

Using GNU Fortran

For GCC version 16.0.0 (pre-release)

(GCC)

The gfortran team

Published by the Free Software Foundation
51 Franklin Street, Fifth Floor
Boston, MA 02110-1301, USA

Copyright © 1999-2025 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with the Invariant Sections being “Funding Free Software”, the Front-Cover Texts being (a) (see below), and with the Back-Cover Texts being (b) (see below). A copy of the license is included in the section entitled “GNU Free Documentation License”.

(a) The FSF’s Front-Cover Text is:

A GNU Manual

(b) The FSF’s Back-Cover Text is:

You have freedom to copy and modify this GNU Manual, like GNU software. Copies published by the Free Software Foundation raise funds for GNU development.

Short Contents

1	Introduction	1
	Invoking GNU Fortran	
2	GNU Fortran Command Options	7
3	Runtime: Influencing runtime behavior with environment variables	37
	Language Reference	
4	Compiler Characteristics	43
5	Extensions	49
6	Mixed-Language Programming	73
7	Coarray Programming	89
8	Intrinsic Procedures	119
9	Intrinsic Modules	311
	Contributing	321
	GNU General Public License	323
	GNU Free Documentation License	335
	Funding Free Software	343
	Option Index	345
	Keyword Index	347

Table of Contents

1	Introduction	1
1.1	About GNU Fortran	1
1.2	GNU Fortran and GCC	2
1.3	Standards	3
1.3.1	Fortran 95 status	3
1.3.2	Fortran 2003 status	3
1.3.3	Fortran 2008 status	4
1.3.4	Fortran 2018 status	4
	Part I: Invoking GNU Fortran	5
2	GNU Fortran Command Options	7
2.1	Option summary	7
2.2	Options controlling Fortran dialect	9
2.3	Enable and customize preprocessing	14
2.4	Options to request or suppress errors and warnings	18
2.5	Options for debugging your program	22
2.6	Options for directory search	24
2.7	Influencing the linking step	24
2.8	Influencing runtime behavior	24
2.9	GNU Fortran Developer Options	25
2.10	Options for code generation conventions	26
2.11	Options for interoperability with other languages	34
2.12	Environment variables affecting <code>gfortran</code>	35
3	Runtime: Influencing runtime behavior with environment variables	37
3.1	<code>TMPDIR</code> —Directory for scratch files	37
3.2	<code>GFORTTRAN_STDIN_UNIT</code> —Unit number for standard input	37
3.3	<code>GFORTTRAN_STDOUT_UNIT</code> —Unit number for standard output	37
3.4	<code>GFORTTRAN_STDERR_UNIT</code> —Unit number for standard error	37
3.5	<code>GFORTTRAN_UNBUFFERED_ALL</code> —Do not buffer I/O on all units	37
3.6	<code>GFORTTRAN_UNBUFFERED_PRECONNECTED</code> —Do not buffer I/O on preconnected units	37
3.7	<code>GFORTTRAN_SHOW_LOCUS</code> —Show location for runtime errors	37
3.8	<code>GFORTTRAN_OPTIONAL_PLUS</code> —Print leading + where permitted	38
3.9	<code>GFORTTRAN_LIST_SEPARATOR</code> —Separator for list output	38
3.10	<code>GFORTTRAN_CONVERT_UNIT</code> —Set conversion for unformatted I/O	38
3.11	<code>GFORTTRAN_ERROR_BACKTRACE</code> —Show backtrace on run-time errors	39
3.12	<code>GFORTTRAN_FORMATTED_BUFFER_SIZE</code> —Set buffer size for formatted I/O	39

3.13	GFORTTRAN_UNFORMATTED_BUFFER_SIZE—Set buffer size for unformatted I/O	39
------	--	----

Part II: Language Reference 41

4 Compiler Characteristics 43

4.1	KIND Type Parameters	43
4.2	Internal representation of LOGICAL variables	43
4.3	Evaluation of logical expressions	44
4.4	MAX and MIN intrinsics with REAL NaN arguments	44
4.5	Thread-safety of the runtime library	44
4.6	Data consistency and durability	45
4.7	Files opened without an explicit ACTION= specifier	46
4.8	File operations on symbolic links	46
4.9	File format of unformatted sequential files	46
4.10	Asynchronous I/O	47
4.11	Behavior on integer overflow	47

5 Extensions 49

5.1	Extensions implemented in GNU Fortran	49
5.1.1	Old-style kind specifications	49
5.1.2	Old-style variable initialization	49
5.1.3	Extensions to namelist	50
5.1.4	X format descriptor without count field	51
5.1.5	Commas in FORMAT specifications	51
5.1.6	Missing period in FORMAT specifications	51
5.1.7	Default widths for 'F', 'G' and 'I' format descriptors	51
5.1.8	I/O item lists	51
5.1.9	'Q' exponent-letter	51
5.1.10	BOZ literal constants	52
5.1.11	Real array indices	52
5.1.12	Unary operators	52
5.1.13	Implicitly convert LOGICAL and INTEGER values	52
5.1.14	Hollerith constants support	52
5.1.15	Character conversion	53
5.1.16	Cray pointers	54
5.1.17	CONVERT specifier	55
5.1.18	OpenMP	56
5.1.19	OpenACC	57
5.1.20	Argument list functions %VAL, %REF and %LOC	57
5.1.21	Read/Write after EOF marker	58
5.1.22	STRUCTURE and RECORD	58
5.1.23	UNION and MAP	61
5.1.24	Type variants for integer intrinsics	62
5.1.25	AUTOMATIC and STATIC attributes	64

5.1.26	Form feed as whitespace	64
5.1.27	TYPE as an alias for PRINT	64
5.1.28	%LOC as an rvalue	65
5.1.29	.XOR. operator	65
5.1.30	Bitwise logical operators	65
5.1.31	Extended I/O specifiers	65
5.1.32	Legacy PARAMETER statements	67
5.1.33	Default exponents	67
5.1.34	Unsigned integers	67
5.2	Extensions not implemented in GNU Fortran	69
5.2.1	ENCODE and DECODE statements	70
5.2.2	Variable FORMAT expressions	70
5.2.3	Alternate complex function syntax	71
5.2.4	Volatile COMMON blocks	71
5.2.5	OPEN(... NAME=)	71
5.2.6	Q edit descriptor	71
6	Mixed-Language Programming	73
6.1	Interoperability with C	73
6.1.1	Intrinsic Types	73
6.1.2	Derived Types and struct	73
6.1.3	Interoperable Global Variables	74
6.1.4	Interoperable Subroutines and Functions	74
6.1.5	Working with C Pointers	76
6.1.6	Further Interoperability of Fortran with C	78
6.1.7	Generating C prototypes from Fortran	78
6.2	GNU Fortran Compiler Directives	78
6.2.1	ATTRIBUTES directive	78
6.2.2	UNROLL directive	80
6.2.3	BUILTIN directive	80
6.2.4	IVDEP directive	80
6.2.5	VECTOR directive	80
6.2.6	NOVECTOR directive	81
6.3	Non-Fortran Main Program	81
6.3.1	_gfortran_set_args — Save command-line arguments ...	81
6.3.2	_gfortran_set_options — Set library option flags	82
6.3.3	_gfortran_set_convert — Set endian conversion	83
6.3.4	_gfortran_set_record_marker — Set length of record markers	83
6.3.5	_gfortran_set_fpe — Enable floating point exception traps ..	84
6.3.6	_gfortran_set_max_subrecord_ length — Set subrecord length	84
6.4	Naming and argument-passing conventions	85
6.4.1	Naming conventions	85
6.4.2	Argument passing conventions	85

7	Coarray Programming	89
7.1	Type and enum ABI Documentation	89
7.1.1	caf_token_t	89
7.1.2	caf_register_t	89
7.1.3	caf_deregister_t	89
7.1.4	caf_reference_t	89
7.1.5	caf_team_t	91
7.2	Function ABI Documentation	91
7.2.1	_gfortran_caf_init — Initialization function	91
7.2.2	_gfortran_caf_finish — Finalization function	92
7.2.3	_gfortran_caf_this_image — Querying the image number	92
7.2.4	_gfortran_caf_num_images — Querying the maximal number of images	92
7.2.5	_gfortran_caf_image_status — Query the status of an image	93
7.2.6	_gfortran_caf_failed_images — Get an array of the indexes of the failed images	93
7.2.7	_gfortran_caf_stopped_images — Get an array of the indexes of the stopped images	94
7.2.8	_gfortran_caf_register — Registering coarrays	94
7.2.9	_gfortran_caf_deregister — Deregistering coarrays	95
7.2.10	_gfortran_caf_register_accessor — Register an accessor for remote access	96
7.2.11	_gfortran_caf_register_accessors_finish — Finish registering accessor functions	96
7.2.12	_gfortran_caf_get_remote_function_ index — Get the index of an accessor	97
7.2.13	_gfortran_caf_get_from_remote — Getting data from a remote image using a remote side accessor	97
7.2.14	_gfortran_caf_is_present_on_remote — Check that a coarray or a part of it is allocated on the remote image	99
7.2.15	_gfortran_caf_send_to_remote — Send data to a remote image using a remote side accessor to store it	100
7.2.16	_gfortran_caf_transfer_between_remotes — Initiate data transfer between to remote images	101
7.2.17	_gfortran_caf_sendget_by_ref — Sending data between remote images using enhanced references on both sides	104
7.2.18	_gfortran_caf_lock — Locking a lock variable	105
7.2.19	_gfortran_caf_lock — Unlocking a lock variable	106
7.2.20	_gfortran_caf_event_post — Post an event	106
7.2.21	_gfortran_caf_event_wait — Wait that an event occurred	107
7.2.22	_gfortran_caf_event_query — Query event count	107
7.2.23	_gfortran_caf_sync_all — All-image barrier	108
7.2.24	_gfortran_caf_sync_images — Barrier for selected images	108

7.2.25	<code>_gfortran_caf_sync_memory</code> — Wait for completion of segment-memory operations	109
7.2.26	<code>_gfortran_caf_error_stop</code> — Error termination with exit code	109
7.2.27	<code>_gfortran_caf_error_stop_str</code> — Error termination with string	109
7.2.28	<code>_gfortran_caf_fail_image</code> — Mark the image failed and end its execution	109
7.2.29	<code>_gfortran_caf_atomic_define</code> — Atomic variable assignment	110
7.2.30	<code>_gfortran_caf_atomic_ref</code> — Atomic variable reference ..	110
7.2.31	<code>_gfortran_caf_atomic_cas</code> — Atomic compare and swap ..	110
7.2.32	<code>_gfortran_caf_atomic_op</code> — Atomic operation	111
7.2.33	<code>_gfortran_caf_co_broadcast</code> — Sending data to all images	112
7.2.34	<code>_gfortran_caf_co_max</code> — Collective maximum reduction ..	112
7.2.35	<code>_gfortran_caf_co_min</code> — Collective minimum reduction ..	113
7.2.36	<code>_gfortran_caf_co_sum</code> — Collective summing reduction ..	113
7.2.37	<code>_gfortran_caf_co_reduce</code> — Generic collective reduction ..	114
7.2.38	<code>_gfortran_caf_form_team</code> — Team creation function ..	115
7.2.39	<code>_gfortran_caf_change_team</code> — Team activation function ..	116
7.2.40	<code>_gfortran_caf_end_team</code> — Team termination function ..	116
7.2.41	<code>_gfortran_caf_sync_team</code> — Synchronize all images of a given team	116
7.2.42	<code>_gfortran_caf_get_team</code> — Get the opaque handle of the specified team	117
7.2.43	<code>_gfortran_caf_team_number</code> — Get the unique id of the given team	117
8	Intrinsic Procedures	119
8.1	Introduction to intrinsic procedures	119
8.2	<code>ABORT</code> — Abort the program	119
8.3	<code>ABS</code> — Absolute value	120
8.4	<code>ACCESS</code> — Checks file access modes	120
8.5	<code>ACHAR</code> — Character in ASCII collating sequence	121
8.6	<code>ACOS</code> — Arccosine function	122
8.7	<code>ACOSD</code> — Arccosine function, degrees	122
8.8	<code>ACOSH</code> — Inverse hyperbolic cosine function	123
8.9	<code>ACOSPI</code> — Circular arc cosine function	124
8.10	<code>ADJUSTL</code> — Left adjust a string	124
8.11	<code>ADJUSTR</code> — Right adjust a string	125
8.12	<code>AIMAG</code> — Imaginary part of complex number	125
8.13	<code>AINT</code> — Truncate to a whole number	126
8.14	<code>ALARM</code> — Execute a routine after a given delay	127
8.15	<code>ALL</code> — All values in <i>MASK</i> along <i>DIM</i> are true	127
8.16	<code>ALLOCATED</code> — Status of an allocatable entity	128

8.17	AND — Bitwise logical AND	129
8.18	ANINT — Nearest whole number	130
8.19	ANY — Any value in <i>MASK</i> along <i>DIM</i> is true	130
8.20	ASIN — Arcsine function	131
8.21	ASIND — Arcsine function, degrees	132
8.22	ASINH — Inverse hyperbolic sine function	132
8.23	ASINPI — Circular arc sine function	133
8.24	ASSOCIATED — Status of a pointer or pointer/target pair	134
8.25	ATAN — Arctangent function	135
8.26	ATAN2 — Arctangent function	136
8.27	ATAN2D — Arctangent function, degrees	136
8.28	ATAN2PI — Circular arc tangent function	137
8.29	ATAND — Arctangent function, degrees	138
8.30	ATANH — Inverse hyperbolic tangent function	139
8.31	ATANPI — Circular arc tangent function	139
8.32	ATOMIC_ADD — Atomic ADD operation	140
8.33	ATOMIC_AND — Atomic bitwise AND operation	141
8.34	ATOMIC_CAS — Atomic compare and swap	141
8.35	ATOMIC_DEFINE — Setting a variable atomically	142
8.36	ATOMIC_FETCH_ADD — Atomic ADD operation with prior fetch ..	143
8.37	ATOMIC_FETCH_AND — Atomic bitwise AND operation with prior fetch	144
8.38	ATOMIC_FETCH_OR — Atomic bitwise OR operation with prior fetch	144
8.39	ATOMIC_FETCH_XOR — Atomic bitwise XOR operation with prior fetch	145
8.40	ATOMIC_OR — Atomic bitwise OR operation	146
8.41	ATOMIC_REF — Obtaining the value of a variable atomically ..	147
8.42	ATOMIC_XOR — Atomic bitwise OR operation	148
8.43	BACKTRACE — Show a backtrace	148
8.44	BESSEL_J0 — Bessel function of the first kind of order 0	149
8.45	BESSEL_J1 — Bessel function of the first kind of order 1	149
8.46	BESSEL_JN — Bessel function of the first kind	150
8.47	BESSEL_Y0 — Bessel function of the second kind of order 0 ..	150
8.48	BESSEL_Y1 — Bessel function of the second kind of order 1 ..	151
8.49	BESSEL_YN — Bessel function of the second kind	152
8.50	BGE — Bitwise greater than or equal to	152
8.51	BGT — Bitwise greater than	153
8.52	BIT_SIZE — Bit size inquiry function	153
8.53	BLE — Bitwise less than or equal to	154
8.54	BLT — Bitwise less than	154
8.55	BTEST — Bit test function	155
8.56	C_ASSOCIATED — Status of a C pointer	155
8.57	C_F_POINTER — Convert C into Fortran pointer	156
8.58	C_F_PROCPINTER — Convert C into Fortran procedure pointer ..	157
8.59	C_FUNLOC — Obtain the C address of a procedure	158
8.60	C_LOC — Obtain the C address of an object	158

8.61	C_SIZEOF — Size in bytes of an expression.....	159
8.62	CEILING — Integer ceiling function.....	160
8.63	CHAR — Character conversion function.....	160
8.64	CHDIR — Change working directory.....	161
8.65	CHMOD — Change access permissions of files.....	162
8.66	CMPLX — Complex conversion function.....	162
8.67	CO_BROADCAST — Copy a value to all images the current set of images.....	163
8.68	CO_MAX — Maximal value on the current set of images.....	164
8.69	CO_MIN — Minimal value on the current set of images.....	165
8.70	CO_REDUCE — Reduction of values on the current set of images..	166
8.71	CO_SUM — Sum of values on the current set of images.....	167
8.72	COMMAND_ARGUMENT_COUNT — Get number of command line arguments.....	168
8.73	COMPILER_OPTIONS — Options passed to the compiler.....	168
8.74	COMPILER_VERSION — Compiler version string.....	169
8.75	COMPLEX — Complex conversion function.....	169
8.76	CONJG — Complex conjugate function.....	170
8.77	COS — Cosine function.....	171
8.78	COSD — Cosine function, degrees.....	171
8.79	COSH — Hyperbolic cosine function.....	172
8.80	COSPI — Circular cosine function.....	173
8.81	COTAN — Cotangent function.....	173
8.82	COTAND — Cotangent function, degrees.....	174
8.83	COUNT — Count function.....	174
8.84	CPU_TIME — CPU elapsed time in seconds.....	175
8.85	CSHIFT — Circular shift elements of an array.....	176
8.86	CTIME — Convert a time into a string.....	177
8.87	DATE_AND_TIME — Date and time subroutine.....	178
8.88	DBLE — Double conversion function.....	179
8.89	DCMPLX — Double complex conversion function.....	179
8.90	DIGITS — Significant binary digits function.....	180
8.91	DIM — Positive difference.....	180
8.92	DOT_PRODUCT — Dot product function.....	181
8.93	DPROD — Double product function.....	182
8.94	DREAL — Double real part function.....	182
8.95	DSHIFTL — Combined left shift.....	183
8.96	DSHIFTR — Combined right shift.....	183
8.97	DTIME — Execution time subroutine (or function).....	184
8.98	EOSHIFT — End-off shift elements of an array.....	185
8.99	EPSILON — Epsilon function.....	186
8.100	ERF — Error function.....	187
8.101	ERFC — Error function.....	187
8.102	ERFC_SCALED — Error function.....	188
8.103	ETIME — Execution time subroutine (or function).....	188
8.104	EVENT_QUERY — Query whether a coarray event has occurred..	189
8.105	EXECUTE_COMMAND_LINE — Execute a shell command.....	190

8.106	EXIT	— Exit the program with status.....	191
8.107	EXP	— Exponential function.....	191
8.108	EXPONENT	— Exponent function.....	192
8.109	EXTENDS_TYPE_OF	— Query dynamic type for extension.....	192
8.110	FDATE	— Get the current time as a string.....	193
8.111	FGET	— Read a single character in stream mode from stdin..	194
8.112	FGETC	— Read a single character in stream mode.....	194
8.113	FINDLOC	— Search an array for a value.....	195
8.114	FLOOR	— Integer floor function.....	196
8.115	FLUSH	— Flush I/O unit(s).....	197
8.116	FNUM	— File number function.....	198
8.117	FPUT	— Write a single character in stream mode to stdout..	198
8.118	FPUTC	— Write a single character in stream mode.....	199
8.119	FRACTION	— Fractional part of the model representation....	200
8.120	FREE	— Frees memory.....	200
8.121	FSEEK	— Low level file positioning subroutine.....	201
8.122	FSTAT	— Get file status.....	202
8.123	FTELL	— Current stream position.....	203
8.124	GAMMA	— Gamma function.....	203
8.125	GERROR	— Get last system error message.....	204
8.126	GETARG	— Get command line arguments.....	204
8.127	GET_COMMAND	— Get the entire command line.....	205
8.128	GET_COMMAND_ARGUMENT	— Get command line arguments...	206
8.129	GETCWD	— Get current working directory.....	207
8.130	GETENV	— Get an environmental variable.....	207
8.131	GET_ENVIRONMENT_VARIABLE	— Get an environmental variable..	208
8.132	GETGID	— Group ID function.....	209
8.133	GETLOG	— Get login name.....	209
8.134	GETPID	— Process ID function.....	210
8.135	GET_TEAM	— Get the handle of a team.....	210
8.136	GETUID	— User ID function.....	211
8.137	GMTIME	— Convert time to GMT info.....	211
8.138	HOSTNM	— Get system host name.....	212
8.139	HUGE	— Largest number of a kind.....	212
8.140	HYPOT	— Euclidean distance function.....	213
8.141	IACHAR	— Code in ASCII collating sequence.....	213
8.142	IALL	— Bitwise AND of array elements.....	214
8.143	IAND	— Bitwise logical and.....	215
8.144	IANY	— Bitwise OR of array elements.....	216
8.145	IARGC	— Get the number of command line arguments.....	217
8.146	IBCLR	— Clear bit.....	217
8.147	IBITS	— Bit extraction.....	218
8.148	IBSET	— Set bit.....	218
8.149	ICHAR	— Character-to-integer conversion function.....	219
8.150	IDATE	— Get current local time subroutine (day/month/year) ..	220
8.151	IEOR	— Bitwise logical exclusive or.....	221
8.152	IERRNO	— Get the last system error number.....	222

8.153	IMAGE_INDEX — Function that converts a cosubscript to an image index	222
8.154	INDEX — Position of a substring within a string	222
8.155	INT — Convert to integer type	223
8.156	INT2 — Convert to 16-bit integer type	224
8.157	INT8 — Convert to 64-bit integer type	224
8.158	IOR — Bitwise logical or	225
8.159	IPARITY — Bitwise XOR of array elements	225
8.160	IRAND — Integer pseudo-random number	226
8.161	IS_CONTIGUOUS — Test whether an array is contiguous	227
8.162	IS_IOSTAT_END — Test for end-of-file value	227
8.163	IS_IOSTAT_EOR — Test for end-of-record value	228
8.164	ISATTY — Whether a unit is a terminal device	229
8.165	ISHFT — Shift bits	229
8.166	ISHFTC — Shift bits circularly	230
8.167	ISNAN — Test for a NaN	230
8.168	ITIME — Get current local time subroutine (hour/minutes/seconds)	231
8.169	KILL — Send a signal to a process	231
8.170	KIND — Kind of an entity	232
8.171	LBOUND — Lower dimension bounds of an array	232
8.172	LCOBOUND — Lower codimension bounds of an array	233
8.173	LEADZ — Number of leading zero bits of an integer	234
8.174	LEN — Length of a character entity	234
8.175	LEN_TRIM — Length of a character entity without trailing blank characters	235
8.176	LGE — Lexical greater than or equal	235
8.177	LGT — Lexical greater than	236
8.178	LINK — Create a hard link	237
8.179	LLE — Lexical less than or equal	237
8.180	LLT — Lexical less than	238
8.181	LNBLNK — Index of the last non-blank character in a string ..	238
8.182	LOC — Returns the address of a variable	239
8.183	LOG — Natural logarithm function	239
8.184	LOG10 — Base 10 logarithm function	240
8.185	LOG_GAMMA — Logarithm of the Gamma function	241
8.186	LOGICAL — Convert to logical type	241
8.187	LSHIFT — Left shift bits	242
8.188	LSTAT — Get file status	242
8.189	LTIME — Convert time to local time info	243
8.190	MALLOC — Allocate dynamic memory	244
8.191	MASKL — Left justified mask	244
8.192	MASKR — Right justified mask	245
8.193	MATMUL — matrix multiplication	245
8.194	MAX — Maximum value of an argument list	246
8.195	MAXEXPONENT — Maximum exponent of a real kind	247
8.196	MAXLOC — Location of the maximum value within an array ..	247

8.197	MAXVAL — Maximum value of an array	248
8.198	MCLOCK — Time function	249
8.199	MCLOCK8 — Time function (64-bit)	249
8.200	MERGE — Merge variables	250
8.201	MERGE_BITS — Merge of bits under mask	250
8.202	MIN — Minimum value of an argument list	251
8.203	MINEXPONENT — Minimum exponent of a real kind	251
8.204	MINLOC — Location of the minimum value within an array ..	252
8.205	MINVAL — Minimum value of an array	253
8.206	MOD — Remainder function	253
8.207	MODULO — Modulo function	254
8.208	MOVE_ALLOC — Move allocation from one object to another ..	255
8.209	MVBITS — Move bits from one integer to another	256
8.210	NEAREST — Nearest representable number	257
8.211	NEW_LINE — New line character	257
8.212	NINT — Nearest whole number	258
8.213	NORM2 — Euclidean vector norms	258
8.214	NOT — Logical negation	259
8.215	NULL — Function that returns an disassociated pointer	259
8.216	NUM_IMAGES — Function that returns the number of images ..	260
8.217	OR — Bitwise logical OR	261
8.218	OUT_OF_RANGE — Range check for numerical conversion	262
8.219	PACK — Pack an array into an array of rank one	262
8.220	PARITY — Reduction with exclusive OR	263
8.221	PERROR — Print system error message	264
8.222	POPCNT — Number of bits set	264
8.223	POPPAR — Parity of the number of bits set	265
8.224	PRECISION — Decimal precision of a real kind	265
8.225	PRESENT — Determine whether an optional dummy argument is specified	266
8.226	PRODUCT — Product of array elements	266
8.227	RADIX — Base of a model number	267
8.228	RAN — Real pseudo-random number	268
8.229	RAND — Real pseudo-random number	268
8.230	RANDOM_INIT — Initialize a pseudo-random number generator ..	268
8.231	RANDOM_NUMBER — Pseudo-random number	269
8.232	RANDOM_SEED — Initialize a pseudo-random number sequence ..	270
8.233	RANGE — Decimal exponent range	271
8.234	RANK — Rank of a data object	271
8.235	REAL — Convert to real type	272
8.236	RENAME — Rename a file	273
8.237	REPEAT — Repeated string concatenation	273
8.238	RESHAPE — Function to reshape an array	274
8.239	RRSPACING — Reciprocal of the relative spacing	274
8.240	RSHIFT — Right shift bits	275
8.241	SAME_TYPE_AS — Query dynamic types for equality	275
8.242	SCALE — Scale a real value	276

8.243	SCAN	— Scan a string for the presence of a set of characters ..	276
8.244	SECNDS	— Time function	277
8.245	SECOND	— CPU time function	277
8.246	SELECTED_CHAR_KIND	— Choose character kind	278
8.247	SELECTED_INT_KIND	— Choose integer kind	279
8.248	SELECTED_LOGICAL_KIND	— Choose logical kind	279
8.249	SELECTED_REAL_KIND	— Choose real kind	280
8.250	SELECTED_UNSIGNED_KIND	— Choose unsigned kind	281
8.251	SET_EXPONENT	— Set the exponent of the model	281
8.252	SHAPE	— Determine the shape of an array	282
8.253	SHIFTA	— Right shift with fill	282
8.254	SHIFTL	— Left shift	283
8.255	SHIFTR	— Right shift	283
8.256	SIGN	— Sign copying function	284
8.257	SIGNAL	— Signal handling subroutine (or function)	284
8.258	SIN	— Sine function	285
8.259	SIND	— Sine function, degrees	286
8.260	SINH	— Hyperbolic sine function	286
8.261	SINPI	— Circular sine function	287
8.262	SIZE	— Determine the size of an array	288
8.263	SIZEOF	— Size in bytes of an expression	288
8.264	SLEEP	— Sleep for the specified number of seconds	289
8.265	SPACING	— Smallest distance between two numbers of a given type	289
8.266	SPLIT	— Parse a string into tokens, one at a time	290
8.267	SPREAD	— Add a dimension to an array	291
8.268	SQRT	— Square-root function	291
8.269	SRAND	— Reinitialize the random number generator	292
8.270	STAT	— Get file status	292
8.271	STORAGE_SIZE	— Storage size in bits	294
8.272	SUM	— Sum of array elements	294
8.273	SYMLNK	— Create a symbolic link	295
8.274	SYSTEM	— Execute a shell command	295
8.275	SYSTEM_CLOCK	— Time function	296
8.276	TAN	— Tangent function	297
8.277	TAND	— Tangent function, degrees	298
8.278	TANH	— Hyperbolic tangent function	298
8.279	TANPI	— Circular tangent function	299
8.280	TEAM_NUMBER	— Retrieve team id of given team	299
8.281	THIS_IMAGE	— Function that returns the cosubscript index of this image	300
8.282	TIME	— Time function	301
8.283	TIME8	— Time function (64-bit)	301
8.284	TINY	— Smallest positive number of a real kind	302
8.285	TRAILZ	— Number of trailing zero bits of an integer	302
8.286	TRANSFER	— Transfer bit patterns	303
8.287	TRANSPOSE	— Transpose an array of rank two	304

8.288	TRIM — Remove trailing blank characters of a string	304
8.289	TTYNAM — Get the name of a terminal device	305
8.290	UBOUND — Upper dimension bounds of an array	305
8.291	UCOBOUND — Upper codimension bounds of an array	306
8.292	UINT — Convert to UNSIGNED type	306
8.293	UMASK — Set the file creation mask	306
8.294	UMASKL — Unsigned left justified mask	307
8.295	UMASKR — Unsigned right justified mask	307
8.296	UNLINK — Remove a file from the file system	308
8.297	UNPACK — Unpack an array of rank one into an array	308
8.298	VERIFY — Scan a string for characters not a given set	309
8.299	XOR — Bitwise logical exclusive OR	310
9	Intrinsic Modules	311
9.1	ISO_FORTRAN_ENV	311
9.2	ISO_C_BINDING	313
9.3	IEEE modules: IEEE_EXCEPTIONS, IEEE_ARITHMETIC, and IEEE_FEATURES	315
9.4	OpenMP Modules OMP_LIB and OMP_LIB_KINDS	315
9.5	OpenACC Module OPENACC	318
	Contributing	321
	Contributors to GNU Fortran	321
	Projects	322
	GNU General Public License	323
	GNU Free Documentation License	335
	ADDENDUM: How to use this License for your documents	342
	Funding Free Software	343
	Option Index	345
	Keyword Index	347

1 Introduction

This manual documents the use of `gfortran`, the GNU Fortran compiler. You can find in this manual how to invoke `gfortran`, as well as its features and incompatibilities.

