate on. The reason might be that a library file has not been specified.

La011: file *file* already exists

The file could not be created because a file with the same name already exists.

La013: file confusions, *lib* given as both library and object

The library file was also mentioned in the list of object modules.

La014: module *module* not present in archive *lib*

The specified object module could not be found in the archive.

La015: internal error

The invocation triggered an unexpected error in `iarchive`.

Ms003: could not open file *filename* for writing

`iarchive` failed to open the archive file for writing. Make sure that it is not write protected.

Ms004: problem writing to file *filename*

An error occurred while writing to file *filename*. A possible reason for this is that the volume is full.

Ms005: problem closing file *filename*

An error occurred while closing the file *filename*.

The IAR ELF Tool—ielftool

The IAR ELF Tool, `ielftool`, can generate a checksum on specific ranges of memories. This checksum can be compared with a checksum calculated on your application.

The source code for `ielftool` and a Microsoft VisualStudio 2005 template project are available in the `arm\src\elfutils` directory. If you have specific requirements for how the checksum should be generated or requirements for format conversion, you can modify the source code accordingly.

INVOCATION SYNTAX

The invocation syntax for the IAR ELF Tool is:

```
ielftool [options] inputfile outputfile [options]
```

The ielftool tool will first process all the fill options, then it will process all the checksum options (from left to right).

Parameters

The parameters are:

Parameter	Description
<i>inputfile</i>	An absolute ELF executable image produced by the ILINK linker.
<i>options</i>	Any of the available command line options, see <i>Summary of ielftool options</i> , page 481.
<i>outputfile</i>	An absolute ELF executable image.

Table 44: ielftool parameters

See also *Rules for specifying a filename or directory as parameters*, page 254.

Example

This example fills a memory range with 0xFF and then calculates a checksum on the same range:

```
ielftool my_input.out my_output.out --fill 0xFF;0-0xFF
--checksum __checksum:4,crc32;0-0xFF
```

SUMMARY OF IELFTOOL OPTIONS

This table summarizes the ielftool command line options:

Command line option	Description
--bin	Sets the format of the output file to binary.
--checksum	Generates a checksum.
--fill	Specifies fill requirements.
--ihex	Sets the format of the output file to linear Intel hex.
--offset	Adds (or subtracts) an offset to all addresses in the generated output file.
--parity	Generates parity bits.
--self_reloc	Not for general use.
--silent	Sets silent operation.

Table 45: ielftool options summary

Command line option	Description
--simple	Sets the format of the output file to Simple code.
--simple-ne	As --simple, but without an entry record.
--srec	Sets the format of the output file to Motorola S-records.
--srec-len	Restricts the number of data bytes in each S-record.
--srec-s3only	Restricts the S-record output to contain only a subset of records.
--strip	Removes debug information.
--titxt	Saves as TI-txt format.
--verbose, -V	Prints all performed operations.

Table 45: ielftool options summary (Continued)

For more information, see *Descriptions of options*, page 492.

The IAR ELF Dumper—ielfdump

The IAR ELF Dumper for ARM, `ielfdumparm`, can be used for creating a text representation of the contents of a relocatable or absolute ELF file.

`ielfdumparm` can be used in one of three ways:

- To produce a listing of the general properties of the input file and the ELF segments and ELF sections it contains. This is the default behavior when no command line options are used.
- To also include a textual representation of the contents of each ELF section in the input file. To specify this behavior, use the command line option `--all`.
- To produce a textual representation of selected ELF sections from the input file. To specify this behavior, use the command line option `--section`.

INVOCATION SYNTAX

The invocation syntax for `ielfdumparm` is:

```
ielfdumparm input_file [output_file]
```

Note: `ielfdumparm` is a command line tool which is not primarily intended to be used in the IDE.

Parameters

The parameters are:

Parameter	Description
<i>input_file</i>	An ELF relocatable or executable file to use as input.
<i>output_file</i>	A file or directory where the output is emitted. If absent and no <code>--output</code> option is specified, output is directed to the console.

Table 46: *ielfdumparm* parameters

See also *Rules for specifying a filename or directory as parameters*, page 254.

SUMMARY OF IELFDUMP OPTIONS

This table summarizes the `ielfdumparm` command line options:

Command line option	Description
<code>--a</code>	Generates output for all sections except string table sections.
<code>--all</code>	Generates output for all input sections regardless of their names or numbers.
<code>--code</code>	Dumps all sections that contain executable code.
<code>-f</code>	Extends the command line.
<code>--output, -o</code>	Specifies an output file.
<code>--no_header</code>	Suppresses production of a list header in the output.
<code>--no_rel_section</code>	Suppresses dumping of <code>.rel/.rela</code> sections.
<code>--no_strtab</code>	Suppresses dumping of string table sections.
<code>--range</code>	Disassembles only addresses in the specified range.
<code>--raw</code>	Uses the generic hexadecimal/ASCII output format for the contents of any selected section, instead of any dedicated output format for that section.
<code>--section, -s</code>	Generates output for selected input sections.
<code>--segment, -g</code>	Generates output for segments with specified numbers.
<code>--source</code>	Includes source with disassembled code in executable files.
<code>--use_full_std_t emplate_names</code>	Uses full short full names for some Standard C++ templates.
<code>--version</code>	Sends tool output to the console and then exits.

Table 47: *ielfdumparm* options summary

For more information, see *Descriptions of options*, page 492.

The IAR ELF Object Tool—iobjmanip

Use the IAR ELF Object Tool, `iobjmanip`, to perform low-level manipulation of ELF object files.

INVOCATION SYNTAX

The invocation syntax for the IAR ELF Object Tool is:

```
iobjmanip options inputfile outputfile
```

Parameters

The parameters are:

Parameter	Description
<i>options</i>	Command line options that define actions to be performed. These options can be placed anywhere on the command line. At least one of the options must be specified.
<i>inputfile</i>	A relocatable ELF object file.
<i>outputfile</i>	A relocatable ELF object file with all the requested operations applied.

Table 48: *iobjmanip* parameters

See also *Rules for specifying a filename or directory as parameters*, page 254.

Examples

This example renames the section `.example` in `input.o` to `.example2` and stores the result in `output.o`:

```
iobjmanip --rename_section .example=.example2 input.o output.o
```

SUMMARY OF IOBJMANIP OPTIONS

This table summarizes the `iobjmanip` options:

Command line option	Description
<code>-f</code>	Extends the command line.
<code>--remove_file_path</code>	Removes path information from the file symbol.
<code>--remove_section</code>	Removes one or more section.
<code>--rename_section</code>	Renames a section.
<code>--rename_symbol</code>	Renames a symbol.
<code>--silent</code>	Sets silent operation.
<code>--strip</code>	Removes debug information.

Table 49: *iobjmanip* options summary

Command line option	Description
<code>--version</code>	Sends tool output to the console and then exits.

Table 49: *iobjmanip* options summary (Continued)

For more information, see *Descriptions of options*, page 492.

DIAGNOSTIC MESSAGES

This section lists the messages produced by *iobjmanip*:

Lm001: No operation given

None of the command line parameters specified an operation to perform.

Lm002: Expected *nr* parameters but got *nr*

Too few or too many parameters. Check invocation syntax for *iobjmanip* and for the used command line options.

Lm003: Invalid section/symbol renaming pattern *pattern*

The pattern does not define a valid renaming operation.

Lm004: Could not open file *filename*

iobjmanip failed to open the input file.

Lm005: ELF format error *msg*

The input file is not a valid ELF object file.

Lm006: Unsupported section type *nr*

The object file contains a section that *iobjmanip* cannot handle. This section will be ignored when generating the output file.

Lm007: Unknown section type *nr*

iobjmanip encountered an unrecognized section. *iobjmanip* will try to copy the content as is.

Lm008: Symbol *symbol* has unsupported format

iobjmanip encountered a symbol that cannot be handled. *iobjmanip* will ignore this symbol when generating the output file.

Lm009: Group type *nr* not supported

`iobjmanip` only supports groups of type `GRP_COMDAT`. If any other group type is encountered, the result is undefined.

Lm010: Unsupported ELF feature in file: *msg*

The input file uses a feature that `iobjmanip` does not support.

Lm011: Unsupported ELF file type

The input file is not a relocatable object file.

Lm012: Ambiguous rename for section/symbol name (*alt1* and *alt2*)

An ambiguity was detected while renaming a section or symbol. One of the alternatives will be used.

Lm013: Section *name* removed due to transitive dependency on *name*

A section was removed as it depends on an explicitly removed section.

Lm014: File has no section with index *nr*

A section index, used as a parameter to `--remove_section` or `--rename_section`, did not refer to a section in the input file.

Ms003: could not open file *filename* for writing

`iobjmanip` failed to open the output file for writing. Make sure that it is not write protected.

Ms004: problem writing to file *filename*

An error occurred while writing to file *filename*. A possible reason for this is that the volume is full.

Ms005: problem closing file *filename*

An error occurred while closing the file *filename*.

The IAR Absolute Symbol Exporter—`isymexport`

The IAR Absolute Symbol Exporter, `isymexport`, can export absolute symbols from a ROM image file, so that they can be used when you link an add-on application.

To keep symbols from your symbols file in your final application, the symbols must be referred to, either from your source code or by using the linker option `--keep`.

INVOCATION SYNTAX

The invocation syntax for the IAR Absolute Symbol Exporter is:

```
isymexport [options] inputfile outputfile [options]
```

Parameters

The parameters are:

Parameter	Description
<i>inputfile</i>	A ROM image in the form of an executable ELF file (output from linking).
<i>options</i>	Any of the available command line options, see <i>Summary of isymexport options</i> , page 488.
<i>outputfile</i>	A relocatable ELF file that can be used as input to linking, and which contains all or a selection of the absolute symbols in the input file. The output file contains only the symbols, not the actual code or data sections. A steering file can be used to control which symbols that are included, and also to rename some of the symbols if that is desired.

Table 50: isymexport parameters

See also *Rules for specifying a filename or directory as parameters*, page 254.



In the IDE, to add the export of library symbols, choose **Project>Options>Build Actions** and specify your command line in the **Post-build command line** text field, for example like this:

```
$TOOLKIT_DIR$\bin\isymexport.exe "$TARGET_PATH$"  
"$PROJ_DIR$\const_lib.symbols"
```

SUMMARY OF ISYMEXPORT OPTIONS

This table summarizes the ismexport command line options:

Command line option	Description
--edit	Specifies a steering file.
-f	Extends the command line.
--generate_vfe_header	Declares that the image does not contain any virtual function calls to potentially discarded functions.
--reserve_ranges	Generates symbols to reserve the areas in ROM and RAM that the image uses.
--show_entry_as	Exports the entry point of the application with the given name.
--version	Sends tool output to the console and then exits.

Table 51: ismexport options summary

For more information, see *Descriptions of options*, page 492.

STEERING FILES

A steering file can be used for controlling which symbols that are included, and also to rename some of the symbols if that is desired. In the file, you can use `show` and `hide` directives to select which public symbols from the input file that are to be included in the output file. `rename` directives can be used for changing the names of symbols in the input file.

When you use a steering file, only actively exported symbols will be available in the output file. Thus, a steering file without `show` directives will generate an output file without symbols.

Syntax

The following syntax rules apply:

- Each directive is specified on a separate line.
- C comments (`/* . . . */`) and C++ comments (`// . . .`) can be used.
- Patterns can contain wildcard characters that match more than one possible character in a symbol name.
- The `*` character matches any sequence of zero or more characters in a symbol name.
- The `?` character matches any single character in a symbol name.

Example

```
rename xxx_* as YY_* /*Change symbol prefix from xxx_ to YY_* */
show YY_*           /* Export all symbols from YY package */
hide *_internal     /* But do not export internal symbols */
show zzz?           /* Export zzza, but not zzzaaa */
hide zzzx           /* But do not export zzzx */
```

Show directive

Syntax	<code>show <i>pattern</i></code>
Parameters	<i>pattern</i> A pattern to match against a symbol name.
Description	A symbol with a name that matches the pattern will be included in the output file unless this is overridden by a later <code>hide</code> directive.
Example	<pre>/* Include all public symbols ending in _pub. */ show *_pub</pre>

Hide directive

Syntax	<code>hide <i>pattern</i></code>
Parameters	<i>pattern</i> A pattern to match against a symbol name.
Description	A symbol with a name that matches the pattern will not be included in the output file unless this is overridden by a later <code>show</code> directive.
Example	<pre>/* Do not include public symbols ending in _sys. */ hide *_sys</pre>

Rename directive

Syntax	<code>rename <i>pattern1</i> as <i>pattern2</i></code>
Parameters	<i>pattern1</i> A pattern used for finding symbols to be renamed. The pattern can contain no more than one <code>*</code> or <code>?</code> wildcard character.

pattern2 A pattern used for the new name for a symbol. If the pattern contains a wildcard character, it must be of the same kind as in *pattern1*.

Description

Use this directive to rename symbols from the output file to the input file. No exported symbol is allowed to match more than one `rename` pattern.

`rename` directives can be placed anywhere in the steering file, but they are executed before any `show` and `hide` directives. Thus, if a symbol will be renamed, all `show` and `hide` directives in the steering file must refer to the new name.

If the name of a symbol matches a *pattern1* pattern that contains no wildcard characters, the symbol will be renamed *pattern2* in the output file.

If the name of a symbol matches a *pattern1* pattern that contains a wildcard character, the symbol will be renamed *pattern2* in the output file, with part of the name matching the wildcard character preserved.

Example

```
/* xxx_start will be renamed Y_start_X in the output file,
   xxx_stop will be renamed Y_stop_X in the output file. */
rename xxx_* as Y_*_X
```

DIAGNOSTIC MESSAGES

This section lists the messages produced by `isymexport`:

Es001: could not open file *filename*

`isymexport` failed to open the specified file.

Es002: illegal path *pathname*

The path *pathname* is not a valid path.

Es003: format error: *message*

A problem occurred while reading the input file.

Es004: no input file

No input file was specified.

Es005: no output file

An input file, but no output file was specified.

Es006: too many input files

More than two files were specified.

Es007: input file is not an ELF executable

The input file is not an ELF executable file.

Es008: unknown directive: *directive*

The specified directive in the steering file is not recognized.

Es009: unexpected end of file

The steering file ended when more input was required.

Es010: unexpected end of line

A line in the steering file ended before the directive was complete.

Es011: unexpected text after end of directive

There is more text on the same line after the end of a steering file directive.

Es012: expected text

The specified text was not present in the steering file, but must be present for the directive to be correct.

Es013: pattern can contain at most one * or ?

Each pattern in the current directive can contain at most one * or one ? wildcard character.

Es014: rename patterns have different wildcards

Both patterns in the current directive must contain exactly the same kind of wildcard. That is, both must either contain:

- No wildcards
- Exactly one *
- Exactly one ?

This error occurs if the patterns are not the same in this regard.


Es014: ambiguous pattern match: *symbol* matches more than one rename pattern

A symbol in the input file matches more than one `rename` pattern.


Descriptions of options

This section gives detailed reference information about each command line option available for the different utilities.

--a

Syntax	--a
For use with	ielfdumparm
Description	Use this option as a shortcut for <code>--all --no_strtab</code> .
	 This option is not available in the IDE.

--all

Syntax	--all
For use with	ielfdumparm
Description	Use this option to include the contents of all ELF sections in the output, in addition to the general properties of the input file. Sections are output in index order, except that each relocation section is output immediately after the section it holds relocations for. By default, no section contents are included in the output.
	 This option is not available in the IDE.

--bin

Syntax	--bin
For use with	ielftool

Description Sets the format of the output file to binary.



To set related options, choose:

Project>Options>Output converter

--checksum

Syntax

```
--checksum
{symbol[+offset] | address}:size,algorithm[:[1|2][a|m|z][L|W][r][R]
[o][i|p]]
[,start];range[;range...]
```

Parameters

<i>symbol</i>	The name of the symbol where the checksum value should be stored. Note that it must exist in the symbol table in the input ELF file.
<i>offset</i>	An offset to the symbol.
<i>address</i>	The absolute address where the checksum value should be stored.
<i>size</i>	The number of bytes in the checksum: 1, 2, or 4; must not be larger than the size of the checksum symbol.
<i>algorithm</i>	<div>The checksum algorithm used, one of:<ul style="list-style-type: none">• <code>sum</code>, a byte-wise calculated arithmetic sum. The result is truncated to 8 bits.• <code>sum8wide</code>, a byte-wise calculated arithmetic sum. The result is truncated to the size of the symbol.• <code>sum32</code>, a word-wise (32 bits) calculated arithmetic sum.• <code>crc16</code>, CRC16 (generating polynomial 0x1021); used by default.• <code>crc32</code>, CRC32 (generating polynomial 0x04C11DB7).• <code>crc64iso</code>, CRC64iso (generating polynomial 0x1B).• <code>crc64ecma</code>, CRC64ECMA (generating polynomial 0x42F0E1EBA9EA3693).• <code>crc=n</code>, CRC with a generating polynomial of <i>n</i>.</div>

$1 2$	<p>If specified, can be one of:</p> <ul style="list-style-type: none"> • 1 – Specifies one's complement. • 2 – Specifies two's complement.
$a m z$	<p>Reverses the order of the bits for the checksum; choose between:</p> <p>a, reverses the input bytes (but nothing else).</p> <p>m, reverses the input bytes and the final checksum.</p> <p>z, reverses the final checksum (but nothing else).</p> <p>Note that using a and z in combination has the same effect as m.</p>
$L W$	<p>Specifies the size of the unit for which a checksum should be calculated.</p> <p>Choose between:</p> <p>L, calculates a checksum on 32 bits in every iteration</p> <p>W, calculates a checksum on 16 bits in every iteration.</p> <p>If you do not specify a unit size, 8 bits will be used by default. Using these parameters does not add any additional error detection power to the checksum.</p>
r	<p>Reverses the byte order of the input data. This has no effect unless the number of bits per iteration has been set using the L or W parameters.</p>
R	<p>Traverses the checksum range(s) in reverse order.</p> <p>If the range is for example 0x100–0xFFFF;0x2000–0x2FFF, the checksum calculation will normally start on 0x100 and then calculate every byte up to and including 0xFFFF, followed by calculating the byte on 0x2000 and continue to 0x2FFF.</p> <p>Using the R parameter, the calculation instead starts on 0x2FFF and continues by calculating every byte down to 0x2000, then from 0xFFFF down to and including 0x100.</p>
o	<p>Uses the Rocksoft model specification for the checksum.</p>

<i>i</i> <i>p</i>	<p>Use either <i>i</i> or <i>p</i>, if the <i>start</i> value is bigger than 0. If specified, can be one of:</p> <ul style="list-style-type: none"> • <i>i</i> - Initializes the checksum value with the start value. • <i>p</i> - Prefixes the input data with a word of size <i>size</i> that contains the <i>start</i> value.
<i>start</i>	<p>By default, the initial value of the checksum is 0. If necessary, use <i>start</i> to supply a different initial value. If not 0, then either <i>i</i> or <i>p</i> must be specified.</p>
<i>range</i>	<p>The address range on which the checksum should be calculated. Hexadecimal and decimal notation is allowed (for example, 0x8002-0x8FFF).</p> <p>Symbols that are present in ELF file can be used in the range description (for example, __checksum_begin-__checksum_end).</p>
For use with	<code>ielftool</code>
Description	<p>Use this option to calculate a checksum with the specified algorithm for the specified ranges. If you have an external definition for the checksum (for example, a hardware CRC implementation), use the appropriate parameters to the <code>--checksum</code> option to match the external design. (In this case, learn more about that design in the hardware documentation.) The checksum will then replace the original value in <i>symbol</i>. A new absolute symbol will be generated; with the <i>symbol</i> name suffixed with <i>_value</i> containing the calculated checksum. This symbol can be used for accessing the checksum value later when needed, for example during debugging.</p> <p>If the <code>--checksum</code> option is used more than once on the command line, the options are evaluated from left to right. If a checksum is calculated for a <i>symbol</i> that is specified in a later evaluated <code>--checksum</code> option, an error is issued.</p>
Example	<p>This example shows how to use the <code>crc16</code> algorithm with the start value 0 over the address range 0x8000-0x8FFF:</p> <pre>ielftool --checksum=__checksum:2,crc16;0x8000-0x8FFF sourceFile.out destinationFile.out</pre> <p>The input data <i>i</i> read from <code>sourceFile.out</code>, and the resulting checksum value of size 2 bytes will be stored at the symbol <code>__checksum</code>. The modified ELF file is saved as <code>destinationFile.out</code> leaving <code>sourceFile.out</code> untouched.</p>

In the next example, a symbol is used for specifying the start of the range:

```
ielftool --checksum=__checksum:2,crc16;__checksum_begin-0x8FFF
sourceFile.out destinationFile.out
```

See also

Checksum calculation for verifying image integrity, page 208



To set related options, choose:

Project>Options>Linker>Checksum

--code

Syntax

```
--code
```

For use with

```
ielfdump
```

Description

Use this option to dump all sections that contain executable code (sections with the ELF section attribute `SHF_EXECINSTR`).



This option is not available in the IDE.

--create

Syntax

```
--create libraryfile objectfile1 ... objectfileN
```

Parameters

```
libraryfile      The library file that the command operates on. See Rules for
                  specifying a filename or directory as parameters, page 254.

objectfile1 ...  The object file(s) to build the library from.
objectfileN
```

For use with

```
iarchive
```

Description


Use this command to build a new library from a set of object files (modules). The object files are added to the library in the exact order that they are specified on the command line.

If no command is specified on the command line, `--create` is used by default.




This option is not available in the IDE.

--delete, -d


Syntax	<pre>--delete libraryfile objectfile1 ... objectfileN -d libraryfile objectfile1 ... objectfileN</pre>	
Parameters	<p><i>libraryfile</i> The library file that the command operates on. See <i>Rules for specifying a filename or directory as parameters</i>, page 254.</p> <p><i>objectfile1 ... objectfileN</i> The object file(s) that the command operates on.</p>	
For use with	iarchive	
Description	<p>Use this command to remove object files (modules) from an existing library. All object files that are specified on the command line will be removed from the library.</p> <p> This option is not available in the IDE.</p>	

--edit


Syntax	<pre>--edit steering_file</pre>	
For use with	isymexport	
Description	<p>Use this option to specify a steering file to control which symbols that are included in the isymexport output file, and also to rename some of the symbols if that is desired.</p>	
See also	<p><i>Steering files</i>, page 488.</p> <p> This option is not available in the IDE.</p>	

--extract, -x

Syntax	<pre>--extract libraryfile [objectfile1 ... objectfileN] -x libraryfile [objectfile1 ... objectfileN]</pre>	
Parameters	<p><i>libraryfile</i> The library file that the command operates on. See <i>Rules for specifying a filename or directory as parameters</i>, page 254.</p>	

	<code>objectfile1 ...</code> The object file(s) that the command operates on. <code>objectfileN</code>
For use with	iarchive
Description	Use this command to extract object files (modules) from an existing library. If a list of object files is specified, only these files are extracted. If a list of object files is not specified, all object files in the library are extracted.
	 This option is not available in the IDE.

-f

Syntax	<code>-f filename</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
For use with	iarchive, ielfdumparm, iobjmanip, and isymexport.
Description	<p>Use this option to make the tool read command line options from the named file, with the default filename extension <code>.xcl</code>.</p> <p>In the command file, you format the items exactly as if they were on the command line itself, except that you can use multiple lines, because the newline character acts just as a space or tab character.</p> <p>Both C and C++ style comments are allowed in the file. Double quotes behave in the same way as in the Microsoft Windows command line environment.</p>
	 This option is not available in the IDE.

--fill

Syntax	<code>--fill [v;]pattern;range[;range...]</code>	
Parameters	<code>v</code>	Generates virtual fill for the fill command. Virtual fill is filler bytes that are included in checksumming, but that are not included in the output file. The primary use for this is certain types of hardware where bytes that are not specified by the image have a known value (typically, 0xFF or 0x0).
	<code>pattern</code>	A hexadecimal string with the 0x prefix (for example, 0xEF) interpreted as a sequence of bytes, where each pair of digits corresponds to one byte (for example 0x123456, for the sequence of bytes 0x12, 0x34, and 0x56). This sequence is repeated over the fill area. If the length of the fill pattern is greater than 1 byte, it is repeated as if it started at address 0.
	<code>range</code>	Specifies the address range for the fill. Hexadecimal and decimal notation is allowed (for example, 0x8002-0x8FFF). Note that each address must be 4-byte aligned. Symbols that are present in the ELF file can be used in the range description (for example. __checksum_begin-__checksum_end).

For use with `ielftool`

Description Use this option to fill all gaps in one or more ranges with a pattern, which can be either an expression or a hexadecimal string. The contents will be calculated as if the fill pattern was repeatedly filled from the start address until the end address is passed, and then the real contents will overwrite that pattern.

If the `--fill` option is used more than once on the command line, the fill ranges cannot overlap each other.



To set related options, choose:
Project>Options>Linker>Checksum

--generate_vfe_header

Syntax	<code>--generate_vfe_header</code>
For use with	<code>ismexport</code>

Description

Use this option to declare that the image does not contain any virtual function calls to potentially discarded functions.

When the linker performs virtual function elimination, it discards virtual functions that appear not to be needed. For the optimization to be applied correctly, there must be no virtual function calls in the image that affect the functions that are discarded.

See also

Virtual function elimination, page 217.



To set this options, use:

Project>Options>Linker>Extra Options

--ihex

Syntax

--ihex

For use with

ielftool

Description

Sets the format of the output file to linear Intel hex.



To set related options, choose:

Project>Options>Linker>Output converter

--no_header

Syntax

--no_header

For use with

ielfdumparm

Description

By default, a standard list header is added before the actual file content. Use this option to suppress output of the list header.



This option is not available in the IDE.


--no_rel_section

Syntax


--no_rel_section

For use with


ielfdumparm

Description	By default, whenever the content of a section of a relocatable file is generated as output, the associated section, if any, is also included in the output. Use this option to suppress output of the relocation section.
	 This option is not available in the IDE.


--no_strtab

Syntax	--no_strtab
For use with	ielfdumparm
Description	Use this option to suppress dumping of string table sections (sections of type SHT_STRTAB).
	 This option is not available in the IDE.

--offset


Syntax	--offset [-]offset
Parameters	<div>offset</div> <div>The offset will be added (or subtracted if - is specified) to all addresses in the generated output file.</div>
For use with	ieftool
Description	Use this option to add or subtract an offset to the address of each output record in the generated output file. The option only works on Motorola S-records, Intel Hex, TI-Txt, and Simple-Code. The option has no effect when generating an ELF file or when binary files (--bin contain no address information) are generated. No content, including the entry point, will be changed by using this option, only the addresses in the output format.
Example	<div>--offset 0x30000</div> <div>This will add an offset of 0x30000 to all addresses. As a result, content that was linked at address 0x4000 will be placed at 0x34000.</div> <div> This option is not available in the IDE.</div>

--output, -o

Syntax	<code>-o {filename directory}</code> <code>--output {filename directory}</code>
Parameters	See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
For use with	<code>iarchive</code> and <code>ielfdumparm</code> .
Description	<p><code>iarchive</code></p> <p>By default, <code>iarchive</code> assumes that the first argument after the <code>iarchive</code> command is the name of the destination library. Use this option to explicitly specify a different filename for the library.</p> <p><code>ielfdumparm</code></p> <p>By default, output from the dumper is directed to the console. Use this option to direct the output to a file instead. The default name of the output file is the name of the input file with an added <code>id</code> filename extension</p> <p>You can also specify the output file by specifying a file or directory following the name of the input file.</p> <p> This option is not available in the IDE.</p>

--parity

Syntax	<code>--parity{symbol[+offset] address}:size,algo:flashbase[:flags];range[:range...]</code>								
Parameters	<table><tr><td><i>symbol</i></td><td>The name of the symbol where the parity bytes should be stored. Note that it must exist in the symbol table in the input ELF file.</td></tr><tr><td><i>offset</i></td><td>An offset to the symbol. By default, 0.</td></tr><tr><td><i>address</i></td><td>The absolute address where the parity bytes should be stored.</td></tr><tr><td><i>size</i></td><td>The maximum number of bytes that the parity generation can use. An error will be issued if this value is exceeded. Note that the size must fit in the specified symbol in the ELF file.</td></tr></table>	<i>symbol</i>	The name of the symbol where the parity bytes should be stored. Note that it must exist in the symbol table in the input ELF file.	<i>offset</i>	An offset to the symbol. By default, 0.	<i>address</i>	The absolute address where the parity bytes should be stored.	<i>size</i>	The maximum number of bytes that the parity generation can use. An error will be issued if this value is exceeded. Note that the size must fit in the specified symbol in the ELF file.
<i>symbol</i>	The name of the symbol where the parity bytes should be stored. Note that it must exist in the symbol table in the input ELF file.								
<i>offset</i>	An offset to the symbol. By default, 0.								
<i>address</i>	The absolute address where the parity bytes should be stored.								
<i>size</i>	The maximum number of bytes that the parity generation can use. An error will be issued if this value is exceeded. Note that the size must fit in the specified symbol in the ELF file.								

	<i>algo</i>	Choose between: odd, uses odd parity. even, uses even parity.
	<i>flashbase</i>	The start address of the flash memory. Parity bits will not be generated for the addresses between <i>flashbase</i> and the start address of the range. If <i>flashbase</i> and the start address of the range coincide, parity bits will be generated for all addresses
	<i>flags</i>	Choose between: r, reverses the byte order within each word. L, processes 4 bytes at a time. W, processes 2 bytes at a time. B, processes 1 byte at a time.
	<i>range</i>	The address range over which the parity bytes should be generated. Hexadecimal and decimal notation are allowed (for example, 0x8002–0x8FFF).
For use with	ielftool	
Description	Use this option to generate parity bytes over specified ranges. The range is traversed left to the right and the parity bits are generated using the odd or even algorithm. The parity bits are finally stored in the specified symbol where they can be accessed by your application.	
		This option is not available in the IDE.

--ram_reserve_ranges

Syntax	<code>--ram_reserve_ranges [=symbol_prefix]</code>	
Parameters	<i>symbol_prefix</i>	The prefix of symbols created by this option.
For use with	ismexport	
Description	Use this option to generate symbols for the areas in RAM that the image uses. One symbol will be generated for each such area. The name of each symbol is based on the name of the area and is prefixed by the optional parameter <i>symbol_prefix</i> .	

Generating symbols that cover an area in this way prevents the linker from placing other content at the affected addresses. This can be useful when linking against an existing image.

If `--ram_reserve_ranges` is used together with `--reserve_ranges`, the RAM areas will get their prefix from the `--ram_reserve_ranges` option and the non-RAM areas will get their prefix from the `--reserve_ranges` option.

See also

`--reserve_ranges`, page 507.



This option is not available in the IDE.

--range

Syntax

`--range start-end`

Parameters

start-end

Disassemble code where the start address is greater than or equal to *start*, and where the end address is less than *end*.

For use with

`ielfdumparm`

Description

Use this option to specify a range for which code from an executable will be dumped.



This option is not available in the IDE.

--raw

Syntax

`--raw`

For use with

`ielfdumparm`

Description


By default, many ELF sections will be dumped using a text format specific to a particular kind of section. Use this option to dump each selected ELF section using the generic text format.

The generic text format dumps each byte in the section in hexadecimal format, and where appropriate, as ASCII text.




This option is not available in the IDE.

--remove_file_path


Syntax	<code>--remove_file_path</code>	
For use with	<code>iobjmanip</code>	
Description	Use this option to make <code>iobjmanip</code> remove information about the directory structure of the project source tree from the generated object file, which means that the file symbol in the ELF object file is modified. This option must be used in combination with <code>--remove_section ".comment"</code> .  This option is not available in the IDE.	

--remove_section


Syntax	<code>--remove_section {<i>section</i> <i>number</i>}</code>	
Parameters	<i>section</i>	The section—or sections, if there are more than one section with the same name—to be removed.
	<i>number</i>	The number of the section to be removed. Section numbers can be obtained from an object dump created using <code>ielfdumparm</code> .
For use with	<code>iobjmanip</code>	
Description	Use this option to make <code>iobjmanip</code> omit the specified section when generating the output file.  This option is not available in the IDE.	

--rename_section

Syntax	<code>--rename_section {<i>oldname</i> <i>oldnumber</i>}=<i>newname</i></code>	
Parameters	<i>oldname</i>	The section—or sections, if there are more than one section with the same name—to be renamed.

	<i>oldnumber</i>	The number of the section to be renamed. Section numbers can be obtained from an object dump created using <code>ielfdumparm</code> .
	<i>newname</i>	The new name of the section.
For use with	<code>iobjmanip</code>	
Description	Use this option to make <code>iobjmanip</code> rename the specified section when generating the output file.	
		This option is not available in the IDE.

--rename_symbol

Syntax	<code>--rename_symbol oldname =newname</code>	
Parameters	<i>oldname</i>	The symbol to be renamed.
	<i>newname</i>	The new name of the symbol.
For use with	<code>iobjmanip</code>	
Description	Use this option to make <code>iobjmanip</code> rename the specified symbol when generating the output file.	
		This option is not available in the IDE.

--replace, -r

Syntax	<code>--replace libraryfile objectfile1 ... objectfileN</code> <code>-r libraryfile objectfile1 ... objectfileN</code>	
Parameters	<i>libraryfile</i>	The library file that the command operates on. See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
	<i>objectfile1 ... objectfileN</i>	The object file(s) that the command operates on.
For use with	<code>iarchive</code>	

Description

Use this command to replace or add object files (modules) to an existing library. The object files specified on the command line either replace existing object files in the library (if they have the same name) or are appended to the library.



This option is not available in the IDE.

--reserve_ranges**Syntax**

```
--reserve_ranges [=symbol_prefix]
```

Parameters

symbol_prefix The prefix of symbols created by this option.

For use with

isymexport

Description

Use this option to generate symbols for the areas in ROM and RAM that the image uses. One symbol will be generated for each such area. The name of each symbol is based on the name of the area and is prefixed by the optional parameter *symbol_prefix*.

Generating symbols that cover an area in this way prevents the linker from placing other content at the affected addresses. This can be useful when linking against an existing image.

If `--reserve_ranges` is used together with `--ram_reserve_ranges`, the RAM areas will get their prefix from the `--ram_reserve_ranges` option and the non-RAM areas will get their prefix from the `--reserve_ranges` option.

See also

`--ram_reserve_ranges`, page 503.



This option is not available in the IDE.

--section, -s**Syntax**


```
--section section_number|section_name[,...]  
--s section_number|section_name[,...]
```

Parameters


section_number The number of the section to be dumped.

section_name The name of the section to be dumped.

For use with	<code>ielfdumparm</code>
Description	<p>Use this option to dump the contents of a section with the specified number, or any section with the specified name. If a relocation section is associated with a selected section, its contents are output as well.</p> <p>If you use this option, the general properties of the input file will not be included in the output.</p> <p>You can specify multiple section numbers or names by separating them with commas, or by using this option more than once.</p> <p>By default, no section contents are included in the output.</p>

Example	<pre>-s 3,17 /* Sections #3 and #17 -s .debug_frame,42 /* Any sections named .debug_frame and also section #42 */</pre> <p> This option is not available in the IDE.</p>
---------	---

--segment, -g

Syntax	<code>--segment <i>segment_number</i>[,...]</code> <code>-g <i>segment_number</i>[,...]</code>		
Parameters	<table><tr><td><i>segmnt_number</i></td><td>The number of a segment whose contents will be included in the output.</td></tr></table>	<i>segmnt_number</i>	The number of a segment whose contents will be included in the output.
<i>segmnt_number</i>	The number of a segment whose contents will be included in the output.		
For use with	<code>ielfdumparm</code>		
Description	<p>Use this option to select specific segments (parts of an executable image indicated by program headers) for inclusion in the output.</p> <p> This option is not available in the IDE.</p>		

--self_reloc

Syntax	<code>--self_reloc</code>
For use with	<code>ielftool</code>

Description

This option is intentionally not documented as it is not intended for general use.



This option is not available in the IDE.

--show_entry_as

Syntax

--show_entry_as *name*

Parameters

name The name to give to the program entry point in the output file.

For use with

isymexport

Description

Use this option to export the entry point of the application given as input under the name *name*.



This option is not available in the IDE.

--silent

Syntax

--silent

For use with

iobjmanip and ielftool.

Description

Causes the tool to operate without sending any messages to the standard output stream. By default, the tool sends various messages via the standard output stream. You can use this option to prevent this. The tool sends error and warning messages to the error output stream, so they are displayed regardless of this setting.



This option is not available in the IDE.

--simple

Syntax

--simple

For use with

ielftool

Description

Sets the format of the output file to Simple code.



To set related options, choose:
Project>Options>Output converter

--simple-ne

Syntax	--simple-ne
For use with	ielftool
Description	Sets the format of the output file to Simple code, but no entry record is generated.



To set related options, choose:
Project>Options>Output converter

--source

Syntax	--source
For use with	ielfdumparm
Description	Use this option to make <code>ielftool</code> include source for each statement before the code for that statement, when dumping code from an executable file. To make this work, the executable image must be built with debug information, and the source code must still be accessible in its original location.



This option is not available in the IDE.


--srec

Syntax	--srec
For use with	ielftool
Description	Sets the format of the output file to Motorola S-records.




To set related options, choose:
Project>Options>Output converter


--srec-len

Syntax	<code>--srec-len=length</code>	
Parameters	<i>length</i>	The number of data bytes in each S-record.
For use with	<code>ielftool</code>	
Description	Restricts the number of data bytes in each S-record. This option can be used in combination with the <code>--srec</code> option.	
		This option is not available in the IDE.


--srec-s3only

Syntax	<code>--srec-s3only</code>	
For use with	<code>ielftool</code>	
Description	Restricts the S-record output to contain only a subset of records, that is S0, S3 and S7 records. This option can be used in combination with the <code>--srec</code> option.	
		This option is not available in the IDE.


--strip

Syntax	<code>--strip</code>	
For use with	<code>iobjmanip</code> and <code>ielftool</code> .	
Description	Use this option to remove all sections containing debug information before the output file is written.	
	Note that <code>ielftool</code> needs an unstripped input ELF image. If you use the <code>--strip</code> option in the linker, remove it and use the <code>--strip</code> option in <code>ielftool</code> instead.	
		To set related options, choose: Project>Options>Linker>Output>Include debug information in output

--symbols


Syntax	<code>--symbols libraryfile</code>	
Parameters	<code>libraryfile</code>	The library file that the command operates on. See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
For use with	<code>iarchive</code>	
Description	<p>Use this command to list all external symbols that are defined by any object file (module) in the specified library, together with the name of the object file (module) that defines it.</p> <p>In silent mode (<code>--silent</code>), this command performs symbol table-related syntax checks on the library file and displays only errors and warnings.</p> <div> This option is not available in the IDE.</div>	

--titxt


Syntax	<code>--titxt</code>	
For use with	<code>ielftool</code>	
Description	<p>Sets the format of the output file to TI-txt.</p> <div> To set related options, choose: Project>Options>Output converter</div>	

--toc, -t


Syntax	<code>--toc libraryfile</code> <code>-t libraryfile</code>	
Parameters	<code>libraryfile</code>	The library file that the command operates on. See <i>Rules for specifying a filename or directory as parameters</i> , page 254.
For use with	<code>iarchive</code>	

Description	Use this command to list the names of all object files (modules) in a specified library. In silent mode (<code>--silent</code>), this command performs basic syntax checks on the library file, and displays only errors and warnings.
	 This option is not available in the IDE.

--use_full_std_template_names


Syntax	<code>--use_full_std_template_names</code>
For use with	<code>ielfdumparm</code>
Description	Normally, the names of some standard C++ templates are used in the output in an abbreviated form in the unmangled names of symbols (for example, " <code>std::string</code> " instead of " <code>std::basic_string<char, std::char_traits<char>, std::allocator<char>></code> "). Use this option to make <code>ielfdump</code> use the unabbreviated form.
	 This option is not available in the IDE.

--verbose, -V

Syntax	<code>--verbose</code> <code>-V</code> (<code>iarchive</code> only)
For use with	<code>iarchive</code> and <code>ielftool</code> .
Description	Use this option to make the tool report which operations it performs, in addition to giving diagnostic messages.
	 This option is not available in the IDE because this setting is always enabled.