Warning: This document, and the compiler it describes, are still under development. While efforts are made to keep it up-to-date, it might not accurately reflect the status of the most recent GNU Fortran compiler.

1.1 About GNU Fortran

The GNU Fortran compiler is the successor to `g77`, the Fortran 77 front end included in GCC prior to version 4 (released in 2005). While it is backward-compatible with most `g77` extensions and command-line options, `gfortran` is a completely new implementation designed to support more modern dialects of Fortran. GNU Fortran implements the Fortran 77, 90 and 95 standards completely, most of the Fortran 2003 and 2008 standards, and some features from the 2018 standard. It also implements several extensions including OpenMP and OpenACC support for parallel programming.

The GNU Fortran compiler passes the NIST Fortran 77 Test Suite (http://www.fortran-2000.com/ArnaudRecipes/fcvs21_f95.html), and produces acceptable results on the LAPACK Test Suite (<https://www.netlib.org/lapack/faq.html>). It also provides respectable performance on the Polyhedron Fortran compiler benchmarks (https://polyhedron.com/?page_id=175) and the Livermore Fortran Kernels test (<https://www.netlib.org/benchmark/livermore>). It has been used to compile a number of large real-world programs, including the HARMONIE and HIRLAM weather forecasting code (<http://hirlam.org/>) and the Tonto quantum chemistry package (<https://github.com/dylan-jayatilaka/tonto>); see <https://gcc.gnu.org/wiki/GfortranApps> for an extended list.

GNU Fortran provides the following functionality:

- Read a program, stored in a file and containing *source code* instructions written in Fortran 77.
- Translate the program into instructions a computer can carry out more quickly than it takes to translate the original Fortran instructions. The result after compilation of a program is *machine code*, which is efficiently translated and processed by a machine such as your computer. Humans usually are not as good writing machine code as they are at writing Fortran (or C++, Ada, or Java), because it is easy to make tiny mistakes writing machine code.
- Provide information about the reasons why the compiler may be unable to create a binary from the source code, for example if the source code is flawed. The Fortran language standards require that the compiler can point out mistakes in your code. An incorrect usage of the language causes an *error message*.

The compiler also attempts to diagnose cases where your program contains a correct usage of the language, but instructs the computer to do something questionable. This kind of diagnostic message is called a *warning message*.

- Provide optional information about the translation passes from the source code to machine code. This can help you to find the cause of certain bugs which may not be

obvious in the source code, but may be more easily found at a lower level compiler output. It also helps developers to find bugs in the compiler itself.

- Provide information in the generated machine code that can make it easier to find bugs in the program (using a debugging tool, called a *debugger*, such as the GNU Debugger *gdb*).
- Locate and gather machine code already generated to perform actions requested by statements in the program. This machine code is organized into *modules* and is located and *linked* to the user program.

The GNU Fortran compiler consists of several components:

- A version of the `gcc` command (which also might be installed as the system's `cc` command) that also understands and accepts Fortran source code. The `gcc` command is the *driver* program for all the languages in the GNU Compiler Collection (GCC); With `gcc`, you can compile the source code of any language for which a front end is available in GCC.
- The `gfortran` command itself, which also might be installed as the system's `f95` command. `gfortran` is just another driver program, but specifically for the Fortran compiler only. The primary difference between the `gcc` and `gfortran` commands is that the latter automatically links the correct libraries to your program.
- A collection of run-time libraries. These libraries contain the machine code needed to support capabilities of the Fortran language that are not directly provided by the machine code generated by the `gfortran` compilation phase, such as intrinsic functions and subroutines, and routines for interaction with files and the operating system.
- The Fortran compiler itself, (`f951`). This is the GNU Fortran parser and code generator, linked to and interfaced with the GCC backend library. `f951` “translates” the source code to assembler code. You would typically not use this program directly; instead, the `gcc` or `gfortran` driver programs call it for you.

1.2 GNU Fortran and GCC

GNU Fortran is a part of GCC, the *GNU Compiler Collection*. GCC consists of a collection of front ends for various languages, which translate the source code into a language-independent form called *GENERIC*. This is then processed by a common middle end which provides optimization, and then passed to one of a collection of back ends which generate code for different computer architectures and operating systems.

Functionally, this is implemented with a driver program (`gcc`) which provides the command-line interface for the compiler. It calls the relevant compiler front-end program (e.g., `f951` for Fortran) for each file in the source code, and then calls the assembler and linker as appropriate to produce the compiled output. In a copy of GCC that has been compiled with Fortran language support enabled, `gcc` recognizes files with `.f`, `.for`, `.ftn`, `.f90`, `.f95`, `.f03` and `.f08` extensions as Fortran source code, and compiles it accordingly. A `gfortran` driver program is also provided, which is identical to `gcc` except that it automatically links the Fortran runtime libraries into the compiled program.

Source files with `.f`, `.for`, `.fpp`, `.ftn`, `.F`, `.FOR`, `.FPP`, and `.FTN` extensions are treated as fixed form. Source files with `.f90`, `.f95`, `.f03`, `.f08`, `.F90`, `.F95`, `.F03` and `.F08` extensions are treated as free form. The capitalized versions of either form are run through

preprocessing. Source files with the lower case `.fpp` extension are also run through preprocessing.

This manual specifically documents the Fortran front end, which handles the programming language's syntax and semantics. The aspects of GCC that relate to the optimization passes and the back-end code generation are documented in the GCC manual; see Section "Introduction" in *Using the GNU Compiler Collection (GCC)*. The two manuals together provide a complete reference for the GNU Fortran compiler.

1.3 Standards

Fortran is developed by the Working Group 5 of Sub-Committee 22 of the Joint Technical Committee 1 of the International Organization for Standardization and the International Electrotechnical Commission (IEC). This group is known as WG5 (<http://www.nag.co.uk/sc22wg5/>). Official Fortran standard documents are available for purchase from ISO; a collection of free documents (typically final drafts) are also available on the wiki (<https://gcc.gnu.org/wiki/GFortranStandards>).

The GNU Fortran compiler implements ISO/IEC 1539:1997 (Fortran 95). As such, it can also compile essentially all standard-compliant Fortran 90 and Fortran 77 programs. It also supports the ISO/IEC TR-15581 enhancements to allocatable arrays.

GNU Fortran also supports almost all of ISO/IEC 1539-1:2004 (Fortran 2003) and ISO/IEC 1539-1:2010 (Fortran 2008). It has partial support for features introduced in ISO/IEC 1539:2018 (Fortran 2018), the most recent version of the Fortran language standard, including full support for the Technical Specification **Further Interoperability of Fortran with C** (ISO/IEC TS 29113:2012). More details on support for these standards can be found in the following sections of the documentation.

1.3.1 Fortran 95 status

The Fortran 95 standard specifies in Part 2 (ISO/IEC 1539-2:2000) varying length character strings. While GNU Fortran currently does not support such strings directly, there exist two Fortran implementations for them, which work with GNU Fortran. One can be found at <http://user.astro.wisc.edu/~townsend/static.php?ref=iso-varying-string>.

Deferred-length character strings of Fortran 2003 supports part of the features of `ISO_VARYING_STRING` and should be considered as replacement. (Namely, allocatable or pointers of the type `character(len=:)`.)

Part 3 of the Fortran 95 standard (ISO/IEC 1539-3:1998) defines Conditional Compilation, which is not widely used and not directly supported by the GNU Fortran compiler. You can use the program `coco` to preprocess such files (<http://www.daniellnagle.com/coco.html>).

1.3.2 Fortran 2003 status

GNU Fortran implements the Fortran 2003 (ISO/IEC 1539-1:2004) standard except for finalization support, which is incomplete. See the wiki page (<https://gcc.gnu.org/wiki/Fortran2003>) for a full list of new features introduced by Fortran 2003 and their implementation status.

1.3.3 Fortran 2008 status

The GNU Fortran compiler supports almost all features of Fortran 2008; the wiki (<https://gcc.gnu.org/wiki/Fortran2008Status>) has some information about the current implementation status. In particular, the following are not yet supported:

- `DO CONCURRENT` and `FORALL` do not recognize a type-spec in the loop header.
- The change to permit any constant expression in subscripts and nested implied-do limits in a `DATA` statement has not been implemented.

1.3.4 Fortran 2018 status

Fortran 2018 (ISO/IEC 1539:2018) is the most recent version of the Fortran language standard. GNU Fortran implements some of the new features of this standard:

- All Fortran 2018 features derived from ISO/IEC TS 29113:2012, “Further Interoperability of Fortran with C”, are supported by GNU Fortran. This includes assumed-type and assumed-rank objects and the `SELECT RANK` construct as well as the parts relating to `BIND(C)` functions. See also Section 6.1.6 [Further Interoperability of Fortran with C], page 78.
- GNU Fortran supports a subset of features derived from ISO/IEC TS 18508:2015, “Additional Parallel Features in Fortran”:
 - The new atomic `ADD`, `CAS`, `FETCH` and `ADD/OR/XOR`, `OR` and `XOR` intrinsics.
 - The `CO_MIN` and `CO_MAX` and `SUM` reduction intrinsics, and the `CO_BROADCAST` and `CO_REDUCE` intrinsic, except that those do not support polymorphic types or types with allocatable, pointer or polymorphic components.
 - Events (`EVENT POST`, `EVENT WAIT`, `EVENT_QUERY`).
 - Failed images (`FAIL IMAGE`, `IMAGE_STATUS`, `FAILED_IMAGES`, `STOPPED_IMAGES`).
- An `ERROR STOP` statement is permitted in a `PURE` procedure.
- GNU Fortran supports the `IMPLICIT NONE` statement with an `implicit-none-spec-list`.
- The behavior of the `INQUIRE` statement with the `RECL=` specifier now conforms to Fortran 2018.

Part I: Invoking GNU Fortran

2 GNU Fortran Command Options

The `gfortran` command supports all the options supported by the `gcc` command. Only options specific to GNU Fortran are documented here.

See Section “GCC Command Options” in *Using the GNU Compiler Collection (GCC)*, for information on the non-Fortran-specific aspects of the `gcc` command (and, therefore, the `gfortran` command).

All GCC and GNU Fortran options are accepted both by `gfortran` and by `gcc` (as well as any other drivers built at the same time, such as `g++`), since adding GNU Fortran to the GCC distribution enables acceptance of GNU Fortran options by all of the relevant drivers.

In some cases, options have positive and negative forms; the negative form of `-ffoo` would be `-fno-foo`. This manual documents only one of these two forms, whichever one is not the default.

2.1 Option summary

Here is a summary of all the options specific to GNU Fortran, grouped by type. Explanations are in the following sections.

Fortran Language Options

See Section 2.2 [Options controlling Fortran dialect], page 9.

```
-fall-intrinsics -fallow-argument-mismatch -fallow-invalid-boz
-fbackslash -fcray-pointer -fd-lines-as-code -fd-lines-as-comments
-fdec -fdec-char-conversions -fdec-structure -fdec-intrinsic-ints
-fdec-static -fdec-math -fdec-include -fdec-format-defaults
-fdec-blank-format-item -fdefault-double-8 -fdefault-integer-8
-fdefault-real-8 -fdefault-real-10 -fdefault-real-16 -fdollar-ok
-ffixed-line-length-n -ffixed-line-length-none -fpad-source
-ffree-form -ffree-line-length-n -ffree-line-length-none
-fimplicit-none -finteger-4-integer-8 -fmax-identifier-length
-fmodule-private -ffixed-form -fno-range-check -fopenacc -fopenmp
-fopenmp-allocators -fopenmp-simd -freal-4-real-10 -freal-4-real-16
-freal-4-real-8 -freal-8-real-10 -freal-8-real-16 -freal-8-real-4
-std=std -ftest-forall-temp -funsigned
```

Preprocessing Options

See Section 2.3 [Enable and customize preprocessing], page 14.

```
-A-question[=answer]
-Aquestion=answer -C -CC -Dmacro[=defn]
-H -P
-Umacro -cpp -dD -dI -dM -dN -dU -fworking-directory
-imultilib dir
-iprefix file -iquote -isysroot dir -isystem dir -nocpp
-nostdinc
-undef
```

Error and Warning Options

See Section 2.4 [Options to request or suppress errors and warnings], page 18.

```
-Waliasing -Wall -Wampersand -Warray-bounds
-Wc-binding-type -Wcharacter-truncation -Wconversion
-Wdo-subscript -Wfunction-elimination -Wimplicit-interface
```

```
-Wimplicit-procedure -Wintrinsic-shadow -Wuse-without-only
-Wintrinsic-std -Wline-truncation -Wno-align-commons
-Wno-overwrite-recursive -Wno-tabs -Wreal-q-constant -Wsurprising
-Wunderflow -Wunused-parameter -Wrealloc-lhs -Wrealloc-lhs-all
-Wfrontend-loop-interchange -Wtarget-lifetime -fmax-errors=n
-fsyntax-only -pedantic
-pedantic-errors
```

Debugging Options

See Section 2.5 [Options for debugging your program], page 22.

```
-fbacktrace -fdebug-aux-vars -ffpe-trap=list
-ffpe-summary=list
```

Directory Options

See Section 2.6 [Options for directory search], page 24.

```
-Idir -Jdir -fintrinsic-modules-path dir
```

Link Options

See Section 2.7 [Options for influencing the linking step], page 24.

```
-static-libgfortran -static-libquadmath
```

Runtime Options

See Section 2.8 [Options for influencing runtime behavior], page 24.

```
-fconvert=conversion -fmax-subrecord-length=length
-frecord-marker=length -fsign-zero
```

Interoperability Options

See Section 2.11 [Options for interoperability], page 34.

```
-fc-prototypes -fc-prototypes-external
```

Code Generation Options

See Section 2.10 [Options for code generation conventions], page 26.

```
-faggressive-function-elimination -fbblas-matmul-limit=n
-fbounds-check -ftail-call-workaround -ftail-call-workaround=n
-fcheck-array-temporaries
-fcheck=<all|array-temps|bits|bounds|do|mem|pointer|recursion>
-fcoarray=<none|single|lib> -fexternal-blas -fexternal-blas64 -ff2c
-ffrontend-loop-interchange -ffrontend-optimize
-finit-character=n -finit-integer=n -finit-local-zero
-finit-derived -finit-logical=<true|false>
-finit-real=<zero|inf|-inf|nan|snan>
-finline-intrinsics[=<minloc,maxloc>]
-finline-matmul-limit=n
-finline-arg-packing -fmax-array-constructor=n
-fmax-stack-var-size=n -fno-align-commons -fno-automatic
-fno-protect-parens -fno-underscoring -fsecond-underscore
-fpack-derived -frealloc-lhs -frecursive -frepack-arrays
-fshort-enums -fstack-arrays
```

Developer Options

See Section 2.9 [GNU Fortran Developer Options], page 25.

```
-fdump-fortran-global -fdump-fortran-optimized
-fdump-fortran-original -fdump-parse-tree -save-temps
```

2.2 Options controlling Fortran dialect

The following options control the details of the Fortran dialect accepted by the compiler:

-ffree-form

-ffixed-form

Specify the layout used by the source file. The free form layout was introduced in Fortran 90. Fixed form was traditionally used in older Fortran programs. When neither option is specified, the source form is determined by the file extension.

-fall-intrinsics

This option causes all intrinsic procedures (including the GNU-specific extensions) to be accepted. This can be useful with **-std=** to force standard compliance but get access to the full range of intrinsics available with **gfortran**. As a consequence, **-Wintrinsics-std** is ignored and no user-defined procedure with the same name as any intrinsic is called except when it is explicitly declared **EXTERNAL**.

-fallow-argument-mismatch

Some code contains calls to external procedures with mismatches between the calls and the procedure definition, or with mismatches between different calls. Such code is nonconforming, and is usually flagged with an error. This options degrades the error to a warning that can only be disabled by disabling all warnings via **-w**. Only a single occurrence per argument is flagged by this warning. **-fallow-argument-mismatch** is implied by **-std=legacy**.

Using this option is *strongly* discouraged. It is possible to provide standard-conforming code that allows different types of arguments by using an explicit interface and **TYPE(*)**.

-fallow-invalid-boz

A BOZ literal constant can occur in a limited number of contexts in standard conforming Fortran. This option degrades an error condition to a warning, and allows a BOZ literal constant to appear where the Fortran standard would otherwise prohibit its use.

-fd-lines-as-code

-fd-lines-as-comments

Enable special treatment for lines beginning with **d** or **D** in fixed form sources. If the **-fd-lines-as-code** option is given they are treated as if the first column contained a blank. If the **-fd-lines-as-comments** option is given, they are treated as comment lines.

-fdec

DEC compatibility mode. Enables extensions and other features that mimic the default behavior of older compilers (such as DEC). These features are non-standard and should be avoided at all costs. For details on GNU Fortran's implementation of these extensions see the full documentation.

Other flags enabled by this switch are: **-fdollar-ok** **-fcray-pointer**
-fdec-char-conversions **-fdec-structure** **-fdec-intrinsic-ints**
-fdec-static **-fdec-math** **-fdec-include** **-fdec-blank-format-item**
-fdec-format-defaults

If `-fd-lines-as-code/-fd-lines-as-comments` are unset, then `-fdec` also sets `-fd-lines-as-comments`.

-fdec-char-conversions

Enable the use of character literals in assignments and `DATA` statements for non-character variables.

-fdec-structure

Enable `DEC STRUCTURE` and `RECORD` as well as `UNION`, `MAP`, and dot (‘.’) as a member separator (in addition to ‘%’). This is provided for compatibility only; Fortran 90 derived types should be used instead where possible.

-fdec-intrinsic-ints

Enable B/I/J/K kind variants of existing integer functions (e.g. `BIAND`, `IIAND`, `JRAND`, etc...). For a complete list of intrinsics see Chapter 8 [Intrinsic Procedures], page 119.

-fdec-math

Obsolete flag. The purpose of this option was to enable legacy math intrinsics such as `COTAN` and degree-valued trigonometric functions (e.g. `TAND`, `ATAND`, etc...) for compatibility with older code. This option is no longer operable. The trigonometric functions are now either part of Fortran 2023 or GNU extensions.

-fdec-static

Enable DEC-style `STATIC` and `AUTOMATIC` attributes to explicitly specify the storage of variables and other objects.

-fdec-include

Enable parsing of `INCLUDE` as a statement in addition to parsing it as `INCLUDE` line. When parsed as `INCLUDE` statement, `INCLUDE` does not have to be on a single line and can use line continuations.

-fdec-format-defaults

Enable format specifiers ‘F’, ‘G’ and ‘I’ to be used without width specifiers; default widths are used instead.

-fdec-blank-format-item

Enable a blank format item at the end of a format specification i.e. nothing following the final comma.

-fdollar-ok

Allow ‘\$’ as a valid non-first character in a symbol name. Symbols that start with ‘\$’ are rejected since it is unclear which rules to apply to implicit typing as different vendors implement different rules. Using ‘\$’ in `IMPLICIT` statements is also rejected.

-fbackslash

Change the interpretation of backslashes in string literals from a single backslash character to “C-style” escape characters. The following combinations are expanded: ‘\a’, ‘\b’, ‘\f’, ‘\n’, ‘\r’, ‘\t’, ‘\v’, ‘\\’, and ‘\0’ to the ASCII characters alert, backspace, form feed, newline, carriage return, horizontal tab, vertical tab, backslash, and NUL, respectively. Additionally, ‘\xnn’, ‘\unnnn’

and ‘\Unnnnnnnn’ (where each *n* is a hexadecimal digit) are translated into the Unicode characters corresponding to the specified code points. All other combinations of a character preceded by ‘\’ are unexpanded.

-fmodule-private

Set the default accessibility of module entities to **PRIVATE**. Use-associated entities are not accessible unless they are explicitly declared as **PUBLIC**.

-ffixed-line-length-*n*

Set column after which characters are ignored in typical fixed-form lines in the source file, and, unless **-fno-pad-source**, through which spaces are assumed (as if padded to that length) after the ends of short fixed-form lines.

Popular values for *n* include 72 (the standard and the default), 80 (card image), and 132 (corresponding to “extended-source” options in some popular compilers). *n* may also be ‘none’, meaning that the entire line is meaningful and that continued character constants never have implicit spaces appended to them to fill out the line. **-ffixed-line-length-0** means the same thing as **-ffixed-line-length-none**.

-fno-pad-source

By default fixed-form lines have spaces assumed (as if padded to that length) after the ends of short fixed-form lines. This is not done either if **-ffixed-line-length-0**, **-ffixed-line-length-none** or if **-fno-pad-source** option is used. With any of those options continued character constants never have implicit spaces appended to them to fill out the line.

-ffree-line-length-*n*

Set column after which characters are ignored in typical free-form lines in the source file. The default value is 132. *n* may be ‘none’, meaning that the entire line is meaningful. **-ffree-line-length-0** means the same thing as **-ffree-line-length-none**.

-fmax-identifier-length=*n*

Specify the maximum allowed identifier length. Typical values are 31 (Fortran 95) and 63 (Fortran 2003 and later).

-fimplicit-none

Specify that no implicit typing is allowed, unless overridden by explicit **IMPLICIT** statements. This is the equivalent of adding **implicit none** to the start of every procedure.

-fcray-pointer

Enable the Cray pointer extension, which provides C-like pointer functionality.

-fopenacc

Enable handling of OpenACC directives ‘!\$acc’ in free-form Fortran and ‘!\$acc’, ‘c\$acc’ and ‘*\$acc’ in fixed-form Fortran. When **-fopenacc** is specified, the compiler generates accelerated code according to the OpenACC Application Programming Interface v2.6 <https://www.openacc.org>. This option implies **-pthread**, and thus is only supported on targets that have support for **-pthread**. The option **-fopenacc** implies **-frecursive**.

-fopenmp Enable handling of OpenMP directives ‘`!$omp`’ in Fortran. It additionally enables the conditional compilation sentinel ‘`!$`’ in Fortran. In fixed source form Fortran, the sentinels can also start with ‘`c`’ or ‘`*`’. When **-fopenmp** is specified, the compiler generates parallel code according to the OpenMP Application Program Interface v4.5 <https://www.openmp.org>. This option implies **-pthread**, and thus is only supported on targets that have support for **-pthread**. **-fopenmp** implies **-fopenmp-simd** and **-frecursive**.

-fopenmp-allocators

Enables handling of allocation, reallocation and deallocation of Fortran allocatable and pointer variables that are allocated using the ‘`!$omp allocators`’ and ‘`!$omp allocate`’ constructs. Files containing either directive have to be compiled with this option in addition to **-fopenmp**. Additionally, all files that might deallocate or reallocate a variable that has been allocated with an OpenMP allocator have to be compiled with this option. This includes intrinsic assignment to allocatable variables when reallocation may occur and deallocation due to either of the following: end of scope, explicit deallocation, ‘`intent(out)`’, deallocation of allocatable components etc. Files not changing the allocation status or only for components of a derived type that have not been allocated using those two directives do not need to be compiled with this option. Nor do files that handle such variables after they have been deallocated or allocated by the normal Fortran allocator.

-fopenmp-simd

Enable handling of OpenMP’s **simd**, **declare simd**, **declare reduction**, **assume**, **ordered**, **scan** and **loop** directive, and of combined or composite directives with **simd** as constituent with `!$omp` in Fortran. It additionally enables the conditional compilation sentinel ‘`!$`’ in Fortran. In fixed source form Fortran, the sentinels can also start with ‘`c`’ or ‘`*`’. Other OpenMP directives are ignored. Unless **-fopenmp** is additionally specified, the **loop** region binds to the current task region, independent of the specified **bind** clause.

-fno-range-check

Disable range checking on results of simplification of constant expressions during compilation. For example, GNU Fortran gives an error at compile time when simplifying `a = 1. / 0.` With this option, no error is given and `a` is assigned the value `+Infinity`. If an expression evaluates to a value outside of the relevant range of `[-HUGE():HUGE()]`, then the expression is replaced by `-Inf` or `+Inf` as appropriate. Similarly, `DATA i/Z'FFFFFFFF'/` results in an integer overflow on most systems, but with **-fno-range-check** the value “wraps around” and `i` is initialized to `-1` instead.

-fdefault-integer-8

Set the default integer and logical types to an 8 byte wide type. This option also affects the kind of integer constants like `42`. Unlike **-finteger-4-integer-8**, it does not promote variables with explicit kind declaration.

-fdefault-real-8

Set the default real type to an 8 byte wide type. This option also affects the kind of non-double real constants like 1.0. This option promotes the default width of DOUBLE PRECISION and double real constants like 1.d0 to 16 bytes if possible. If **-fdefault-double-8** is given along with **fdefault-real-8**, DOUBLE PRECISION and double real constants are not promoted. Unlike **-freal-4-real-8**, **fdefault-real-8** does not promote variables with explicit kind declarations.

-fdefault-real-10

Set the default real type to an 10 byte wide type. This option also affects the kind of non-double real constants like 1.0. This option promotes the default width of DOUBLE PRECISION and double real constants like 1.d0 to 16 bytes if possible. If **-fdefault-double-8** is given along with **fdefault-real-10**, DOUBLE PRECISION and double real constants are not promoted. Unlike **-freal-4-real-10**, **fdefault-real-10** does not promote variables with explicit kind declarations.