--version

Syntax	<code>--version</code>
For use with	<code>iarchive</code> , <code>ielfdump</code> , <code>iobjmanip</code> , <code>isymexport</code>

Description	Use this option to make the tool send version information to the console and then exit.
	This option is not available in the IDE.

Implementation-defined behavior for Standard C

- Descriptions of implementation-defined behavior

If you are using C89 instead of Standard C, see *Implementation-defined behavior for C89*, page 531. For a short overview of the differences between Standard C and C89, see *C language overview*, page 181.

Descriptions of implementation-defined behavior

This section follows the same order as the C standard. Each item includes references to the ISO chapter and section (in parenthesis) that explains the implementation-defined behavior.

Note: The IAR Systems implementation adheres to a freestanding implementation of Standard C. This means that parts of a standard library can be excluded in the implementation.

J.3.1 TRANSLATION

Diagnostics (3.10, 5.1.1.3)

Diagnostics are produced in the form:

filename,linenumber level[tag]: message

where *filename* is the name of the source file in which the error was encountered, *linenumber* is the line number at which the compiler detected the error, *level* is the level of seriousness of the message (remark, warning, error, or fatal error), *tag* is a unique tag that identifies the message, and *message* is an explanatory message, possibly several lines.

White-space characters (5.1.1.2)

At translation phase three, each non-empty sequence of white-space characters is retained.

J.3.2 ENVIRONMENT

The character set (5.1.1.2)

The source character set is the same as the physical source file multibyte character set. By default, the standard ASCII character set is used. However, if you use the `--enable_multibytes` compiler option, the host character set is used instead.

Main (5.1.2.1)

The function called at program startup is called `main`. No prototype is declared for `main`, and the only definition supported for `main` is:

```
int main(void)
```

To change this behavior, see *System initialization*, page 137.

The effect of program termination (5.1.2.1)

Terminating the application returns the execution to the startup code (just after the call to `main`).

Alternative ways to define main (5.1.2.2.1)

There is no alternative ways to define the `main` function.

The argv argument to main (5.1.2.2.1)

The `argv` argument is not supported.

Streams as interactive devices (5.1.2.3)

The streams `stdin`, `stdout`, and `stderr` are treated as interactive devices.

Signals, their semantics, and the default handling (7.14)

In the DLIB runtime environment, the set of supported signals is the same as in Standard C. A raised signal will do nothing, unless the `signal` function is customized to fit the application.

Signal values for computational exceptions (7.14.1.1)

In the DLIB runtime environment, there are no implementation-defined values that correspond to a computational exception.

Signals at system startup (7.14.1.1)

In the DLIB runtime environment, there are no implementation-defined signals that are executed at system startup.

Environment names (7.20.4.5)

In the DLIB runtime environment, there are no implementation-defined environment names that are used by the `getenv` function.

The system function (7.20.4.6)

The `system` function is not supported.

J.3.3 IDENTIFIERS

Multibyte characters in identifiers (6.4.2)

Additional multibyte characters may not appear in identifiers.

Significant characters in identifiers (5.2.4.1, 6.1.2)

The number of significant initial characters in an identifier with or without external linkage is guaranteed to be no less than 200.

J.3.4 CHARACTERS

Number of bits in a byte (3.6)

A byte contains 8 bits.

Execution character set member values (5.2.1)

The values of the members of the execution character set are the values of the ASCII character set, which can be augmented by the values of the extra characters in the host character set.

Alphabetic escape sequences (5.2.2)

The standard alphabetic escape sequences have the values `\a-7`, `\b-8`, `\f-12`, `\n-10`, `\r-13`, `\t-9`, and `\v-11`.

Characters outside of the basic executive character set (6.2.5)

A character outside of the basic executive character set that is stored in a `char` is not transformed.

Plain char (6.2.5, 6.3.1.1)

A plain `char` is treated as an unsigned `char`.

Source and execution character sets (6.4.4.4, 5.1.1.2)

The source character set is the set of legal characters that can appear in source files. By default, the source character set is the standard ASCII character set. However, if you use the command line option `--enable_multibytes`, the source character set will be the host computer's default character set.

The execution character set is the set of legal characters that can appear in the execution environment. By default, the execution character set is the standard ASCII character set.

However, if you use the command line option `--enable_multibytes`, the execution character set will be the host computer's default character set. The DLIB runtime environment needs a multibyte character scanner to support a multibyte execution character set. See *Locale*, page 148.

Integer character constants with more than one character (6.4.4.4)

An integer character constant that contains more than one character will be treated as an integer constant. The value will be calculated by treating the leftmost character as the most significant character, and the rightmost character as the least significant character, in an integer constant. A diagnostic message will be issued if the value cannot be represented in an integer constant.

Wide character constants with more than one character (6.4.4.4)

A wide character constant that contains more than one multibyte character generates a diagnostic message.

Locale used for wide character constants (6.4.4.4)

By default, the C locale is used. If the `--enable_multibytes` compiler option is used, the default host locale is used instead.

Locale used for wide string literals (6.4.5)

By default, the C locale is used. If the `--enable_multibytes` compiler option is used, the default host locale is used instead.

Source characters as executive characters (6.4.5)

All source characters can be represented as executive characters.

J.3.5 INTEGERS

Extended integer types (6.2.5)

There are no extended integer types.

Range of integer values (6.2.6.2)

The representation of integer values are in the two's complement form. The most significant bit holds the sign; 1 for negative, 0 for positive and zero.

For information about the ranges for the different integer types, see *Basic data types—integer types*, page 331.

The rank of extended integer types (6.3.1.1)

There are no extended integer types.

Signals when converting to a signed integer type (6.3.1.3)

No signal is raised when an integer is converted to a signed integer type.

Signed bitwise operations (6.5)

Bitwise operations on signed integers work the same way as bitwise operations on unsigned integers; in other words, the sign-bit will be treated as any other bit.

J.3.6 FLOATING POINT

Accuracy of floating-point operations (5.2.4.2.2)

The accuracy of floating-point operations is unknown.

Rounding behaviors (5.2.4.2.2)

There are no non-standard values of `FLT_ROUNDS`.

Evaluation methods (5.2.4.2.2)

There are no non-standard values of `FLT_EVAL_METHOD`.

Converting integer values to floating-point values (6.3.1.4)

When an integral value is converted to a floating-point value that cannot exactly represent the source value, the round-to-nearest rounding mode is used (`FLT_ROUNDS` is defined to 1).

Converting floating-point values to floating-point values (6.3.1.5)

When a floating-point value is converted to a floating-point value that cannot exactly represent the source value, the round-to-nearest rounding mode is used (`FLT_ROUNDS` is defined to 1).

Denoting the value of floating-point constants (6.4.4.2)

The round-to-nearest rounding mode is used (`FLT_ROUNDS` is defined to 1).

Contraction of floating-point values (6.5)

Floating-point values are contracted. However, there is no loss in precision and because signaling is not supported, this does not matter.

Default state of `FENV_ACCESS` (7.6.1)

The default state of the pragma directive `FENV_ACCESS` is `OFF`.

Additional floating-point mechanisms (7.6, 7.12)

There are no additional floating-point exceptions, rounding-modes, environments, and classifications.

Default state of `FP_CONTRACT` (7.12.2)

The default state of the pragma directive `FP_CONTRACT` is `OFF`.

J.3.7 ARRAYS AND POINTERS

Conversion from/to pointers (6.3.2.3)

For information about casting of data pointers and function pointers, see *Casting*, page 338.

`ptrdiff_t` (6.5.6)

For information about `ptrdiff_t`, see *ptrdiff_t*, page 338.

J.3.8 HINTS

Honoring the register keyword (6.7.1)

User requests for register variables are not honored.

Inlining functions (6.7.4)

User requests for inlining functions increases the chance, but does not make it certain, that the function will actually be inlined into another function. See *Inlining functions*, page 79.

J.3.9 STRUCTURES, UNIONS, ENUMERATIONS, AND BITFIELDS**Sign of 'plain' bitfields (6.7.2, 6.7.2.1)**

For information about how a 'plain' `int` bitfield is treated, see *Bitfields*, page 332.

Possible types for bitfields (6.7.2.1)

All integer types can be used as bitfields in the compiler's extended mode, see *-e*, page 269.

Bitfields straddling a storage-unit boundary (6.7.2.1)

A bitfield is always placed in one—and one only—storage unit, which means that the bitfield cannot straddle a storage-unit boundary.

Allocation order of bitfields within a unit (6.7.2.1)

For information about how bitfields are allocated within a storage unit, see *Bitfields*, page 332.

Alignment of non-bitfield structure members (6.7.2.1)

The alignment of non-bitfield members of structures is the same as for the member types, see *Alignment*, page 329.

Integer type used for representing enumeration types (6.7.2.2)

The chosen integer type for a specific enumeration type depends on the enumeration constants defined for the enumeration type. The chosen integer type is the smallest possible.

J.3.10 QUALIFIERS**Access to volatile objects (6.7.3)**

Any reference to an object with `volatile` qualified type is an access, see *Declaring objects volatile*, page 341.

J.3.11 PREPROCESSING DIRECTIVES

Mapping of header names (6.4.7)

Sequences in header names are mapped to source file names verbatim. A backslash '\' is not treated as an escape sequence. See *Overview of the preprocessor*, page 419.

Character constants in constant expressions (6.10.1)

A character constant in a constant expression that controls conditional inclusion matches the value of the same character constant in the execution character set.

The value of a single-character constant (6.10.1)

A single-character constant may only have a negative value if a plain character (`char`) is treated as a signed character, see *--char_is_signed*, page 262.

Including bracketed filenames (6.10.2)

For information about the search algorithm used for file specifications in angle brackets `<>`, see *Include file search procedure*, page 245.

Including quoted filenames (6.10.2)

For information about the search algorithm used for file specifications enclosed in quotes, see *Include file search procedure*, page 245.

Preprocessing tokens in #include directives (6.10.2)

Preprocessing tokens in an `#include` directive are combined in the same way as outside an `#include` directive.

Nesting limits for #include directives (6.10.2)

There is no explicit nesting limit for `#include` processing.

Universal character names (6.10.3.2)

Universal character names (UCN) are not supported.

Recognized pragma directives (6.10.6)

In addition to the pragma directives described in the chapter *Pragma directives*, the following directives are recognized and will have an indeterminate effect. If a pragma

directive is listed both in the chapter *Pragma directives* and here, the information provided in the chapter *Pragma directives* overrides the information here.

alignment
baseaddr
basic_template_matching
building_runtime
can_instantiate
codeseg
constseg
cspy_support
dataseg
define_type_info
do_not_instantiate
early_dynamic_initialization
function
function_effects
hdrstop
important_typedef
instantiate
keep_definition
library_default_requirements
library_provides
library_requirement_override
memory
module_name
no_pch
once
system_include
vector

warnings

Default `__DATE__` and `__TIME__` (6.10.8)

The definitions for `__TIME__` and `__DATE__` are always available.

J.3.12 LIBRARY FUNCTIONS

Additional library facilities (5.1.2.1)

Most of the standard library facilities are supported. Some of them—the ones that need an operating system—require a low-level implementation in the application. For more information, see *The DLIB runtime environment*, page 115.

Diagnostic printed by the `assert` function (7.2.1.1)

The `assert()` function prints:

```
filename:linenr expression -- assertion failed
```

when the parameter evaluates to zero.

Representation of the floating-point status flags (7.6.2.2)

For information about the floating-point status flags, see *fenv.h*, page 434.

`feraiseexcept` raising floating-point exception (7.6.2.3)

For information about the `feraiseexcept` function raising floating-point exceptions, see *Floating-point environment*, page 336.

Strings passed to the `setlocale` function (7.11.1.1)

For information about strings passed to the `setlocale` function, see *Locale*, page 148.

Types defined for `float_t` and `double_t` (7.12)

The `FLT_EVAL_METHOD` macro can only have the value 0.

Domain errors (7.12.1)

No function generates other domain errors than what the standard requires.

Return values on domain errors (7.12.1)

Mathematic functions return a floating-point NaN (not a number) for domain errors.

Underflow errors (7.12.1)

Mathematic functions set `errno` to the macro `ERANGE` (a macro in `errno.h`) and return zero for underflow errors.

`fmod` return value (7.12.10.1)

The `fmod` function returns a floating-point NaN when the second argument is zero.

The magnitude of `remquo` (7.12.10.3)

The magnitude is congruent modulo `INT_MAX`.

`signal()` (7.14.1.1)

The signal part of the library is not supported.

Note: The default implementation of `signal` does not perform anything. Use the template source code to implement application-specific signal handling. See *signal*, page 144 and *raise*, page 142, respectively.

NULL macro (7.17)

The `NULL` macro is defined to 0.

Terminating newline character (7.19.2)

`stdout` stream functions recognize either `newline` or end of file (EOF) as the terminating character for a line.

Space characters before a newline character (7.19.2)

Space characters written to a stream immediately before a newline character are preserved.

Null characters appended to data written to binary streams (7.19.2)

No null characters are appended to data written to binary streams.

File position in append mode (7.19.3)

The file position is initially placed at the beginning of the file when it is opened in append-mode.

Truncation of files (7.19.3)

Whether a write operation on a text stream causes the associated file to be truncated beyond that point, depends on the application-specific implementation of the low-level file routines. See *Briefly about input and output (I/O)*, page 116.

File buffering (7.19.3)

An open file can be either block-buffered, line-buffered, or unbuffered.

A zero-length file (7.19.3)

Whether a zero-length file exists depends on the application-specific implementation of the low-level file routines.

Legal file names (7.19.3)

The legality of a filename depends on the application-specific implementation of the low-level file routines.

Number of times a file can be opened (7.19.3)

Whether a file can be opened more than once depends on the application-specific implementation of the low-level file routines.

Multibyte characters in a file (7.19.3)

The encoding of multibyte characters in a file depends on the application-specific implementation of the low-level file routines.

remove() (7.19.4.1)

The effect of a remove operation on an open file depends on the application-specific implementation of the low-level file routines. See *Briefly about input and output (I/O)*, page 116.

rename() (7.19.4.2)

The effect of renaming a file to an already existing filename depends on the application-specific implementation of the low-level file routines. See *Briefly about input and output (I/O)*, page 116.

Removal of open temporary files (7.19.4.3)

Whether an open temporary file is removed depends on the application-specific implementation of the low-level file routines.

Mode changing (7.19.5.4)

`freopen` closes the named stream, then reopens it in the new mode. The streams `stdin`, `stdout`, and `stderr` can be reopened in any new mode.

Style for printing infinity or NaN (7.19.6.1, 7.24.2.1)

The style used for printing infinity or NaN for a floating-point constant is `inf` and `nan` (`INF` and `NAN` for the `F` conversion specifier), respectively. The `n`-char-sequence is not used for `nan`.

%p in printf() (7.19.6.1, 7.24.2.1)

The argument to a `%p` conversion specifier, print pointer, to `printf()` is treated as having the type `void *`. The value will be printed as a hexadecimal number, similar to using the `%x` conversion specifier.

Reading ranges in scanf (7.19.6.2, 7.24.2.1)

A `-` (dash) character is always treated as a range symbol.

%p in scanf (7.19.6.2, 7.24.2.2)

The `%p` conversion specifier, scan pointer, to `scanf()` reads a hexadecimal number and converts it into a value with the type `void *`.

File position errors (7.19.9.1, 7.19.9.3, 7.19.9.4)

On file position errors, the functions `fgetpos`, `ftell`, and `fsetpos` store `EFPOS` in `errno`.

An n-char-sequence after nan (7.20.1.3, 7.24.4.1.1)

An `n`-char-sequence after a NaN is read and ignored.

errno value at underflow (7.20.1.3, 7.24.4.1.1)

`errno` is set to `ERANGE` if an underflow is encountered.

Zero-sized heap objects (7.20.3)

A request for a zero-sized heap object will return a valid pointer and not a null pointer.

Behavior of abort and exit (7.20.4.1, 7.20.4.4)

A call to `abort()` or `_Exit()` will not flush stream buffers, not close open streams, and not remove temporary files.

Termination status (7.20.4.1, 7.20.4.3, 7.20.4.4)

The termination status will be propagated to `__exit()` as a parameter. `exit()` and `_Exit()` use the input parameter, whereas `abort` uses `EXIT_FAILURE`.

The system function return value (7.20.4.6)

The `system` function is not supported.

The time zone (7.23.1)

The local time zone and daylight savings time must be defined by the application. For more information, see `__time32`, `__time64`, page 145.

Range and precision of time (7.23)

For information about range and precision, see *time.h*, page 435. The application must supply the actual implementation for the functions `time` and `clock`. See `__time32`, `__time64`, page 145 and `clock`, page 140, respectively.

clock() (7.23.2.1)

The application must supply an implementation of the `clock` function. See *clock*, page 140.

%Z replacement string (7.23.3.5, 7.24.5.1)

By default, ":" is used as a replacement for %Z. Your application should implement the time zone handling. See `__time32`, `__time64`, page 145.

Math functions rounding mode (F.9)

The functions in *math.h* honor the rounding direction mode in `FLT-ROUNDS`.

J.3.13 ARCHITECTURE

Values and expressions assigned to some macros (5.2.4.2, 7.18.2, 7.18.3)

There are always 8 bits in a byte.

`MB_LEN_MAX` is at the most 6 bytes depending on the library configuration that is used.

For information about sizes, ranges, etc for all basic types, see *Data representation*, page 329.

The limit macros for the exact-width, minimum-width, and fastest minimum-width integer types defined in `stdint.h` have the same ranges as `char`, `short`, `int`, `long`, and `long long`.

The floating-point constant `FLT_ROUNDS` has the value 1 (to nearest) and the floating-point constant `FLT_EVAL_METHOD` has the value 0 (treat as is).

The number, order, and encoding of bytes (6.2.6.1)

See *Data representation*, page 329.

The value of the result of the sizeof operator (6.5.3.4)

See *Data representation*, page 329.

J.4 LOCALE

Members of the source and execution character set (5.2.1)

By default, the compiler accepts all one-byte characters in the host's default character set. If the compiler option `--enable_multibytes` is used, the host multibyte characters are accepted in comments and string literals as well.

The meaning of the additional character set (5.2.1.2)

Any multibyte characters in the extended source character set is translated verbatim into the extended execution character set. It is up to your application with the support of the library configuration to handle the characters correctly.

Shift states for encoding multibyte characters (5.2.1.2)

Using the compiler option `--enable_multibytes` enables the use of the host's default multibyte characters as extended source characters.

Direction of successive printing characters (5.2.2)

The application defines the characteristics of a display device.

The decimal point character (7.1.1)

The default decimal-point character is a `'.'`. You can redefine it by defining the library configuration symbol `_LOCALE_DECIMAL_POINT`.

Printing characters (7.4, 7.25.2)

The set of printing characters is determined by the chosen locale.

Control characters (7.4, 7.25.2)

The set of control characters is determined by the chosen locale.

Characters tested for (7.4.1.2, 7.4.1.3, 7.4.1.7, 7.4.1.9, 7.4.1.10, 7.4.1.11, 7.25.2.1.2, 7.25.5.1.3, 7.25.2.1.7, 7.25.2.1.9, 7.25.2.1.10, 7.25.2.1.11)

The sets of characters tested are determined by the chosen locale.

The native environment (7.1.1.1)

The native environment is the same as the "C" locale.

Subject sequences for numeric conversion functions (7.20.1, 7.24.4.1)

There are no additional subject sequences that can be accepted by the numeric conversion functions.

The collation of the execution character set (7.21.4.3, 7.24.4.4.2)

The collation of the execution character set is determined by the chosen locale.

Message returned by strerror (7.21.6.2)

The messages returned by the `strerror` function depending on the argument is:

Argument	Message
EZERO	no error
EDOM	domain error
ERANGE	range error
EFPOS	file positioning error
EILSEQ	multi-byte encoding error
<0 >99	unknown error
all others	error <i>nnn</i>

Table 52: Message returned by `strerror()`—DLIB runtime environment

Implementation-defined behavior for C89

- Descriptions of implementation-defined behavior

If you are using Standard C instead of C89, see *Implementation-defined behavior for Standard C*, page 515. For a short overview of the differences between Standard C and C89, see *C language overview*, page 181.