-fdefault-real-16

Set the default real type to an 16 byte wide type. This option also affects the kind of non-double real constants like 1.0. This option promotes the default width of DOUBLE PRECISION and double real constants like 1.d0 to 16 bytes if possible. If **-fdefault-double-8** is given along with **fdefault-real-16**, DOUBLE PRECISION and double real constants are not promoted. Unlike **-freal-4-real-16**, **fdefault-real-16** does not promote variables with explicit kind declarations.

-fdefault-double-8

Set the DOUBLE PRECISION type and double real constants like 1.d0 to an 8 byte wide type. Do nothing if this is already the default. This option prevents **-fdefault-real-8**, **-fdefault-real-10**, and **-fdefault-real-16**, from promoting DOUBLE PRECISION and double real constants like 1.d0 to 16 bytes.

-finteger-4-integer-8

Promote all INTEGER(KIND=4) entities to an INTEGER(KIND=8) entities. If KIND=8 is unavailable, then an error is issued. This option should be used with care and may not be suitable for your codes. Areas of possible concern include calls to external procedures, alignment in EQUIVALENCE and/or COMMON, generic interfaces, BOZ literal constant conversion, and I/O. Inspection of the intermediate representation of the translated Fortran code, produced by **-fdump-tree-original**, is suggested.

-freal-4-real-8**-freal-4-real-10****-freal-4-real-16****-freal-8-real-4****-freal-8-real-10****-freal-8-real-16**

Promote all REAL(KIND=M) entities to REAL(KIND=N) entities. If REAL(KIND=N) is unavailable, then an error is issued. The **-freal-4-** flags also affect the de-

fault real kind and the `-freal-8-` flags also the double-precision real kind. All other real-kind types are unaffected by this option. The promotion is also applied to real literal constants of default and double-precision kind and a specified kind number of 4 or 8, respectively. However, `-fdefault-real-8`, `-fdefault-real-10`, `-fdefault-real-10`, and `-fdefault-double-8` take precedence for the default and double-precision real kinds, both for real literal constants and for declarations without a kind number. Note that for `REAL(KIND=KIND(1.0))` the literal may get promoted and then the result may get promoted again. These options should be used with care and may not be suitable for your codes. Areas of possible concern include calls to external procedures, alignment in `EQUIVALENCE` and/or `COMMON`, generic interfaces, `BOZ` literal constant conversion, and I/O and calls to intrinsic procedures when passing a value to the `kind=` dummy argument. Inspection of the intermediate representation of the translated Fortran code, produced by `-fdump-fortran-original` or `-fdump-tree-original`, is suggested.

-std=std Specify the standard to which the program is expected to conform, which may be one of 'f95', 'f2003', 'f2008', 'f2018', 'f2023', 'gnu', or 'legacy'. The default value for *std* is 'gnu', which specifies a superset of the latest Fortran standard that includes all of the extensions supported by GNU Fortran, although warnings are given for obsolete extensions not recommended for use in new code. The 'legacy' value is equivalent but without the warnings for obsolete extensions, and may be useful for old nonstandard programs. The 'f95', 'f2003', 'f2008', 'f2018', and 'f2023' values specify strict conformance to the Fortran 95, Fortran 2003, Fortran 2008, Fortran 2018 and Fortran 2023 standards, respectively; errors are given for all extensions beyond the relevant language standard, and warnings are given for the Fortran 77 features that are permitted but obsolescent in later standards. The deprecated option '`-std=f2008ts`' acts as an alias for '`-std=f2018`'. It is only present for backwards compatibility with earlier gfortran versions and should not be used any more. '`-std=f202y`' acts as an alias for '`-std=f2023`' and enables proposed features for testing Fortran 202y. As the Fortran 202y standard develops, implementation might change or the experimental new features might be removed.

-ftest-forall-temp

Enhance test coverage by forcing most forall assignments to use temporary.

-funsigned

Allow the experimental unsigned extension.

2.3 Enable and customize preprocessing

Many Fortran compilers including GNU Fortran allow passing the source code through a C preprocessor (CPP; sometimes also called the Fortran preprocessor, FPP) to allow for conditional compilation. In the case of GNU Fortran, this is the GNU C Preprocessor in the traditional mode. On systems with case-preserving file names, the preprocessor is automatically invoked if the filename extension is `.F`, `.FOR`, `.FTN`, `.fpp`, `.FPP`, `.F90`, `.F95`, `.F03` or `.F08`. To manually invoke the preprocessor on any file, use `-cpp`, to disable preprocessing on files where the preprocessor is run automatically, use `-nocpp`.

If a preprocessed file includes another file with the Fortran `INCLUDE` statement, the included file is not preprocessed. To preprocess included files, use the equivalent preprocessor statement `#include`.

If GNU Fortran invokes the preprocessor, `__GFORTRAN__` is defined. The macros `__GNUC__`, `__GNUC_MINOR__` and `__GNUC_PATCHLEVEL__` can be used to determine the version of the compiler. See Section “Overview” in *The C Preprocessor* for details.

GNU Fortran supports a number of `INTEGER` and `REAL` kind types in addition to the kind types required by the Fortran standard. The availability of any given kind type is architecture dependent. The following predefined preprocessor macros can be used to conditionally include code for these additional kind types: `__GFC_INT_1__`, `__GFC_INT_2__`, `__GFC_INT_8__`, `__GFC_INT_16__`, `__GFC_REAL_10__`, and `__GFC_REAL_16__`.

While CPP is the de facto standard for preprocessing Fortran code, Part 3 of the Fortran 95 standard (ISO/IEC 1539-3:1998) defines Conditional Compilation, which is not widely used and not directly supported by the GNU Fortran compiler.

The following options control preprocessing of Fortran code:

- cpp**
- nocpp** Enable preprocessing. The preprocessor is automatically invoked if the file extension is `.fpp`, `.FPP`, `.F`, `.FOR`, `.FTN`, `.F90`, `.F95`, `.F03` or `.F08`. Use this option to manually enable preprocessing of any kind of Fortran file.
 To disable preprocessing of files with any of the above listed extensions, use the negative form: `-nocpp`.
 The preprocessor is run in traditional mode. Any restrictions of the file format, especially the limits on line length, apply for preprocessed output as well, so it might be advisable to use the `-ffree-line-length-none` or `-ffixed-line-length-none` options.
- dM** Instead of the normal output, generate a list of `'#define'` directives for all the macros defined during the execution of the preprocessor, including predefined macros. This gives you a way of finding out what is predefined in your version of the preprocessor. Assuming you have no file `foo.f90`, the command
 `touch foo.f90; gfortran -cpp -E -dM foo.f90`
 shows all the predefined macros.
- dD** Like `-dM` except in two respects: it does not include the predefined macros, and it outputs both the `#define` directives and the result of preprocessing. Both kinds of output go to the standard output file.
- dN** Like `-dD`, but emit only the macro names, not their expansions.
- dU** Like `dD` except that only macros that are expanded, or whose definedness is tested in preprocessor directives, are output; the output is delayed until the use or test of the macro; and `'#undef'` directives are also output for macros tested but undefined at the time.
- dI** Output `'#include'` directives in addition to the result of preprocessing.
- fworking-directory**
 Enable generation of linemarkers in the preprocessor output that let the compiler know the current working directory at the time of preprocessing. When

this option is enabled, the preprocessor emits, after the initial linemarker, a second linemarker with the current working directory followed by two slashes. GCC uses this directory, when it is present in the preprocessed input, as the directory emitted as the current working directory in some debugging information formats. This option is implicitly enabled if debugging information is enabled, but this can be inhibited with the negated form `-fno-working-directory`. If the `-P` flag is present in the command line, this option has no effect, since no `#line` directives are emitted whatsoever.

-idirafter *dir*

Search *dir* for include files, but do it after all directories specified with `-I` and the standard system directories have been exhausted. *dir* is treated as a system include directory. If *dir* begins with `=`, then the `=` is replaced by the sysroot prefix; see `--sysroot` and `-isysroot`.

-imultilib *dir*

Use *dir* as a subdirectory of the directory containing target-specific C++ headers.

-iprefix *prefix*

Specify *prefix* as the prefix for subsequent `-iwithprefix` options. If the *prefix* represents a directory, you should include the final `'/'`.

-isysroot *dir*

This option is like the `--sysroot` option, but applies only to header files. See the `--sysroot` option for more information.

-iquote *dir*

Search *dir* only for header files requested with `#include "file"`; they are not searched for `#include <file>`, before all directories specified by `-I` and before the standard system directories. If *dir* begins with `=`, then the `=` is replaced by the sysroot prefix; see `--sysroot` and `-isysroot`.

-isystem *dir*

Search *dir* for header files, after all directories specified by `-I` but before the standard system directories. Mark it as a system directory, so that it gets the same special treatment as is applied to the standard system directories. If *dir* begins with `=`, then the `=` is replaced by the sysroot prefix; see `--sysroot` and `-isysroot`.

-nostdinc

Do not search the standard system directories for header files. Only the directories you have specified with `-I` options (and the directory of the current file, if appropriate) are searched.

-undef

Do not predefine any system-specific or GCC-specific macros. The standard predefined macros remain defined.

-Apredicate=*answer*

Make an assertion with the predicate *predicate* and answer *answer*. This form is preferred to the older form `-A predicate(answer)`, which is still supported, because it does not use shell special characters.

-A-predicate=answer

Cancel an assertion with the predicate *predicate* and answer *answer*.

-C

Do not discard comments. All comments are passed through to the output file, except for comments in processed directives, which are deleted along with the directive.

You should be prepared for side effects when using **-C**; it causes the preprocessor to treat comments as tokens in their own right. For example, comments appearing at the start of what would be a directive line have the effect of turning that line into an ordinary source line, since the first token on the line is no longer a '#'.
Warning: this currently handles C-Style comments only. The preprocessor does not yet recognize Fortran-style comments.

-CC

Do not discard comments, including during macro expansion. This is like **-C**, except that comments contained within macros are also passed through to the output file where the macro is expanded.

In addition to the side-effects of the **-C** option, the **-CC** option causes all C++-style comments inside a macro to be converted to C-style comments. This is to prevent later use of that macro from inadvertently commenting out the remainder of the source line. The **-CC** option is generally used to support lint comments.

Warning: this currently handles C- and C++-Style comments only. The preprocessor does not yet recognize Fortran-style comments.

-Dname

Predefine name as a macro, with definition 1.

-Dname=definition

The contents of *definition* are tokenized and processed as if they appeared during translation phase three in a '#define' directive. In particular, the definition is truncated by embedded newline characters.

If you are invoking the preprocessor from a shell or shell-like program you may need to use the shell's quoting syntax to protect characters such as spaces that have a meaning in the shell syntax.

If you wish to define a function-like macro on the command line, write its argument list with surrounding parentheses before the equals sign (if any). Parentheses are meaningful to most shells, so you need to quote the option. With sh and csh, **-D'name(args...)=definition'** works.

-D and **-U** options are processed in the order they are given on the command line. All **-imacros** file and **-include** file options are processed after all **-D** and **-U** options.

-H

Print the name of each header file used, in addition to other normal activities. Each name is indented to show how deep in the '#include' stack it is.

-P

Inhibit generation of linemarkers in the output from the preprocessor. This might be useful when running the preprocessor on something that is not C code, and is sent to a program that might be confused by the linemarkers.

-Uname Cancel any previous definition of *name*, either built in or provided with a **-D** option.

2.4 Options to request or suppress errors and warnings

Errors are diagnostic messages that report that the GNU Fortran compiler cannot compile the relevant piece of source code. The compiler continues to process the program in an attempt to report further errors to aid in debugging, but does not produce any compiled output.

Warnings are diagnostic messages that report constructions that are not inherently erroneous but that are risky or suggest there is likely to be a bug in the program. Unless **-Werror** is specified, they do not prevent compilation of the program.

You can request many specific warnings with options beginning **-W**, for example **-Wimplicit** to request warnings on implicit declarations. Each of these specific warning options also has a negative form beginning **-Wno-** to turn off warnings; for example, **-Wno-implicit**. This manual lists only one of the two forms, whichever is not the default.

These options control the amount and kinds of errors and warnings produced by GNU Fortran:

-fmax-errors=n

Limits the maximum number of error messages to *n*, at which point GNU Fortran bails out rather than attempting to continue processing the source code. If *n* is 0, there is no limit on the number of error messages produced.

-fsyntax-only

Check the code for syntax errors, but do not actually compile it. This generates module files for each module present in the code, but no other output file.

-Wpedantic

-pedantic

Issue warnings for uses of extensions to Fortran. **-pedantic** also applies to C-language constructs where they occur in GNU Fortran source files, such as use of `'\e'` in a character constant within a directive like `#include`.

Valid Fortran programs should compile properly with or without this option. However, without this option, certain GNU extensions and traditional Fortran features are supported as well. With this option, many of them are rejected.

Some users try to use **-pedantic** to check programs for conformance. They soon find that it does not do quite what they want—it finds some nonstandard practices, but not all. However, improvements to GNU Fortran in this area are welcome.

This should be used in conjunction with **-std=f95**, **-std=f2003**, **-std=f2008**, **-std=f2018** or **-std=f2023**.

-pedantic-errors

Like **-pedantic**, except that errors are produced rather than warnings.

-Wall

Enables commonly used warning options pertaining to usage that we recommend avoiding and that we believe are easy to avoid. This currently includes **-Waliasing**, **-Wampersand**, **-Wconversion**, **-Wsurprising**,

`-Wc-binding-type`, `-Wintrinsics-std`, `-Wtabs`, `-Wintrinsic-shadow`,
`-Wline-truncation`, `-Wtarget-lifetime`, `-Winteger-division`,
`-Wreal-q-constant`, `-Wunused` and `-Wundefined-do-loop`.

`-Waliasing`

Warn about possible aliasing of dummy arguments. Specifically, it warns if the same actual argument is associated with a dummy argument with `INTENT(IN)` and a dummy argument with `INTENT(OUT)` in a call with an explicit interface.

The following example triggers the warning.

```
interface
  subroutine bar(a,b)
    integer, intent(in) :: a
    integer, intent(out) :: b
  end subroutine
end interface
integer :: a

call bar(a,a)
```

`-Wampersand`

Warn about missing ampersand in continued character constants. The warning is given with `-Wampersand`, `-pedantic`, `-std=f95`, `-std=f2003`, `-std=f2008`, `-std=f2018` and `-std=f2023`. Note: With no ampersand given in a continued character constant, GNU Fortran assumes continuation at the first non-comment, non-whitespace character after the ampersand that initiated the continuation.

`-Warray-temporaries`

Warn about array temporaries generated by the compiler. The information generated by this warning is sometimes useful in optimization, in order to avoid such temporaries.

`-Wc-binding-type`

Warn if the a variable might not be C interoperable. In particular, warn if the variable has been declared using an intrinsic type with default kind instead of using a kind parameter defined for C interoperability in the intrinsic `ISO_C_Binding` module. This option is implied by `-Wall`.

`-Wcharacter-truncation`

Warn when a character assignment truncates the assigned string.

`-Wline-truncation`

Warn when a source code line is truncated. This option is implied by `-Wall`. For free-form source code, the default is `-Werror=line-truncation` such that truncations are reported as error.

`-Wconversion`

Warn about implicit conversions that are likely to change the value of the expression after conversion. Implied by `-Wall`.

`-Wconversion-extra`

Warn about implicit conversions between different types and kinds. This option does *not* imply `-Wconversion`.

-Wexternal-argument-mismatch

Warn about argument mismatches for dummy external procedures. This is implied by `-fc-prototypes-external` because generation of a valid C23 interface is not possible in such a case. Also implied by `-Wall`.

-Wextra Enables some warning options for usages of language features that may be problematic. This currently includes `-Wcompare-reals`, `-Wunused-parameter` and `-Wdo-subscript`.

-Wfrontend-loop-interchange

Warn when using `-ffrontend-loop-interchange` for performing loop interchanges.

-Wimplicit-interface

Warn if a procedure is called without an explicit interface. Note this only checks that an explicit interface is present. It does not check that the declared interfaces are consistent across program units.

-Wimplicit-procedure

Warn if a procedure is called that has neither an explicit interface nor has been declared as `EXTERNAL`.

-Winteger-division

Warn if a constant integer division truncates its result. As an example, `3/5` evaluates to 0.

-Wintrinsics-std

Warn if `gfortran` finds a procedure named like an intrinsic not available in the currently selected standard (with `-std`) and treats it as `EXTERNAL` procedure because of this. `-fall-intrinsics` can be used to never trigger this behavior and always link to the intrinsic regardless of the selected standard.

-Wno-overwrite-recursive

Do not warn when `-fno-automatic` is used with `-frecursive`. Recursion is broken if the relevant local variables do not have the attribute `AUTOMATIC` explicitly declared. This option can be used to suppress the warning when it is known that recursion is not broken. Useful for build environments that use `-Werror`.

-Wreal-q-constant

Produce a warning if a real-literal-constant contains a `q` exponent-letter.

-Wsurprising

Produce a warning when “suspicious” code constructs are encountered. While technically legal these usually indicate that an error has been made.

This currently produces a warning under the following circumstances:

- An `INTEGER`-typed `SELECT CASE` construct has a `CASE` that can never be matched as its lower value is greater than its upper value.
- A `LOGICAL`-typed `SELECT CASE` construct has three `CASE` statements.
- A `TRANSFER` specifies a source that is shorter than the destination.

- The type of a function result is declared more than once with the same type. If `-pedantic` or standard-conforming mode is enabled, this is an error.
 - A `CHARACTER` variable is declared with negative length.
 - With `-fopenmp`, for fixed-form source code, when an `omx` vendor-extension sentinel is encountered. (The equivalent `omp`, used in free-form source code, is diagnosed by default.)
 - With `-fopenacc`, when using named constances with clauses that take a variable as doing so has no effect.
- Wtabs** By default, tabs are accepted as whitespace, but tabs are not members of the Fortran Character Set. For continuation lines, a tab followed by a digit between 1 and 9 is supported. `-Wtabs` causes a warning to be issued if a tab is encountered. Note, `-Wtabs` is active for `-pedantic`, `-std=f95`, `-std=f2003`, `-std=f2008`, `-std=f2018`, `-std=f2023` and `-Wall`.
- Wundefined-do-loop**
Warn if a `DO` loop with step either 1 or -1 yields an underflow or an overflow during iteration of an induction variable of the loop. This option is implied by `-Wall`.
- Wunderflow**
Produce a warning when numerical constant expressions that yield an underflow are encountered during compilation. Enabled by default.
- Wintrinsic-shadow**
Warn if a user-defined procedure or module procedure has the same name as an intrinsic; in this case, an explicit interface or `EXTERNAL` or `INTRINSIC` declaration might be needed to get calls later resolved to the desired intrinsic/procedure. This option is implied by `-Wall`.
- Wuse-without-only**
Warn if a `USE` statement has no `ONLY` qualifier and thus implicitly imports all public entities of the used module.
- Wunused-dummy-argument**
Warn about unused dummy arguments. This option is implied by `-Wall`.
- Wunused-parameter**
Contrary to `gcc`'s meaning of `-Wunused-parameter`, `gfortran`'s implementation of this option does not warn about unused dummy arguments (see `-Wunused-dummy-argument`), but about unused `PARAMETER` values. `-Wunused-parameter` is implied by `-Wextra` if also `-Wunused` or `-Wall` is used.
- Walign-commons**
By default, `gfortran` warns about any occasion of variables being padded for proper alignment inside a `COMMON` block. This warning can be turned off via `-Wno-align-commons`. See also `-falign-commons`.
- Wfunction-elimination**
Warn if any calls to impure functions are eliminated by the optimizations enabled by the `-ffrontend-optimize` option. This option is implied by `-Wextra`.

-Wrealloc-lhs

Warn when the compiler might insert code to for allocation or reallocation of an allocatable array variable of intrinsic type in intrinsic assignments. In hot loops, the Fortran 2003 reallocation feature may reduce the performance. If the array is already allocated with the correct shape, consider using a whole-array array-spec (e.g. `(:,:,:)`) for the variable on the left-hand side to prevent the reallocation check. Note that in some cases the warning is shown, even if the compiler optimizes reallocation checks away. For instance, when the right-hand side contains the same variable multiplied by a scalar. See also **-frealloc-lhs**.

-Wrealloc-lhs-all

Warn when the compiler inserts code to for allocation or reallocation of an allocatable variable; this includes scalars and derived types.

-Wcompare-reals

Warn when comparing real or complex types for equality or inequality. This option is implied by **-Wextra**.

-Wtarget-lifetime

Warn if the pointer in a pointer assignment might be longer than the its target. This option is implied by **-Wall**.

-Wzerotrip

Warn if a DO loop is known to execute zero times at compile time. This option is implied by **-Wall**.

-Wdo-subscript

Warn if an array subscript inside a DO loop could lead to an out-of-bounds access even if the compiler cannot prove that the statement is actually executed, in cases like

```

      real a(3)
      do i=1,4
        if (condition(i)) then
          a(i) = 1.2
        end if
      end do

```

This option is implied by **-Wextra**.

-Werror Turns all warnings into errors.

See Section “Options to Request or Suppress Errors and Warnings” in *Using the GNU Compiler Collection (GCC)*, for information on more options offered by the back end shared by **gfortran**, **gcc** and other GNU compilers.

Some of these have no effect when compiling programs written in Fortran.

2.5 Options for debugging your program

GNU Fortran has various special options that are used for debugging your program.

-fdebug-aux-vars

Renames internal variables created by the **gfortran** front end and makes them accessible to a debugger. The name of the internal variables then start with

uppercase letters followed by an underscore. This option is useful for debugging the compiler's code generation together with `-fdump-tree-original` and enabling debugging of the executable program by using `-g` or `-ggdb3`.

`-ffpe-trap=list`

Specify a list of floating point exception traps to enable. On most systems, if a floating point exception occurs and the trap for that exception is enabled, a SIGFPE signal is sent and the program being aborted, producing a core file useful for debugging. *list* is a (possibly empty) comma-separated list of either `'none'` (to clear the set of exceptions to be trapped), or of the following exceptions: `'invalid'` (invalid floating point operation, such as `SQRT(-1.0)`), `'zero'` (division by zero), `'overflow'` (overflow in a floating point operation), `'underflow'` (underflow in a floating point operation), `'inexact'` (loss of precision during operation), and `'denormal'` (operation performed on a denormal value). The first five exceptions correspond to the five IEEE 754 exceptions, whereas the last one (`'denormal'`) is not part of the IEEE 754 standard but is available on some common architectures such as x86.

The first three exceptions (`'invalid'`, `'zero'`, and `'overflow'`) often indicate serious errors, and unless the program has provisions for dealing with these exceptions, enabling traps for these three exceptions is probably a good idea.

If the option is used more than once in the command line, the lists are joined: `'ffpe-trap=list1 ffpe-trap=list2'` is equivalent to `ffpe-trap=list1,list2`.

Note that once enabled an exception cannot be disabled (no negative form), except by clearing all traps by specifying `'none'`.

Many, if not most, floating point operations incur loss of precision due to rounding, and hence the `ffpe-trap=inexact` is likely to be uninteresting in practice.

By default no exception traps are enabled.

`-ffpe-summary=list`

Specify a list of floating-point exceptions, whose flag status is printed to `ERROR_UNIT` when invoking `STOP` and `ERROR STOP`. *list* can be either `'none'`, `'all'` or a comma-separated list of the following exceptions: `'invalid'`, `'zero'`, `'overflow'`, `'underflow'`, `'inexact'` and `'denormal'`. (See `-ffpe-trap` for a description of the exceptions.)

If the option is used more than once in the command line, only the last one is used.

By default, a summary for all exceptions but `'inexact'` is shown.

`-fno-backtrace`

When a serious runtime error is encountered or a deadly signal is emitted (segmentation fault, illegal instruction, bus error, floating-point exception, and the other POSIX signals that have the action `'core'`), the Fortran runtime library tries to output a backtrace of the error. `-fno-backtrace` disables the backtrace generation. This option only has influence for compilation of the Fortran main program.

See Section "Options for Debugging Your Program or GCC" in *Using the GNU Compiler Collection (GCC)*, for more information on debugging options.

2.6 Options for directory search

These options affect how GNU Fortran searches for files specified by the `INCLUDE` directive and where it searches for previously compiled modules.

It also affects the search paths used by `cpp` when used to preprocess Fortran source.

-I*dir* These affect interpretation of the `INCLUDE` directive (as well as of the `#include` directive of the `cpp` preprocessor).

Also note that the general behavior of `-I` and `INCLUDE` is pretty much the same as of `-I` with `#include` in the `cpp` preprocessor, with regard to looking for `header.gcc` files and other such things.

This path is also used to search for `.mod` files when previously compiled modules are required by a `USE` statement.

See Section “Options for Directory Search” in *Using the GNU Compiler Collection (GCC)*, for information on the `-I` option.

-J*dir* This option specifies where to put `.mod` files for compiled modules. It is also added to the list of directories to searched by an `USE` statement.

The default is the current directory.

-fintrinsic-modules-path *dir*

This option specifies the location of pre-compiled intrinsic modules, if they are not in the default location expected by the compiler.

2.7 Influencing the linking step

These options come into play when the compiler links object files into an executable output file. They are meaningless if the compiler is not doing a link step.

-static-libgfortran

On systems that provide `libgfortran` as a shared and a static library, this option forces the use of the static version. If no shared version of `libgfortran` was built when the compiler was configured, this option has no effect.

-static-libquadmath

On systems that provide `libquadmath` as a shared and a static library, this option forces the use of the static version. If no shared version of `libquadmath` was built when the compiler was configured, this option has no effect.

Please note that the `libquadmath` runtime library is licensed under the GNU Lesser General Public License (LGPL), and linking it statically introduces requirements when redistributing the resulting binaries.

2.8 Influencing runtime behavior

These options affect the runtime behavior of programs compiled with GNU Fortran.

-fconvert=*conversion*

Specify the representation of data for unformatted files. Valid values for *conversion* on most systems are: `'native'`, the default; `'swap'`, swap between big- and little-endian; `'big-endian'`, use big-endian representation for unformatted files; `'little-endian'`, use little-endian representation for unformatted files.

On POWER systems that support `-mabi=ieeelongdouble`, there are additional options, which can be combined with others with commas. Those are

`-fconvert=r16_ieee` Use IEEE 128-bit format for `REAL(KIND=16)`.

`-fconvert=r16_ibm` Use IBM long double format for `REAL(KIND=16)`.

This option has an effect only when used in the main program. The `CONVERT` specifier and the `GFORTRAN_CONVERT_UNIT` environment variable override the default specified by `-fconvert`.

`-frecord-marker=length`

Specify the length of record markers for unformatted files. Valid values for *length* are 4 and 8. Default is 4. *This is different from previous versions of gfortran*, which specified a default record marker length of 8 on most systems. If you want to read or write files compatible with earlier versions of `gfortran`, use `-frecord-marker=8`.

`-fmax-subrecord-length=length`

Specify the maximum length for a subrecord. The maximum permitted value for *length* is 2147483639, which is also the default. Only really useful for use by the `gfortran` testsuite.

`-fsign-zero`

When enabled, floating point numbers of value zero with the sign bit set are written as negative number in formatted output and treated as negative in the `SIGN` intrinsic. `-fno-sign-zero` does not print the negative sign of zero values (or values rounded to zero for I/O) and regards zero as positive number in the `SIGN` intrinsic for compatibility with Fortran 77. The default is `-fsign-zero`.

2.9 GNU Fortran Developer Options

GNU Fortran has various special options that are used for debugging the GNU Fortran compiler.