Descriptions of implementation-defined behavior

The descriptions follow the same order as the ISO appendix. Each item covered includes references to the ISO chapter and section (in parenthesis) that explains the implementation-defined behavior.

TRANSLATION

Diagnostics (5.1.1.3)

Diagnostics are produced in the form:

filename, *linenumber* *level*[*tag*]: *message*

where *filename* is the name of the source file in which the error was encountered, *linenumber* is the line number at which the compiler detected the error, *level* is the level of seriousness of the message (remark, warning, error, or fatal error), *tag* is a unique tag that identifies the message, and *message* is an explanatory message, possibly several lines.

ENVIRONMENT

Arguments to main (5.1.2.2.1)

The function called at program startup is called `main`. No prototype was declared for `main`, and the only definition supported for `main` is:

```
int main(void)
```

To change this behavior for the DLIB runtime environment, see *System initialization*, page 137.

Interactive devices (5.1.2.3)

The streams `stdin` and `stdout` are treated as interactive devices.

IDENTIFIERS

Significant characters without external linkage (6.1.2)

The number of significant initial characters in an identifier without external linkage is 200.

Significant characters with external linkage (6.1.2)

The number of significant initial characters in an identifier with external linkage is 200.

Case distinctions are significant (6.1.2)

Identifiers with external linkage are treated as case-sensitive.

CHARACTERS

Source and execution character sets (5.2.1)

The source character set is the set of legal characters that can appear in source files. The default source character set is the standard ASCII character set. However, if you use the command line option `--enable_multibytes`, the source character set will be the host computer's default character set.

The execution character set is the set of legal characters that can appear in the execution environment. The default execution character set is the standard ASCII character set. However, if you use the command line option `--enable_multibytes`, the execution character set will be the host computer's default character set. The DLIB runtime environment needs a multibyte character scanner to support a multibyte execution character set.

See *Locale*, page 148.

Bits per character in execution character set (5.2.4.2.1)

The number of bits in a character is represented by the manifest constant `CHAR_BIT`. The standard include file `limits.h` defines `CHAR_BIT` as 8.

Mapping of characters (6.1.3.4)

The mapping of members of the source character set (in character and string literals) to members of the execution character set is made in a one-to-one way. In other words, the

same representation value is used for each member in the character sets except for the escape sequences listed in the ISO standard.

Unrepresented character constants (6.1.3.4)

The value of an integer character constant that contains a character or escape sequence not represented in the basic execution character set or in the extended character set for a wide character constant generates a diagnostic message, and will be truncated to fit the execution character set.

Character constant with more than one character (6.1.3.4)

An integer character constant that contains more than one character will be treated as an integer constant. The value will be calculated by treating the leftmost character as the most significant character, and the rightmost character as the least significant character, in an integer constant. A diagnostic message will be issued if the value cannot be represented in an integer constant.

A wide character constant that contains more than one multibyte character generates a diagnostic message.

Converting multibyte characters (6.1.3.4)

The only locale supported—that is, the only locale supplied with the IAR C/C++ Compiler—is the ‘C’ locale. If you use the command line option `--enable_multibytes`, the DLIB runtime environment will support multibyte characters if you add a locale with multibyte support or a multibyte character scanner to the library.

See *Locale*, page 148.

Range of ‘plain’ char (6.2.1.1)

A ‘plain’ char has the same range as an unsigned char.

INTEGERS

Range of integer values (6.1.2.5)

The representation of integer values are in the two's complement form. The most significant bit holds the sign; 1 for negative, 0 for positive and zero.

See *Basic data types—integer types*, page 331, for information about the ranges for the different integer types.

Demotion of integers (6.2.1.2)

Converting an integer to a shorter signed integer is made by truncation. If the value cannot be represented when converting an unsigned integer to a signed integer of equal length, the bit-pattern remains the same. In other words, a large enough value will be converted into a negative value.

Signed bitwise operations (6.3)

Bitwise operations on signed integers work the same way as bitwise operations on unsigned integers; in other words, the sign-bit will be treated as any other bit.

Sign of the remainder on integer division (6.3.5)

The sign of the remainder on integer division is the same as the sign of the dividend.

Negative valued signed right shifts (6.3.7)

The result of a right-shift of a negative-valued signed integral type preserves the sign-bit. For example, shifting `0xFF00` down one step yields `0xFF80`.

FLOATING POINT**Representation of floating-point values (6.1.2.5)**

The representation and sets of the various floating-point numbers adheres to IEEE 854–1987. A typical floating-point number is built up of a sign-bit (*s*), a biased exponent (*e*), and a mantissa (*m*).

See *Basic data types—floating-point types*, page 336, for information about the ranges and sizes for the different floating-point types: `float` and `double`.

Converting integer values to floating-point values (6.2.1.3)

When an integral number is cast to a floating-point value that cannot exactly represent the value, the value is rounded (up or down) to the nearest suitable value.

Demoting floating-point values (6.2.1.4)

When a floating-point value is converted to a floating-point value of narrower type that cannot exactly represent the value, the value is rounded (up or down) to the nearest suitable value.

ARRAYS AND POINTERS

size_t (6.3.3.4, 7.1.1)

See *size_t*, page 338, for information about *size_t*.

Conversion from/to pointers (6.3.4)

See *Casting*, page 338, for information about casting of data pointers and function pointers.

ptrdiff_t (6.3.6, 7.1.1)

See *ptrdiff_t*, page 338, for information about the *ptrdiff_t*.

REGISTERS

Honoring the register keyword (6.5.1)

User requests for register variables are not honored.

STRUCTURES, UNIONS, ENUMERATIONS, AND BITFIELDS

Improper access to a union (6.3.2.3)

If a union gets its value stored through a member and is then accessed using a member of a different type, the result is solely dependent on the internal storage of the first member.

Padding and alignment of structure members (6.5.2.1)

See the section *Basic data types—integer types*, page 331, for information about the alignment requirement for data objects.

Sign of 'plain' bitfields (6.5.2.1)

A 'plain' *int* bitfield is treated as an *unsigned int* bitfield. All integer types are allowed as bitfields.

Allocation order of bitfields within a unit (6.5.2.1)

Bitfields are allocated within an integer from least-significant to most-significant bit.

Can bitfields straddle a storage-unit boundary (6.5.2.1)

Bitfields cannot straddle a storage-unit boundary for the chosen bitfield integer type.

Integer type chosen to represent enumeration types (6.5.2.2)

The chosen integer type for a specific enumeration type depends on the enumeration constants defined for the enumeration type. The chosen integer type is the smallest possible.

QUALIFIERS

Access to volatile objects (6.5.3)

Any reference to an object with volatile qualified type is an access.

DECLARATORS

Maximum numbers of declarators (6.5.4)

The number of declarators is not limited. The number is limited only by the available memory.

STATEMENTS

Maximum number of case statements (6.6.4.2)

The number of case statements (case values) in a switch statement is not limited. The number is limited only by the available memory.

PREPROCESSING DIRECTIVES

Character constants and conditional inclusion (6.8.1)

The character set used in the preprocessor directives is the same as the execution character set. The preprocessor recognizes negative character values if a 'plain' character is treated as a `signed` character.

Including bracketed filenames (6.8.2)

For file specifications enclosed in angle brackets, the preprocessor does not search directories of the parent files. A parent file is the file that contains the `#include` directive. Instead, it begins by searching for the file in the directories specified on the compiler command line.

Including quoted filenames (6.8.2)

For file specifications enclosed in quotes, the preprocessor directory search begins with the directories of the parent file, then proceeds through the directories of any grandparent files. Thus, searching begins relative to the directory containing the source

file currently being processed. If there is no grandparent file and the file is not found, the search continues as if the filename was enclosed in angle brackets.

Character sequences (6.8.2)

Preprocessor directives use the source character set, except for escape sequences. Thus, to specify a path for an include file, use only one backslash:

```
#include "mydirectory\myfile"
```

Within source code, two backslashes are necessary:

```
file = fopen("mydirectory\\myfile", "rt");
```

Recognized pragma directives (6.8.6)

In addition to the pragma directives described in the chapter *Pragma directives*, the following directives are recognized and will have an indeterminate effect. If a pragma directive is listed both in the chapter *Pragma directives* and here, the information provided in the chapter *Pragma directives* overrides the information here.

```
alignment
baseaddr
basic_template_matching
building_runtime
can_instantiate
codeseg
constseg
cspy_support
dataseg
define_type_info
do_not_instantiate
early_dynamic_initialization
function
function_effects
hdrstop
important_typedef
instantiate
```

`keep_definition`
`library_default_requirements`
`library_provides`
`library_requirement_override`
`memory`
`module_name`
`no_pch`
`once`
`system_include`
`vector`
`warnings`

Default `__DATE__` and `__TIME__` (6.8.8)

The definitions for `__TIME__` and `__DATE__` are always available.

LIBRARY FUNCTIONS FOR THE IAR DLIB RUNTIME ENVIRONMENT

Note that some items in this list only apply when file descriptors are supported by the library configuration. See the chapter *The DLIB runtime environment* for more information about runtime library configurations.

NULL macro (7.1.6)

The `NULL` macro is defined to 0.

Diagnostic printed by the `assert` function (7.2)

The `assert()` function prints:

filename:linenr expression -- assertion failed

when the parameter evaluates to zero.

Domain errors (7.5.1)

NaN (Not a Number) will be returned by the mathematic functions on domain errors.

Underflow of floating-point values sets `errno` to `ERANGE` (7.5.1)

The mathematics functions set the integer expression `errno` to `ERANGE` (a macro in `errno.h`) on underflow range errors.

`fmod()` functionality (7.5.6.4)

If the second argument to `fmod()` is zero, the function returns NaN; `errno` is set to `EDOM`.

`signal()` (7.7.1.1)

The signal part of the library is not supported.

Note: The default implementation of `signal` does not perform anything. Use the template source code to implement application-specific signal handling. See *signal*, page 144 and *raise*, page 142, respectively.

Terminating newline character (7.9.2)

`stdout` stream functions recognize either `newline` or end of file (EOF) as the terminating character for a line.

Blank lines (7.9.2)

Space characters written to the `stdout` stream immediately before a newline character are preserved. There is no way to read the line through the `stdin` stream that was written through the `stdout` stream.

Null characters appended to data written to binary streams (7.9.2)

No null characters are appended to data written to binary streams.

Files (7.9.3)

Whether the file position indicator of an append-mode stream is initially positioned at the beginning or the end of the file, depends on the application-specific implementation of the low-level file routines.

Whether a write operation on a text stream causes the associated file to be truncated beyond that point, depends on the application-specific implementation of the low-level file routines. See *Briefly about input and output (I/O)*, page 116.

The characteristics of the file buffering is that the implementation supports files that are unbuffered, line buffered, or fully buffered.

Whether a zero-length file actually exists depends on the application-specific implementation of the low-level file routines.

Rules for composing valid file names depends on the application-specific implementation of the low-level file routines.

Whether the same file can be simultaneously open multiple times depends on the application-specific implementation of the low-level file routines.

remove() (7.9.4.1)

The effect of a remove operation on an open file depends on the application-specific implementation of the low-level file routines. See *Briefly about input and output (I/O)*, page 116.

rename() (7.9.4.2)

The effect of renaming a file to an already existing filename depends on the application-specific implementation of the low-level file routines. See *Briefly about input and output (I/O)*, page 116.

%p in printf() (7.9.6.1)

The argument to a `%p` conversion specifier, print pointer, to `printf()` is treated as having the type `void *`. The value will be printed as a hexadecimal number, similar to using the `%x` conversion specifier.

%p in scanf() (7.9.6.2)

The `%p` conversion specifier, scan pointer, to `scanf()` reads a hexadecimal number and converts it into a value with the type `void *`.

Reading ranges in scanf() (7.9.6.2)

A – (dash) character is always treated as a range symbol.

File position errors (7.9.9.1, 7.9.9.4)

On file position errors, the functions `fgetpos` and `ftell` store `EFPOS` in `errno`.

Message generated by perror() (7.9.10.4)

The generated message is:

usersuppliedprefix:errormessage

Allocating zero bytes of memory (7.10.3)

The `calloc()`, `malloc()`, and `realloc()` functions accept zero as an argument. Memory will be allocated, a valid pointer to that memory is returned, and the memory block can be modified later by `realloc`.

Behavior of abort() (7.10.4.1)

The `abort()` function does not flush stream buffers, and it does not handle files, because this is an unsupported feature.

Behavior of exit() (7.10.4.3)

The argument passed to the `exit` function will be the return value returned by the `main` function to `cstartup`.

Environment (7.10.4.4)

The set of available environment names and the method for altering the environment list is described in *getenv*, page 141.

system() (7.10.4.5)

How the command processor works depends on how you have implemented the `system` function. See *system*, page 145.

Message returned by strerror() (7.11.6.2)

The messages returned by `strerror()` depending on the argument is:

Argument	Message
EZERO	no error
EDOM	domain error
ERANGE	range error
EFPOS	file positioning error
EILSEQ	multi-byte encoding error
<0 >99	unknown error
all others	error <i>nnn</i>

Table 53: Message returned by `strerror()`—DLIB runtime environment

The time zone (7.12.1)

The local time zone and daylight savings time implementation is described in *__time32*, *__time64*, page 145.

clock() (7.12.2.1)

From where the system clock starts counting depends on how you have implemented the `clock` function. See *clock*, page 140.

A

- a (ielfdump option) 492
- __AAPCS__ (predefined symbol) 420
- aapcs (compiler option) 260
- __AAPCS_VFP__ (predefined symbol) 420
- ABI, AEABI and IA64 217
- abort
 - implementation-defined behavior 527
 - implementation-defined behavior in C89 (DLIB) . . . 541
 - system termination (DLIB) 136
- __absolute (extended keyword) 347
- absolute location
 - data, placing at (@) 225
 - language support for 184
 - placing data in registers (@) 227
 - #pragma location 370
- advanced_heap (linker option) 301
- aeabi (compiler option) 260
- _AEABI_PORTABILITY_LEVEL (preprocessor symbol) 219
- _AEABI_PORTABLE (preprocessor symbol) 219
- algorithm (STL header file) 432
- alignment 329
 - forcing stricter (#pragma data_alignment) 363
 - in structures (#pragma pack) 374
 - in structures, causing problems 222
 - of an object (__ALIGNOF__) 184
 - of data types 330
 - restrictions for inline assembler 159
- alignment (pragma directive) 523, 537
- __ALIGNOF__ (operator) 184
- align_sp_on_irq (compiler option) 261
- all (ielfdump option) 492
- anonymous structures 223
- anonymous symbols, creating 181
- ANSI C. *See* C89
- application
 - building, overview of 63
 - execution, overview of 58
 - startup and termination (DLIB) 133
- argv (argument), implementation-defined behavior 516
- ARM
 - and Thumb code, overview 71
 - supported devices 50
- __arm (extended keyword) 347
- arm (compiler option) 261
- __ARMVFP__ (predefined symbol) 421
- __ARMVFPV2__ (predefined symbol) 421
- __ARMVFPV3__ (predefined symbol) 421
- __ARMVFPV4__ (predefined symbol) 421
- __ARMVFP_D16__ (predefined symbol) 421
- __ARMVFP_FP16__ (predefined symbol) 421
- __ARMVFP_SP__ (predefined symbol) 421
- __ARM_ADVANCED_SIMD__ (predefined symbol) . . 420
- __ARM_MEDIA__ (predefined symbol) 420
- __ARM_PROFILE_M__ (predefined symbol) 421
- __ARM4TM__ (predefined symbol) 422
- __ARM5__ (predefined symbol) 422
- __ARM5E__ (predefined symbol) 422
- __ARM6__ (predefined symbol) 422
- __ARM6M__ (predefined symbol) 422
- __ARM6SM__ (predefined symbol) 422
- __ARM7A__ (predefined symbol) 422
- __ARM7EM__ (predefined symbol) 422
- __ARM7M__ (predefined symbol) 422
- __ARM7R__ (predefined symbol) 422
- arrays
 - designated initializers in 181
 - implementation-defined behavior 520
 - implementation-defined behavior in C89 535
 - incomplete at end of structs 181
 - non-lvalue 187
 - of incomplete types 186
 - single-value initialization 188
- asm, __asm (language extension) 160
- assembler code
 - calling from C 166

calling from C++	168
inserting inline	158
assembler directives	
for call frame information	176
using in inline assembler code	159
assembler instructions	
for software interrupts	78
assembler instructions, inserting inline	158
assembler labels	
default for application startup	63, 105
making public (<code>--public_equ</code>)	290
assembler language interface	157
calling convention. <i>See</i> assembler code	
assembler list file, generating	275
assembler output file	168
asserts	
implementation-defined behavior of	524
implementation-defined behavior of in C89, (DLIB)	538
including in application	426
assert.h (DLIB header file)	430
__assignment_by_bitwise_copy_allowed, symbol used	
in library	435
@ (operator)	
placing at absolute address	225
placing in sections	226
atexit limit, setting up	106
atexit, reserving space for calls	106
attributes	
object	345
type	343
auto variables	68
at function entrance	172
programming hints for efficient code	234
using in inline assembler statements	159
auto, packing algorithm for initializers	447

B

backtrace information *See* call frame information

Barr, Michael	41
baseaddr (pragma directive)	523, 537
__BASE_FILE__ (predefined symbol)	422
--basic_heap (linker option)	301
basic_template_matching (pragma directive)	523, 537
batch files	
error return codes	247
none for building library from command line	124
--BE32 (linker option)	302, 304
--BE8 (linker option)	301
__big_endian (extended keyword)	347
big-endian (byte order)	65
--bin (ielftool option)	492
binary streams	525
binary streams in C89 (DLIB)	539
bit negation	236
bitfields	
data representation of	332
hints	221
implementation-defined behavior	521
implementation-defined behavior in C89	535
non-standard types in	184
bitfields (pragma directive)	361
bits in a byte, implementation-defined behavior	517
bitset (library header file)	432
bold style, in this guide	43
bool (data type)	331
adding support for in DLIB	430, 433
--bounds_table_size (linker option)	297
.bss (ELF section)	464
building_runtime (pragma directive)	523, 537
__BUILD_NUMBER__ (predefined symbol)	422
byte order	65
identifying	424