`-fdump-fortran-global`

Output a list of the global identifiers after translating into middle-end representation. Mostly useful for debugging the GNU Fortran compiler itself. The output generated by this option might change between releases. This option may also generate internal compiler errors for features that have only recently been added.

`-fdump-fortran-optimized`

Output the parse tree after front-end optimization. Mostly useful for debugging the GNU Fortran compiler itself. The output generated by this option might change between releases. This option may also generate internal compiler errors for features that have only recently been added.

`-fdump-fortran-original`

Output the internal parse tree after translating the source program into internal representation. This option is mostly useful for debugging the GNU Fortran compiler itself. The output generated by this option might change between

releases. This option may also generate internal compiler errors for features that have only recently been added.

-fdump-parse-tree

Output the internal parse tree after translating the source program into internal representation. Mostly useful for debugging the GNU Fortran compiler itself. The output generated by this option might change between releases. This option may also generate internal compiler errors for features that have only recently been added. This option is deprecated; use **-fdump-fortran-original** instead.

-save-temps

Store the usual “temporary” intermediate files permanently; name them as auxiliary output files, as specified described under GCC **-dumpbase** and **-dumpdir**.

```
gfortran -save-temps -c foo.F90
```

preprocesses input file `foo.F90` to `foo.fii`, compiles to an intermediate `foo.s`, and then assembles to the (implied) output file `foo.o`, whereas:

```
gfortran -save-temps -S foo.F
```

saves the preprocessor output in `foo.fi`, and then compiles to the (implied) output file `foo.s`.

2.10 Options for code generation conventions

These machine-independent options control the interface conventions used in code generation.

Most of them have both positive and negative forms; the negative form of **-ffoo** would be **-fno-foo**. In the table below, only one of the forms is listed—the one that is not the default. You can figure out the other form by either removing **no-** or adding it.

-fno-automatic

Treat each program unit (except those marked as **RECURSIVE**) as if the **SAVE** statement were specified for every local variable and array referenced in it. Does not affect common blocks. (Some Fortran compilers provide this option under the name **-static** or **-save**.) The default, which is **-fautomatic**, uses the stack for local variables smaller than the value given by **-fmax-stack-var-size**. Use the option **-frecursive** to use no static memory.

Local variables or arrays having an explicit **SAVE** attribute are silently ignored unless the **-pedantic** option is added.

-ff2c

Generate code designed to be compatible with code generated by **g77** and **f2c**. The calling conventions used by **g77** (originally implemented in **f2c**) require functions that return type default **REAL** to actually return the C type **double**, and functions that return type **COMPLEX** to return the values via an extra argument in the calling sequence that points to where to store the return value. Under the default GNU calling conventions, such functions simply return their results as they would in GNU C—default **REAL** functions return the C type **float**, and **COMPLEX** functions return the GNU C type **complex**. Additionally, this option implies the **-fsecond-underscore** option, unless **-fno-second-underscore** is explicitly requested.

This does not affect the generation of code that interfaces with the `libgfortran` library.

Caution: It is not a good idea to mix Fortran code compiled with `-ff2c` with code compiled with the default `-fno-f2c` calling conventions as, calling `COMPLEX` or default `REAL` functions between program parts that were compiled with different calling conventions will break at execution time.

Caution: This breaks code that passes intrinsic functions of type default `REAL` or `COMPLEX` as actual arguments, as the library implementations use the `-fno-f2c` calling conventions.

`-fno-underscoring`

Do not transform names of entities specified in the Fortran source file by appending underscores to them.

With `-funderscoring` in effect, GNU Fortran appends one underscore to external names. This is done to ensure compatibility with code produced by many UNIX Fortran compilers. Note this does not apply to names declared with `C` binding, or within a module.

Caution: The default behavior of GNU Fortran is incompatible with `f2c` and `g77`, please use the `-ff2c` option if you want object files compiled with GNU Fortran to be compatible with object code created with these tools.

Use of `-fno-underscoring` is not recommended unless you are experimenting with issues such as integration of GNU Fortran into existing system environments (vis-à-vis existing libraries, tools, and so on).

For example, with `-funderscoring`, and assuming that `j()` and `max_count()` are external functions while `my_var` and `lvar` are local variables, a Fortran statement like

```
I = J() + MAX_COUNT (MY_VAR, LVAR)
```

is implemented as something akin to the C code:

```
i = j_() + max_count_(&my_var, &lvar);
```

With `-fno-underscoring`, the same statement is implemented as:

```
i = j() + max_count(&my_var, &lvar);
```

Use of `-fno-underscoring` allows direct specification of user-defined names while debugging and when interfacing GNU Fortran code with other languages.

Note that just because the names match does *not* mean that the interface implemented by GNU Fortran for an external name matches the interface implemented by some other language for that same name. That is, getting code produced by GNU Fortran to link to code produced by some other compiler using this or any other method can be only a small part of the overall solution—getting the code generated by both compilers to agree on issues other than naming can require significant effort, and, unlike naming disagreements, linkers normally cannot detect disagreements in these other areas.

Also, note that with `-fno-underscoring`, the lack of appended underscores introduces the very real possibility that a user-defined external name conflicts with a name in a system library, which could make finding unresolved-reference

bugs quite difficult in some cases—they might occur at program run time, and show up only as buggy behavior at run time.

See Section 6.4 [Naming and argument-passing conventions], page 85, for more information. Also note that declaring symbols as `bind(C)` is a more robust way to interface with code written in other languages or compiled with different Fortran compilers than the command-line options documented in this section.

`-fsecond-underscore`

By default, GNU Fortran appends an underscore to external names. If this option is used, GNU Fortran appends two underscores to names with underscores and one underscore to names with no underscores.

For example, an external name such as `MAX_COUNT` is implemented as a reference to the link-time external symbol `max_count__`, instead of `max_count_`. This is required for compatibility with `g77` and `f2c`, and is implied by use of the `-ff2c` option.

This option has no effect if `-fno-underscoring` is in effect. It is implied by the `-ff2c` option.

`-fcoarray=<keyword>`

- `'none'` Disable coarray support; using coarray declarations and image-control statements produces a compile-time error. (Default)
- `'single'` Single-image mode, i.e. `num_images()` is always one.
- `'lib'` Library-based coarray parallelization; a suitable GNU Fortran coarray library such as <http://opencoarrays.org> needs to be linked. Alternatively, GCC's `libcaf_single` library can be linked, albeit it only supports a single image.

`-fcheck=<keyword>`

Enable the generation of run-time checks; the argument shall be a comma-delimited list of the following keywords. Prefixing a check with `no-` disables it if it was activated by a previous specification.

- `'all'` Enable all run-time test of `-fcheck`.
- `'array-temps'` Warns at run time when for passing an actual argument a temporary array had to be generated. The information generated by this warning is sometimes useful in optimization, in order to avoid such temporaries.
Note: The warning is only printed once per location.
- `'bits'` Enable generation of run-time checks for invalid arguments to the bit manipulation intrinsics.
- `'bounds'` Enable generation of run-time checks for array subscripts and against the declared minimum and maximum values. It also checks array indices for assumed and deferred shape arrays against the actual allocated bounds and ensures that all string lengths

are equal for character array constructors without an explicit typespec.

Some checks require that `-fcheck=bounds` is set for the compilation of the main program.

Note: In the future this may also include other forms of checking, e.g., checking substring references.

- ‘do’ Enable generation of run-time checks for invalid modification of loop iteration variables.
- ‘mem’ Enable generation of run-time checks for memory allocation. Note: This option does not affect explicit allocations using the `ALLOCATE` statement, which are always checked.
- ‘pointer’ Enable generation of run-time checks for pointers and allocatables.
- ‘recursion’
 Enable generation of run-time checks for recursively called subroutines and functions that are not marked as recursive. See also `-frecursive`. Note: This check does not work for OpenMP programs and is disabled if used together with `-frecursive` and `-fopenmp`.

Example: Assuming you have a file `foo.f90`, the command

```
gfortran -fcheck=all,no-array-temps foo.f90
```

compiles the file with all checks enabled as specified above except warnings for generated array temporaries.

`-fbounds-check`

Deprecated alias for `-fcheck=bounds`.

`-ftail-call-workaround`

`-ftail-call-workaround=n`

Some C interfaces to Fortran codes violate the gfortran ABI by omitting the hidden character length arguments as described in See Section 6.4.2 [Argument passing conventions], page 85. This can lead to crashes because pushing arguments for tail calls can overflow the stack.

To provide a workaround for existing binary packages, this option disables tail call optimization for gfortran procedures with character arguments. With `-ftail-call-workaround=2` tail call optimization is disabled in all gfortran procedures with character arguments, with `-ftail-call-workaround=1` or equivalent `-ftail-call-workaround` only in gfortran procedures with character arguments that call implicitly prototyped procedures.

Using this option can lead to problems including crashes due to insufficient stack space.

It is *very strongly* recommended to fix the code in question. The `-fc-prototypes-external` option can be used to generate prototypes that conform to gfortran’s ABI, for inclusion in the source code.

Support for this option will likely be withdrawn in a future release of gfortran.

The negative form, `-fno-tail-call-workaround` or equivalent `-ftail-call-workaround=0`, can be used to disable this option.

Default is currently `-ftail-call-workaround`, this will change in future releases.

`-fcheck-array-temporaries`

Deprecated alias for `-fcheck=array-temps`.

`-fmax-array-structor=n`

This option can be used to increase the upper limit permitted in array constructors. The code below requires this option to expand the array at compile time.

```
program test
implicit none
integer j
integer, parameter :: n = 100000
integer, parameter :: i(n) = (/ (2*j, j = 1, n) /)
print '(10(I0,1X))', i
end program test
```

Caution: This option can lead to long compile times and excessively large object files.

The default value for *n* is 65535.

`-fmax-stack-var-size=n`

This option specifies the size in bytes of the largest array that is put on the stack; if the size is exceeded static memory is used (except in procedures marked as `RECURSIVE`). Use the option `-frecursive` to allow for recursive procedures that do not have a `RECURSIVE` attribute or for parallel programs. Use `-fno-automatic` to never use the stack.

This option currently only affects local arrays declared with constant bounds, and may not apply to all character variables. Future versions of GNU Fortran may improve this behavior.

The default value for *n* is 65536.

`-fstack-arrays`

Adding this option makes the Fortran compiler put all arrays of unknown size and array temporaries onto stack memory. If your program uses very large local arrays it is possible that you have to extend your runtime limits for stack memory on some operating systems. This flag is enabled by default at optimization level `-Ofast` unless `-fmax-stack-var-size` is specified.

`-fpack-derived`

This option tells GNU Fortran to pack derived type members as closely as possible. Code compiled with this option is likely to be incompatible with code compiled without this option, and may execute slower.

`-frepack-arrays`

In some circumstances GNU Fortran may pass assumed shape array sections via a descriptor describing a noncontiguous area of memory. This option adds code to the function prologue to repack the data into a contiguous block at runtime.

This should result in faster accesses to the array. However it can introduce significant overhead to the function call, especially when the passed data is noncontiguous.

-fshort-enums

This option is provided for interoperability with C code that was compiled with the `-fshort-enums` option. It makes GNU Fortran choose the smallest `INTEGER` kind a given enumerator set fits in, and give all its enumerators this kind.

-finline-arg-packing

When passing an assumed-shape argument of a procedure as actual argument to an assumed-size or explicit size or as argument to a procedure that does not have an explicit interface, the argument may have to be packed; that is, put into contiguous memory. An example is the call to `foo` in

```
subroutine foo(a)
  real, dimension(*) :: a
end subroutine foo
subroutine bar(b)
  real, dimension(:) :: b
  call foo(b)
end subroutine bar
```

When `-finline-arg-packing` is in effect, this packing is performed by inline code. This allows for more optimization while increasing code size.

`-finline-arg-packing` is implied by any of the `-O` options except when optimizing for size via `-Os`. If the code contains a very large number of argument that have to be packed, code size and also compilation time may become excessive. If that is the case, it may be better to disable this option. Instances of packing can be found by using `-Warray-temporaries`.

-fexternal-blas

This option makes `gfortran` generate calls to BLAS functions for some matrix operations like `MATMUL`, instead of using our own algorithms, if the size of the matrices involved is larger than a given limit (see `-fblas-matmul-limit`). This may be profitable if an optimized vendor BLAS library is available. The BLAS library has to be specified at link time. This option specifies a BLAS library with integer arguments of default kind (32 bits). It cannot be used together with `-fexternal-blas64`.

-fexternal-blas64

makes `gfortran` generate calls to BLAS functions for some matrix operations like `MATMUL`, instead of using our own algorithms, if the size of the matrices involved is larger than a given limit (see `-fblas-matmul-limit`). This may be profitable if an optimized vendor BLAS library is available. The BLAS library has to be specified at link time. This option specifies a BLAS library with integer arguments of `KIND=8` (64 bits). It cannot be used together with `-fexternal-blas`, and requires a 64-bit system. This option also requires `-ffrontend-optimize`.

-fblas-matmul-limit=n

Only significant when `-fexternal-blas` or `-fexternal-blas64` are in effect. Matrix multiplication of matrices with size larger than or equal to `n` is performed

by calls to BLAS functions, while others are handled by `gfortran` internal algorithms. If the matrices involved are not square, the size comparison is performed using the geometric mean of the dimensions of the argument and result matrices.

The default value for n is 30.

`-finline-intrinsics`

`-finline-intrinsics=intr1,intr2,...`

Prefer generating inline code over calls to libgfortran functions to implement intrinsics.

Usage of intrinsics can be implemented either by generating a call to the libgfortran library function or by directly generating inline code. For most intrinsics, only a single variant is available, and there is no choice of implementation. However, some intrinsics can use a library function or inline code, where inline code typically offers opportunities for additional optimization over a library function. With `-finline-intrinsics=...` or `-fno-inline-intrinsics=...`, the choice applies only to the intrinsics present in the comma-separated list provided as argument.

For each intrinsic, if no choice of implementation was made through either of the flag variants, a default behavior is chosen depending on optimization: library calls are generated when not optimizing or when optimizing for size; otherwise inline code is preferred.

The set of intrinsics allowed as argument to `-finline-intrinsics=` is currently limited to `MAXLOC` and `MINLOC`. The effect of the flag is moreover limited to calls of those intrinsics without `DIM` argument and with `ARRAY` of a non-`CHARACTER` type. The case of rank-1 argument and `DIM` argument present, i.e. `MAXLOC(A(:),DIM=1)` or `MINLOC(A(:),DIM=1)` is inlined unconditionally for numeric rank-1 array argument `A`.

`-finline-matmul-limit=n`

When front-end optimization is active, some calls to the `MATMUL` intrinsic function are inlined. This may result in code size increase if the size of the matrix cannot be determined at compile time, as code for both cases is generated. Setting `-finline-matmul-limit=0` disables inlining in all cases. Setting this option with a value of n produces inline code for matrices with size up to n . If the matrices involved are not square, the size comparison is performed using the geometric mean of the dimensions of the argument and result matrices.

The default value for n is 30. The `-fblas-matmul-limit` can be used to change this value.

`-frecursive`

Allow indirect recursion by forcing all local arrays to be allocated on the stack. This flag cannot be used together with `-fmax-stack-var-size=` or `-fno-automatic`.

```

-finit-local-zero
-finit-derived
-finit-integer=n
-finit-real=<zero|inf|-inf|nan|snan>
-finit-logical=<true|false>
-finit-character=n

```

The `-finit-local-zero` option instructs the compiler to initialize local `INTEGER`, `REAL`, and `COMPLEX` variables to zero, `LOGICAL` variables to false, and `CHARACTER` variables to a string of null bytes. Finer-grained initialization options are provided by the `-finit-integer=n`, `-finit-real=<zero|inf|-inf|nan|snan>` (which also initializes the real and imaginary parts of local `COMPLEX` variables), `-finit-logical=<true|false>`, and `-finit-character=n` (where *n* is an ASCII character value) options.

With `-finit-derived`, components of derived type variables are initialized according to these flags. Components whose type is not covered by an explicit `-finit-*` flag are treated as described above with `-finit-local-zero`.

These options do not initialize

- objects with the `POINTER` attribute
- allocatable arrays
- variables that appear in an `EQUIVALENCE` statement.

(These limitations may be removed in future releases).

Note that the `-finit-real=nan` option initializes `REAL` and `COMPLEX` variables with a quiet NaN. For a signalling NaN use `-finit-real=snan`; note, however, that compile-time optimizations may convert them into quiet NaN and that trapping needs to be enabled (e.g. via `-ffpe-trap`).

The `-finit-integer` option parses the value into an integer of type `INTEGER(kind=C_LONG)` on the host. Said value is then assigned to the integer variables in the Fortran code, which might result in wraparound if the value is too large for the kind.

Finally, note that enabling any of the `-finit-*` options silences warnings that would have been emitted by `-Wuninitialized` for the affected local variables.

```

-falign-commons

```

By default, `gfortran` enforces proper alignment of all variables in a `COMMON` block by padding them as needed. On certain platforms this is mandatory, on others it increases performance. If a `COMMON` block is not declared with consistent data types everywhere, this padding can cause trouble, and `-fno-align-commons` can be used to disable automatic alignment. The same form of this option should be used for all files that share a `COMMON` block. To avoid potential alignment issues in `COMMON` blocks, it is recommended to order objects from largest to smallest.

```

-fno-protect-parens

```

By default the parentheses in expression are honored for all optimization levels such that the compiler does not do any reassociation. Using `-fno-protect-parens` allows the compiler to reorder `REAL` and `COMPLEX` expressions to produce

faster code. Note that for the reassociation optimization `-fno-signed-zeros` and `-fno-trapping-math` need to be in effect. The parentheses protection is enabled by default, unless `-Ofast` is given.

`-frealloc-lhs`

An allocatable left-hand side of an intrinsic assignment is automatically (re)allocated if it is either unallocated or has a different shape. The option is enabled by default except when `-std=f95` is given. See also `-wrealloc-lhs`.

`-faggressive-function-elimination`

Functions with identical argument lists are eliminated within statements, regardless of whether these functions are marked `PURE` or not. For example, in

```
a = f(b,c) + f(b,c)
```

there is only a single call to `f`. This option only works if `-ffrontend-optimize` is in effect.

`-ffrontend-optimize`

This option performs front-end optimization, based on manipulating parts of the Fortran parse tree. Enabled by default by any `-O` option except `-O0` and `-Og`. Optimizations enabled by this option include:

- inlining calls to `MATMUL`,
- elimination of identical function calls within expressions,
- removing unnecessary calls to `TRIM` in comparisons and assignments,
- replacing `TRIM(a)` with `a(1:LEN_TRIM(a))` and
- short-circuiting of logical operators (`.AND.` and `.OR.`).

It can be deselected by specifying `-fno-frontend-optimize`.

`-ffrontend-loop-interchange`

Attempt to interchange loops in the Fortran front end where profitable. Enabled by default by any `-O` option. At the moment, this option only affects `FORALL` and `DO CONCURRENT` statements with several forall triplets.

See Section “Options for Code Generation Conventions” in *Using the GNU Compiler Collection (GCC)*, for information on more options offered by the back end shared by `gfortran`, `gcc`, and other GNU compilers.

2.11 Options for interoperability with other languages

`-fc-prototypes`

This option generates C prototypes from `BIND(C)` variable declarations, types and procedure interfaces and writes them to standard output. `ENUM` is not yet supported.

The generated prototypes may need inclusion of an appropriate header, such as `<stdint.h>` or `<stdlib.h>`. For types that are not specified using the appropriate kind from the `iso_c_binding` module, a warning is added as a comment to the code.

For function pointers, a pointer to a function returning `int` without an explicit argument list is generated.

Example of use:

```
$ gfortran -fc-prototypes -fsyntax-only foo.f90 > foo.h
```

where the C code intended for interoperating with the Fortran code then uses `#include "foo.h"`.

-fc-prototypes-external

This option generates C prototypes from external functions and subroutines and writes them to standard output. This may be useful for making sure that C bindings to Fortran code are correct. This option does not generate prototypes for `BIND(C)` procedures; use `-fc-prototypes` for that.

The generated prototypes may need inclusion of an appropriate header, such as `<stdint.h>` or `<stdlib.h>`.

This is primarily meant for legacy code to ensure that existing C bindings match what `gfortran` emits. The generated C prototypes should be correct for the current version of the compiler, but may not match what other compilers or earlier versions of `gfortran` need. For new development, use of the `BIND(C)` features is recommended.

Example of use:

```
$ gfortran -fc-prototypes-external -fsyntax-only foo.f > foo.h
```

where the C code intended for interoperating with the Fortran code then uses `#include "foo.h"`.

2.12 Environment variables affecting gfortran

The `gfortran` compiler currently does not make use of any environment variables to control its operation above and beyond those that affect the operation of `gcc`.

See Section “Environment Variables Affecting GCC” in *Using the GNU Compiler Collection (GCC)*, for information on environment variables.

See Chapter 3 [Runtime], page 37, for environment variables that affect the run-time behavior of programs compiled with GNU Fortran.

3 Runtime: Influencing runtime behavior with environment variables

The behavior of the `gfortran` can be influenced by environment variables.

Malformed environment variables are silently ignored.

3.1 TMPDIR—Directory for scratch files

When opening a file with `STATUS='SCRATCH'`, GNU Fortran tries to create the file in one of the potential directories by testing each directory in the order below.

1. The environment variable `TMPDIR`, if it exists.
2. On the MinGW target, the directory returned by the `GetTempPath` function. Alternatively, on the Cygwin target, the `TMP` and `TEMP` environment variables, if they exist, in that order.
3. The `P_tmpdir` macro if it is defined, otherwise the directory `/tmp`.

3.2 GFORTRAN_STDIN_UNIT—Unit number for standard input

This environment variable can be used to select the unit number preconnected to standard input. This must be a positive integer. The default value is 5.

3.3 GFORTRAN_STDOUT_UNIT—Unit number for standard output

This environment variable can be used to select the unit number preconnected to standard output. This must be a positive integer. The default value is 6.

3.4 GFORTRAN_STDERR_UNIT—Unit number for standard error

This environment variable can be used to select the unit number preconnected to standard error. This must be a positive integer. The default value is 0.

3.5 GFORTRAN_UNBUFFERED_ALL—Do not buffer I/O on all units

This environment variable controls whether all I/O is unbuffered. If the first letter is 'y', 'Y' or '1', all I/O is unbuffered. This slows down small sequential reads and writes. If the first letter is 'n', 'N' or '0', I/O is buffered. This is the default.

3.6 GFORTRAN_UNBUFFERED_PRECONNECTED—Do not buffer I/O on preconnected units

The environment variable named `GFORTRAN_UNBUFFERED_PRECONNECTED` controls whether I/O on a preconnected unit (i.e. `STDOUT` or `STDERR`) is unbuffered. If the first letter is 'y', 'Y' or '1', I/O is unbuffered. This slows down small sequential reads and writes. If the first letter is 'n', 'N' or '0', I/O is buffered. This is the default.

3.7 GFORTRAN_SHOW_LOCUS—Show location for runtime errors

If the first letter is 'y', 'Y' or '1', filename and line numbers for runtime errors are printed. If the first letter is 'n', 'N' or '0', do not print filename and line numbers for runtime errors. The default is to print the location.

3.8 GFORTRAN_OPTIONAL_PLUS—Print leading + where permitted

If the first letter is ‘y’, ‘Y’ or ‘1’, a plus sign is printed where permitted by the Fortran standard. If the first letter is ‘n’, ‘N’ or ‘O’, a plus sign is not printed in most cases. Default is not to print plus signs.

3.9 GFORTRAN_LIST_SEPARATOR—Separator for list output

This environment variable specifies the separator when writing list-directed output. It may contain any number of spaces and at most one comma. If you specify this on the command line, be sure to quote spaces, as in

```
$ GFORTRAN_LIST_SEPARATOR=' , ' ./a.out
```

when `a.out` is the compiled Fortran program that you want to run. Default is a single space.

3.10 GFORTRAN_CONVERT_UNIT—Set conversion for unformatted I/O

By setting the `GFORTRAN_CONVERT_UNIT` variable, it is possible to change the representation of data for unformatted files. The syntax for the `GFORTRAN_CONVERT_UNIT` variable for most systems is:

```
GFORTRAN_CONVERT_UNIT: mode | mode ';' exception | exception ;
mode: 'native' | 'swap' | 'big_endian' | 'little_endian' ;
exception: mode ':' unit_list | unit_list ;
unit_list: unit_spec | unit_list unit_spec ;
unit_spec: INTEGER | INTEGER '-' INTEGER ;
```

The variable consists of an optional default mode, followed by a list of optional exceptions, which are separated by semicolons from the preceding default and each other. Each exception consists of a format and a comma-separated list of units. Valid values for the modes are the same as for the `CONVERT` specifier:

NATIVE Use the native format. This is the default.

SWAP Swap between little- and big-endian.

LITTLE_ENDIAN Use the little-endian format for unformatted files.

BIG_ENDIAN Use the big-endian format for unformatted files.

For POWER systems that support `-mabi=ieeelongdouble`, there are additional options, which can be combined with the others with commas. Those are

R16_IEEE Use IEEE 128-bit format for `REAL(KIND=16)`.

R16_IBM Use IBM long double format for `REAL(KIND=16)`.

A missing mode for an exception is taken to mean **BIG_ENDIAN**. Examples of values for `GFORTRAN_CONVERT_UNIT` are:

`'big_endian'` Do all unformatted I/O in big-endian mode.

`'little_endian;native:10-20,25'` Do all unformatted I/O in little-endian mode, except for units 10 to 20 and 25, which are in native format.

`'10-20'` Units 10 to 20 are big-endian, the rest is native.

`'big_endian,r16_ibm'` Do all unformatted I/O in big-endian mode and use IBM long double for output of `REAL(KIND=16)` values.

Setting the environment variables should be done on the command line or via the `export` command for `sh`-compatible shells and via `setenv` for `csh`-compatible shells.

Example for `sh`:

```
$ gfortran foo.f90
$ GFORTRAN_CONVERT_UNIT='big_endian;native:10-20' ./a.out
```

Example code for `csh`:

```
% gfortran foo.f90
% setenv GFORTRAN_CONVERT_UNIT 'big_endian;native:10-20'
% ./a.out
```

Using anything but the native representation for unformatted data carries a significant speed overhead. If speed in this area matters to you, it is best if you use this only for data that needs to be portable.

See Section 5.1.17 [CONVERT specifier], page 55, for an alternative way to specify the data representation for unformatted files. See Section 2.8 [Runtime Options], page 24, for setting a default data representation for the whole program. The `CONVERT` specifier overrides the `-fconvert` compile options.

Note that the values specified via the `GFORTRAN_CONVERT_UNIT` environment variable override the `CONVERT` specifier in the `OPEN` statement. This is to give control over data formats to users who do not have the source code of their program available.

3.11 GFORTRAN_ERROR_BACKTRACE—Show backtrace on run-time errors

If the `GFORTRAN_ERROR_BACKTRACE` variable is set to 'y', 'Y' or '1' (only the first letter is relevant) then a backtrace is printed when a serious run-time error occurs. To disable the backtracing, set the variable to 'n', 'N', '0'. Default is to print a backtrace unless the `-fno-backtrace` compile option was used.

3.12 GFORTRAN_FORMATTED_BUFFER_SIZE—Set buffer size for formatted I/O

The `GFORTRAN_FORMATTED_BUFFER_SIZE` environment variable specifies buffer size in bytes to be used for formatted output. The default value is 8192.

3.13 GFORTRAN_UNFORMATTED_BUFFER_SIZE—Set buffer size for unformatted I/O

The `GFORTRAN_UNFORMATTED_BUFFER_SIZE` environment variable specifies buffer size in bytes to be used for unformatted output. The default value is 131072.

Part II: Language Reference

4 Compiler Characteristics

This chapter describes certain characteristics of the GNU Fortran compiler that are not specified by the Fortran standard, but which might in some way or another become visible to the programmer.

4.1 KIND Type Parameters

The KIND type parameters supported by GNU Fortran for the primitive data types are:

INTEGER 1, 2, 4, 8*, 16*, default: 4**

LOGICAL 1, 2, 4, 8*, 16*, default: 4**

REAL 4, 8, 10*, 16*, default: 4***

COMPLEX 4, 8, 10*, 16*, default: 4***

DOUBLE PRECISION
4, 8, 10*, 16*, default: 8***

CHARACTER
1, 4, default: 1

* not available on all systems

** unless `-fdefault-integer-8` is used

*** unless `-fdefault-real-8` is used (see Section 2.2 [Fortran Dialect Options], page 9)

The KIND value matches the storage size in bytes, except for **COMPLEX** where the storage size is twice as much (or both real and imaginary part are a real value of the given size). It is recommended to use the Section 8.246 [SELECTED_CHAR_KIND], page 278, Section 8.247 [SELECTED_INT_KIND], page 279, Section 8.248 [SELECTED_LOGICAL_KIND], page 279, and Section 8.249 [SELECTED_REAL_KIND], page 280, intrinsics or the INT8, INT16, INT32, INT64, REAL32, REAL64, and REAL128 parameters of the ISO_FORTRAN_ENV module instead of the concrete values. The available kind parameters can be found in the constant arrays CHARACTER_KINDS, INTEGER_KINDS, LOGICAL_KINDS and REAL_KINDS in the Section 9.1 [ISO_FORTRAN_ENV], page 311, module. For C interoperability, the kind parameters of the Section 9.2 [ISO_C_BINDING], page 313, module should be used.

4.2 Internal representation of LOGICAL variables

The Fortran standard does not specify how variables of **LOGICAL** type are represented, beyond requiring that **LOGICAL** variables of default kind have the same storage size as default **INTEGER** and **REAL** variables. The GNU Fortran internal representation is as follows.

A **LOGICAL**(KIND=N) variable is represented as an **INTEGER**(KIND=N) variable, however, with only two permissible values: 1 for **.TRUE.** and 0 for **.FALSE.** Any other integer value results in undefined behavior.

See also Section 6.4.2 [Argument passing conventions], page 85, and Section 6.1 [Interoperability with C], page 73.

4.3 Evaluation of logical expressions

The Fortran standard does not require the compiler to evaluate all parts of an expression, if they do not contribute to the final result. For logical expressions with `.AND.` or `.OR.` operators, in particular, GNU Fortran optimizes out function calls (even to impure functions) if the result of the expression can be established without them. However, since not all compilers do that, and such an optimization can potentially modify the program flow and subsequent results, GNU Fortran throws warnings for such situations with the `-Wfunction-elimination` flag.

4.4 MAX and MIN intrinsics with REAL NaN arguments

The Fortran standard does not specify what the result of the `MAX` and `MIN` intrinsics are if one of the arguments is a `NaN`. Accordingly, the GNU Fortran compiler does not specify that either, as this allows for faster and more compact code to be generated. If the programmer wishes to take some specific action in case one of the arguments is a `NaN`, it is necessary to explicitly test the arguments before calling `MAX` or `MIN`, e.g. with the `IEEE_IS_NAN` function from the intrinsic module `IEEE_ARITHMETIC`.

4.5 Thread-safety of the runtime library

GNU Fortran can be used in programs with multiple threads, e.g. by using OpenMP, by calling OS thread handling functions via the `ISO_C_BINDING` facility, or by GNU Fortran compiled library code being called from a multi-threaded program.

The GNU Fortran runtime library, (`libgfortran`), supports being called concurrently from multiple threads with the following exceptions.

During library initialization, the C `getenv` function is used, which need not be thread-safe. Similarly, the `getenv` function is used to implement the `GET_ENVIRONMENT_VARIABLE` and `GETENV` intrinsics. It is the responsibility of the user to ensure that the environment is not being updated concurrently when any of these actions are taking place.

The `EXECUTE_COMMAND_LINE` and `SYSTEM` intrinsics are implemented with the `system` function, which need not be thread-safe. It is the responsibility of the user to ensure that `system` is not called concurrently.

For platforms not supporting thread-safe POSIX functions, further functionality might not be thread-safe. For details, please consult the documentation for your operating system.

The GNU Fortran runtime library uses various C library functions that depend on the locale, such as `strtod` and `snprintf`. In order to work correctly in locale-aware programs that set the locale using `setlocale`, the locale is reset to the default “C” locale while executing a formatted `READ` or `WRITE` statement. On targets supporting the POSIX 2008 per-thread locale functions (e.g. `newlocale`, `uselocale`, `freelocale`), these are used and thus the global locale set using `setlocale` or the per-thread locales in other threads are not affected. However, on targets lacking this functionality, the global `LC_NUMERIC` locale is set to “C” during the formatted I/O. Thus, on such targets it’s not safe to call `setlocale` concurrently from another thread while a Fortran formatted I/O operation is in progress. Also, other threads doing something dependent on the `LC_NUMERIC` locale might not work correctly if a formatted I/O operation is in progress in another thread.

4.6 Data consistency and durability

This section contains a brief overview of data and metadata consistency and durability issues when doing I/O.

With respect to durability, GNU Fortran makes no effort to ensure that data is committed to stable storage. If this is required, the GNU Fortran programmer can use the intrinsic `FNUM` to retrieve the low level file descriptor corresponding to an open Fortran unit. Then, using e.g. the `ISO_C_BINDING` feature, one can call the underlying system call to flush dirty data to stable storage, such as `fsync` on POSIX, `_commit` on MinGW, or `fcntl(fd, F_FULLSYNC, 0)` on macOS. The following example shows how to call `fsync`:

```
! Declare the interface for POSIX fsync function
interface
  function fsync (fd) bind(c,name="fsync")
    use iso_c_binding, only: c_int
    integer(c_int), value :: fd
    integer(c_int) :: fsync
  end function fsync
end interface

! Variable declaration
integer :: ret

! Opening unit 10
open (10,file="foo")

! ...
! Perform I/O on unit 10
! ...

! Flush and sync
flush(10)
ret = fsync(fnum(10))

! Handle possible error
if (ret /= 0) stop "Error calling FSYNC"
```

With respect to consistency, for regular files GNU Fortran uses buffered I/O in order to improve performance. This buffer is flushed automatically when full and in some other situations, e.g. when closing a unit. It can also be explicitly flushed with the `FLUSH` statement. Also, the buffering can be turned off with the `GFORTRAN_UNBUFFERED_ALL` and `GFORTRAN_UNBUFFERED_PRECONNECTED` environment variables. Special files, such as terminals and pipes, are always unbuffered. Sometimes, however, further things may need to be done in order to allow other processes to see data that GNU Fortran has written, as follows.

The Windows platform supports a relaxed metadata consistency model, where file metadata is written to the directory lazily. This means that, for instance, the `dir` command can show a stale size for a file. One can force a directory metadata update by closing the unit, or by calling `_commit` on the file descriptor. Note, though, that `_commit` forces all dirty data to stable storage, which is often a very slow operation.

The Network File System (NFS) implements a relaxed consistency model called open-to-close consistency. Closing a file forces dirty data and metadata to be flushed to the server, and opening a file forces the client to contact the server in order to revalidate cached data. `fsync` also forces a flush of dirty data and metadata to the server. Similar to `open`

and `close`, acquiring and releasing `fcntl` file locks, if the server supports them, also forces cache validation and flushing dirty data and metadata.

4.7 Files opened without an explicit `ACTION=` specifier

The Fortran standard says that if an `OPEN` statement is executed without an explicit `ACTION=` specifier, the default value is processor dependent. GNU Fortran behaves as follows:

1. Attempt to open the file with `ACTION='READWRITE'`
2. If that fails, try to open with `ACTION='READ'`
3. If that fails, try to open with `ACTION='WRITE'`
4. If that fails, generate an error

4.8 File operations on symbolic links

This section documents the behavior of GNU Fortran for file operations on symbolic links, on systems that support them.

- Results of `INQUIRE` statements of the “inquire by file” form relate to the target of the symbolic link. For example, `INQUIRE(FILE="foo",EXIST=ex)` sets `ex` to *.true.* if *foo* is a symbolic link pointing to an existing file, and *.false.* if *foo* points to a non-existing file (“dangling” symbolic link).
- Using the `OPEN` statement with a `STATUS="NEW"` specifier on a symbolic link results in an error condition, whether the symbolic link points to an existing target or is dangling.
- If a symbolic link was connected, using the `CLOSE` statement with a `STATUS="DELETE"` specifier causes the symbolic link itself to be deleted, not its target.

4.9 File format of unformatted sequential files

Unformatted sequential files are stored as logical records using record markers. Each logical record consists of one or more subrecords.

Each subrecord consists of a leading record marker, the data written by the user program, and a trailing record marker. The record markers are four-byte integers by default, and eight-byte integers if the `-fmax-subrecord-length=8` option (which exists for backwards compatibility only) is in effect.

The representation of the record markers is that of unformatted files given with the `-fconvert` option, the Section 5.1.17 [`CONVERT` specifier], page 55, in an `OPEN` statement or the Section 3.10 [`GFORTRAN_CONVERT_UNIT`], page 38, environment variable.

The maximum number of bytes of user data in a subrecord is 2147483639 (2 GiB - 9) for a four-byte record marker. This limit can be lowered with the `-fmax-subrecord-length` option, although this is rarely useful. If the length of a logical record exceeds this limit, the data is distributed among several subrecords.

The absolute of the number stored in the record markers is the number of bytes of user data in the corresponding subrecord. If the leading record marker of a subrecord contains a negative number, another subrecord follows the current one. If the trailing record marker contains a negative number, then there is a preceding subrecord.

In the most simple case, with only one subrecord per logical record, both record markers contain the number of bytes of user data in the record.

The format for unformatted sequential data can be duplicated using unformatted stream, as shown in the example program for an unformatted record containing a single subrecord:

```

program main
  use iso_fortran_env, only: int32
  implicit none
  integer(int32) :: i
  real, dimension(10) :: a, b
  call random_number(a)
  open (10,file='test.dat',form='unformatted',access='stream')
  inquire (iolength=i) a
  write (10) i, a, i
  close (10)
  open (10,file='test.dat',form='unformatted')
  read (10) b
  if (all (a == b)) print *, 'success!'
end program main

```

4.10 Asynchronous I/O

Asynchronous I/O is supported if the program is linked against the POSIX thread library. If that is not the case, all I/O is performed as synchronous. On systems that do not support pthread condition variables, such as AIX, I/O is also performed as synchronous.

On some systems, such as Darwin or Solaris, the POSIX thread library is always linked in, so asynchronous I/O is always performed. On other systems, such as Linux, it is necessary to specify `-pthread`, `-lpthread` or `-fopenmp` during the linking step.

4.11 Behavior on integer overflow

Integer overflow is prohibited by the Fortran standard. The behavior of gfortran on integer overflow is undefined by default. Traditional code, like linear congruential pseudo-random number generators in old programs that rely on specific, nonstandard behavior may generate unexpected results. The `-fsanitize=undefined` option can be used to detect such code at runtime.

It is recommended to use the intrinsic subroutine `RANDOM_NUMBER` for random number generators or, if the old behavior is desired, to use the `-fwrapv` option. Note that this option can impact performance.

5 Extensions

The two sections below detail the extensions to standard Fortran that are implemented in GNU Fortran, as well as some of the popular or historically important extensions that are not (or not yet) implemented. For the latter case, we explain the alternatives available to GNU Fortran users, including replacement by standard-conforming code or GNU extensions.

5.1 Extensions implemented in GNU Fortran

GNU Fortran implements a number of extensions over standard Fortran. This chapter contains information on their syntax and meaning. There are currently two categories of GNU Fortran extensions, those that provide functionality beyond that provided by any standard, and those that are supported by GNU Fortran purely for backward compatibility with legacy compilers. By default, `-std=gnu` allows the compiler to accept both types of extensions, but to warn about the use of the latter. Specifying either `-std=f95`, `-std=f2003`, `-std=f2008`, or `-std=f2018` disables both types of extensions, and `-std=legacy` allows both without warning. The special compile flag `-fdec` enables additional compatibility extensions along with those enabled by `-std=legacy`.

5.1.1 Old-style kind specifications

GNU Fortran allows old-style kind specifications in declarations. These look like:

```
TYPESPEC*size x,y,z
```

where `TYPESPEC` is a basic type (`INTEGER`, `REAL`, etc.), and where *size* is a byte count corresponding to the storage size of a valid kind for that type. (For `COMPLEX` variables, *size* is the total size of the real and imaginary parts.) The statement then declares `x`, `y` and `z` to be of type `TYPESPEC` with the appropriate kind. This is equivalent to the standard-conforming declaration

```
TYPESPEC(k) x,y,z
```

where `k` is the kind parameter suitable for the intended precision. As kind parameters are implementation-dependent, use the `KIND`, `SELECTED_INT_KIND`, `SELECTED_LOGICAL_KIND` and `SELECTED_REAL_KIND` intrinsics to retrieve the correct value, for instance `REAL*8 x` can be replaced by:

```
INTEGER, PARAMETER :: dbl = KIND(1.0d0)
REAL(KIND=dbl) :: x
```

5.1.2 Old-style variable initialization

GNU Fortran allows old-style initialization of variables of the form:

```
INTEGER i/1/,j/2/
REAL x(2,2) /3*0.,1./
```

The syntax for the initializers is as for the `DATA` statement, but unlike in a `DATA` statement, an initializer only applies to the variable immediately preceding the initialization. In other words, something like `INTEGER I,J/2,3/` is not valid. This style of initialization is only allowed in declarations without double colons (`::`); the double colons were introduced in Fortran 90, which also introduced a standard syntax for initializing variables in type declarations.

Examples of standard-conforming code equivalent to the above example are:

```
! Fortran 90
```

```

      INTEGER :: i = 1, j = 2
      REAL :: x(2,2) = RESHAPE((/0.,0.,0.,1./),SHAPE(x))
! Fortran 77
      INTEGER i, j
      REAL x(2,2)
      DATA i/1/, j/2/, x/3*0.,1./

```

Note that variables that are explicitly initialized in declarations or in DATA statements automatically acquire the SAVE attribute.

5.1.3 Extensions to namelist

GNU Fortran fully supports the Fortran 95 standard for namelist I/O including array qualifiers, substrings and fully qualified derived types. The output from a namelist write is compatible with namelist read. The output has all names in upper case and indentation to column 1 after the namelist name. The following extensions are permitted:

- Old-style use of '\$' instead of '&'

```

$MYNML
  X(:)%Y(2) = 1.0 2.0 3.0
  CH(1:4) = "abcd"
$END

```

It should be noted that the default terminator is '/' rather than '&END'.

- Querying of the namelist when inputting from stdin. After at least one space, entering '?' sends to stdout the namelist name and the names of the variables in the namelist:

```

?

&mynml
  x
  x%y
  ch
&end

```

Entering '=?' outputs the namelist to stdout, as if WRITE(*,NML = mynml) had been called:

```

=?

&MYNML
  X(1)%Y=  0.000000      ,  1.000000      ,  0.000000      ,
  X(2)%Y=  0.000000      ,  2.000000      ,  0.000000      ,
  X(3)%Y=  0.000000      ,  3.000000      ,  0.000000      ,
  CH=abcd,  /

```

To aid this dialog, when input is from stdin, errors send their messages to stderr and execution continues, even if IOSTAT is set.

- PRINT namelist is permitted. This causes an error if -std=f95 is used.

```

PROGRAM test_print
  REAL, dimension (4)  :: x = (/1.0, 2.0, 3.0, 4.0/)
  NAMELIST /mynml/ x
  PRINT mynml
END PROGRAM test_print

```

- Expanded namelist reads are permitted. This causes an error if -std=f95 is used. In the following example, the first element of the array is given the value 0.00 and the two succeeding elements are given the values 1.00 and 2.00.

```

&MYNML

```

```

      X(1,1) = 0.00 , 1.00 , 2.00
/

```

When writing a namelist, if no `DELIM=` is specified, by default a double quote is used to delimit character strings. With `-std=f95` or later, the `delim` status is set to `'none'`. Defaulting to quotes ensures that namelists with character strings can be subsequently read back in accurately.

5.1.4 X format descriptor without count field

To support legacy codes, GNU Fortran permits the count field of the `X` edit descriptor in `FORMAT` statements to be omitted. When omitted, the count is implicitly assumed to be one.

```

      PRINT 10, 2, 3
10      FORMAT (I1, X, I1)

```

5.1.5 Commas in FORMAT specifications

To support legacy codes, GNU Fortran allows the comma separator to be omitted immediately before and after character string edit descriptors in `FORMAT` statements. A comma with no following format descriptor is permitted if the `-fdec-blank-format-item` is given on the command line. This is considered non-conforming code and is discouraged.

```

      PRINT 10, 2, 3
10      FORMAT ('FOO=' I1' BAR=' I2)
      print 20, 5, 6
20      FORMAT (I3, I3,)

```

5.1.6 Missing period in FORMAT specifications

To support legacy codes, GNU Fortran allows missing periods in format specifications if and only if `-std=legacy` is given on the command line. This is considered non-conforming code and is discouraged.

```

      REAL :: value
      READ(*,10) value
10      FORMAT ('F4')

```

5.1.7 Default widths for ‘F’, ‘G’ and ‘I’ format descriptors

To support legacy codes, GNU Fortran allows width to be omitted from format specifications if and only if `-fdec-format-defaults` is given on the command line. Default widths are used. This is considered non-conforming code and is discouraged.

```

      REAL :: value1
      INTEGER :: value2
      WRITE(*,10) value1, value1, value2
10      FORMAT ('F, G, I')

```

5.1.8 I/O item lists

To support legacy codes, GNU Fortran allows the input item list of the `READ` statement, and the output item lists of the `WRITE` and `PRINT` statements, to start with a comma.

5.1.9 ‘Q’ exponent-letter

GNU Fortran accepts real literal constants with an exponent-letter of `‘Q’`, for example, `1.23Q45`. The constant is interpreted as a `REAL(16)` entity on targets that support this

type. If the target does not support `REAL(16)` but has a `REAL(10)` type, then the real-literal-constant is interpreted as a `REAL(10)` entity. In the absence of `REAL(16)` and `REAL(10)`, an error occurs.

5.1.10 BOZ literal constants

Besides decimal constants, Fortran also supports binary ('b'), octal ('o') and hexadecimal ('z') integer constants. The syntax is: `'prefix quote digits quote'`, where the prefix is either 'b', 'o' or 'z', quote is either ' or ' and the digits are 0 or 1 for binary, between 0 and 7 for octal, and between 0 and F for hexadecimal. (Example: `b'01011101'`.)

Up to Fortran 95, BOZ literal constants were only allowed to initialize integer variables in `DATA` statements. Since Fortran 2003 BOZ literal constants are also allowed as actual arguments to the `REAL`, `DBLE`, `INT` and `CMPLX` intrinsic functions. The BOZ literal constant is simply a string of bits, which is padded or truncated as needed, during conversion to a numeric type. The Fortran standard states that the treatment of the sign bit is processor dependent. Gfortran interprets the sign bit as a user would expect.

As a deprecated extension, GNU Fortran allows hexadecimal BOZ literal constants to be specified using the 'X' prefix. That the BOZ literal constant can also be specified by adding a suffix to the string, for example, `Z'ABC'` and `'ABC'X` are equivalent. Additionally, as extension, BOZ literals are permitted in some contexts outside of `DATA` and the intrinsic functions listed in the Fortran standard. Use `-fallow-invalid-boz` to enable the extension.

5.1.11 Real array indices

As an extension, GNU Fortran allows the use of `REAL` expressions or variables as array indices.

5.1.12 Unary operators

As an extension, GNU Fortran allows unary plus and unary minus operators to appear as the second operand of binary arithmetic operators without the need for parenthesis.

```
X = Y * -Z
```

5.1.13 Implicitly convert LOGICAL and INTEGER values

As an extension for backwards compatibility with other compilers, GNU Fortran allows the implicit conversion of `LOGICAL` values to `INTEGER` values and vice versa. When converting from a `LOGICAL` to an `INTEGER`, `.FALSE.` is interpreted as zero, and `.TRUE.` is interpreted as one. When converting from `INTEGER` to `LOGICAL`, the value zero is interpreted as `.FALSE.` and any nonzero value is interpreted as `.TRUE.`.

```
LOGICAL :: l
l = 1
INTEGER :: i
i = .TRUE.
```

However, there is no implicit conversion of `INTEGER` values in `if`-statements, nor of `LOGICAL` or `INTEGER` values in I/O operations.

5.1.14 Hollerith constants support

GNU Fortran supports Hollerith constants in assignments, `DATA` statements, function and subroutine arguments. A Hollerith constant is written as a string of characters preceded

by an integer constant indicating the character count, and the letter H or h, and stored in bitwise fashion in a numeric (INTEGER, REAL, or COMPLEX), LOGICAL or CHARACTER variable. The constant is padded with spaces or truncated to fit the size of the variable in which it is stored.

Examples of valid uses of Hollerith constants:

```
complex*16 x(2)
data x /16Habcdefghijklmnop, 16Hqrstuvwxyz012345/
x(1) = 16HABCDEFGHJKLMNOP
call foo (4h abc)
```

Examples of Hollerith constants:

```
integer*4 a
a = 0H          ! Invalid, at least one character is needed.
a = 4HAB12      ! Valid
a = 8H12345678 ! Valid, but the Hollerith constant is truncated.
a = 3Hxyz       ! Valid, but the Hollerith constant is padded.
```

In general, Hollerith constants were used to provide a rudimentary facility for handling character strings in early Fortran compilers, prior to the introduction of CHARACTER variables in Fortran 77; in those cases, the standard-compliant equivalent is to convert the program to use proper character strings. On occasion, there may be a case where the intent is specifically to initialize a numeric variable with a given byte sequence. In these cases, the same result can be obtained by using the TRANSFER statement, as in this example.

```
integer(kind=4) :: a
a = transfer ("abcd", a)      ! equivalent to: a = 4Habcd
```

The use of the -fdec option extends support of Hollerith constants to comparisons:

```
integer*4 a
a = 4hABCD
if (a .ne. 4habcd) then
  write(*,*) "no match"
end if
```

Supported types are numeric (INTEGER, REAL, or COMPLEX), and CHARACTER.

5.1.15 Character conversion

Allowing character literals to be used in a similar way to Hollerith constants is a nonstandard extension. This feature is enabled using -fdec-char-conversions and only applies to character literals of kind=1.

Character literals can be used in DATA statements and assignments with numeric (INTEGER, REAL, or COMPLEX) or LOGICAL variables. Like Hollerith constants they are copied bitwise fashion. The constant is padded with spaces or truncated to fit the size of the variable in which it is stored.

Examples:

```
integer*4 x
data x / 'abcd' /

x = 'A'          ! Is padded.
x = 'ab1234'     ! Is truncated.
```

5.1.16 Cray pointers

Cray pointers are part of a nonstandard extension that provides a C-like pointer in Fortran. This is accomplished through a pair of variables: an integer “pointer” that holds a memory address, and a “pointee” that is used to dereference the pointer.

Pointer/pointee pairs are declared in statements of the form:

```
pointer ( <pointer> , <pointee> )
```

or,

```
pointer ( <pointer1> , <pointee1> ), ( <pointer2> , <pointee2> ), ...
```

The pointer is an integer that is intended to hold a memory address. The pointee may be an array or scalar. If an assumed-size array is permitted within the scoping unit, a pointee can be an assumed-size array. That is, the last dimension may be left unspecified by using a `*` in place of a value. A pointee cannot be an assumed shape array. No space is allocated for the pointee.

The pointee may have its type declared before or after the pointer statement, and its array specification (if any) may be declared before, during, or after the pointer statement. The pointer may be declared as an integer prior to the pointer statement. However, some machines have default integer sizes that are different than the size of a pointer, and so the following code is not portable:

```
integer ipt
pointer (ipt, iarr)
```

If a pointer is declared with a kind that is too small, the compiler issues a warning; the resulting binary will probably not work correctly, because the memory addresses stored in the pointers may be truncated. It is safer to omit the first line of the above example; if explicit declaration of `ipt`’s type is omitted, then the compiler ensures that `ipt` is an integer variable large enough to hold a pointer.

Pointer arithmetic is valid with Cray pointers, but it is not the same as C pointer arithmetic. Cray pointers are just ordinary integers, so the user is responsible for determining how many bytes to add to a pointer in order to increment it. Consider the following example:

```
real target(10)
real pointee(10)
pointer (ipt, pointee)
ipt = loc (target)
ipt = ipt + 1
```

The last statement does not set `ipt` to the address of `target(1)`, as it would in C pointer arithmetic. Adding 1 to `ipt` just adds one byte to the address stored in `ipt`.

Any expression involving the pointee is translated to use the value stored in the pointer as the base address.

To get the address of elements, this extension provides an intrinsic function `LOC()`. The `LOC()` function is equivalent to the `&` operator in C, except the address is cast to an integer type:

```
real ar(10)
pointer(ipt, arpte(10))
real arpte
ipt = loc(ar) ! Makes arpte is an alias for ar
arpte(1) = 1.0 ! Sets ar(1) to 1.0
```

The pointer can also be set by a call to the `MALLOC` intrinsic (see Section 8.190 [MALLOC], page 244).

Cray pointees often are used to alias an existing variable. For example:

```
integer target(10)
integer iarr(10)
pointer (ipt, iarr)
ipt = loc(target)
```

As long as `ipt` remains unchanged, `iarr` is now an alias for `target`. The optimizer, however, does not detect this aliasing, so it is unsafe to use `iarr` and `target` simultaneously. Using a pointee in any way that violates the Fortran aliasing rules or assumptions is invalid. It is the user's responsibility to avoid doing this; the compiler works under the assumption that no such aliasing occurs.

Cray pointers work correctly when there is no aliasing (i.e., when they are used to access a dynamically allocated block of memory), and also in any routine where a pointee is used, but any variable with which it shares storage is not used. Code that violates these rules may not run as the user intends. This is not a bug in the optimizer; any code that violates the aliasing rules is invalid. (Note that this is not unique to GNU Fortran; any Fortran compiler that supports Cray pointers "incorrectly" optimizes code with invalid aliasing.)

There are a number of restrictions on the attributes that can be applied to Cray pointers and pointees. Pointees may not have the `ALLOCATABLE`, `INTENT`, `OPTIONAL`, `DUMMY`, `TARGET`, `INTRINSIC`, or `POINTER` attributes. Pointers may not have the `DIMENSION`, `POINTER`, `TARGET`, `ALLOCATABLE`, `EXTERNAL`, or `INTRINSIC` attributes, nor may they be function results. Pointees may not occur in more than one pointer statement. A pointee cannot be a pointer. Pointees cannot occur in equivalence, common, or data statements.

A Cray pointer may also point to a function or a subroutine. For example, the following excerpt is valid:

```
implicit none
external sub
pointer (subptr,subpte)
external subpte
subptr = loc(sub)
call subpte()
[...]
subroutine sub
[...]
end subroutine sub
```

A pointer may be modified during the course of a program, and this changes the location to which the pointee refers. However, when pointees are passed as arguments, they are treated as ordinary variables in the invoked function. Subsequent changes to the pointer do not change the base address of the array that was passed.