C

C and C++ linkage	170
C/C++ calling convention. <i>See</i> calling convention	

- C header files 429
- C language, overview 181
- call frame information 176
 - in assembler list file 168
 - in assembler list file (-IA) 276
- call frame information, disabling (--no_call_frame_info) . 278
- call graph root (stack usage control directive) 470
- call stack 176
- callee-save registers, stored on stack 68
- calling convention
 - C++, requiring C linkage 168
 - in compiler 169
- calloc (library function) 69
 - See also* heap
 - implementation-defined behavior in C89 (DLIB) . . . 540
- calls (pragma directive) 362
- call_graph (linker option) 302
- call_graph_root (pragma directive) 363
- call-info (in stack usage control file) 474
- can_instantiate (pragma directive) 523, 537
- cassert (library header file) 432
- cast operators
 - in Extended EC++ 192, 199
 - missing from Embedded C++ 192
- casting
 - of pointers and integers 338
 - pointers to integers, language extension 187
- category (in stack usage control file) 473
- ccomplex (library header file) 432
- cctype (DLIB header file) 433
- cerrno (DLIB header file) 433
- cexit (system termination code)
 - customizing system termination 137
 - in DLIB 133
- cfnv (library header file) 433
- CFI (assembler directive) 176
- CFI_COMMON_ARM (call frame information macro) . . 179
- CFI_COMMON_Thumb (call frame information macro) . 179
- CFI_NAMES_BLOCK (call frame information macro) . 179
- cfloat (DLIB header file) 433
- char (data type) 331
 - changing default representation (--char_is_signed) . . 262
 - changing representation (--char_is_unsigned) 262
 - implementation-defined behavior 518
 - signed and unsigned 332
- character set, implementation-defined behavior 516
- characters
 - implementation-defined behavior 517
 - implementation-defined behavior in C89 532
- char_is_signed (compiler option) 262
- char_is_unsigned (compiler option) 262
- check that (linker directive) 457
- checksum
 - calculation of 208
 - display format in C-SPY for symbol 216
- checksum (ielftool option) 493
- cinttypes (DLIB header file) 433
- ciso646 (library header file) 433
- climits (DLIB header file) 433
- clobber 159
- locale (DLIB header file) 433
- clock (DLIB library function),
 - implementation-defined behavior in C89 541
- clock (library function)
 - implementation-defined behavior 528
- __CLREX (intrinsic function) 390
- clustering (compiler transformation) 233
 - disabling (--no_clustering) 279
- __CLZ (intrinsic function) 390
- cmain (system initialization code)
 - in DLIB 133
- cmath (DLIB header file) 433
- CMSIS integration 219
- code
 - ARM and Thumb, overview 71
 - facilitating for good generation of 234
 - interruption of execution 74
- code (ielfdump option) 496

code motion (compiler transformation)	232
disabling (<code>--no_code_motion</code>)	279
codeseg (pragma directive)	523, 537
<code>__code</code> , symbol used in library	435
command line options	
<i>See also</i> compiler options	
<i>See also</i> linker options	
part of compiler invocation syntax	243
part of linker invocation syntax	244
passing	244
typographic convention	42
command prompt icon, in this guide	43
.comment (ELF section)	464
comments	
after preprocessor directives	187
C++ style, using in C code	181
common block (call frame information)	177
common subexpr elimination (compiler transformation)	232
disabling (<code>--no_cse</code>)	280
Common.i (CFI header example file)	179
compilation date	
exact time of (<code>__TIME__</code>)	425
identifying (<code>__DATE__</code>)	423
compiler	
environment variables	245
invocation syntax	243
output from	246
compiler listing, generating (<code>-l</code>)	275
compiler object file	56
including debug information in (<code>--debug, -r</code>)	265
output from compiler	246
compiler optimization levels	230
compiler options	253
passing to compiler	244
reading from file (<code>-f</code>)	273
specifying parameters	255
summary	255
syntax	253
for creating skeleton code	167
instruction scheduling	234
<code>--warnings_affect_exit_code</code>	247
compiler platform, identifying	424
compiler transformations	228
compiler version number	425
compiling	
from the command line	63
syntax	243
complex numbers, supported in Embedded C++	192
complex (library header file)	431
complex.h (library header file)	430
compound literals	181
computer style, typographic convention	42
<code>--config</code> (linker option)	302
configuration	
basic project settings	63
<code>__low_level_init</code>	137
configuration file for linker. <i>See</i> linker configuration file	
configuration symbols	
for file input and output	148
for locale	149
for printf and scanf	147
in library configuration files	123
in linker configuration files	457
specifying for linker	303
<code>--config_def</code> (linker option)	303
<code>--config_search</code> (linker option)	303
consistency, module	113
const	
declaring objects	342
non-top level	187
<code>__constrange()</code> , symbol used in library	435
<code>__construction_by_bitwise_copy_allowed</code> , symbol used in library	435
constseg (pragma directive)	523, 537
const_cast (cast operator)	192
contents, of this guide	38
control characters,	
implementation-defined behavior	530
conventions, used in this guide	42

copyright notice 2

__CORE__ (predefined symbol). 422

core

- identifying 422
- selecting. 64

Cortex-M7 220

Cortex, special considerations for interrupt functions. 73

cos (library function) 428

cos (library routine) 132–133

cosf (library routine). 132–133

cosl (library routine). 132–133

__COUNTER__ (predefined symbol). 422

__cplusplus (predefined symbol) 422

--cpp_init_routine (linker option) 304

--cpu (compiler option) 262

__CPU_MODE__ (predefined symbol) 422

--cpu_mode (compiler option) 263

CPU, specifying on command line for compiler 262

--create (iarchive option). 496

csetjmp (DLIB header file). 433

csignal (DLIB header file) 433

cspy_support (pragma directive). 523, 537

CSTACK (ELF block) 464

See also stack

 setting up size for. 105

cstartup (system startup code)

- customizing system initialization. 137
- source files for (DLIB). 133

cstdarg (DLIB header file) 433

cstdbool (DLIB header file) 433

cstddef (DLIB header file) 433

cstdio (DLIB header file) 433

cstdlib (DLIB header file). 433

cstring (DLIB header file). 433

ctgmath (library header file) 433

ctime (DLIB header file). 433

ctype.h (library header file). 430

cwctype.h (library header file) 433

C_INCLUDE (environment variable). 245

C-SPY

- debug support for C++. 199
- interface to system termination 137

C-STAT for static analysis, documentation for. 40

C++

- See also* Embedded C++ and Extended Embedded C++
- absolute location 226
- calling convention 168
- header files. 430
- language extensions. 200
- standard template library (STL). 432
- static member variables 226
- support for 49

--c++ (compiler option) 264

C++ header files 431

C++ terminology. 42

C++-style comments. 181

C89

- implementation-defined behavior 531
- support for 181

--c89 (compiler option). 261

C99. *See* Standard C

D

-D (compiler option). 264

-d (iarchive option). 497

data

- alignment of. 329
- different ways of storing 67
- located, declaring extern 226
- placing 224, 292
 - at absolute location 225
- placing in registers. 227
- representation of 329
- storage 67

data block (call frame information). 177

data pointers 338

data types	331	deque (STL header file)	432
floating point	336	destructors and interrupts, using	198
in C++	342	device description files, preconfigured for C-SPY	50
integer types	331	diagnostic messages	248
dataseg (pragma directive)	523, 537	classifying as compilation errors	266
data_alignment (pragma directive)	363	classifying as compilation remarks	267
.data_init (ELF section)	465	classifying as compiler warnings	267
__data, symbol used in library	435	classifying as errors	281, 317
__DATE__ (predefined symbol)	423	classifying as linker warnings	307
date (library function), configuring support for	121	classifying as linking errors	306
DC32 (assembler directive)	159	classifying as linking remarks	306
--debug (compiler option)	265	disabling compiler warnings	287
debug information, including in object file	265	disabling linker warnings	320
.debug (ELF section)	464	disabling wrapping of in compiler	287
--debug_heap (linker option)	297	disabling wrapping of in linker	321
decimal point, implementation-defined behavior	529	enabling compiler remarks	290
declarations		enabling linker remarks	323
empty	188	listing all used by compiler	268
in for loops	181	listing all used by linker	307
Kernighan & Ritchie	236	suppressing in compiler	267
of functions	170	suppressing in linker	307
declarations and statements, mixing	181	diagnostics	
declarators, implementation-defined behavior in C89	536	iarchive	479
default_no_bounds (pragma directive)	359	iobjmanip	485
--default_to_complex_ranges (linker option)	304	isymexport	490
define block (linker directive)	444	--diagnostics_tables (compiler option)	268
define memory (linker directive)	439	--diagnostics_tables (linker option)	307
define overlay (linker directive)	445	diagnostics, implementation-defined behavior	515
define region (linker directive)	439	diag_default (pragma directive)	366
define symbol (linker directive)	457	--diag_error (compiler option)	266
--define_symbol (linker option)	305	--diag_error (linker option)	306
define_type_info (pragma directive)	523, 537	--no_fragments (compiler option)	281
define_without_bounds (pragma directive)	360	--no_fragments (linker option)	317
define_with_bounds (pragma directive)	359	diag_error (pragma directive)	366
--delete (iarchive option)	497	--diag_remark (compiler option)	267
delete (keyword)	69	--diag_remark (linker option)	306
denormalized numbers. <i>See</i> subnormal numbers		diag_remark (pragma directive)	366
--dependencies (compiler option)	265	--diag_suppress (compiler option)	267
--dependencies (linker option)	305	--diag_suppress (linker option)	307

diag_suppress (pragma directive) 367
 --diag_warning (compiler option) 267
 --diag_warning (linker option) 307
 diag_warning (pragma directive) 367
 directives
 pragma 51, 359
 to the linker 437
 directory, specifying as parameter 254
 disable_check (pragma directive) 360
 __disable_fiq (intrinsic function) 390
 __disable_interrupt (intrinsic function) 391
 __disable_irq (intrinsic function) 391
 --discard_unused_publics (compiler option) 268
 disclaimer 2
 DLIB 429
 configurations 125
 configuring 123, 269
 documentation 40
 naming convention 43
 reference information. *See* the online help system . . . 427
 runtime environment 115
 --dlib_config (compiler option) 269
 DLib_Defaults.h (library configuration file) 123
 __DLIB_FILE_DESCRIPTOR (configuration symbol) . . 148
 __DMB (intrinsic function) 391
 do not initialize (linker directive) 449
 document conventions 42
 documentation
 contents of this 38
 how to use this 37
 overview of guides 39
 who should read this 37
 domain errors, implementation-defined behavior 524
 domain errors, implementation-defined behavior in C89
 (DLIB) 538
 double (data type) 336
 do_not_instantiate (pragma directive) 523, 537
 --do_segment_pad (linker option) 308
 __DSB (intrinsic function) 391

dynamic initialization 133
 and C++ 92
 dynamic memory 69
 dynamic RTTI data, including in the image 316

E

-e (compiler option) 269
 early_initialization (pragma directive) 523, 537
 --ec++ (compiler option) 270
 --edit (isymexport option) 497
 edition, of this guide 2
 --eec++ (compiler option) 270
 ELF utilities 477
 Embedded C++ 191
 differences from C++ 192
 enabling 270
 function linkage 170
 language extensions 191
 overview 191
 Embedded C++ Technical Committee 42
 embedded systems, IAR special support for 51
 __embedded_cplusplus (predefined symbol) 423
 empty region (in linker configuration file) 442
 __enable_fiq (intrinsic function) 391
 --enable_hardware_workaround (compiler option) 270
 --enable_hardware_workaround (linker option) 308
 __enable_interrupt (intrinsic function) 392
 __enable_irq (intrinsic function) 392
 --enable_multibytes (compiler option) 271
 --enable_restrict (compiler option) 271
 enabling restrict keyword 271
 endianness. *See* byte order
 --entry (linker option) 309
 entry label, program 134
 enumerations
 implementation-defined behavior 521
 implementation-defined behavior in C89 535

enums	
data representation	331
forward declarations of	186
--enum_is_int (compiler option)	272
environment	
implementation-defined behavior	516
implementation-defined behavior in C89	531
runtime (DLIB)	115
environment names, implementation-defined behavior	517
environment variables	
C_INCLUDE	245
ILINKARM_CMD_LINE	245
QCCARM	245
environment (native), implementation-defined behavior	530
EQU (assembler directive)	290
ERANGE	525
ERANGE (C89)	539
errno value at underflow, implementation-defined behavior	527
errno.h (library header file)	430
error messages	250
classifying	281, 317
classifying for compiler	266
classifying for linker	306
range	111
error return codes	247
error (linker directive)	460
error (pragma directive)	367
--error_limit (compiler option)	272
--error_limit (linker option)	309
escape sequences, implementation-defined behavior	517
exception flags, for floating-point values	336
exception handling, missing from Embedded C++	192
exception (library header file)	431
exceptions, code for in section	465
--exception_tables (linker option)	310
exclude (stack usage control directive)	470
.exc.text (ELF section)	465
_Exit (library function)	136
exit (library function)	136
implementation-defined behavior	527
implementation-defined behavior in C89	541
_exit (library function)	136
_exit (library function)	136
exp (library routine)	132
expf (library routine)	132
expl (library routine)	132
export keyword, missing from Extended EC++	199
export (linker directive)	458
--export_builtin_config (linker option)	310
expressions (in linker configuration file)	459
extended command line file	
for compiler	273
for linker	311
passing options	244
Extended Embedded C++	192
enabling	270
extended keywords	343
enabling (-e)	269
overview	51
summary	346
syntax	
object attributes	345
type attributes on data objects	344
type attributes on functions	344
extended-selectors (in linker configuration file)	455
extern "C" linkage	197
--extract (iarchive option)	497
--extra_init (linker option)	311

F

-f (compiler option)	273
-f (IAR utility option)	498
-f (linker option)	311
fast interrupts	76
fatal error messages	250
fdopen, in stdio.h	434

- fegettrapdisable. 434
- fegettrappable. 434
- FENV_ACCESS, implementation-defined behavior . . . 520
- fenv.h (library header file) 430, 433
 - additional C functionality. 434
- fgetpos (library function), implementation-defined behavior 527
- fgetpos (library function), implementation-defined behavior in C89 540
- __FILE__ (predefined symbol). 423
- file buffering, implementation-defined behavior. 526
- file dependencies, tracking 265
- file input and output
 - configuration symbols for 148
- file paths, specifying for #include files 275
- file position, implementation-defined behavior. 525
- file streams lock interface 153
- file (zero-length), implementation-defined behavior. . . . 526
- filename
 - extension for device description files. 50
 - extension for header files. 50
 - of object executable image. 321
 - of object file. 288, 321
 - search procedure for. 245
 - specifying as parameter 254
- filenames (legal), implementation-defined behavior. . . . 526
- fileno, in stdio.h 434
- files, implementation-defined behavior
 - handling of temporary 526
 - multibyte characters in. 526
 - opening 526
- fill (ielftool option). 499
- __fiq (extended keyword). 348
- float (data type). 336
- floating-point constants
 - hexadecimal notation 181
 - hints 222
- floating-point environment, accessing or not 378
- floating-point expressions
 - contracting or not. 378
- floating-point format. 336
 - hints 221–222
 - implementation-defined behavior 519
 - implementation-defined behavior in C89. 534
 - special cases. 337
 - 32-bits 336
 - 64-bits 337
- floating-point status flags 434
- floating-point unit. 274
- float.h (library header file) 430
- FLT_EVAL_METHOD, implementation-defined behavior 519, 524, 529
- FLT_ROUNDS, implementation-defined behavior 519, 528–529
- fmod (library function),
 - implementation-defined behavior in C89 539
- for loops, declarations in. 181
- force_exceptions (linker option). 312
- force_output (linker option). 312
- formats
 - floating-point values 336
 - standard IEEE (floating point) 336
- fpu (compiler option) 273
- FP_CONTRACT, implementation-defined behavior . . . 520
- fragmentation, of heap memory 70
- free (library function). *See also* heap 69
- fsetpos (library function), implementation-defined behavior 527
- fstream (library header file) 431
- ftell (library function), implementation-defined behavior . 527
 - in C89 540
- Full DLIB (library configuration) 125
- __func__ (predefined symbol) 188, 423
- __FUNCTION__ (predefined symbol) 188, 423
- function calls
 - calling convention 169
 - eliminating overhead of by inlining 80
 - preserved registers across. 171
- function declarations, Kernighan & Ritchie 236
- function execution, in RAM 72

function inlining (compiler transformation)	232
disabling (--no_inline)	281
function pointers	337
function prototypes	235
enforcing	291
function (pragma directive)	523, 537
function (stack usage control directive)	470
functional (STL header file)	432
functions	71
declaring	170, 235
inlining	181, 232, 234, 368
interrupt	74
intrinsic	157, 235
parameters	172
placing in memory	224, 226, 292
recursive	
avoiding	235
storing data on stack	68
reentrancy (DLIB)	428
related extensions	71
return values from	173
function_effects (pragma directive)	523, 537
function-spec (in stack usage control file)	473

G

-g (ielfdump option)	508
generate_entry_without_bounds (pragma directive)	360
--generate_vfe_header (isymexport option)	499
getw, in stdio.h	434
getzone (library function), configuring support for	121
__get_BASEPRI (intrinsic function)	392
__get_CONTROL (intrinsic function)	392
__get_CPSR (intrinsic function)	392
__get_FAULTMASK (intrinsic function)	393
__get_FPSCR (intrinsic function)	393
__get_interrupt_state (intrinsic function)	393
__get_IPSR (intrinsic function)	394
__get_LR (intrinsic function)	394

__get_MSP (intrinsic function)	394
__get_PRIMASK (intrinsic function)	394
__get_PSP (intrinsic function)	394
__get_PSR (intrinsic function)	395
__get_SB (intrinsic function)	395
__get_SP (intrinsic function)	395
global variables	
affected by static clustering	233
handled during system termination	136
hints for not using	234
initialized during system startup	135
GRP_COMDAT, group type	486
--guard_calls (compiler option)	274
guidelines, reading	37

H

Harbison, Samuel P.	41
hardware support in compiler	115
hash_map (STL header file)	432
hash_set (STL header file)	432
__has_constructor, symbol used in library	436
__has_destructor, symbol used in library	436
hdrstop (pragma directive)	523, 537
header files	
C	429
C++	430–431
library	427
special function registers	237
STL	432
DLib_Defaults.h	123
including stdbool.h for bool	331
including stddef.h for wchar_t	332
header names, implementation-defined behavior	522
--header_context (compiler option)	274
heap	
advanced versus basic	205
dynamic memory	69
storing data	67

VLA allocated on. 295

heap sections

 placing. 106

heap size

 and standard I/O. 206

 changing default. 106

HEAP (ELF section). 465

heap (zero-sized), implementation-defined behavior . . . 527

hide (isymexport directive). 489

hints

 for good code generation 234

 implementation-defined behavior 520

 using efficient data types 221

I

-I (compiler option). 275

IAR Command Line Build Utility. 124

IAR Systems Technical Support. 250

iarbuild.exe (utility) 124

iarchive. 477

 commands summary 478

 options summary 479

__iar_cos_accurate (library routine) 133

__iar_cos_accuratef (library routine) 133

__iar_cos_accuratef (library function) 428

__iar_cos_accuratel (library routine) 133

__iar_cos_accuratel (library function) 428

__iar_cos_small (library routine) 132

__iar_cos_smallf (library routine) 132

__iar_cos_smallll (library routine) 132

__IAR_DLIB_PERTHREAD_INIT_SIZE (macro) 154

__IAR_DLIB_PERTHREAD_SIZE (macro) 154

__IAR_DLIB_PERTHREAD_SYMBOL_OFFSET (symbolptr) 154

iar_dlmalloc.h (library header file)

 additional C functionality. 434

__iar_exp_small (library routine) 132

__iar_exp_smallf (library routine) 132

__iar_exp_smallll (library routine) 132

__iar_FPow (library routine) 133

__iar_FSin (library routine) 132

__iar_log_small (library routine) 132

__iar_log_smallf (library routine) 132

__iar_log_smallll (library routine) 132

__iar_log10_small (library routine) 132

__iar_log10_smallf (library routine) 132

__iar_log10_smallll (library routine) 132

__iar_LPow (library routine) 133

__iar_LSin (library routine) 132–133

__iar_maximum_atexit_calls 106

__iar_Pow (library routine) 133

__iar_Pow_accurate (library routine) 133

__iar_pow_accurate (library routine) 133

__iar_Pow_accuratef (library routine) 133

__iar_pow_accuratef (library routine) 133

__iar_pow_accuratef (library function) 428

__iar_Pow_accuratel (library routine) 133

__iar_pow_accuratel (library routine) 133

__iar_pow_accuratel (library function) 428

__iar_pow_small (library routine) 132

__iar_pow_smallf (library routine) 132

__iar_pow_smallll (library routine) 132

__iar_program_start (label). 134

__iar_ReportAssert (library function) 140

__iar_Sin (library routine) 132–133

__iar_Sin_accurate (library routine) 133

__iar_sin_accurate (library routine) 133

__iar_Sin_accuratef (library routine) 133

__iar_sin_accuratef (library routine) 133

__iar_sin_accuratef (library function) 428

__iar_Sin_accuratel (library routine) 133

__iar_sin_accuratel (library routine) 133

__iar_sin_accuratel (library function) 428

__iar_Sin_small (library routine) 132

__iar_sin_small (library routine) 132

__iar_Sin_smallf (library routine) 132

__iar_sin_smallf (library routine) 132

__iar_Sin_smallf (library routine)	132
__iar_sin_smallf (library routine)	132
__IAR_SYSTEMS_ICC__ (predefined symbol)	424
__iar_tan_accurate (library routine)	133
__iar_tan_accuratef (library routine)	133
__iar_tan_accuratef (library function)	428
__iar_tan_accuratel (library routine)	133
__iar_tan_accuratel (library function)	428
__iar_tan_small (library routine)	132
__iar_tan_smallf (library routine)	132
__iar_tan_smallf (library routine)	132
.iar.debug (ELF section)	464
.iar.dynexit (ELF section)	466
IA64 ABI	217
__ICCARM__ (predefined symbol)	424
icons, in this guide	43
IDE	
building a library from	123
overview of build tools	47
identifiers, implementation-defined behavior	517
identifiers, implementation-defined behavior in C89	532
IEEE format, floating-point values	336
ielfdump	482
options summary	483
ielftool	480
options summary	481
if (linker directive)	461
--ignore_uninstrumented_pointers (linker option)	298
--ihex (ielftool option)	500
ILINK options. <i>See</i> linker options	
ILINKARM_CMD_LINE (environment variable)	245
ILINK. <i>See</i> linker	
--image_input (linker option)	312
important_typedef (pragma directive)	523, 537
include files	
including before source files	289
specifying	245
include (linker directive)	461
include_alias (pragma directive)	368

infinity	337
infinity (style for printing), implementation-defined	
behavior	527
inheritance, in Embedded C++	191
initialization	
changing default.	106
C++ dynamic	92
dynamic	133
manual	107
packing algorithm for.	107
single-value	188
suppressing	106
initialize (linker directive)	446
initializers, static	187
.init_array (section)	466
--inline (linker option)	313
inline assembler	158
avoiding.	235
for passing values between C and assembler	239
<i>See also</i> assembler language interface	
inline functions	181
in compiler.	232
inline (pragma directive)	368
inlining functions	80
implementation-defined behavior	521
installation directory	42
instantiate (pragma directive)	523, 537
instruction scheduling (compiler option)	234
int (data type) signed and unsigned.	331
integer types	331
casting	338
implementation-defined behavior	519
intptr_t	338
ptrdiff_t	338
size_t	338
uintptr_t	338
integers, implementation-defined behavior in C89	533
integral promotion	236
Intel hex	203
Intel IA64 ABI	217

internal error 250

interrupt functions 74

 fast interrupts 76

 in Cortex 73

 nested interrupts 77

 operations 79

 software interrupts 78

interrupt handler. *See* interrupt service routine

interrupt service routine 74

interrupt state, restoring 405

interrupt vector table 79

 start address for 75

interrupts

 processor state 68

 using with EC++ destructors 198

__interwork (extended keyword) 348

--interwork (compiler option) 275

intptr_t (integer type) 338

__intrinsic (extended keyword) 348

intrinsic functions 235

 for Neon 389

 overview 157

 summary 383

intrinsics.h (header file) 383

inttypes.h (library header file) 430

.intvec (ELF section) 466

invocation syntax 243

iobjmanip 484

 options summary 484

iomanip (library header file) 431

ios (library header file) 431

iosfwd (library header file) 431

iostream (library header file) 431

__irq (extended keyword) 348

IRQ_STACK (section) 466

__ISB (intrinsic function) 395

iso646.h (library header file) 430

istream (library header file) 431

isymexport 486

 options summary 488

italic style, in this guide 42–43

iterator (STL header file) 432

I/O register. *See* SFR

J

Josuttis, Nicolai M. 41

K

--keep (linker option) 313

keep (linker directive) 449

keep_definition (pragma directive) 523, 538

Kernighan & Ritchie function declarations 236

 disallowing 291

Kernighan, Brian W. 41

keywords 343

 extended, overview of 51

L

-L (linker option) 323

-l (compiler option) 275

 for creating skeleton code 167

labels 188

 assembler, making public 290

 __iar_program_start 134

 __program_start 134

Labrosse, Jean J. 41

Lajoie, Josée 41

language extensions

 Embedded C++ 191

 enabling using pragma 369

 enabling (-e) 269

language overview 49

language (pragma directive) 369

__LDC (intrinsic function) 395

__LDCL (intrinsic function)	395	__LINE__ (predefined symbol)	424
__LDCL_noidx (intrinsic function)	396	linkage, C and C++	170
__LDC_noidx (intrinsic function)	396	linker	83
__LDC2 (intrinsic function)	395	output from	248
__LDC2L (intrinsic function)	395	linker configuration file	
__LDC2L_noidx (intrinsic function)	396	for placing code and data	87
__LDC2_noidx (intrinsic function)	396	in depth	437
__LDREX (intrinsic function)	397	overview of	437
__LDREXB (intrinsic function)	397	selecting	101
__LDREXD (intrinsic function)	397	linker object executable image	
__LDREXH (intrinsic function)	397	specifying filename of (-o)	321
--legacy (compiler option)	277	linker options	297
libraries	220	reading from file (-f)	311
reason for using	56	summary	297
standard template library	432	typographic convention	42
using a prebuilt	126	linking	
library configuration files		from the command line	63
DLIB	125	in the build process	56
DLib_Defaults.h	123	introduction	83
modifying	123	process for	85
specifying	269	Lippman, Stanley B.	41
library documentation	427	list (STL header file)	432
library features, missing from Embedded C++	192	listing, generating	275
library files, linker search path to (--search)	323	literals, compound	181
library functions		literature, recommended	41
summary, DLIB	429	__LITTLE_ENDIAN__ (predefined symbol)	424
online help for	41	__little_endian (extended keyword)	349
library header files	427	little-endian (byte order)	65
library modules		local symbols, removing from ELF image	319
introduction	84	local variables, <i>See</i> auto variables	
overriding	122	locale	
library object files	428	adding support for in library	149
library project, building using a template	123	changing at runtime	150
library_default_requirements (pragma directive)	523, 538	implementation-defined behavior	518, 529
library_provides (pragma directive)	523, 538	library header file	431
library_requirement_override (pragma directive)	523, 538	removing support for	149
lightbulb icon, in this guide	43	support for	148
limits (library header file)	431	locale.h (library header file)	430
limits.h (library header file)	430	located data, declaring extern	226
		location (pragma directive)	225, 370

--lock_regs (compiler option) 276
 --log (linker option) 314
 log (library routine) 132
 logf (library routine) 132
 logl (library routine) 132
 --log_file (linker option) 315
 log10 (library routine) 132
 log10f (library routine) 132
 log10l (library routine) 132
 long double (data type) 336
 long float (data type), synonym for double 187
 long long (data type) signed and unsigned 331
 long (data type) signed and unsigned 331
 longjmp, restrictions for using 429
 loop unrolling (compiler transformation) 232
 disabling 286
 #pragma unroll 379
 loop-invariant expressions 232
 __low_level_init 134
 customizing 137
 initialization phase 59
 low_level_init.c 133
 low-level processor operations 182
 accessing 157
 lz77, packing algorithm for initializers 447

M

macros

embedded in #pragma optimize 373
 ERANGE (in errno.h) 525, 539
 inclusion of assert 426
 NULL, implementation-defined behavior 525
 in C89 for DLIB 538
 substituted in #pragma directives 182
 variadic 181
 --macro_positions_in_diagnostics (compiler option) 277
 main (function)
 definition (C89) 531

implementation-defined behavior 516
 --make_all_definitions_weak (compiler option) 277
 malloc (library function)
 See also heap 69
 implementation-defined behavior in C89 540
 --mangled_names_in_messages (linker option) 315
 Mann, Bernhard 41
 -map (linker option) 315
 map file, producing 315
 map (STL header file) 432
 math functions rounding mode,
 implementation-defined behavior 528
 math functions (library functions) 131
 math.h (library header file) 430
 max recursion depth (stack usage control directive) 471
 MB_LEN_MAX, implementation-defined behavior 528
 __MCR (intrinsic function) 397
 __MCR2 (intrinsic function) 397
 memory
 allocating in C++ 69
 dynamic 69
 heap 69
 non-initialized 239
 RAM, saving 235
 releasing in C++ 69
 stack 68
 saving 235
 used by global or static variables 67
 memory clobber 159
 memory management, type-safe 191
 memory map
 initializing SFRs 137
 linker configuration for 102
 output from linker 248
 producing (--map) 315
 memory (pragma directive) 523, 538
 memory (STL header file) 432
 __memory_of
 symbol used in library 436
 -merge_duplicate_sections (linker option) 316

message (pragma directive)	371
messages	
disabling	293, 324
forcing	371
Meyers, Scott	41
--mfc (compiler option)	278
migration, from earlier IAR compilers	40
MISRA C	
documentation	40
--misrac (compiler option)	257
--misrac (linker option)	299
--misrac_verbose (compiler option)	257
--misrac_verbose (linker option)	299
--misrac1998 (compiler option)	257
--misrac1998 (linker option)	299
--misrac2004 (compiler option)	257
--misrac2004 (linker option)	299
mode changing, implementation-defined behavior	527
module consistency	113
rtmodel	375
modules, introduction	84
module_name (pragma directive)	523, 538
module-spec (in stack usage control file)	473
Motorola S-records	203
__MRC (intrinsic function)	398
__MRC2 (intrinsic function)	398
multibyte character support	271
multibyte characters, implementation-defined behavior	517, 529
multiple inheritance	
in Extended EC++	192
missing from Embedded C++	192
multithreaded environment	150
multi-file compilation	229
mutable attribute, in Extended EC++	192, 199

N

name (in stack usage control file)	474
--	-----

names block (call frame information)	176
namespace support	
in Extended EC++	192, 199
missing from Embedded C++	192
naming conventions	43
NaN	
implementation of	337
implementation-defined behavior	527
native environment,	
implementation-defined behavior	530
NDEBUG (preprocessor symbol)	426
Neon intrinsic functions	389
__nested (extended keyword)	349
nested interrupts	77
new (keyword)	69
new (library header file)	431
no calls from (stack usage control directive)	471
.noinit (ELF section)	467
non-initialized variables, hints for	239
non-scalar parameters, avoiding	235
NOP (assembler instruction)	399
__noreturn (extended keyword)	351
Normal DLIB (library configuration)	125
Not a number (NaN)	337
--no_alignment_reduction (compiler option)	278
__no_alloc (extended keyword)	350
__no_alloc_str (operator)	350
__no_alloc_str16 (operator)	350
__no_alloc16 (extended keyword)	350
no_bounds (pragma directive)	360
--no_call_frame_info (compiler option)	278
--no_clustering (compiler option)	279
--no_code_motion (compiler option)	279
--no_const_align (compiler option)	279
--no_cse (compiler option)	280
--no_dynamic_rtti_elimination (linker option)	316
--no_entry (linker option)	317
--no_exceptions (compiler option)	280
--no_exceptions (linker option)	317
--no_header (ielfdump option)	500

__no_init (extended keyword) 239, 351
 --no_inline (compiler option) 281
 --no_library_search (linker option) 318
 --no_literal_pool (compiler option) 281
 --no_literal_pool (linker option) 318
 --no_locals (linker option) 319
 --no_loop_align (compiler option) 282
 --no_mem_idioms (compiler option) 282
 __no_operation (intrinsic function) 399
 --no_path_in_file_macros (compiler option) 282
 no_pch (pragma directive) 523, 538
 --no_range_reservations (linker option) 319
 --no_rel_section (ielfdump option) 500
 --no_remove (linker option) 319
 --no_rtti (compiler option) 283
 --no_rw_dynamic_init (compiler option) 283
 --no_scheduling (compiler option) 283
 --no_size_constraints (compiler option) 284
 --no_static_destruction (compiler option) 284
 --no_strtab (ielfdump option) 501
 --no_system_include (compiler option) 284
 --no_tbaa (compiler option) 285
 --no_typedefs_in_diagnostics (compiler option) 285
 --no_unaligned_access (compiler option) 285
 --no_unroll (compiler option) 286
 --no_var_align (compiler option) 286
 --no_veneers (linker option) 320
 --no_vfe (linker option) 320
 --no_warnings (compiler option) 287
 --no_warnings (linker option) 320
 --no_wrap_diagnostics (compiler option) 287
 --no_wrap_diagnostics (linker option) 321
 NULL
 implementation-defined behavior 525
 implementation-defined behavior in C89 (DLIB) 538
 pointer constant, relaxation to Standard C 186
 numbers (in linker configuration file) 460
 numeric conversion functions,
 implementation-defined behavior 530
 numeric (STL header file) 432

O

-O (compiler option) 287
 -o (compiler option) 288
 -o (iarchive option) 502
 -o (ielfdump option) 502
 -o (linker option) 321
 object attributes 345
 object filename, specifying (-o) 288, 321
 object files, linker search path to (--search) 323
 object_attribute (pragma directive) 239, 371
 --offset (ielftool option) 501
 once (pragma directive) 523, 538
 --only_stdout (compiler option) 288
 --only_stdout (linker option) 321
 operators
 See also @ (operator)
 for cast
 in Extended EC++. 192
 missing from Embedded C++. 192
 for region expressions 442
 for section control 185
 precision for 32-bit float 336
 precision for 64-bit float 337
 sizeof, implementation-defined behavior 529
 variants for cast 199
 __Pragma (preprocessor) 181
 __ALIGNOF__, for alignment control 184
 ?, language extensions for 201
 optimization
 clustering, disabling 279
 code motion, disabling 279
 common sub-expression elimination, disabling 280
 configuration 65
 disabling 231
 function inlining, disabling (--no_inline) 281
 hints 234
 loop unrolling, disabling 286
 scheduling, disabling 283

specifying (-O)	287
techniques	231
type-based alias analysis, disabling (--tbaa)	285
using inline assembler code	159
using pragma directive	372
optimization levels	230
optimize (pragma directive)	372
option parameters	253
options, compiler. <i>See</i> compiler options	
options, iarchive. <i>See</i> iarchive options	
options, ielfdump. <i>See</i> ielfdump options	
options, ielftool. <i>See</i> ielftool options	
options, iobjmanip. <i>See</i> iobjmanip options	
options, isymexport. <i>See</i> isymexport options	
options, linker. <i>See</i> linker options	
--option_name (compiler option)	309
Oram, Andy	41
ostream (library header file)	431
output	
from preprocessor	289
specifying for linker	63
--output (compiler option)	288
--output (iarchive option)	502
--output (ielfdump option)	502
--output (linker option)	321
overhead, reducing	232

P

pack (pragma directive)	373
packbits, packing algorithm for initializers	447
__packed (extended keyword)	352
packed structure types	339
packing, algorithms for initializers	447
parameters	
function	172
hidden	172
non-scalar, avoiding	235
register	172

rules for specifying a file or directory	254
specifying	255
stack	172–173
typographic convention	42
--parity (ielftool option)	502
part number, of this guide	2
__pcrel (extended keyword)	346
permanent registers	171
perorr (library function), implementation-defined behavior in C89	540
--pi_veneers (linker option)	321
__PKHBT (intrinsic function)	399
__PKHTB (intrinsic function)	399
place at (linker directive)	450
place in (linker directive)	451
placement	
in named sections	226
of code and data, introduction to	87
--place_holder (linker option)	322
plain char, implementation-defined behavior	518
__PLD (intrinsic function)	400
__PLDW (intrinsic function)	400
__PLI (intrinsic function)	400
pointer types	337
mixing	187
pointers	
casting	338
data	338
function	337
implementation-defined behavior	520
implementation-defined behavior in C89	535
polymorphic RTTI data, including in the image	316
polymorphism, in Embedded C++	191
porting, code containing pragma directives	361
possible calls (stack usage control directive)	472
pow (library routine)	132–133
alternative implementation of	428
powf (library routine)	132–133
powl (library routine)	132–133

- pragma directives 51
 - summary 359
 - for absolute located data 225
 - list of all recognized 522
 - list of all recognized (C89). 537
 - pack 373
- _Pragma (preprocessor operator) 181
- predefined symbols
 - overview 51
 - summary 420
- predef_macro (compiler option). 288
- Prefetch_Handler (exception function). 76
- preinclude (compiler option) 289
- .preinit_array (section) 467
- .prepreinit_array (section). 467
- preprocess (compiler option) 289
- preprocessor
 - operator (_Pragma) 181
 - output. 289
- preprocessor directives
 - comments at the end of 187
 - implementation-defined behavior 522
 - implementation-defined behavior in C89. 536
 - #pragma 359
- preprocessor extensions
 - __VA_ARGS__ 181
 - #warning message 426
- preprocessor symbols 420
 - defining 264, 305
- preserved registers 171
- __PRETTY_FUNCTION__ (predefined symbol). 424
- print formatter, selecting. 129
- printf (library function). 128
 - choosing formatter 128
 - configuration symbols 147
 - implementation-defined behavior 527
 - implementation-defined behavior in C89. 540
- __printf_args (pragma directive). 374
- printing characters, implementation-defined behavior . . . 529

- processor configuration. 64
- processor operations
 - accessing 157
 - low-level 182
- program entry label. 134
- program termination, implementation-defined behavior . . 516
- programming hints 234
- __program_start (label). 134
- projects
 - basic settings for 63
 - setting up for a library 123
- prototypes, enforcing 291
- ptrdiff_t (integer type) 338
- PUBLIC (assembler directive) 290
- publication date, of this guide. 2
- public_equ (compiler option) 289
- public_equ (pragma directive) 374
- putenv (library function), absent from DLIB 141
- putw, in stdio.h 434

Q

- __QADD (intrinsic function) 400
- __QADD8 (intrinsic function) 401
- __QADD16 (intrinsic function) 401
- __QASX (intrinsic function). 401
- QCCARM (environment variable) 245
- __QCFlag (intrinsic function). 401
- __QDADD (intrinsic function). 400
- __QDOUBLE (intrinsic function). 401
- __QDSUB (intrinsic function) 400
- __QFlag (intrinsic function) 402
- __QSAX (intrinsic function). 401
- __QSUB (intrinsic function). 400
- __QSUB16 (intrinsic function). 401
- __QSUB8 (intrinsic function). 401
- qualifiers
 - const and volatile 340
 - implementation-defined behavior 521

implementation-defined behavior in C89	536
queue (STL header file)	432

R

-r (compiler option)	265
-r (iarchive option)	506

RAM

example of declaring region	88
execution	72
initializers copied from ROM	61
running code from	109
saving memory	235
__ramfunc (extended keyword)	353
--ram_reserve_ranges (ismexport option)	503
--range (ielfdump option)	504
range errors	111
--raw (ielfdump option)	504
__RBIT (intrinsic function)	402
read formatter, selecting	130
reading guidelines	37
reading, recommended	41
realloc (library function)	69
implementation-defined behavior in C89	540
<i>See also</i> heap	
recursive functions	
avoiding	235
storing data on stack	68
--redirect (linker option)	322
reentrancy (DLIB)	428
reference information, typographic convention	43
region expression (in linker configuration file)	441
region literal (in linker configuration file)	440
register keyword, implementation-defined behavior	520
register parameters	172
registered trademarks	2
registers	
assigning to parameters	173
callee-save, stored on stack	68

for function returns	174
implementation-defined behavior in C89	535
in assembler-level routines	169
preserved	171
scratch	171
reinterpret_cast (cast operator)	192
.rel (ELF section)	464
.rela (ELF section)	464
--relaxed_fp (compiler option)	290
relay, <i>see</i> veneers	111
relocation errors, resolving	111
remark (diagnostic message)	249
classifying for compiler	267
classifying for linker	306
enabling in compiler	290
enabling in linker	323
--remarks (compiler option)	290
--remarks (linker option)	323
remove (library function)	
implementation-defined behavior	526
implementation-defined behavior in C89 (DLIB)	540
--remove_file_path (iobjmanip option)	505
--remove_section (iobjmanip option)	505
remquo, magnitude of	525
rename (ismexport directive)	489
rename (library function)	
implementation-defined behavior	526
implementation-defined behavior in C89 (DLIB)	540
--rename_section (iobjmanip option)	505
--rename_symbol (iobjmanip option)	506
--replace (iarchive option)	506
required (pragma directive)	375
--require_prototypes (compiler option)	291
--reserve_ranges (ismexport option)	507
reset vector table	466
__reset_QC_flag (intrinsic function)	402
__reset_Q_flag (intrinsic function)	402
restrict keyword, enabling	271
return values, from functions	173

__REV (intrinsic function) 403
 __REVSH (intrinsic function) 403
 __REV16 (intrinsic function) 403
 Ritchie, Dennis M. 41
 .rodata (ELF section) 467
 ROM to RAM, copying 109
 __root (extended keyword) 354
 __ROPI__ (predefined symbol) 424
 --ropi (compiler option) 291
 routines, time-critical 157, 182
 __ro_placement (extended keyword) 353
 rtmodel (assembler directive) 114
 rtmodel (pragma directive) 375
 RTTI data (dynamic), including in the image 316
 rtti support, missing from STL 193
 runtime environment
 DLIB 115
 setting up (DLIB) 120
 runtime libraries (DLIB)
 introduction 427
 customizing system startup code 137
 filename syntax 126
 overriding modules in 122
 using prebuilt 126
 runtime model attributes 113
 runtime model definitions 376
 runtime type information, missing from Embedded C++ . 192
 __RWPI__ (predefined symbol) 425
 --rwpi (compiler option) 292