5.1.17 CONVERT specifier

GNU Fortran allows the conversion of unformatted data between little- and big-endian representation to facilitate moving of data between different systems. The conversion can be indicated with the `CONVERT` specifier on the `OPEN` statement. See Section 3.10 [GFORTRAN_CONVERT_UNIT], page 38, for an alternative way of specifying the data format via an environment variable.

Valid values for `CONVERT` on most systems are:

`CONVERT='NATIVE'` Use the native format. This is the default.

`CONVERT='SWAP'` Swap between little- and big-endian.

`CONVERT='LITTLE_ENDIAN'` Use the little-endian representation for unformatted files.

`CONVERT='BIG_ENDIAN'` Use the big-endian representation for unformatted files.

On POWER systems that support `-mabi=ieeelongdouble`, there are additional options, which can be combined with the others with commas. Those are

`CONVERT='R16_IEEE'` Use IEEE 128-bit format for `REAL(KIND=16)`.

`CONVERT='R16_IBM'` Use IBM long double format for `REAL(KIND=16)`.

Using the option could look like this:

```
open(file='big.dat',form='unformatted',access='sequential', &
      convert='big_endian')
```

The value of the conversion can be queried by using `INQUIRE(CONVERT=ch)`. The values returned are `'BIG_ENDIAN'` and `'LITTLE_ENDIAN'`.

`CONVERT` works between big- and little-endian for `INTEGER` values of all supported kinds and for `REAL` on IEEE systems of kinds 4 and 8. Conversion between different “extended double” types on different architectures such as m68k and x86_64, which GNU Fortran supports as `REAL(KIND=10)` and `REAL(KIND=16)`, probably does not work.

Note that the values specified via the `GFORTTRAN_CONVERT_UNIT` environment variable overrides the `CONVERT` specifier in the `OPEN` statement. This is to give control over data formats to users who do not have the source code of their program available.

Using anything but the native representation for unformatted data carries a significant speed overhead. If speed in this area matters to you, it is best if you use this only for data that needs to be portable.

5.1.18 OpenMP

OpenMP (Open Multi-Processing) is an application programming interface (API) that supports multi-platform shared memory multiprocessing programming in C/C++ and Fortran on many architectures, including Unix and Microsoft Windows platforms. It consists of a set of compiler directives, library routines, and environment variables that influence run-time behavior.

GNU Fortran implements all of the OpenMP Application Program Interface v4.5 (<https://openmp.org/specifications/>), and many features from later versions of the OpenMP specification. See Section “OpenMP Implementation Status” in *GNU Offloading and Multi Processing Runtime Library*, for more details about currently supported OpenMP features.

To enable the processing of the OpenMP directive `!$omp` in free-form source code; the `c$omp`, `*$omp` and `!$omp` directives in fixed form; the `!$` conditional compilation sentinels in free form; and the `c$`, `*$` and `!$` sentinels in fixed form, `gfortran` needs to be invoked with the `-fopenmp` option. This option also arranges for automatic linking of the OpenMP runtime library. See *GNU Offloading and Multi Processing Runtime Library*.

The OpenMP Fortran runtime library routines are provided both in a form of a Fortran 90 module named `omp_lib` and in a form of a Fortran `include` file named `omp_lib.h`.

An example of a parallelized loop taken from Appendix A.1 of the OpenMP Application Program Interface v2.5:

```
SUBROUTINE A1(N, A, B)
  INTEGER I, N
  REAL B(N), A(N)
  !$OMP PARALLEL DO !I is private by default
  DO I=2,N
    B(I) = (A(I) + A(I-1)) / 2.0
  ENDDO
  !$OMP END PARALLEL DO
END SUBROUTINE A1
```

See Section “OpenMP and OpenACC Options” in *Using the GNU Compiler Collection (GCC)*, for additional options useful with `-fopenmp`.

Please note:

- `-fopenmp` implies `-frecursive`, i.e., all local arrays are allocated on the stack. When porting existing code to OpenMP, this may lead to surprising results, especially segmentation faults if the stack size is limited.
- On glibc-based systems, OpenMP-enabled applications cannot be statically linked due to limitations of the underlying pthreads implementation. It might be possible to get a working solution if `-Wl,--whole-archive -lpthread -Wl,--no-whole-archive` is added to the command line. However, this is not supported by GCC and thus not recommended.

5.1.19 OpenACC

OpenACC is an application programming interface (API) that supports offloading of code to accelerator devices. It consists of a set of compiler directives, library routines, and environment variables that influence run-time behavior.

GNU Fortran strives to be compatible with the OpenACC Application Programming Interface v2.6 (<https://www.openacc.org/>).

To enable the processing of the OpenACC directive `!$acc` in free-form source code; the `c$acc`, `*$acc` and `!$acc` directives in fixed form; the `!$` conditional compilation sentinels in free form; and the `c$`, `*$` and `!$` sentinels in fixed form, `gfortran` needs to be invoked with the `-fopenacc` option. This option also arranges for automatic linking of the OpenACC runtime library. See *GNU Offloading and Multi Processing Runtime Library*.

The OpenACC Fortran runtime library routines are provided both in a form of a Fortran 90 module named `openacc` and in a form of a Fortran `include` file named `openacc_lib.h`.

See Section “OpenMP and OpenACC Options” in *Using the GNU Compiler Collection (GCC)*, for additional options useful with `-fopenacc`.

5.1.20 Argument list functions %VAL, %REF and %LOC

GNU Fortran supports argument list functions `%VAL`, `%REF` and `%LOC` statements, for backward compatibility with g77. It is recommended that these should be used only for code that is accessing facilities outside of GNU Fortran, such as operating system or windowing facilities. It is best to constrain such uses to isolated portions of a program—portions that deal specifically and exclusively with low-level, system-dependent facilities. Such portions

might well provide a portable interface for use by the program as a whole, but are themselves not portable, and should be thoroughly tested each time they are rebuilt using a new compiler or version of a compiler.

%VAL passes a scalar argument by value, **%REF** passes it by reference and **%LOC** passes its memory location. Since gfortran already passes scalar arguments by reference, **%REF** is in effect a do-nothing. **%LOC** has the same effect as a Fortran pointer.

An example of passing an argument by value to a C subroutine foo.:

```
C
C prototype      void foo_ (float x);
C
      external foo
      real*4 x
      x = 3.14159
      call foo (%VAL (x))
      end
```

For details refer to the g77 manual <https://gcc.gnu.org/onlinedocs/gcc-3.4.6/g77/index.html#Top>.

Also, `c_by_val.f` and its partner `c_by_val.c` of the GNU Fortran testsuite are worth a look.

5.1.21 Read/Write after EOF marker

Some legacy codes rely on allowing **READ** or **WRITE** after the EOF file marker in order to find the end of a file. GNU Fortran normally rejects these codes with a run-time error message and suggests the user consider **BACKSPACE** or **REWIND** to properly position the file before the EOF marker. As an extension, the run-time error may be disabled using `-std=legacy`.

5.1.22 STRUCTURE and RECORD

Record structures are a pre-Fortran-90 vendor extension to create user-defined aggregate data types. Support for record structures in GNU Fortran can be enabled with the `-fdec-structure` compile flag. If you have a choice, you should instead use Fortran 90's "derived types", which have a different syntax.

In many cases, record structures can easily be converted to derived types. To convert, replace **STRUCTURE** */structure-name/* by **TYPE** *type-name*. Additionally, replace **RECORD** */structure-name/* by **TYPE**(*type-name*). Finally, in the component access, replace the period (.) by the percent sign (%).

Here is an example of code using the non portable record structure syntax:

```
! Declaring a structure named ``item'' and containing three fields:
! an integer ID, an description string and a floating-point price.
STRUCTURE /item/
  INTEGER id
  CHARACTER(LEN=200) description
  REAL price
END STRUCTURE

! Define two variables, an single record of type ``item''
! named ``pear'', and an array of items named ``store_catalog''
```

```

RECORD /item/ pear, store_catalog(100)

! We can directly access the fields of both variables
pear.id = 92316
pear.description = "juicy D'Anjou pear"
pear.price = 0.15
store_catalog(7).id = 7831
store_catalog(7).description = "milk bottle"
store_catalog(7).price = 1.2

! We can also manipulate the whole structure
store_catalog(12) = pear
print *, store_catalog(12)

```

This code can easily be rewritten in the Fortran 90 syntax as following:

```

! ``STRUCTURE /name/ ... END STRUCTURE'' becomes
! ``TYPE name ... END TYPE''
TYPE item
  INTEGER id
  CHARACTER(LEN=200) description
  REAL price
END TYPE

! ``RECORD /name/ variable'' becomes ``TYPE(name) variable''
TYPE(item) pear, store_catalog(100)

! Instead of using a dot (.) to access fields of a record, the
! standard syntax uses a percent sign (%)
pear%id = 92316
pear%description = "juicy D'Anjou pear"
pear%price = 0.15
store_catalog(7)%id = 7831
store_catalog(7)%description = "milk bottle"
store_catalog(7)%price = 1.2

! Assignments of a whole variable do not change
store_catalog(12) = pear
print *, store_catalog(12)

```

GNU Fortran implements structures like derived types with the following rules and exceptions:

- Structures act like derived types with the `SEQUENCE` attribute. Otherwise they may contain no specifiers.
- Structures may contain a special field with the name `%FILL`. This creates an anonymous component that cannot be accessed but occupies space just as if a component of the same type was declared in its place, useful for alignment purposes. As an example, the following structure consists of at least sixteen bytes:

```

structure /padded/

```

```

        character(4) start
        character(8) %FILL
        character(4) end
    end structure

```

- Structures may share names with other symbols. For example, the following is invalid for derived types, but valid for structures:

```

    structure /header/
    ! ...
    end structure
    record /header/ header

```

- Structure types may be declared nested within another parent structure. The syntax is:

```

    structure /type-name/
    ...
    structure [/<type-name>/] <field-list>
    ...

```

The type name may be omitted, in which case the structure type itself is anonymous, and other structures of the same type cannot be instantiated. The following shows some examples:

```

    structure /appointment/
    ! nested structure definition: app_time is an array of two 'time'
    structure /time/ app_time (2)
        integer(1) hour, minute
    end structure
    character(10) memo
end structure

```

```

! The 'time' structure is still usable
record /time/ now
now = time(5, 30)

```

```

...

```

```

    structure /appointment/
    ! anonymous nested structure definition
    structure start, end
        integer(1) hour, minute
    end structure
    character(10) memo
end structure

```

- Structures may contain UNION blocks. For more detail see the section on Section 5.1.23 [UNION and MAP], page 61.
- Structures support old-style initialization of components, like those described in Section 5.1.2 [Old-style variable initialization], page 49. For array initializers, an initializer may contain a repeat specification of the form `<literal-integer> * <constant-initializer>`. The value of the integer indicates the number of times to repeat the constant initializer when expanding the initializer list.


```

      map
        character(8) rl      ! ral
      end map
      map
        character(8) ex      ! eax
      end map
      map
        character(4) eh      ! eah
        union ! U2
          map
            character(4) el ! eal
          end map
          map
            character(4) x  ! ax
          end map
          map
            character(2) h  ! ah
            character(2) l  ! al
          end map
        end union
      end map
    end union
  end map
end union
end structure
record /reg/ a

! After this assignment...
a.rx    =    'AAAAAAAA.BBB.C.D'

! The following is true:
a.rx == 'AAAAAAAA.BBB.C.D'
a.rh == 'AAAAAAAA'
a.rl == '.BBB.C.D'
a.ex == '.BBB.C.D'
a.eh == '.BBB'
a.el == '.C.D'
a.x  == '.C.D'
a.h  == '.C'
a.l  == '.D'

```

5.1.24 Type variants for integer intrinsics

Similar to the D/C prefixes to real functions to specify the input/output types, GNU Fortran offers B/I/J/K prefixes to integer functions for compatibility with DEC programs. The types implied by each are:

B - INTEGER(kind=1)

I - INTEGER(kind=2)
 J - INTEGER(kind=4)
 K - INTEGER(kind=8)

GNU Fortran supports these with the flag `-fdec-intrinsic-ints`. Intrinsic for which prefixed versions are available and in what form are noted in Chapter 8 [Intrinsic Procedures], page 119. The complete list of supported intrinsics is here:

Intrinsic	B	I	J	K
Section 8.3 [ABS], page 120	BABS	IIABS	JIABS	KIABS
Section 8.55 [BTEST], page 155	BBTEST	BITEST	BJTEST	BKTEST
Section 8.143 [IAND], page 215	BIAND	IIAND	JIAND	KIAND
Section 8.146 [IBCLR], page 217	BBCLR	IIBCLR	JIBCLR	KIBCLR
Section 8.147 [IBITS], page 218	BBITS	IIBITS	JIBITS	KIBITS
Section 8.148 [IBSET], page 218	BBSET	IIBSET	JIBSET	KIBSET
Section 8.151 [IEOR], page 221	BIEOR	IIEOR	JIEOR	KIEOR
Section 8.158 [IOR], page 225	BIOR	IIOR	JIOR	KIOR
Section 8.165 [ISHFT], page 229	BSHFT	IISHFT	JISHFT	KISHFT
Section 8.166 [ISHFTC], page 230	BSHFTC	IISHFTC	JISHFTC	KISHFTC
Section 8.206 [MOD], page 253	BMOD	IMOD	JMOD	KMOD
Section 8.214 [NOT], page 259	BNOT	INOT	JNOT	KNOT
Section 8.235 [REAL], page 272	--	FLOATI	FLOATJ	FLOATK

5.1.25 AUTOMATIC and STATIC attributes

With `-fdec-static` GNU Fortran supports the DEC extended attributes `STATIC` and `AUTOMATIC` to provide explicit specification of entity storage. These follow the syntax of the Fortran standard `SAVE` attribute.

`STATIC` is exactly equivalent to `SAVE`, and specifies that an entity should be allocated in static memory. As an example, `STATIC` local variables retain their values across multiple calls to a function.

Entities marked `AUTOMATIC` are stack automatic whenever possible. `AUTOMATIC` is the default for local variables smaller than `-fmax-stack-var-size`, unless `-fno-automatic` is given. This attribute overrides `-fno-automatic`, `-fmax-stack-var-size`, and blanket `SAVE` statements.

Examples:

```
subroutine f
  integer, automatic :: i  ! automatic variable
  integer x, y             ! static variables
  save
  ...
endsubroutine

subroutine f
  integer a, b, c, x, y, z
  static :: x
  save y
  automatic z, c
  ! a, b, c, and z are automatic
  ! x and y are static
endsubroutine

! Compiled with -fno-automatic
subroutine f
  integer a, b, c, d
  automatic :: a
  ! a is automatic; b, c, and d are static
endsubroutine
```

5.1.26 Form feed as whitespace

Historically, legacy compilers allowed insertion of form feed characters (`'\f'`, ASCII 0xC) at the beginning of lines for formatted output to line printers, though the Fortran standard does not mention this. GNU Fortran supports the interpretation of form feed characters in source as whitespace for compatibility.

5.1.27 TYPE as an alias for PRINT

For compatibility, GNU Fortran interprets `TYPE` statements as `PRINT` statements with the flag `-fdec`. With this flag asserted, the following two examples are equivalent:

```
TYPE *, 'hello world'
PRINT *, 'hello world'
```

5.1.28 %LOC as an rvalue

Normally %LOC is allowed only in parameter lists. However the intrinsic function LOC does the same thing, and is usable as the right-hand-side of assignments. For compatibility, GNU Fortran supports the use of %LOC as an alias for the builtin LOC with `-std=legacy`. With this feature enabled the following two examples are equivalent:

```
integer :: i, l
l = %loc(i)
call sub(l)

integer :: i
call sub(%loc(i))
```

5.1.29 .XOR. operator

GNU Fortran supports `.XOR.` as a logical operator with `-std=legacy` for compatibility with legacy code. `.XOR.` is equivalent to `.NEQV..` That is, the output is true if and only if the inputs differ.

5.1.30 Bitwise logical operators

With `-fdec`, GNU Fortran relaxes the type constraints on logical operators to allow integer operands, and performs the corresponding bitwise operation instead. This flag is for compatibility only, and should be avoided in new code. Consider:

```
INTEGER :: i, j
i = z'33'
j = z'cc'
print *, i .AND. j
```

In this example, compiled with `-fdec`, GNU Fortran replaces the `.AND.` operation with a call to the intrinsic [Section 8.143 \[IAND\]](#), [page 215](#) function, yielding the bitwise-and of `i` and `j`.

Note that this conversion occurs if at least one operand is of integral type. As a result, a logical operand is converted to an integer when the other operand is an integer in a logical operation. In this case, `.TRUE.` is converted to 1 and `.FALSE.` to 0.

Here is the mapping of logical operator to bitwise intrinsic used with `-fdec`:

Operator	Intrinsic	Bitwise operation
<code>.NOT.</code>	NOT	complement (see Section 8.214 [NOT] , page 259)
<code>.AND.</code>	IAND	intersection (see Section 8.143 [IAND] , page 215)
<code>.OR.</code>	IOR	union (see Section 8.158 [IOR] , page 225)
<code>.NEQV.</code>	IEOR	exclusive or (see Section 8.151 [IEOR] , page 221)
<code>.EQV.</code>	NOT IEO	complement of exclusive or (see Section 8.151 [IEOR] , page 221)

5.1.31 Extended I/O specifiers

GNU Fortran supports the additional legacy I/O specifiers `CARRIAGECONTROL`, `READONLY`, and `SHARE` with the compile flag `-fdec`, for compatibility.

CARRIAGECONTROL

The `CARRIAGECONTROL` specifier allows a user to control line termination settings between output records for an I/O unit. The specifier has no meaning for

readonly files. When `CARRIAGECONTROL` is specified upon opening a unit for formatted writing, the exact `CARRIAGECONTROL` setting determines what characters to write between output records. The syntax is:

```
OPEN(..., CARRIAGECONTROL=cc)
```

where `cc` is a character expression that evaluates to one of the following values:

'LIST'	One line feed between records (default)
'FORTRAN'	Legacy interpretation of the first character (see below)
'NONE'	No separator between records

With `CARRIAGECONTROL='FORTRAN'`, when a record is written, the first character of the input record is not written, and instead determines the output record separator as follows:

Leading character	Meaning	Output character(s)	separating
'+'	Overprinting	Carriage return only	
'-'	New line	Line feed and carriage return	
'0'	Skip line	Two line feeds and carriage return	
'1'	New page	Form feed and carriage return	
'\$'	Prompting	Line feed (no carriage return)	
CHAR(0)	Overprinting (no advance)	None	

READONLY The `READONLY` specifier may be given upon opening a unit, and is equivalent to specifying `ACTION='READ'`, except that the file may not be deleted on close (i.e. `CLOSE` with `STATUS="DELETE"`). The syntax is:

```
OPEN(..., READONLY)
```

SHARE The `SHARE` specifier allows system-level locking on a unit upon opening it for controlled access from multiple processes/threads. The `SHARE` specifier has several forms:

```
OPEN(..., SHARE=sh)
OPEN(..., SHARED)
OPEN(..., NOSHARED)
```

Where `sh` in the first form is a character expression that evaluates to a value as seen in the table below. The latter two forms are aliases for particular values of `sh`:

Explicit form	Short form	Meaning
SHARE='DENYRW'	NOSHARED	Exclusive (write) lock
SHARE='DENYNONE'	SHARED	Shared (read) lock

In general only one process may hold an exclusive (write) lock for a given file at a time, whereas many processes may hold shared (read) locks for the same file.

The behavior of locking may vary with your operating system. On POSIX systems, locking is implemented with `fcntl`. Consult your corresponding operating system's manual pages for further details. Locking via `SHARE=` is not supported on other systems.

5.1.32 Legacy PARAMETER statements

For compatibility, GNU Fortran supports legacy `PARAMETER` statements without parentheses with `-std=legacy`. A warning is emitted if used with `-std=gnu`, and an error is acknowledged with a real Fortran standard flag (`-std=f95`, etc...). These statements take the following form:

```
implicit real (E)
parameter e = 2.718282
real c
parameter c = 3.0e8
```

5.1.33 Default exponents

For compatibility, GNU Fortran supports a default exponent of zero in real constants with `-fdec`. For example, `9e` would be interpreted as `9e0`, rather than an error.

5.1.34 Unsigned integers

If the `-funsigned` option is given, GNU Fortran supports unsigned integers according to J3/24-116 (<https://j3-fortran.org/doc/year/24/24-116.txt>). The data type is called `UNSIGNED`. For an unsigned type with `n` bits, it implements integer arithmetic modulo 2^{**n} , comparable to the `unsigned` data type in C.

The data type has `KIND` numbers comparable to other Fortran data types, which can be selected via the `SELECTED_UNSIGNED_KIND` function.

Mixed arithmetic, comparisons and assignment between `UNSIGNED` and other types are only possible via explicit conversion. Conversion from `UNSIGNED` to other types is done via type conversion functions like `INT` or `REAL`. Conversion from other types to `UNSIGNED` is done via `UINT`. Unsigned variables cannot be used as index variables in `DO` loops or as array indices.

Unsigned numbers have a trailing `u` as suffix, optionally followed by a `KIND` number separated by an underscore.

Input and output can be done using the ‘I’, ‘B’, ‘O’ and ‘Z’ descriptors, plus unformatted I/O.

Unsigned integers as implemented in gfortran are compatible with `flang`.

Here is a small, somewhat contrived example of their use:

```
program main
  use iso_fortran_env, only : uint64
  unsigned(kind=uint64) :: v
  v = huge(v) - 32u_uint64
  print *,v
end program main
```

which outputs the number 18446744073709551583.

Arithmetic operations work on unsigned integers, also for exponentiation. As an extension to J3/24-116.txt, unary minus and exponentiation of unsigned integers are permitted unless `-pedantic` is in force.

In intrinsic procedures, unsigned arguments are typically permitted for arguments for the data to be processed, analogous to the use of `REAL` arguments. Unsigned values are prohibited as index variables in `DO` loops and as array indices.

Unsigned numbers can be read and written using list-directed, formatted and unformatted I/O. For formatted I/O, the ‘B’, ‘I’, ‘O’ and ‘Z’ descriptors are valid. Negative values and values that would overflow are rejected with `-pedantic`.

SELECT CASE is supported for unsigned integers.

The following intrinsics take unsigned arguments:

- BGE, see Section 8.50 [BGE], page 152,
- BGT, see Section 8.51 [BGT], page 153,
- BIT_SIZE, see Section 8.52 [BIT_SIZE], page 153,
- BLE, see Section 8.53 [BLE], page 154,
- BLT, see Section 8.54 [BLT], page 154,
- CMPLX, see Section 8.66 [CMPLX], page 162,
- CSHIFT, see Section 8.85 [CSHIFT], page 176,
- DIGITS, see Section 8.90 [DIGITS], page 180,
- DOT_PRODUCT, see Section 8.92 [DOT_PRODUCT], page 181,
- DSHIFTL, see Section 8.95 [DSHIFTL], page 183,
- DSHIFTR, see Section 8.96 [DSHIFTR], page 183,
- EOSHIFT, see Section 8.98 [EOSHIFT], page 185,
- FINDLOC, see Section 8.113 [FINDLOC], page 195,
- HUGE, see Section 8.139 [HUGE], page 212,
- IALL, see Section 8.142 [IALL], page 214,
- IAND, see Section 8.143 [IAND], page 215,
- IANY, see Section 8.144 [IANY], page 216,
- IBCLR, see Section 8.146 [IBCLR], page 217,
- IBITS, see Section 8.147 [IBITS], page 218,
- IBSET, see Section 8.148 [IBSET], page 218,
- IEOR, see Section 8.151 [IEOR], page 221,
- INT, see Section 8.155 [INT], page 223,
- IOR, see Section 8.158 [IOR], page 225,
- IPARITY, see Section 8.159 [IPARITY], page 225,
- ISHFT, see Section 8.165 [ISHFT], page 229,
- ISHFTC, see Section 8.166 [ISHFTC], page 230,
- MATMUL, see Section 8.193 [MATMUL], page 245,
- MAX, see Section 8.194 [MAX], page 246,
- MAXLOC, see Section 8.196 [MAXLOC], page 247,
- MAXVAL, see Section 8.197 [MAXVAL], page 248,
- MERGE, see Section 8.200 [MERGE], page 250,
- MERGE_BITS, see Section 8.201 [MERGE_BITS], page 250,
- MIN, see Section 8.202 [MIN], page 251,
- MINLOC, see Section 8.204 [MINLOC], page 252,

- MINVAL, see Section 8.205 [MINVAL], page 253,
- MOD, see Section 8.206 [MOD], page 253,
- MODULO, see Section 8.207 [MODULO], page 254,
- MVBITS, see Section 8.209 [MVBITS], page 256,
- NOT, see Section 8.214 [NOT], page 259,
- OUT_OF_RANGE, see Section 8.218 [OUT_OF_RANGE], page 262,
- PRODUCT, see Section 8.226 [PRODUCT], page 266,
- RANDOM_NUMBER, see Section 8.231 [RANDOM_NUMBER], page 269,
- RANGE, see Section 8.233 [RANGE], page 271,
- REAL, see Section 8.235 [REAL], page 272,
- SHIFTA, see Section 8.253 [SHIFTA], page 282,
- SHIFTL, see Section 8.254 [SHIFTL], page 283,
- SHIFTR, see Section 8.255 [SHIFTR], page 283,
- SUM, see Section 8.272 [SUM], page 294,
- TRANSPOSE, see Section 8.287 [TRANSPOSE], page 304,
- TRANSFER, see Section 8.286 [TRANSFER], page 303,

The following intrinsics are enabled with `-funsigned`:

- UINT, see Section 8.292 [UINT], page 306,
- UMASKL, see Section 8.294 [UMASKL], page 307,
- UMASKR, see Section 8.295 [UMASKR], page 307,
- SELECTED_UNSIGNED_KIND, see Section 8.250 [SELECTED_UNSIGNED_KIND], page 281,

The following constants have been added to the intrinsic `ISO_C_BINDING` module: `c_unsigned`, `c_unsigned_short`, `c_unsigned_char`, `c_unsigned_long`, `c_unsigned_long_long`, `c_uintmax_t`, `c_uint8_t`, `c_uint16_t`, `c_uint32_t`, `c_uint64_t`, `c_uint128_t`, `c_uint_fast8_t`, `c_uint_fast16_t`, `c_uint_fast32_t`, `c_uint_fast64_t`, `c_uint_fast128_t`, `c_uint_least8_t`, `c_uint_least16_t`, `c_uint_least32_t`, `c_uint_least64_t` and `c_uint_least128_t`.

The following constants have been added to the intrinsic `ISO_FORTRAN_ENV` module: `uint8`, `uint16`, `uint32` and `uint64`.

5.2 Extensions not implemented in GNU Fortran

The long history of the Fortran language, its wide use and broad userbase, the large number of different compiler vendors and the lack of some features crucial to users in the first standards have lead to the existence of a number of important extensions to the language. While some of the most useful or popular extensions are supported by the GNU Fortran compiler, not all existing extensions are supported. This section aims at listing these extensions and offering advice on how best make code that uses them running with the GNU Fortran compiler.