S

-s (ielfdump option) 507
 __SADD8 (intrinsic function) 403
 __SADD16 (intrinsic function) 403
 __SASX (intrinsic function) 403
 __sbrel (extended keyword) 346
 scanf (library function)
 choosing formatter (DLIB) 129
 configuration symbols 147
 implementation-defined behavior 527
 implementation-defined behavior in C89 (DLIB) . . . 540
 __scanf_args (pragma directive) 376
 scheduling (compiler transformation) 234
 disabling 283
 scratch registers 171
 --search (linker option) 323
 search directory, for linker configuration files (--config_
 search) 303
 search path to library files (--search) 323
 search path to object files (--search) 323
 --section (ielfdump option) 507
 --section (compiler option) 292
 sections
 summary 463
 allocation of 87
 declaring (#pragma section) 376
 introduction 84
 specifying (--section) 292
 __section_begin (extended operator) 185
 __section_end (extended operator) 185
 __section_size (extended operator) 185
 section-selectors (in linker configuration file) 452
 --segment (ielfdump option) 508
 segment (pragma directive) 376
 __SEL (intrinsic function) 403
 --self_reloc (ielftool option) 508
 --semihosting (linker option) 323
 semihosting, overview 130
 --separate_cluster_for_initialized_variables
 (compiler option) 293
 set (STL header file) 432
 setjmp.h (library header file) 430
 setlocale (library function) 150
 settings, basic for project configuration 63
 __set_BASEPRI (intrinsic function) 404
 __set_CONTROL (intrinsic function) 404
 __set_CPSR (intrinsic function) 404
 __set_FAULTMASK (intrinsic function) 404

__set_FPSCR (intrinsic function)	404	--silent (ielftool option)	509
__set_interrupt_state (intrinsic function)	405	--silent (linker option)	324
__set_LR (intrinsic function)	405	silent operation	
__set_MSP (intrinsic function)	405	specifying in compiler	293
__set_PRIMASK (intrinsic function)	405	specifying in linker	324
__set_PSP (intrinsic function)	405	--simple (ielftool option)	509
__set_SB (intrinsic function)	406	--simple-ne (ielftool option)	510
__set_SP (intrinsic function)	406	sin (library function)	428
__SEV (intrinsic function)	406	sin (library routine)	132–133
severity level, of diagnostic messages	249	sinf (library routine)	132–133
specifying	250	sinl (library routine)	132–133
SFR		64-bits (floating-point format)	337
accessing special function registers	237	size (in stack usage control file)	475
declaring extern special function registers	226	size_t (integer type)	338
__SHADD8 (intrinsic function)	406	skeleton code, creating for assembler language interface .	167
__SHADD16 (intrinsic function)	406	--skip_dynamic_initialization (linker option)	324
shared object	247, 316	slist (STL header file)	432
__SHASX (intrinsic function)	406	smallest, packing algorithm for initializers	447
short (data type)	331	__SMLABB (intrinsic function)	407
show (isymexport directive)	489	__SMLABT (intrinsic function)	407
--show_entry_as (isymexport option)	509	__SMLAD (intrinsic function)	407
__SHSAX (intrinsic function)	406	__SMLADX (intrinsic function)	407
.shstrtab (ELF section)	464	__SMLALBB (intrinsic function)	408
__SHSUB16 (intrinsic function)	406	__SMLALBT (intrinsic function)	408
__SHSUB8 (intrinsic function)	406	__SMLALD (intrinsic function)	408
signal (library function)		__SMLALDX (intrinsic function)	408
implementation-defined behavior	525	__SMLALTB (intrinsic function)	408
implementation-defined behavior in C89	539	__SMLALTT (intrinsic function)	408
signals, implementation-defined behavior	516	__SMLATB (intrinsic function)	407
at system startup	517	__SMLATT (intrinsic function)	407
signal.h (library header file)	430	__SMLAWB (intrinsic function)	407
signed char (data type)	331–332	__SMLAWT (intrinsic function)	407
specifying	262	__SMLSD (intrinsic function)	407
signed int (data type)	331	__SMLSDX (intrinsic function)	407
signed long long (data type)	331	__SMLSLD (intrinsic function)	408
signed long (data type)	331	__SMLSLDX (intrinsic function)	408
signed short (data type)	331	__SMMLA (intrinsic function)	409
--silent (compiler option)	293	__SMMLAR (intrinsic function)	409
--silent (iarchive option)	509	__SMMLS (intrinsic function)	409

- __SMMLSR (intrinsic function) 409
- __SMMUL (intrinsic function) 409
- __SMMULR (intrinsic function) 409
- __SMUAD (intrinsic function) 409
- __SMUL (intrinsic function) 410
- __SMULBB (intrinsic function) 410
- __SMULBT (intrinsic function) 410
- __SMULTB (intrinsic function) 410
- __SMULTT (intrinsic function) 410
- __SMULWB (intrinsic function) 410
- __SMULWT (intrinsic function) 410
- __SMUSD (intrinsic function) 409
- __SMUSDX (intrinsic function) 409
- software interrupts 78
- source (ielfdump option) 510
- source files, list all referred 274
- space characters, implementation-defined behavior 525
- special function registers (SFR) 237
- sprintf (library function) 128
 - choosing formatter 128
- srec (ielftool option) 510
- srec-len (ielftool option) 511
- srec-s3only (ielftool option) 511
- __SSAT (intrinsic function) 410
- __SSAT16 (intrinsic function) 411
- __SSAX (intrinsic function) 403
- sscanf (library function)
 - choosing formatter (DLIB) 129
- sstream (library header file) 431
- __SSUB16 (intrinsic function) 403
- __SSUB8 (intrinsic function) 403
- stack 68
 - advantages and problems using 68
 - block for holding 464
 - cleaning after function return 174
 - contents of 68
 - layout 173
 - saving space 235
 - setting up size for 105
 - size 204
- stack parameters 172–173
- stack pointer 68
- stack usage control file
 - in depth 469
 - overview of 469
- stack (STL header file) 432
- __stackless (extended keyword) 354
- stack_usage_control (linker option) 325
- stack-size (in stack usage control file) 474
- Standard C 271
 - library compliance with 427
 - specifying strict usage 293
- standard error
 - redirecting in compiler 288
 - redirecting in linker 321
 - See also diagnostic messages 247
- standard output
 - specifying in compiler 288
 - specifying in linker 321
- standard template library (STL)
 - in C++ 432
 - in Extended EC++ 192, 199
 - missing from Embedded C++ 192
- startup code
 - cstartup 137
- startup system. *See* system startup
- statements, implementation-defined behavior in C89 . . . 536
- static analysis
 - documentation for 40
- static clustering (compiler transformation) 233
- static variables 67
 - taking the address of 234
- static_assert() 184
- static_cast (cast operator) 192
- status flags for floating-point 434
- __STC (intrinsic function) 411
- __STCL (intrinsic function) 411
- __STCL_noidx (intrinsic function) 412

__STC_noidx (intrinsic function)	412
__STC2 (intrinsic function)	411
__STC2L (intrinsic function)	411
__STC2L_noidx (intrinsic function)	412
__STC2_noidx (intrinsic function)	412
std namespace, missing from EC++	
and Extended EC++	200
stdarg.h (library header file)	430
stdbool.h (library header file)	331, 430
__STDC__ (predefined symbol)	425
STDC CX_LIMITED_RANGE (pragma directive)	377
STDC FENV_ACCESS (pragma directive)	377
STDC FP_CONTRACT (pragma directive)	378
__STDC_VERSION__ (predefined symbol)	425
stddef.h (library header file)	332, 430
stderr	121, 288, 321
stdexcept (library header file)	431
stdin	121
implementation-defined behavior in C89 (DLIB)	539
stdint.h (library header file)	430, 433
stdio.h (library header file)	430
stdio.h, additional C functionality	434
stdlib.h (library header file)	430
stdout	121, 288, 321
implementation-defined behavior	525
implementation-defined behavior in C89 (DLIB)	539
Steele, Guy L.	41
steering file, input to isymexport	488
STL	199
strcasecmp, in string.h	434
strdup, in string.h	434
streambuf (library header file)	431
streams	
implementation-defined behavior	516
supported in Embedded C++	192
strerror (library function), implementation-defined	
behavior	530
strerror (library function),	
implementation-defined behavior in C89 (DLIB)	541
__STREX (intrinsic function)	413
__STREXB (intrinsic function)	413
__STREXD (intrinsic function)	413
__STREXH (intrinsic function)	413
--strict (compiler option)	293
string (library header file)	431
strings, supported in Embedded C++	192
string.h (library header file)	430
string.h, additional C functionality	434
--strip (ielftool option)	511
--strip (iobjmanip option)	511
--strip (linker option)	325
strncasecmp, in string.h	434
strlen, in string.h	434
Stroustrup, Bjarne	41
strstream (library header file)	431
.strtab (ELF section)	464
structure types	
alignment	339
layout of	339
packed	339
structures	
aligning	374
anonymous	184, 223
implementation-defined behavior	521
implementation-defined behavior in C89	535
packing and unpacking	223
subnormal numbers	336–337
support, technical	250
Sutter, Herb	41
SVC #immed, for software interrupts	78
__swi (extended keyword)	355
SWI_Handler (exception function)	76
swi_number (pragma directive)	378
SWO, directing stdout/stderr via	120
__SWP (intrinsic function)	413
__SWPB (intrinsic function)	413
__SXTAB (intrinsic function)	414
__SXTAB16 (intrinsic function)	414
__SXTAH (intrinsic function)	414

__SXTB16 (intrinsic function) 414
 symbols
 anonymous, creating 181
 directing from one to another 322
 including in output 375
 local, removing from ELF image 319
 overview of predefined 51
 preprocessor, defining 264, 305
 --symbols (iarchive option) 512
 .symtab (ELF section) 464
 syntax
 command line options 253
 extended keywords 344–345
 invoking compiler and linker 243
 system function, implementation-defined behavior . . 517, 528
 system locks interface 152
 system startup
 customizing 137
 DLIB 134
 initialization phase 59
 system termination
 C-SPY interface to 137
 DLIB 136
 system (library function)
 implementation-defined behavior in C89 (DLIB) . . . 541
 system_include (pragma directive) 523, 538
 --system_include_dir (compiler option) 294

T

-t (iarchive option) 512
 tan (library function) 428
 tan (library routine) 132–133
 tanf (library routine) 132–133
 tanl (library routine) 132–133
 __task (extended keyword) 356
 technical support, IAR Systems 250
 template support
 in Extended EC++ 192, 199

 missing from Embedded C++ 192
 Terminal I/O window
 not supported when 122–123
 termination of system. *See* system termination
 termination status, implementation-defined behavior . . . 528
 terminology 42
 .text (ELF section) 467
 tgmth.h (library header file) 430
 32-bits (floating-point format) 336
 this (pointer) 168
 threaded environment 150
 --threaded_lib (linker option) 325
 thread-local storage 153
 __thumb (extended keyword) 356
 --thumb (compiler option) 294
 __TIME__ (predefined symbol) 425
 time zone (library function)
 implementation-defined behavior in C89 541
 time zone (library function), implementation-defined
 behavior 528
 __TIMESTAMP__ (predefined symbol) 425
 time-critical routines 157, 182
 time.h (library header file) 430
 additional C functionality 435
 time32 (library function), configuring support for 121
 time64 (library function), configuring support for 121
 tips, programming 234
 --tixt (ielftool option) 512
 TLS handling 153
 --toc (iarchive option) 512
 tools icon, in this guide 43
 trademarks 2
 transformations, compiler 228
 translation
 implementation-defined behavior 515
 implementation-defined behavior in C89 531
 type attributes 343
 specifying 379
 type qualifiers
 const and volatile 340

implementation-defined behavior	521
implementation-defined behavior in C89	536
typedefs	
excluding from diagnostics	285
repeated	187
typeinfo (library header file)	431
type_attribute (pragma directive)	378
type-based alias analysis (compiler transformation)	233
disabling	285
type-safe memory management	191
typographic conventions	42

U

__UADD8 (intrinsic function)	414
__UADD16 (intrinsic function)	414
__UASX (intrinsic function)	414
uchar.h (library header file)	430
__UHADD8 (intrinsic function)	415
__UHADD16 (intrinsic function)	415
__UHASX (intrinsic function)	415
__UHSAX (intrinsic function)	415
__UHSUB16 (intrinsic function)	415
__UHSUB8 (intrinsic function)	415
uintptr_t (integer type)	338
__UMAAL (intrinsic function)	415
underflow errors, implementation-defined behavior	525
underflow range errors,	
implementation-defined behavior in C89	539
__ungetchar, in stdio.h	434
unions	
anonymous	184, 223
implementation-defined behavior	521
implementation-defined behavior in C89	535
universal character names, implementation-defined	
behavior	522
unroll (pragma directive)	379
unsigned char (data type)	331–332
changing to signed char	262

unsigned int (data type)	331
unsigned long long (data type)	331
unsigned long (data type)	331
unsigned short (data type)	331
__UQADD8 (intrinsic function)	416
__UQADD16 (intrinsic function)	416
__UQASX (intrinsic function)	416
__UQSAX (intrinsic function)	416–417
__UQSUB16 (intrinsic function)	416–417
__UQSUB8 (intrinsic function)	416–417
__USADA8 (intrinsic function)	416
__USAD8 (intrinsic function)	416
__USAT (intrinsic function)	416
__USAT16 (intrinsic function)	417
__USAX (intrinsic function)	414
use init table (linker directive)	451
--use_c++_inline (compiler option)	294
--use_full_std_template_names (ielfdump option)	513
--use_full_std_template_names (linker option)	326
--use_unix_directory_separators (compiler option)	295
__USUB16 (intrinsic function)	414
__USUB8 (intrinsic function)	414
utilities (ELF)	477
utility (STL header file)	432
__UXTAB (intrinsic function)	417

V

-V (iarchive option)	513
valarray (library header file)	432
variables	
auto	68
defined inside a function	68
global	
placement in memory	67
hints for choosing	234
local. <i>See</i> auto variables	
non-initialized	239
placing at absolute addresses	226

- placing in named sections 226
- static
 - placement in memory 67
 - taking the address of 234
- variadic macros 185
- vector floating-point unit 274
- vector (pragma directive) 523, 538
- vector (STL header file) 432
- vectorize (compiler option) 295
- vectorize (pragma directive) 380
- __vector_table, array holding vector table 73
- veneers 111
- verbose (iarchive option) 513
- verbose (ielftool option) 513
- version
 - identifying C standard in use (__STDC_VERSION__) 425
 - of compiler (__VER__) 425
 - of this guide 2
- version (utilities option) 513
- vfe (linker option) 326
- VFP 274
- vla (compiler option) 295
- void, pointers to 186
- volatile
 - and const, declaring objects 342
 - declaring objects 341
 - protecting simultaneously accesses variables 237
 - rules for access 341

W

- #warning message (preprocessor extension) 426
- warnings 250
 - classifying in compiler 267
 - classifying in linker 307
 - disabling in compiler 287
 - disabling in linker 320
 - exit code in compiler 296
 - exit code in linker 327

- warnings icon, in this guide 43
- warnings (pragma directive) 524, 538
- warnings_affect_exit_code (compiler option) 247, 296
- warnings_affect_exit_code (linker option) 327
- warnings_are_errors (compiler option) 296
- warnings_are_errors (linker option) 327
- warn_about_c_style_casts (compiler option) 296
- wchar_t (data type), adding support for in C 332
- wchar.h (library header file) 430, 433
- wctype.h (library header file) 430
- __weak (extended keyword) 356
- weak (pragma directive) 380
- web sites, recommended 42
- __WFE (intrinsic function) 417
- __WFI (intrinsic function) 417
- white-space characters, implementation-defined behavior 515
- whole_archive (linker option) 327
- __write_array, in stdio.h 434
- __write_buffered (DLIB library function) 120

X

- x (iarchive option) 497

Y

- __YIELD (intrinsic function) 417

Z

- zeros, packing algorithm for initializers 447

Symbols

- _AEABI_PORTABILITY_LEVEL (preprocessor symbol) 219
- _AEABI_PORTABLE (preprocessor symbol) 219
- _Exit (library function) 136

__exit (library function)	136
__AAPCS_VFP__ (predefined symbol)	420
__AAPCS__ (predefined symbol)	420
__absolute (extended keyword)	347
__ALIGNOF__ (operator)	184
__arm (extended keyword)	347
__ARMVFPV2__ (predefined symbol)	421
__ARMVFPV3__ (predefined symbol)	421
__ARMVFPV4__ (predefined symbol)	421
__ARMVFP_D16__ (predefined symbol)	421
__ARMVFP_FP16__ (predefined symbol)	421
__ARMVFP_SP__ (predefined symbol)	421
__ARMVFP__ (predefined symbol)	421
__ARM_ADVANCED_SIMD__ (predefined symbol)	420
__ARM_MEDIA__ (predefined symbol)	420
__ARM_PROFILE_M__ (predefined symbol)	421
__ARM4TM__ (predefined symbol)	422
__ARM5E__ (predefined symbol)	422
__ARM5__ (predefined symbol)	422
__ARM6M__ (predefined symbol)	422
__ARM6SM__ (predefined symbol)	422
__ARM6__ (predefined symbol)	422
__ARM7A__ (predefined symbol)	422
__ARM7EM__ (predefined symbol)	422
__ARM7M__ (predefined symbol)	422
__ARM7R__ (predefined symbol)	422
__asm (language extension)	160
__assignment_by_bitwise_copy_allowed, symbol used in library	435
__as_get_base (C-RUN operator)	383
__as_get_bounds (C-RUN operator)	383
__as_make_bounds (C-RUN operator)	383
__BASE_FILE__ (predefined symbol)	422
__big_endian (extended keyword)	347
__BUILD_NUMBER__ (predefined symbol)	422
__CLREX (intrinsic function)	390
__CLZ (intrinsic function)	390
__code, symbol used in library	435
__constrange(), symbol used in library	435
__construction_by_bitwise_copy_allowed, symbol used in library	435
__CORE__ (predefined symbol)	422
__COUNTER__ (predefined symbol)	422
__cplusplus (predefined symbol)	422
__CPU_MODE__ (predefined symbol)	422
__data, symbol used in library	435
__DATE__ (predefined symbol)	423
__disable_fiq (intrinsic function)	390
__disable_interrupt (intrinsic function)	391
__disable_irq (intrinsic function)	391
__DLIB_FILE_DESCRIPTOR (configuration symbol)	148
__DLIB_PERTHREAD (ELF section)	465
__DMB (intrinsic function)	391
__DSB (intrinsic function)	391
__embedded_cplusplus (predefined symbol)	423
__enable_fiq (intrinsic function)	391
__enable_interrupt (intrinsic function)	392
__enable_irq (intrinsic function)	392
__exit (library function)	136
__FILE__ (predefined symbol)	423
__fiq (extended keyword)	348
__FUNCTION__ (predefined symbol)	188, 423
__func__ (predefined symbol)	188, 423
__gets, in stdio.h.	434
__get_BASEPRI (intrinsic function)	392
__get_CONTROL (intrinsic function)	392
__get_CPSR (intrinsic function)	392
__get_FAULTMASK (intrinsic function)	393
__get_FPSCR (intrinsic function)	393
__get_interrupt_state (intrinsic function)	393
__get_IPSR (intrinsic function)	394
__get_LR (intrinsic function)	394
__get_MSP (intrinsic function)	394
__get_PRIMASK (intrinsic function)	394
__get_PSP (intrinsic function)	394
__get_PSR (intrinsic function)	395
__get_SB (intrinsic function)	395
__get_SP (intrinsic function)	395
__has_constructor, symbol used in library	436