5.2.1 ENCODE and DECODE statements

GNU Fortran does not support the `ENCODE` and `DECODE` statements. These statements are best replaced by `READ` and `WRITE` statements involving internal files (`CHARACTER` variables and arrays), which have been part of the Fortran standard since Fortran 77. For example, replace a code fragment like

```

      INTEGER*1 LINE(80)
      REAL A, B, C
c      ... Code that sets LINE
      DECODE (80, 9000, LINE) A, B, C
      9000 FORMAT (1X, 3(F10.5))

```

with the following:

```

      CHARACTER(LEN=80) LINE
      REAL A, B, C
c      ... Code that sets LINE
      READ (UNIT=LINE, FMT=9000) A, B, C
      9000 FORMAT (1X, 3(F10.5))

```

Similarly, replace a code fragment like

```

      INTEGER*1 LINE(80)
      REAL A, B, C
c      ... Code that sets A, B and C
      ENCODE (80, 9000, LINE) A, B, C
      9000 FORMAT (1X, 'OUTPUT IS ', 3(F10.5))

```

with the following:

```

      CHARACTER(LEN=80) LINE
      REAL A, B, C
c      ... Code that sets A, B and C
      WRITE (UNIT=LINE, FMT=9000) A, B, C
      9000 FORMAT (1X, 'OUTPUT IS ', 3(F10.5))

```

5.2.2 Variable FORMAT expressions

A variable `FORMAT` expression is format statement that includes angle brackets enclosing a Fortran expression: `FORMAT(I<N>)`. GNU Fortran does not support this legacy extension. The effect of variable format expressions can be reproduced by using the more powerful (and standard) combination of internal output and string formats. For example, replace a code fragment like this:

```

      WRITE(6,20) INT1
20    FORMAT(I<N+1>)

```

with the following:

```

c      Variable declaration
      CHARACTER(LEN=20) FMT
c
c      Other code here...
c
      WRITE(FMT,'("(I", I0, ")")') N+1
      WRITE(6,FMT) INT1

```

or with:

```

c      Variable declaration
      CHARACTER(LEN=20) FMT
c
c      Other code here...
c

```

```
WRITE(FMT,*) N+1  
WRITE(6,"(I" // ADJUSTL(FMT) // ")") INT1
```

5.2.3 Alternate complex function syntax

Some Fortran compilers, including `g77`, let the user declare complex functions with the syntax `COMPLEX FUNCTION name*16()`, as well as `COMPLEX*16 FUNCTION name()`. Both are nonstandard legacy extensions. `gfortran` accepts the latter form, which is more common, but not the former.

5.2.4 Volatile COMMON blocks

Some Fortran compilers, including `g77`, let the user declare `COMMON` with the `VOLATILE` attribute. This is invalid standard Fortran syntax and is not supported by `gfortran`. Note that `gfortran` accepts `VOLATILE` variables in `COMMON` blocks since revision 4.3.

5.2.5 OPEN(... NAME=)

Some Fortran compilers, including `g77`, let the user declare `OPEN(... NAME=)`. This is invalid standard Fortran syntax and is not supported by `gfortran`. `OPEN(... NAME=)` should be replaced with `OPEN(... FILE=)`.

5.2.6 Q edit descriptor

Some Fortran compilers provide the `Q` edit descriptor, which transfers the number of characters left within an input record into an integer variable.

A direct replacement of the `Q` edit descriptor is not available in `gfortran`. How to replicate its functionality using standard-conforming code depends on what the intent of the original code is.

Options to replace `Q` may be to read the whole line into a character variable and then counting the number of non-blank characters left using `LEN_TRIM`. Another method may be to use formatted stream, read the data up to the position where the `Q` descriptor occurred, use `INQUIRE` to get the file position, count the characters up to the next `NEW_LINE` and then start reading from the position marked previously.

6 Mixed-Language Programming

This chapter is about mixed-language interoperability, but also applies if you link Fortran code compiled by different compilers. In most cases, use of the C Binding features of the Fortran 2003 and later standards is sufficient.

For example, it is possible to mix Fortran code with C++ code as well as C, if you declare the interface functions as `extern "C"` on the C++ side and `BIND(C)` on the Fortran side, and follow the rules for interoperability with C. Note that you cannot manipulate C++ class objects in Fortran or vice versa except as opaque pointers.

You can use the `gfortran` command to link both Fortran and non-Fortran code into the same program, or you can use `gcc` or `g++` if you also add an explicit `-lgfortran` option to link with the Fortran library. If your main program is written in C or some other language instead of Fortran, see Section 6.3 [Non-Fortran Main Program], page 81, below.

6.1 Interoperability with C

Since Fortran 2003 (ISO/IEC 1539-1:2004(E)) there is a standardized way to generate procedure and derived-type declarations and global variables that are interoperable with C (ISO/IEC 9899:1999). The `BIND(C)` attribute has been added to inform the compiler that a symbol shall be interoperable with C; also, some constraints are added. Note, however, that not all C features have a Fortran equivalent or vice versa. For instance, neither C's unsigned integers nor C's functions with variable number of arguments have an equivalent in Fortran.

Note that array dimensions are reversely ordered in C and that arrays in C always start with index 0 while in Fortran they start by default with 1. Thus, an array declaration `A(n,m)` in Fortran matches `A[m][n]` in C and accessing the element `A(i,j)` matches `A[j-1][i-1]`. The element following `A(i,j)` (C: `A[j-1][i-1]`; assuming $i < n$) in memory is `A(i+1,j)` (C: `A[j-1][i]`).

6.1.1 Intrinsic Types

In order to ensure that exactly the same variable type and kind is used in C and Fortran, you should use the named constants for kind parameters that are defined in the `ISO_C_BINDING` intrinsic module. That module contains named constants of character type representing the escaped special characters in C, such as newline. For a list of the constants, see Section 9.2 [ISO_C_BINDING], page 313.

For logical types, please note that the Fortran standard only guarantees interoperability between C99's `_Bool` and Fortran's `C_Bool`-kind logicals and C99 defines that `true` has the value 1 and `false` the value 0. Using any other integer value with GNU Fortran's `LOGICAL` (with any kind parameter) gives an undefined result. (Passing other integer values than 0 and 1 to GCC's `_Bool` is also undefined, unless the integer is explicitly or implicitly casted to `_Bool`.)

6.1.2 Derived Types and struct

For compatibility of derived types with `struct`, use the `BIND(C)` attribute in the type declaration. For instance, the following type declaration

```
USE ISO_C_BINDING
```

```

TYPE, BIND(C) :: myType
  INTEGER(C_INT) :: i1, i2
  INTEGER(C_SIGNED_CHAR) :: i3
  REAL(C_DOUBLE) :: d1
  COMPLEX(C_FLOAT_COMPLEX) :: c1
  CHARACTER(KIND=C_CHAR) :: str(5)
END TYPE

```

matches the following `struct` declaration in C

```

struct {
  int i1, i2;
  /* Note: "char" might be signed or unsigned. */
  signed char i3;
  double d1;
  float _Complex c1;
  char str[5];
} myType;

```

Derived types with the C binding attribute shall not have the `sequence` attribute, type parameters, the `extends` attribute, nor type-bound procedures. Every component must be of interoperable type and kind and may not have the `pointer` or `allocatable` attribute. The names of the components are irrelevant for interoperability.

As there exist no direct Fortran equivalents, neither unions nor structs with bit field or variable-length array members are interoperable.

6.1.3 Interoperable Global Variables

Variables can be made accessible from C using the C binding attribute, optionally together with specifying a binding name. Those variables have to be declared in the declaration part of a `MODULE`, be of interoperable type, and have neither the `pointer` nor the `allocatable` attribute.

```

MODULE m
  USE myType_module
  USE ISO_C_BINDING
  integer(C_INT), bind(C, name="_MyProject_flags") :: global_flag
  type(myType), bind(C) :: tp
END MODULE

```

Here, `_MyProject_flags` is the case-sensitive name of the variable as seen from C programs while `global_flag` is the case-insensitive name as seen from Fortran. If no binding name is specified, as for `tp`, the C binding name is the (lowercase) Fortran binding name. If a binding name is specified, only a single variable may be after the double colon. Note of warning: You cannot use a global variable to access `errno` of the C library as the C standard allows it to be a macro. Use the `IERRNO` intrinsic (GNU extension) instead.

6.1.4 Interoperable Subroutines and Functions

Subroutines and functions have to have the `BIND(C)` attribute to be compatible with C. The dummy argument declaration is relatively straightforward. However, one needs to be careful because C uses call-by-value by default while Fortran behaves usually similar to call-by-reference. Furthermore, strings and pointers are handled differently.

To pass a variable by value, use the `VALUE` attribute. Thus, the following C prototype

```
int func(int i, int *j)
```

matches the Fortran declaration

```
integer(c_int) function func(i,j)
  use iso_c_binding, only: c_int
  integer(c_int), VALUE :: i
  integer(c_int) :: j
```

Note that pointer arguments also frequently need the `VALUE` attribute, see Section 6.1.5 [Working with C Pointers], page 76.

Strings are handled quite differently in C and Fortran. In C a string is a NUL-terminated array of characters while in Fortran each string has a length associated with it and is thus not terminated (by e.g. NUL). For example, if you want to use the following C function,

```
#include <stdio.h>
void print_C(char *string) /* equivalent: char string[] */
{
  printf("%s\n", string);
}
```

to print “Hello World” from Fortran, you can call it using

```
use iso_c_binding, only: C_CHAR, C_NULL_CHAR
interface
  subroutine print_c(string) bind(C, name="print_C")
    use iso_c_binding, only: c_char
    character(kind=c_char) :: string(*)
  end subroutine print_c
end interface
call print_c(C_CHAR_"Hello World"//C_NULL_CHAR)
```

As the example shows, you need to ensure that the string is NUL terminated. Additionally, the dummy argument *string* of `print_C` is a length-one assumed-size array; using `character(len=*)` is not allowed. The example above uses `c_char_"Hello World"` to ensure the string literal has the right type; typically the default character kind and `c_char` are the same and thus `"Hello World"` is equivalent. However, the standard does not guarantee this.

The use of strings is now further illustrated using the C library function `strncpy`, whose prototype is

```
char *strncpy(char *restrict s1, const char *restrict s2, size_t n);
```

The function `strncpy` copies at most *n* characters from string *s2* to *s1* and returns *s1*. In the following example, we ignore the return value:

```
use iso_c_binding
implicit none
character(len=30) :: str, str2
interface
  ! Ignore the return value of strncpy -> subroutine
  ! "restrict" is always assumed if we do not pass a pointer
  subroutine strncpy(dest, src, n) bind(C)
    import
    character(kind=c_char), intent(out) :: dest(*)
    character(kind=c_char), intent(in)  :: src(*)
    integer(c_size_t), value, intent(in) :: n
  end subroutine strncpy
end interface
str = repeat('X',30) ! Initialize whole string with 'X'
call strncpy(str, c_char_"Hello World"//C_NULL_CHAR, &
             len(c_char_"Hello World",kind=c_size_t))
```

```
print '(a)', str ! prints: "Hello WorldXXXXXXXXXXXXXXXXXXXX"
end
```

The intrinsic procedures are described in Chapter 8 [Intrinsic Procedures], page 119.

6.1.5 Working with C Pointers

C pointers are represented in Fortran via the special opaque derived type `type(c_ptr)` (with private components). C pointers are distinct from Fortran objects with the `POINTER` attribute. Thus one needs to use intrinsic conversion procedures to convert from or to C pointers. For some applications, using an assumed type (`TYPE(*)`) can be an alternative to a C pointer, and you can also use library routines to access Fortran pointers from C. See Section 6.1.6 [Further Interoperability of Fortran with C], page 78.

Here is an example of using C pointers in Fortran:

```
use iso_c_binding
type(c_ptr) :: cptr1, cptr2
integer, target :: array(7), scalar
integer, pointer :: pa(:), ps
cptr1 = c_loc(array(1)) ! The programmer needs to ensure that the
                        ! array is contiguous if required by the C
                        ! procedure
cptr2 = c_loc(scalar)
call c_f_pointer(cptr2, ps)
call c_f_pointer(cptr2, pa, shape=[7])
```

When converting C to Fortran arrays, the one-dimensional `SHAPE` argument has to be passed.

If a pointer is a dummy argument of an interoperable procedure, it usually has to be declared using the `VALUE` attribute. `void*` matches `TYPE(C_PTR)`, `VALUE`, while `TYPE(C_PTR)` alone matches `void**`.

Procedure pointers are handled analogously to pointers; the C type is `TYPE(C_FUNPTR)` and the intrinsic conversion procedures are `C_F_PROCPTR` and `C_FUNLOC`.

Let us consider two examples of actually passing a procedure pointer from C to Fortran and vice versa. Note that these examples are also very similar to passing ordinary pointers between both languages. First, consider this code in C:

```
/* Procedure implemented in Fortran. */
void get_values (void (*)(double));

/* Call-back routine we want called from Fortran. */
void
print_it (double x)
{
    printf ("Number is %f.\n", x);
}

/* Call Fortran routine and pass call-back to it. */
void
foobar ()
{
    get_values (&print_it);
}
```

A matching implementation for `get_values` in Fortran that correctly receives the procedure pointer from C and is able to call it, is given in the following `MODULE`:

```
MODULE m
```

```

IMPLICIT NONE

! Define interface of call-back routine.
ABSTRACT INTERFACE
  SUBROUTINE callback (x)
    USE, INTRINSIC :: ISO_C_BINDING
    REAL(KIND=C_DOUBLE), INTENT(IN), VALUE :: x
  END SUBROUTINE callback
END INTERFACE

CONTAINS

! Define C-bound procedure.
SUBROUTINE get_values (cproc) BIND(C)
  USE, INTRINSIC :: ISO_C_BINDING
  TYPE(C_FUNPTR), INTENT(IN), VALUE :: cproc

  PROCEDURE(callback), POINTER :: proc

  ! Convert C to Fortran procedure pointer.
  CALL C_F_PROCPTR (cproc, proc)

  ! Call it.
  CALL proc (1.0_C_DOUBLE)
  CALL proc (-42.0_C_DOUBLE)
  CALL proc (18.12_C_DOUBLE)
END SUBROUTINE get_values

END MODULE m

```

Next, we want to call a C routine that expects a procedure pointer argument and pass it a Fortran procedure (that clearly must be interoperable!). Again, the C function may be:

```

int
call_it (int (*func)(int), int arg)
{
  return func (arg);
}

```

It can be used as in the following Fortran code:

```

MODULE m
  USE, INTRINSIC :: ISO_C_BINDING
  IMPLICIT NONE

  ! Define interface of C function.
  INTERFACE
    INTEGER(KIND=C_INT) FUNCTION call_it (func, arg) BIND(C)
      USE, INTRINSIC :: ISO_C_BINDING
      TYPE(C_FUNPTR), INTENT(IN), VALUE :: func
      INTEGER(KIND=C_INT), INTENT(IN), VALUE :: arg
    END FUNCTION call_it
  END INTERFACE

CONTAINS

! Define procedure passed to C function.
! It must be interoperable!
INTEGER(KIND=C_INT) FUNCTION double_it (arg) BIND(C)
  INTEGER(KIND=C_INT), INTENT(IN), VALUE :: arg
  double_it = arg + arg

```

```

END FUNCTION double_it

! Call C function.
SUBROUTINE foobar ()
  TYPE(C_FUNPTR) :: cproc
  INTEGER(KIND=C_INT) :: i

  ! Get C procedure pointer.
  cproc = C_FUNLOC (double_it)

  ! Use it.
  DO i = 1_C_INT, 10_C_INT
    PRINT *, call_it (cproc, i)
  END DO
END SUBROUTINE foobar

END MODULE m

```

6.1.6 Further Interoperability of Fortran with C

GNU Fortran implements the Technical Specification ISO/IEC TS 29113:2012, which extends the interoperability support of Fortran 2003 and Fortran 2008 and is now part of the 2018 Fortran standard. Besides removing some restrictions and constraints, the Technical Specification adds assumed-type (`TYPE(*)`) and assumed-rank (`DIMENSION(...)`) variables and allows for interoperability of assumed-shape, assumed-rank, and deferred-shape arrays, as well as allocatables and pointers. Objects of these types are passed to `BIND(C)` functions as descriptors with a standard interface, declared in the header file `<ISO_Fortran_binding.h>`.

Note: Currently, GNU Fortran does not use internally the array descriptor (dope vector) as specified in the Technical Specification, but uses an array descriptor with different fields in functions without the `BIND(C)` attribute. Arguments to functions marked `BIND(C)` are converted to the specified form. If you need to access GNU Fortran's internal array descriptor, you can use the Chasm Language Interoperability Tools, <http://chasm-interop.sourceforge.net/>.

6.1.7 Generating C prototypes from Fortran

The options `-fc-prototypes` can be used to write out C declarations and function prototypes for `BIND(C)` entities. The same can be done for writing out prototypes for external procedures using `-fc-prototypes-external`, see Section 2.11 [Interoperability Options], page 34.

Standard Fortran does not specify an interoperable type for C's `unsigned` integer types. For interoperability with unsigned types, GNU Fortran provides unsigned integers, see Section 5.1.34 [Unsigned integers], page 67.

6.2 GNU Fortran Compiler Directives

6.2.1 ATTRIBUTES directive

The Fortran standard describes how a conforming program shall behave; however, the exact implementation is not standardized. In order to allow the user to choose specific implementation details, compiler directives can be used to set attributes of variables and

procedures that are not part of the standard. Whether a given attribute is supported and its exact effects depend on both the operating system and on the processor; see Section “C Extensions” in *Using the GNU Compiler Collection (GCC)* for details.

For procedures and procedure pointers, the following attributes can be used to change the calling convention:

- **CDECL** – standard C calling convention
- **STDCALL** – convention where the called procedure pops the stack
- **FASTCALL** – part of the arguments are passed via registers instead using the stack

Besides changing the calling convention, the attributes also influence the decoration of the symbol name, e.g., by a leading underscore or by a trailing at-sign followed by the number of bytes on the stack. When assigning a procedure to a procedure pointer, both should use the same calling convention.

On some systems, procedures and global variables (module variables and **COMMON** blocks) need special handling to be accessible when they are in a shared library. The following attributes are available:

- **DLEXPOR** – provide a global pointer to a pointer in the DLL
- **DLLIMPORT** – reference the function or variable using a global pointer

For dummy arguments, the **NO_ARG_CHECK** attribute can be used; in other compilers, it is also known as **IGNORE_TKR**. For dummy arguments with this attribute actual arguments of any type and kind (similar to **TYPE(*)**), scalars and arrays of any rank (no equivalent in Fortran standard) are accepted. As with **TYPE(*)**, the argument is unlimited polymorphic and no type information is available. Additionally, the argument may only be passed to dummy arguments with the **NO_ARG_CHECK** attribute and as argument to the **PRESENT** intrinsic function and to **C_LOC** of the **ISO_C_BINDING** module.

Variables with **NO_ARG_CHECK** attribute shall be of assumed-type (**TYPE(*)**; recommended) or of type **INTEGER**, **LOGICAL**, **REAL** or **COMPLEX**. They shall not have the **ALLOCATE**, **CODIMENSION**, **INTENT(OUT)**, **POINTER** or **VALUE** attribute; furthermore, they shall be either scalar or of assumed-size (**dimension(*)**). As **TYPE(*)**, the **NO_ARG_CHECK** attribute requires an explicit interface.

- **NO_ARG_CHECK** – disable the type, kind and rank checking
- **DEPRECATED** – print a warning when using a such-tagged deprecated procedure, variable or parameter; the warning can be suppressed with **-Wno-deprecated-declarations**.
- **NOINLINE** – prevent inlining given function.
- **NORETURN** – add a hint that a given function cannot return.
- **WEAK** – emit the declaration of an external symbol as a weak symbol rather than a global. This is primarily useful in defining library functions that can be overridden in user code, though it can also be used with non-function declarations. The overriding symbol must have the same type as the weak symbol.

The attributes are specified using the syntax

```
!GCC$ ATTRIBUTES attribute-list :: variable-list
```

where in free-form source code only whitespace is allowed before **!GCC\$** and in fixed-form source code **!GCC\$**, **cGCC\$** or ***GCC\$** shall start in the first column.

For procedures, the compiler directives shall be placed into the body of the procedure; for variables and procedure pointers, they shall be in the same declaration part as the variable or procedure pointer.

6.2.2 UNROLL directive

The syntax of the directive is

```
!GCC$ unroll N
```

You can use this directive to control how many times a loop should be unrolled. It must be placed immediately before a **DO** loop and applies only to the loop that follows. *N* is an integer constant specifying the unrolling factor. The values of 0 and 1 block any unrolling of the loop.

For **DO CONCURRENT** constructs the unrolling specification applies only to the first loop control variable.

6.2.3 BUILTIN directive

The syntax of the directive is

```
!GCC$ BUILTIN (B) attributes simd FLAGS IF('target')
```

You can use this directive to define which middle-end built-ins provide vector implementations. *B* is name of the middle-end built-in. *FLAGS* are optional and must be either (**inbranch**) or (**notinbranch**). **IF** statement is optional and is used to filter multilib ABIs for the built-in that should be vectorized. Example usage:

```
!GCC$ builtin (sinf) attributes simd (notinbranch) if('x86_64')
```

The purpose of the directive is to provide an API among the GCC compiler and the GNU C Library which would define vector implementations of math routines.

6.2.4 IVDEP directive

The syntax of the directive is

```
!GCC$ ivdep
```

This directive tells the compiler to ignore vector dependencies in the following loop. It must be placed immediately before a **DO** loop and applies only to the loop that follows.

Sometimes the compiler may not have sufficient information to decide whether a particular loop is vectorizable due to potential dependencies between iterations. The purpose of the directive is to tell the compiler that vectorization is safe.

For **DO CONCURRENT** constructs this annotation is implicit to all loop control variables.

This directive is intended for annotation of existing code. For new code it is recommended to consider OpenMP SIMD directives as potential alternative.

6.2.5 VECTOR directive

The syntax of the directive is

```
!GCC$ vector
```

This directive tells the compiler to vectorize the following loop. It must be placed immediately before a **DO** loop and applies only to the loop that follows.

For **DO CONCURRENT** constructs this annotation applies to all loops specified in the concurrent header.

6.2.6 NOVECTOR directive

The syntax of the directive is

```
!GCC$ novector
```

This directive tells the compiler to not vectorize the following loop. It must be placed immediately before a DO loop and applies only to the loop that follows.

For DO CONCURRENT constructs this annotation applies to all loops specified in the concurrent header.

6.3 Non-Fortran Main Program

Even if you are doing mixed-language programming, it is very likely that you do not need to know or use the information in this section. Since it is about the internal structure of GNU Fortran, it may also change in GCC minor releases.

When you compile a PROGRAM with GNU Fortran, a function with the name `main` (in the symbol table of the object file) is generated, which initializes the libgfortran library and then calls the actual program that uses the name `MAIN__`, for historic reasons. If you link GNU Fortran compiled procedures to, e.g., a C or C++ program or to a Fortran program compiled by a different compiler, the libgfortran library is not initialized and thus a few intrinsic procedures do not work properly, e.g. those for obtaining the command-line arguments.

Therefore, if your PROGRAM is not compiled with GNU Fortran and the GNU Fortran compiled procedures require intrinsics relying on the library initialization, you need to initialize the library yourself. Using the default options, gfortran calls `_gfortran_set_args` and `_gfortran_set_options`. The initialization of the former is needed if the called procedures access the command line (and for backtracing); the latter sets some flags based on the standard chosen or to enable backtracing. In typical programs, it is not necessary to call any initialization function.

If your PROGRAM is compiled with GNU Fortran, you shall not call any of the following functions. The libgfortran initialization functions are shown in C syntax but using C bindings they are also accessible from Fortran.

6.3.1 `_gfortran_set_args` — Save command-line arguments

Synopsis: `void _gfortran_set_args (int argc, char *argv[])`

Description:

`_gfortran_set_args` saves the command-line arguments; this initialization is required if any of the command-line intrinsics is called. Additionally, it shall be called if backtracing is enabled (see `_gfortran_set_options`).

Arguments:

<code>argc</code>	number of command line argument strings
<code>argv</code>	the command-line argument strings; <code>argv[0]</code> is the pathname of the executable itself.

Example:

```
int main (int argc, char *argv[])
{
```

```

    /* Initialize libgfortran. */
    _gfortran_set_args (argc, argv);
    return 0;
}

```

6.3.2 `_gfortran_set_options` — Set library option flags

Synopsis: `void _gfortran_set_options (int num, int options[])`

Description:

`_gfortran_set_options` sets several flags related to the Fortran standard to be used, whether backtracing should be enabled and whether range checks should be performed. The syntax allows for upward compatibility since the number of passed flags is specified; for non-passed flags, the default value is used. See also see Section 2.10 [Code Gen Options], page 26. Please note that not all flags are actually used.

Arguments:

<code>num</code>	number of options passed
<code>argv</code>	The list of flag values

option flag list:

<code>option[0]</code>	Allowed standard; can give run-time errors if e.g. an input-output edit descriptor is invalid in a given standard. Possible values are (bitwise or-ed) <code>GFC_STD_F77</code> (1), <code>GFC_STD_F95_OBS</code> (2), <code>GFC_STD_F95_DEL</code> (4), <code>GFC_STD_F95</code> (8), <code>GFC_STD_F2003</code> (16), <code>GFC_STD_GNU</code> (32), <code>GFC_STD_LEGACY</code> (64), <code>GFC_STD_F2008</code> (128), <code>GFC_STD_F2008_OBS</code> (256), <code>GFC_STD_F2018</code> (512), <code>GFC_STD_F2018_OBS</code> (1024), <code>GFC_STD_F2018_DEL</code> (2048), <code>GFC_STD_F2023</code> (4096), and <code>GFC_STD_F2023_DEL</code> (8192). Default: <code>GFC_STD_F95_OBS GFC_STD_F95_DEL GFC_STD_F95 GFC_STD_F2003 GFC_STD_F2008 GFC_STD_F2008_OBS GFC_STD_F77 GFC_STD_F2018 GFC_STD_F2018_OBS GFC_STD_F2018_DEL GFC_STD_F2023 GFC_STD_F2023_DEL GFC_STD_GNU GFC_STD_LEGACY</code> .
<code>option[1]</code>	Standard-warning flag; prints a warning to standard error. Default: <code>GFC_STD_F95_DEL GFC_STD_LEGACY</code> .
<code>option[2]</code>	If non zero, enable pedantic checking. Default: off.
<code>option[3]</code>	Unused.
<code>option[4]</code>	If non zero, enable backtracing on run-time errors. Default: off. (Default in the compiler: on.) Note: Installs a signal handler and requires command-line initialization using <code>_gfortran_set_args</code> .
<code>option[5]</code>	If non zero, supports signed zeros. Default: enabled.

option[6] Enables run-time checking. Possible values are (bitwise or-ed): GFC_RTCKECK_BOUNDS (1), GFC_RTCKECK_ARRAY_TEMPS (2), GFC_RTCKECK_RECURSION (4), GFC_RTCKECK_DO (8), GFC_RTCKECK_POINTER (16), GFC_RTCKECK_MEM (32), GFC_RTCKECK_BITS (64). Default: disabled.

option[7] Unused.

option[8] Show a warning when invoking STOP and ERROR STOP if a floating-point exception occurred. Possible values are (bitwise or-ed) GFC_FPE_INVALID (1), GFC_FPE_DENORMAL (2), GFC_FPE_ZERO (4), GFC_FPE_OVERFLOW (8), GFC_FPE_UNDERFLOW (16), GFC_FPE_INEXACT (32). Default: None (0). (Default in the compiler: GFC_FPE_INVALID | GFC_FPE_DENORMAL | GFC_FPE_ZERO | GFC_FPE_OVERFLOW | GFC_FPE_UNDERFLOW.)