<code>__has_destructor</code> , symbol used in library	436	<code>__iar_sin_accuratef</code> (library routine)	133
<code>__iar_cos_accurate</code> (library routine)	133	<code>__iar_Sin_accuratel</code> (library routine)	133
<code>__iar_cos_accuratef</code> (library routine)	133	<code>__iar_sin_accuratel</code> (library routine)	133
<code>__iar_cos_accuratel</code> (library routine)	133	<code>__iar_Sin_small</code> (library routine)	132
<code>__iar_cos_small</code> (library routine)	132	<code>__iar_sin_small</code> (library routine)	132
<code>__iar_cos_smallf</code> (library routine)	132	<code>__iar_Sin_smallf</code> (library routine)	132
<code>__iar_cos_smallll</code> (library routine)	132	<code>__iar_sin_smallf</code> (library routine)	132
<code>__IAR_DLIB_PERTHREAD_INIT_SIZE</code> (macro)	154	<code>__iar_Sin_smallll</code> (library routine)	132
<code>__IAR_DLIB_PERTHREAD_SIZE</code> (macro)	154	<code>__iar_sin_smallll</code> (library routine)	132
<code>__IAR_DLIB_PERTHREAD_SYMBOL_OFFSET</code> (symbolptr)	154	<code>__IAR_SYSTEMS_ICC__</code> (predefined symbol)	424
<code>__iar_exp_small</code> (library routine)	132	<code>__iar_tan_accurate</code> (library routine)	133
<code>__iar_exp_smallf</code> (library routine)	132	<code>__iar_tan_accuratef</code> (library routine)	133
<code>__iar_exp_smallll</code> (library routine)	132	<code>__iar_tan_accuratel</code> (library routine)	133
<code>__iar_FPow</code> (library routine)	133	<code>__iar_tan_small</code> (library routine)	132
<code>__iar_FSin</code> (library routine)	132–133	<code>__iar_tan_smallf</code> (library routine)	132
<code>__iar_log_small</code> (library routine)	132	<code>__iar_tan_smallll</code> (library routine)	132
<code>__iar_log_smallf</code> (library routine)	132	<code>__ICCARM__</code> (predefined symbol)	424
<code>__iar_log_smallll</code> (library routine)	132	<code>__interwork</code> (extended keyword)	348
<code>__iar_log10_small</code> (library routine)	132	<code>__intrinsic</code> (extended keyword)	348
<code>__iar_log10_smallf</code> (library routine)	132	<code>__irq</code> (extended keyword)	348
<code>__iar_log10_smallll</code> (library routine)	132	<code>__ISB</code> (intrinsic function)	395
<code>__iar_LPow</code> (library routine)	133	<code>__LDC</code> (intrinsic function)	395
<code>__iar_LSin</code> (library routine)	132–133	<code>__LDCL</code> (intrinsic function)	395
<code>__iar_maximum_atexit_calls</code>	106	<code>__LDCL_noidx</code> (intrinsic function)	396
<code>__iar_Pow</code> (library routine)	133	<code>__LDC_noidx</code> (intrinsic function)	396
<code>__iar_Pow_accurate</code> (library routine)	133	<code>__LDC2</code> (intrinsic function)	395
<code>__iar_pow_accurate</code> (library routine)	133	<code>__LDC2L</code> (intrinsic function)	395
<code>__iar_pow_accuratef</code> (library routine)	133	<code>__LDC2L_noidx</code> (intrinsic function)	396
<code>__iar_pow_accuratel</code> (library routine)	133	<code>__LDC2_noidx</code> (intrinsic function)	396
<code>__iar_pow_small</code> (library routine)	132	<code>__LDREX</code> (intrinsic function)	397
<code>__iar_pow_smallf</code> (library routine)	132	<code>__LDREXB</code> (intrinsic function)	397
<code>__iar_pow_smallll</code> (library routine)	132	<code>__LDREXD</code> (intrinsic function)	397
<code>__iar_program_start</code> (label)	134	<code>__LDREXH</code> (intrinsic function)	397
<code>__iar_ReportAssert</code> (library function)	140	<code>__LINE__</code> (predefined symbol)	424
<code>__iar_Sin</code> (library routine)	132–133	<code>__little_endian</code> (extended keyword)	349
<code>__iar_Sin_accurate</code> (library routine)	133	<code>__LITTLE_ENDIAN__</code> (predefined symbol)	424
<code>__iar_sin_accurate</code> (library routine)	133	<code>__low_level_init</code>	134
<code>__iar_Sin_accuratef</code> (library routine)	133	initialization phase	59
		<code>__low_level_init</code> , customizing	137

__MCR (intrinsic function)	397	__reset_QC_flag (intrinsic function)	402
__MCR2 (intrinsic function)	397	__reset_Q_flag (intrinsic function)	402
__memory_of		__REV (intrinsic function)	403
symbol used in library	436	__REVSH (intrinsic function)	403
__MRC (intrinsic function)	398	__REV16 (intrinsic function)	403
__MRC2 (intrinsic function)	398	__root (extended keyword)	354
__nested (extended keyword)	349	__ROPI__ (predefined symbol)	424
__noreturn (extended keyword)	351	__ro_placement (extended keyword)	353
__no_alloc (extended keyword)	350	__RWPI__ (predefined symbol)	425
__no_alloc_str (operator)	350	__SADD8 (intrinsic function)	403
__no_alloc_str16 (operator)	350	__SADD16 (intrinsic function)	403
__no_alloc16 (extended keyword)	350	__SASX (intrinsic function)	403
__no_init (extended keyword)	239, 351	__sbrel (extended keyword)	346
__no_operation (intrinsic function)	399	__scanf_args (pragma directive)	376
__packed (extended keyword)	352	__section_begin (extended operator)	185
__pcrel (extended keyword)	346	__section_end (extended operator)	185
__PKHBT (intrinsic function)	399	__section_size (extended operator)	185
__PKHTB (intrinsic function)	399	__SEL (intrinsic function)	403
__PLD (intrinsic function)	400	__set_BASEPRI (intrinsic function)	404
__PLDW (intrinsic function)	400	__set_CONTROL (intrinsic function)	404
__PLI (intrinsic function)	400	__set_CPSR (intrinsic function)	404
__PRETTY_FUNCTION__ (predefined symbol)	424	__set_FAULTMASK (intrinsic function)	404
__printf_args (pragma directive)	374	__set_FPSCR (intrinsic function)	404
__program_start (label)	134	__set_interrupt_state (intrinsic function)	405
__QADD (intrinsic function)	400	__set_LR (intrinsic function)	405
__QADD8 (intrinsic function)	401	__set_MSP (intrinsic function)	405
__QADD16 (intrinsic function)	401	__set_PRIMASK (intrinsic function)	405
__QASX (intrinsic function)	401	__set_PSP (intrinsic function)	405
__QCFlag (intrinsic function)	401	__set_SB (intrinsic function)	406
__QDADD (intrinsic function)	400	__set_SP (intrinsic function)	406
__QDOUBLE (intrinsic function)	401	__SEV (intrinsic function)	406
__QDSUB (intrinsic function)	400	__SHADD8 (intrinsic function)	406
__QFlag (intrinsic function)	402	__SHADD16 (intrinsic function)	406
__QSAX (intrinsic function)	401	__SHASX (intrinsic function)	406
__QSUB (intrinsic function)	400	__SHSAX (intrinsic function)	406
__QSUB16 (intrinsic function)	401	__SHSUB16 (intrinsic function)	406
__QSUB8 (intrinsic function)	401	__SHSUB8 (intrinsic function)	406
__ramfunc (extended keyword)	353	__SMLABB (intrinsic function)	407
__RBIT (intrinsic function)	402	__SMLABT (intrinsic function)	407

__SMLAD (intrinsic function)	407	__STCL (intrinsic function)	411
__SMLADX (intrinsic function)	407	__STCL_noidx (intrinsic function)	412
__SMLALBB (intrinsic function)	408	__STC_noidx (intrinsic function)	412
__SMLALBT (intrinsic function)	408	__STC2 (intrinsic function)	411
__SMLALD (intrinsic function)	408	__STC2L (intrinsic function)	411
__SMLALDX (intrinsic function)	408	__STC2L_noidx (intrinsic function)	412
__SMLALTB (intrinsic function)	408	__STC2_noidx (intrinsic function)	412
__SMLALTT (intrinsic function)	408	__STDC_VERSION__ (predefined symbol)	425
__SMLATB (intrinsic function)	407	__STDC__ (predefined symbol)	425
__SMLATT (intrinsic function)	407	__STREX (intrinsic function)	413
__SMLAWB (intrinsic function)	407	__STREXB (intrinsic function)	413
__SMLAWT (intrinsic function)	407	__STREXD (intrinsic function)	413
__SMLS (intrinsic function)	407	__STREXH (intrinsic function)	413
__SMLS (intrinsic function)	407	__swi (extended keyword)	355
__SMLS (intrinsic function)	407	__SWP (intrinsic function)	413
__SMLS (intrinsic function)	408	__SWPB (intrinsic function)	413
__SMLS (intrinsic function)	408	__SXTAB (intrinsic function)	414
__SMLS (intrinsic function)	409	__SXTAB16 (intrinsic function)	414
__SMLS (intrinsic function)	409	__SXTAH (intrinsic function)	414
__SMLS (intrinsic function)	409	__SXTB16 (intrinsic function)	414
__SMLS (intrinsic function)	409	__task (extended keyword)	356
__SMLS (intrinsic function)	409	__thumb (extended keyword)	356
__SMUAD (intrinsic function)	409	__TIMESTAMP__ (predefined symbol)	425
__SMUL (intrinsic function)	410	__TIME__ (predefined symbol)	425
__SMULBB (intrinsic function)	410	__UADD8 (intrinsic function)	414
__SMULBT (intrinsic function)	410	__UADD16 (intrinsic function)	414
__SMULTB (intrinsic function)	410	__UASX (intrinsic function)	414
__SMULTT (intrinsic function)	410	__UHADD8 (intrinsic function)	415
__SMULWB (intrinsic function)	410	__UHADD16 (intrinsic function)	415
__SMULWT (intrinsic function)	410	__UHASX (intrinsic function)	415
__SMUSD (intrinsic function)	409	__UHSAX (intrinsic function)	415
__SMUSD (intrinsic function)	409	__UHSUB16 (intrinsic function)	415
__SSAT (intrinsic function)	410	__UHSUB8 (intrinsic function)	415
__SSAT16 (intrinsic function)	411	__UMAAL (intrinsic function)	415
__SSAX (intrinsic function)	403	__ungetchar, in stdio.h	434
__SSUB16 (intrinsic function)	403	__UQADD8 (intrinsic function)	416
__SSUB8 (intrinsic function)	403	__UQADD16 (intrinsic function)	416
__stackless (extended keyword)	354	__UQASX (intrinsic function)	416
__STC (intrinsic function)	411	__UQSAX (intrinsic function)	416–417

__UQSUB16 (intrinsic function)	416–417	--a (ielfdump option).	492
__UQSUB8 (intrinsic function)	416–417	--aapcs (compiler option)	260
__USADA8 (intrinsic function)	416	--advanced_heap (linker option)	301
__USAD8 (intrinsic function).	416	--aeabi (compiler option)	260
__USAT (intrinsic function)	416	--align_sp_on_irq (compiler option)	261
__USAT16 (intrinsic function)	417	--all (ielfdump option)	492
__USAX (intrinsic function).	414	--arm (compiler option).	261
__USUB16 (intrinsic function).	414	--basic_heap (linker option)	301
__USUB8 (intrinsic function).	414	--BE32 (linker option)	302, 304
__UXTAB (intrinsic function)	417	--BE8 (linker option)	301
__VA_ARGS__ (preprocessor extension)	181	--bin (ielftool option)	492
__weak (extended keyword)	356	--bounds_table_size (linker option)	297
__WFE (intrinsic function).	417	--call_graph (linker option).	302
__WFI (intrinsic function)	417	--char_is_signed (compiler option).	262
__write_array, in stdio.h	434	--char_is_unsigned (compiler option)	262
__write_buffered (DLIB library function)	120	--checksum (ielftool option)	493
__YIELD (intrinsic function)	417	--code (ielfdump option).	496
-D (compiler option)	264	--config (linker option)	302
-d (iarchive option)	497	--config_def (linker option)	303
-e (compiler option)	269	--config_search (linker option)	303
-f (compiler option).	273	--cpp_init_routine (linker option)	304
-f (IAR utility option)	498	--cpu (compiler option).	262
-f (linker option)	311	--cpu_mode (compiler option)	263
-g (ielfdump option)	508	--create (iarchive option).	496
-I (compiler option).	275	--c++ (compiler option)	264
-l (compiler option).	275	--c89 (compiler option).	261
for creating skeleton code	167	--debug (compiler option)	265
-L (linker option)	323	--debug_heap (linker option)	297
-O (compiler option)	287	--default_to_complex_ranges (linker option)	304
-o (compiler option)	288	--define_symbol (linker option)	305
-o (iarchive option)	502	--delete (iarchive option).	497
-o (ielfdump option)	502	--dependencies (compiler option)	265
-o (linker option).	321	--dependencies (linker option)	305
-r (compiler option).	265	--diagnostics_tables (compiler option)	268
-r (iarchive option)	506	--diagnostics_tables (linker option).	307
-s (ielfdump option)	507	--diag_error (compiler option)	266
-t (iarchive option)	512	--diag_error (linker option).	306
-V (iarchive option)	513	--diag_remark (compiler option).	267
-x (iarchive option)	497	--diag_remark (linker option)	306

--diag_suppress (compiler option)	267	--macro_positions_in_diagnostics (compiler option)	277
--diag_suppress (linker option)	307	--make_all_definitions_weak (compiler option)	277
--diag_warning (compiler option)	267	--mangled_names_in_messages (linker option)	315
--diag_warning (linker option)	307	--map (linker option)	315
--discard_unused_publics (compiler option)	268	--merge_duplicate_sections (linker option)	316
--dlib_config (compiler option)	269	--mfc (compiler option)	278
--do_segment_pad (linker option)	308	--misrac (compiler option)	257
--ec++ (compiler option)	270	--misrac (linker option)	299
--edit (ismyexport option)	497	--misrac_verbose (compiler option)	257
--eec++ (compiler option)	270	--misrac_verbose (linker option)	299
--enable_hardware_workaround (compiler option)	270	--misrac1998 (compiler option)	257
--enable_hardware_workaround (linker option)	308	--misrac1998 (linker option)	299
--enable_multibytes (compiler option)	271	--misrac2004 (compiler option)	257
--enable_restrict (compiler option)	271	--misrac2004 (linker option)	299
--entry (linker option)	309	--no_alignment_reduction (compiler option)	278
--enum_is_int (compiler option)	272	--no_call_frame_info (compiler option)	278
--error_limit (compiler option)	272	--no_clustering (compiler option)	279
--error_limit (linker option)	309	--no_code_motion (compiler option)	279
--exception_tables (linker option)	310	--no_const_align (compiler option)	279
--export_builtin_config (linker option)	310	--no_cse (compiler option)	280
--extract (iarchive option)	497	--no_dynamic_rtti_elimination (linker option)	316
--extra_init (linker option)	311	--no_entry (linker option)	317
--fill (ielftool option)	499	--no_exceptions (compiler option)	280
--force_exceptions (linker option)	312	--no_exceptions (linker option)	317
--force_output (linker option)	312	--no_fragments (compiler option)	281
--fpu (compiler option)	273	--no_fragments (linker option)	317
--generate_vfe_header (ismyexport option)	499	--no_header (ielfdump option)	500
--guard_calls (compiler option)	274	--no_inline (compiler option)	281
--header_context (compiler option)	274	--no_library_search (linker option)	318
--ignore_uninstrumented_pointers (linker option)	298	--no_literal_pool (compiler option)	281
--ihex (ielftool option)	500	--no_literal_pool (linker option)	318
--image_input (linker option)	312	--no_locals (linker option)	319
--inline (linker option)	313	--no_loop_align (compiler option)	282
--interwork (compiler option)	275	--no_mem_idioms (compiler option)	282
--keep (linker option)	313	--no_path_in_file_macros (compiler option)	282
--legacy (compiler option)	277	--no_range_reservations (linker option)	319
--lock_regs (compiler option)	276	--no_rel_section (ielfdump option)	500
--log (linker option)	314	--no_remove (linker option)	319
--log_file (linker option)	315	--no_rtti (compiler option)	283

--no_rw_dynamic_init (compiler option)	283	--rename_section (iobjmanip option)	505
--no_scheduling (compiler option)	283	--rename_symbol (iobjmanip option)	506
--no_size_constraints (compiler option)	284	--replace (iarchive option)	506
--no_static_destruction (compiler option)	284	--require_prototypes (compiler option)	291
--no_strtab (ielfdump option)	501	--reserve_ranges (ismexport option)	507
--no_system_include (compiler option)	284	--ropi (compiler option)	291
--no_typedefs_in_diagnostics (compiler option)	285	--rwp (compiler option)	292
--no_unaligned_access (compiler option)	285	--search (linker option)	323
--no_unroll (compiler option)	286	--section (compiler option)	292
--no_var_align (compiler option)	286	--section (ielfdump option)	507
--no_veneers (linker option)	320	--segment (ielfdump option)	508
--no_vfe (linker option)	320	--self_reloc (ielftool option)	508
--no_warnings (compiler option)	287	--seminhosting (linker option)	323
--no_warnings (linker option)	320	--separate_cluster_for_initialized_variables (compiler option)	293
--no_wrap_diagnostics (compiler option)	287	--show_entry_as (ismexport option)	509
--no_wrap_diagnostics (linker option)	321	--silent (compiler option)	293
--offset (ielftool option)	501	--silent (iarchive option)	509
--only_stdout (compiler option)	288	--silent (ielftool option)	509
--only_stdout (linker option)	321	--silent (linker option)	324
--option_name (compiler option)	309	--simple (ielftool option)	509
--output (compiler option)	288	--simple-ne (ielftool option)	510
--output (iarchive option)	502	--skip_dynamic_initialization (linker option)	324
--output (ielfdump option)	502	--source (ielfdump option)	510
--output (linker option)	321	--srec (ielftool option)	510
--parity (ielftool option)	502	--srec-len (ielftool option)	511
--pi_veneers (linker option)	321	--srec-s3only (ielftool option)	511
--place_holder (linker option)	322	--stack_usage_control (linker option)	325
--predef_macro (compiler option)	288	--strict (compiler option)	293
--preinclude (compiler option)	289	--strip (ielftool option)	511
--preprocess (compiler option)	289	--strip (iobjmanip option)	511
--ram_reserve_ranges (ismexport option)	503	--strip (linker option)	325
--range (ielfdump option)	504	--symbols (iarchive option)	512
--raw (ielfdump] option)	504	--system_include_dir (compiler option)	294
--redirect (linker option)	322	--threaded_lib (linker option)	325
--relaxed_fp (compiler option)	290	--thumb (compiler option)	294
--remarks (compiler option)	290	--titxt (ielftool option)	512
--remarks (linker option)	323	--toc (iarchive option)	512
--remove_file_path (iobjmanip option)	505	--use_c++_inline (compiler option)	294
--remove_section (iobjmanip option)	505	--use_full_std_template_names (ielfdump option)	513

- use_full_std_template_names (linker option) 326
- use_unix_directory_separators (compiler option) 295
- vectorize (compiler option) 295
- verbose (iarchive option) 513
- verbose (ielftool option) 513
- version (utilities option) 513
- vfe (linker option) 326
- vla (compiler option) 295
- warnings_affect_exit_code (compiler option) . . . 247, 296
- warnings_affect_exit_code (linker option) 327
- warnings_are_errors (compiler option) 296
- warnings_are_errors (linker option) 327
- warn_about_c_style_casts (compiler option) 296
- whole_archive (linker option) 327
- .bss (ELF section) 464
- .comment (ELF section) 464
- .data (ELF section) 465
- .data_init (ELF section) 465
- .debug (ELF section) 464
- .exc.text (ELF section) 465
- .iar.debug (ELF section) 464
- .iar.dynexit (ELF section) 466
- .init_array (section) 466
- .intvec (ELF section) 466
- .noinit (ELF section) 467
- .preinit_array (section) 467
- .prepreinit_array (section) 467
- .rel (ELF section) 464
- .rela (ELF section) 464
- .rodata (ELF section) 467
- .shstrtab (ELF section) 464
- .strtab (ELF section) 464
- .symtab (ELF section) 464
- .text (ELF section) 467
- .textrw (ELF section) 468
- .textrw_init (ELF section) 468
- @ (operator)
 - placing at absolute address 225
 - placing in sections 226

- #include files, specifying 245, 275
- #warning message (preprocessor extension) 426
- %Z replacement string,
 - implementation-defined behavior 528

Numerics

- 32-bits (floating-point format) 336
- 64-bits (floating-point format) 337