Example:

```
/* Use gfortran 4.9 default options. */
static int options[] = {68, 511, 0, 0, 1, 1, 0, 0, 31};
_gfortran_set_options (9, &options);
```

6.3.3 `_gfortran_set_convert` — Set endian conversion

Synopsis: `void _gfortran_set_convert (int conv)`

Description:

`_gfortran_set_convert` set the representation of data for unformatted files.

Arguments:

conv Endian conversion, possible values:
 GFC_CONVERT_NATIVE (0, default),
 GFC_CONVERT_SWAP (1), GFC_CONVERT_BIG
 (2), GFC_CONVERT_LITTLE (3).

Example:

```
int main (int argc, char *argv[])
{
  /* Initialize libgfortran. */
  _gfortran_set_args (argc, argv);
  _gfortran_set_convert (1);
  return 0;
}
```

6.3.4 `_gfortran_set_record_marker` — Set length of record markers

Synopsis: `void _gfortran_set_record_marker (int val)`

Description:

`_gfortran_set_record_marker` sets the length of record markers for unformatted files.

Arguments:

`val` Length of the record marker; valid values are 4 and 8. Default is 4.

Example:

```
int main (int argc, char *argv[])
{
    /* Initialize libgfortran. */
    _gfortran_set_args (argc, argv);
    _gfortran_set_record_marker (8);
    return 0;
}
```

6.3.5 `_gfortran_set_fpe` — Enable floating point exception traps

Synopsis: `void _gfortran_set_fpe (int val)`

Description:

`_gfortran_set_fpe` enables floating point exception traps for the specified exceptions. On most systems, this results in a SIGFPE signal being sent and the program being aborted.

Arguments:

`option[0]` IEEE exceptions. Possible values are (bitwise or-ed) zero (0, default) no trapping, `GFC_FPE_INVALID` (1), `GFC_FPE_DENORMAL` (2), `GFC_FPE_ZERO` (4), `GFC_FPE_OVERFLOW` (8), `GFC_FPE_UNDERFLOW` (16), and `GFC_FPE_INEXACT` (32).

Example:

```
int main (int argc, char *argv[])
{
    /* Initialize libgfortran. */
    _gfortran_set_args (argc, argv);
    /* FPE for invalid operations such as SQRT(-1.0). */
    _gfortran_set_fpe (1);
    return 0;
}
```

6.3.6 `_gfortran_set_max_subrecord_length` — Set subrecord length

Synopsis: `void _gfortran_set_max_subrecord_length (int val)`

Description:

`_gfortran_set_max_subrecord_length` set the maximum length for a subrecord. This option only makes sense for testing and debugging of unformatted I/O.

Arguments:

`val` the maximum length for a subrecord; the maximum permitted value is 2147483639, which is also the default.

Example:

```
int main (int argc, char *argv[])
{
    /* Initialize libgfortran. */
    _gfortran_set_args (argc, argv);
    _gfortran_set_max_subrecord_length (8);
    return 0;
}
```

6.4 Naming and argument-passing conventions

This section gives an overview about the naming convention of procedures and global variables and about the argument passing conventions used by GNU Fortran. If a C binding has been specified, the naming convention and some of the argument-passing conventions change. If possible, mixed-language and mixed-compiler projects should use the better defined C binding for interoperability. See Section 6.1 [Interoperability with C], page 73.

6.4.1 Naming conventions

According the Fortran standard, valid Fortran names consist of a letter between A to Z, a to z, digits 0, 1 to 9 and underscores (_) with the restriction that names may only start with a letter. As vendor extension, the dollar sign (\$) is additionally permitted with the option `-fdollar-ok`, but not as first character and only if the target system supports it.

By default, the procedure name is the lowercased Fortran name with an appended underscore (_); using `-fno-underscoring` no underscore is appended while `-fsecond-underscore` appends two underscores. Depending on the target system and the calling convention, the procedure might be additionally dressed; for instance, on 32bit Windows with `stdcall`, an at-sign @ followed by an integer number is appended. For the changing the calling convention, see Section 6.2 [GNU Fortran Compiler Directives], page 78.

For common blocks, the same convention is used, i.e. by default an underscore is appended to the lowercased Fortran name. Blank commons have the name `__BLNK__`.

For procedures and variables declared in the specification space of a module, the name is formed by `__`, followed by the lowercased module name, `_MOD_`, and the lowercased Fortran name. Note that no underscore is appended.

6.4.2 Argument passing conventions

Subroutines do not return a value (matching C99's `void`) while functions either return a value as specified in the platform ABI or the result variable is passed as hidden argument to the function and no result is returned. A hidden result variable is used when the result variable is an array or of type `CHARACTER`.

Arguments are passed according to the platform ABI. In particular, complex arguments might not be compatible to a struct with two real components for the real and imaginary part. The argument passing matches the one of C99's `_Complex`. Functions with scalar complex result variables return their value and do not use a by-reference argument. Note that with the `-ff2c` option, the argument passing is modified and no longer completely matches the platform ABI. Some other Fortran compilers use `f2c` semantic by default; this might cause problems with interoperability.

GNU Fortran passes most arguments by reference, i.e. by passing a pointer to the data. Note that the compiler might use a temporary variable into which the actual argument has been copied, if required semantically (copy-in/copy-out).

For arguments with `ALLOCATABLE` and `POINTER` attribute (including procedure pointers), a pointer to the pointer is passed such that the pointer address can be modified in the procedure.

For dummy arguments with the `VALUE` attribute: Scalar arguments of the type `INTEGER`, `LOGICAL`, `REAL` and `COMPLEX` are passed by value according to the platform ABI. (As vendor extension and not recommended, using `%VAL()` in the call to a procedure has the same effect.) For `TYPE(C_PTR)` and procedure pointers, the pointer itself is passed such that it can be modified without affecting the caller.

For Boolean (`LOGICAL`) arguments, please note that GCC expects only the integer value 0 and 1. If a GNU Fortran `LOGICAL` variable contains another integer value, the result is undefined. As some other Fortran compilers use `-1` for `.TRUE.`, extra care has to be taken – such as passing the value as `INTEGER`. (The same value restriction also applies to other front ends of GCC, e.g. to GCC's C99 compiler for `_Bool` or GCC's Ada compiler for `Boolean`.)

For arguments of `CHARACTER` type, the character length is passed as a hidden argument at the end of the argument list, except when the corresponding dummy argument is declared as `TYPE(*)`. For deferred-length strings, the value is passed by reference, otherwise by value. The character length has the C type `size_t` (or `INTEGER(kind=C_SIZE_T)` in Fortran). Note that this is different to older versions of the GNU Fortran compiler, where the type of the hidden character length argument was a C `int`. In order to retain compatibility with older versions, one can e.g. for the following Fortran procedure

```
subroutine fstrlen (s, a)
  character(len=*) :: s
  integer :: a
  print*, len(s)
end subroutine fstrlen
```

define the corresponding C prototype as follows:

```
#if __GNUC__ > 7
typedef size_t fortran_charlen_t;
#else
typedef int fortran_charlen_t;
#endif

void fstrlen_ (char*, int*, fortran_charlen_t);
```

In order to avoid such compiler-specific details, for new code it is instead recommended to use the `ISO_C_BINDING` feature.

Note with C binding, `CHARACTER(len=1)` result variables are returned according to the platform ABI and no hidden length argument is used for dummy arguments; with `VALUE`, those variables are passed by value.

For `OPTIONAL` dummy arguments, an absent argument is denoted by a `NULL` pointer, except for scalar dummy arguments of intrinsic type or derived type (but not `CLASS`) and that have the `VALUE` attribute. For those, a hidden Boolean argument (`logical(kind=C_bool), value`) is used to indicate whether the argument is present.

Arguments that are assumed-shape, assumed-rank or deferred-rank arrays or, with `-fcoarray=lib`, allocatable scalar coarrays use an array descriptor. All other arrays pass

the address of the first element of the array. With `-fcoarray=lib`, the token and the offset belonging to nonallocatable coarrays dummy arguments are passed as hidden argument along the character length hidden arguments. The token is an opaque pointer identifying the coarray and the offset is a passed-by-value integer of kind `C_PTRDIFF_T`, denoting the byte offset between the base address of the coarray and the passed scalar or first element of the passed array.

The arguments are passed in the following order

- Result variable, when the function result is passed by reference
- Character length of the function result, if it is a of type `CHARACTER` and no C binding is used
- The arguments in the order in which they appear in the Fortran declaration
- The present status for optional arguments with value attribute, which are internally passed by value
- The character length and/or coarray token and offset for the first argument which is a `CHARACTER` or a nonallocatable coarray dummy argument, followed by the hidden arguments of the next dummy argument of such a type

7 Coarray Programming

7.1 Type and enum ABI Documentation

7.1.1 caf_token_t

Typedef of type `void *` on the compiler side. Can be any data type on the library side.

7.1.2 caf_register_t

Indicates which kind of coarray variable should be registered.

```
typedef enum caf_register_t {
    CAF_REGTYPE_COARRAY_STATIC,
    CAF_REGTYPE_COARRAY_ALLOC,
    CAF_REGTYPE_LOCK_STATIC,
    CAF_REGTYPE_LOCK_ALLOC,
    CAF_REGTYPE_CRITICAL,
    CAF_REGTYPE_EVENT_STATIC,
    CAF_REGTYPE_EVENT_ALLOC,
    CAF_REGTYPE_COARRAY_ALLOC_REGISTER_ONLY,
    CAF_REGTYPE_COARRAY_ALLOC_ALLOCATE_ONLY
}
caf_register_t;
```

The values `CAF_REGTYPE_COARRAY_ALLOC_REGISTER_ONLY` and `CAF_REGTYPE_COARRAY_ALLOC_ALLOCATE_ONLY` are for allocatable components in derived type coarrays only. The first one sets up the token without allocating memory for allocatable component. The latter one only allocates the memory for an allocatable component in a derived type coarray. The token needs to be set up previously by the `REGISTER_ONLY`. This allows having allocatable components unallocated on some images. The status of whether an allocatable component is allocated on a remote image can be queried by `_caf_is_present` which used internally by the `ALLOCATED` intrinsic.

7.1.3 caf_deregister_t

```
typedef enum caf_deregister_t {
    CAF_DEREGTYPE_COARRAY_DEREGISTER,
    CAF_DEREGTYPE_COARRAY DEALLOCATE_ONLY
}
caf_deregister_t;
```

Allows to specify the type of deregistration of a coarray object. The `CAF_DEREGTYPE_COARRAY DEALLOCATE_ONLY` flag is only allowed for allocatable components in derived type coarrays.

7.1.4 caf_reference_t

The structure used for implementing arbitrary reference chains. A `CAF_REFERENCE_T` allows to specify a component reference or any kind of array reference of any rank supported by gfortran. For array references all kinds as known by the compiler/Fortran standard are supported indicated by a `MODE`.

```
typedef enum caf_ref_type_t {
    /* Reference a component of a derived type, either regular one or an
       allocatable or pointer type. For regular ones idx in caf_reference_t is
```

```

        set to -1. */
CAF_REF_COMPONENT,
/* Reference an allocatable array. */
CAF_REF_ARRAY,
/* Reference a non-allocatable/non-pointer array. I.e., the coarray object
   has no array descriptor associated and the addressing is done
   completely using the ref. */
CAF_REF_STATIC_ARRAY
} caf_ref_type_t;

typedef enum caf_array_ref_t {
/* No array ref. This terminates the array ref. */
CAF_ARR_REF_NONE = 0,
/* Reference array elements given by a vector. Only for this mode
   caf_reference_t.u.a.dim[i].v is valid. */
CAF_ARR_REF_VECTOR,
/* A full array ref (:). */
CAF_ARR_REF_FULL,
/* Reference a range on elements given by start, end and stride. */
CAF_ARR_REF_RANGE,
/* Only a single item is referenced given in the start member. */
CAF_ARR_REF_SINGLE,
/* An array ref of the kind (i:), where i is an arbitrary valid index in the
   array. The index i is given in the start member. */
CAF_ARR_REF_OPEN_END,
/* An array ref of the kind (:i), where the lower bound of the array ref
   is given by the remote side. The index i is given in the end member. */
CAF_ARR_REF_OPEN_START
} caf_array_ref_t;

/* References to remote components of a derived type. */
typedef struct caf_reference_t {
/* A pointer to the next ref or NULL. */
struct caf_reference_t *next;
/* The type of the reference. */
/* caf_ref_type_t, replaced by int to allow specification in fortran FE. */
int type;
/* The size of an item referenced in bytes. I.e. in an array ref this is
   the factor to advance the array pointer with to get to the next item.
   For component refs this gives just the size of the element referenced. */
size_t item_size;
union {
    struct {
        /* The offset (in bytes) of the component in the derived type.
           Unused for allocatable or pointer components. */
        ptrdiff_t offset;
        /* The offset (in bytes) to the caf_token associated with this
           component. NULL, when not allocatable/pointer ref. */
        ptrdiff_t caf_token_offset;
    } c;
    struct {
        /* The mode of the array ref. See CAF_ARR_REF*. */
        /* caf_array_ref_t, replaced by unsigned char to allow specification in
           fortran FE. */
        unsigned char mode[GFC_MAX_DIMENSIONS];
        /* The type of a static array. Unset for array's with descriptors. */
        int static_array_type;
        /* Subscript refs (s) or vector refs (v). */
        union {
            struct {

```

```

        /* The start and end boundary of the ref and the stride. */
        index_type start, end, stride;
    } s;
    struct {
        /* nvec entries of kind giving the elements to reference. */
        void *vector;
        /* The number of entries in vector. */
        size_t nvec;
        /* The integer kind used for the elements in vector. */
        int kind;
    } v;
    } dim[GFC_MAX_DIMENSIONS];
    } a;
    } u;
} caf_reference_t;

```

The references make up a single linked list of reference operations. The `NEXT` member links to the next reference or `NULL` to indicate the end of the chain. Component and array refs can be arbitrarily mixed as long as they comply to the Fortran standard.

Notes: The member `STATIC_ARRAY_TYPE` is used only when the `TYPE` is `CAF_REF_STATIC_ARRAY`. The member gives the type of the data referenced. Because no array descriptor is available for a descriptorless array and type conversion still needs to take place the type is transported here.

At the moment `CAF_ARR_REF_VECTOR` is not implemented in the front end for descriptorless arrays. The library `caf_single` has untested support for it.

7.1.5 `caf_team_t`

Opaque pointer to represent a team-handle. This type is a stand-in for the future implementation of teams. It is about to change without further notice.

7.2 Function ABI Documentation

7.2.1 `_gfortran_caf_init` — Initialization function

Synopsis: `void _gfortran_caf_init (int *argc, char ***argv)`

Description:

This function is called at startup of the program before the Fortran main program, if the latter has been compiled with `-fcoarray=lib`. It takes as arguments the command-line arguments of the program. It is permitted to pass two `NULL` pointers as argument; if non-`NULL`, the library is permitted to modify the arguments.

Arguments:

<code>argc</code>	intent(inout) An integer pointer with the number of arguments passed to the program or <code>NULL</code> .
<code>argv</code>	intent(inout) A pointer to an array of strings with the command-line arguments or <code>NULL</code> .

Notes: The function is modelled after the initialization function of the Message Passing Interface (MPI) specification. Due to the way coarray registration works, it might not be the first call to the library. If the main program is not written

in Fortran and only a library uses coarrays, it can happen that this function is never called. Therefore, it is recommended that the library does not rely on the passed arguments and whether the call has been done.

7.2.2 `_gfortran_caf_finish` — Finalization function

Synopsis: `void _gfortran_caf_finish (void)`

Description:

This function is called at the end of the Fortran main program, if it has been compiled with the `-fcoarray=lib` option.

Notes: For non-Fortran programs, it is recommended to call the function at the end of the main program. To ensure that the shutdown is also performed for programs where this function is not explicitly invoked, for instance non-Fortran programs or calls to the system's `exit()` function, the library can use a destructor function. Note that programs can also be terminated using the `STOP` and `ERROR STOP` statements; those use different library calls.

7.2.3 `_gfortran_caf_this_image` — Querying the image number

Synopsis: `int _gfortran_caf_this_image (caf_team_t team)`

Description:

Return the current image number in the *team*, or in the current team, if no *team* is given.

Arguments:

team intent(in), optional; The team this image's number is requested for. If null, the image number in the current team is returned.

Notes: Available since Fortran 2008 without argument; Since Fortran 2018 with optional team argument. Fortran 2008 uses 0 as argument for team, which is permissible, because a team handle is always an opaque pointer, which as a special case can be null here.

7.2.4 `_gfortran_caf_num_images` — Querying the maximal number of images

Synopsis: `int _gfortran_caf_num_images (caf_team_t team, int32_t *team_number)`

Description:

This function returns the number of images in the team given by *team* or *team_number*, if either one is present. If both are null, then the number of images in the current team is returned.

Arguments:

team intent(in), optional; The team the number of images is requested for. If null, the number of images in the current team is returned.

VALUES and *TIME* are *INTENT(OUT)* and provide the following:

VALUES(1): User time in seconds.
VALUES(2): System time in seconds.
TIME: Run time since start in seconds.

Class: Subroutine, function

Arguments:

VALUES The type shall be *REAL*(4), *DIMENSION*(2).
TIME The type shall be *REAL*(4).

Return value:

Elapsed time in seconds since the last invocation or since the start of program execution if not called before.

Example:

```
program test_dtime
  integer(8) :: i, j
  real, dimension(2) :: tarray
  real :: result
  call dtime(tarray, result)
  print *, result
  print *, tarray(1)
  print *, tarray(2)
  do i=1,100000000 ! Just a delay
    j = i * i - i
  end do
  call dtime(tarray, result)
  print *, result
  print *, tarray(1)
  print *, tarray(2)
end program test_dtime
```

Standard: GNU extension

See also: Section 8.84 [CPU_TIME], page 175,

8.98 EOSHIFT — End-off shift elements of an array

Synopsis: *RESULT* = *EOSHIFT*(*ARRAY*, *SHIFT* [, *BOUNDARY*, *DIM*])

Description:

EOSHIFT(*ARRAY*, *SHIFT* [, *BOUNDARY*, *DIM*]) performs an end-off shift on elements of *ARRAY* along the dimension of *DIM*. If *DIM* is omitted it is taken to be 1. *DIM* is a scalar of type *INTEGER* in the range of $1 \leq DIM \leq n$ where *n* is the rank of *ARRAY*. If the rank of *ARRAY* is one, then all elements of *ARRAY* are shifted by *SHIFT* places. If rank is greater than one, then all complete rank one sections of *ARRAY* along the given dimension are shifted. Elements shifted out one end of each rank one section are dropped. If *BOUNDARY* is present then the corresponding value of from *BOUNDARY* is copied back in the other end. If *BOUNDARY* is not present then the following are copied in depending on the type of *ARRAY*.

<i>Array Type</i>	<i>Boundary Value</i>
Numeric	0 of the type and kind of <i>ARRAY</i> .

Logical .FALSE..
 Character(*len*)*len* blanks.

Class: Transformational function

Arguments:

ARRAY May be any type, not scalar.
SHIFT The type shall be INTEGER.
BOUNDARY Same type as *ARRAY*.
DIM The type shall be INTEGER.

Notes: *ARRAY* can also be UNSIGNED.

Return value:

Returns an array of same type and rank as the *ARRAY* argument.

Example:

```
program test_eoshift
  integer, dimension(3,3) :: a
  a = reshape( (/ 1, 2, 3, 4, 5, 6, 7, 8, 9 /), (/ 3, 3 /))
  print '(3i3)', a(1,:)
  print '(3i3)', a(2,:)
  print '(3i3)', a(3,:)
  a = EOSHIFT(a, SHIFT=(/1, 2, 1/), BOUNDARY=-5, DIM=2)
  print *
  print '(3i3)', a(1,:)
  print '(3i3)', a(2,:)
  print '(3i3)', a(3,:)
end program test_eoshift
```

Standard: Fortran 90 and later, extension for UNSIGNED (see Section 5.1.34 [Unsigned integers], page 67)

8.99 EPSILON — Epsilon function

Synopsis: RESULT = EPSILON(X)

Description:

EPSILON(X) returns the smallest number *E* of the same kind as *X* such that $1 + E > 1$.

Class: Inquiry function

Arguments:

X The type shall be REAL.

Return value:

The return value is of same type as the argument.

Example:

```
program test_epsilon
  real :: x = 3.143
  real(8) :: y = 2.33
  print *, EPSILON(x)
  print *, EPSILON(y)
end program test_epsilon
```

Standard: Fortran 90 and later

8.100 ERF — Error function

Synopsis: `RESULT = ERF(X)`

Description:

`ERF(X)` computes the error function of X .

Class: Elemental function

Arguments:

X The type shall be `REAL`.

Return value:

The return value is of type `REAL`, of the same kind as X and lies in the range $-1 \leq \text{erf}(x) \leq 1$.

Example:

```
program test_erf
  real(8) :: x = 0.17_8
  x = erf(x)
end program test_erf
```

Specific names:

Name	Argument	Return type	Standard
<code>DERF(X)</code>	<code>REAL(8) X</code>	<code>REAL(8)</code>	GNU extension

Standard: Fortran 2008 and later

8.101 ERFC — Error function

Synopsis: `RESULT = ERFC(X)`

Description:

`ERFC(X)` computes the complementary error function of X .

Class: Elemental function

Arguments:

X The type shall be `REAL`.

Return value:

The return value is of type `REAL` and of the same kind as X . It lies in the range $0 \leq \text{erfc}(x) \leq 2$.

Example:

```
program test_erfc
  real(8) :: x = 0.17_8
  x = erfc(x)
end program test_erfc
```

Specific names:

Name	Argument	Return type	Standard
<code>DERFC(X)</code>	<code>REAL(8) X</code>	<code>REAL(8)</code>	GNU extension

Standard: Fortran 2008 and later


```

      INTEGER, PARAMETER :: SGL = SELECTED_REAL_KIND(p=6, r=37)
      INTEGER, PARAMETER :: DBL = SELECTED_REAL_KIND(p=13, r=200)

      WRITE(*,*) spacing(1.0_SGL)      ! "1.1920929E-07"      on i686
      WRITE(*,*) spacing(1.0_DBL)      ! "2.220446049250313E-016" on i686
END PROGRAM

```

Standard: Fortran 90 and later

See also: Section 8.239 [RRSPACING], page 274,

8.266 SPLIT — Parse a string into tokens, one at a time

Synopsis: RESULT = SPLIT(STRING, SET, POS [, BACK])

Description:

Updates the integer *POS* to the position of the next (or previous) separator in *STRING*.

If *BACK* is absent or is present with the value false, *POS* is assigned the position of the leftmost token delimiter in *STRING* whose position is greater than *POS*, or if there is no such character, it is assigned a value one greater than the length of *STRING*. This identifies a token with starting position one greater than the value of *POS* on invocation, and ending position one less than the value of *POS* on return.

If *BACK* is present with the value true, *POS* is assigned the position of the rightmost token delimiter in *STRING* whose position is less than *POS*, or if there is no such character, it is assigned the value zero. This identifies a token with ending position one less than the value of *POS* on invocation, and starting position one greater than the value of *POS* on return.

Class: Subroutine

Arguments:

<i>STRING</i>	Shall be of type CHARACTER.
<i>SET</i>	Shall be of type CHARACTER.
<i>POS</i>	Shall be of type INTEGER.
<i>BACK</i>	(Optional) Shall be of type LOGICAL.

Example:

```

character(len=:), allocatable :: input
character(len=2) :: set = ', '
integer :: p
input = "one,last example"
p = 0
do
  if (p > len(input)) exit
  istart = p + 1
  call split(input, set, p)
  iend = p - 1
  print '(t7, a)', input(istart:iend)
end do

```

Standard: Fortran 2023

See also: Section 8.243 [SCAN], page 276,


```
omp_proc_bind_false  
omp_proc_bind_true  
omp_proc_bind_primary  
omp_proc_bind_master  
omp_proc_bind_close  
omp_proc_bind_spread
```

The following scalar integer named constants are of the kind `omp_lock_hint_kind`:

```
omp_lock_hint_none  
omp_lock_hint_uncontended  
omp_lock_hint_contended  
omp_lock_hint_nonspeculative  
omp_lock_hint_speculative  
omp_sync_hint_none  
omp_sync_hint_uncontended  
omp_sync_hint_contended  
omp_sync_hint_nonspeculative  
omp_sync_hint_speculative
```

And the following two scalar integer named constants are of the kind `omp_pause_resource_kind`:

```
omp_pause_soft  
omp_pause_hard
```

The following scalar integer named constants are of the kind `omp_alloctrail_key_kind`:

```
omp_atk_sync_hint  
omp_atk_alignment  
omp_atk_access  
omp_atk_pool_size  
omp_atk_fallback  
omp_atk_fb_data  
omp_atk_pinned  
omp_atk_partition
```

The following scalar integer named constants are of the kind `omp_alloctrail_val_kind`:

```
omp_alloctrail_key_kind:
```

```

omp_atv_default
omp_atv_false
omp_atv_true
omp_atv_contended
omp_atv_uncontended
omp_atv_serialized
omp_atv_sequential
omp_atv_private
omp_atv_all
omp_atv_thread
omp_atv_ptesteam
omp_atv_cgroup
omp_atv_default_mem_fb
omp_atv_null_fb
omp_atv_abort_fb
omp_atv_allocator_fb
omp_atv_environment
omp_atv_nearest
omp_atv_blocked

```

The following scalar integer named constants are of the kind `omp_allocator_handle_kind`:

```

omp_null_allocator
omp_default_mem_alloc
omp_large_cap_mem_alloc
omp_const_mem_alloc
omp_high_bw_mem_alloc
omp_low_lat_mem_alloc
omp_cgroup_mem_alloc
omp_ptesteam_mem_alloc
omp_thread_mem_alloc

```

The following scalar integer named constants are of the kind `omp_memspace_handle_kind`:

```

omp_default_mem_space
omp_large_cap_mem_space
omp_const_mem_space
omp_high_bw_mem_space
omp_low_lat_mem_space

```

9.5 OpenACC Module OPENACC

Standard: OpenACC Application Programming Interface v2.6

The OpenACC Fortran runtime library routines are provided both in a form of a Fortran 90 module, named `OPENACC`, and in form of a Fortran `include` file named `openacc_lib.h`. The procedures provided by `OPENACC` can be found in the Section “Introduction” in *GNU Offloading and Multi Processing Runtime Library* manual, the named constants defined in the modules are listed below.

For details refer to the actual OpenACC Application Programming Interface v2.6 (<https://www.openacc.org/>).

OPENACC provides the scalar default-integer named constant `openacc_version` with a value of the form `yyyymm`, where `yyyy` is the year and `mm` the month of the OpenACC version; for OpenACC v2.6 the value is 201711.

Contributing

Free software is only possible if people contribute to efforts to create it. We're always in need of more people helping out with ideas and comments, writing documentation and contributing code.

If you want to contribute to GNU Fortran, have a look at the long lists of projects you can take on. Some of these projects are small, some of them are large; some are completely orthogonal to the rest of what is happening on GNU Fortran, but others are “mainstream” projects in need of enthusiastic hackers. All of these projects are important! We will eventually get around to the things here, but they are also things doable by someone who is willing and able.

Contributors to GNU Fortran

Most of the parser was hand-crafted by *Andy Vaught*, who is also the initiator of the whole project. Thanks Andy! Most of the interface with GCC was written by *Paul Brook*.

The following individuals have contributed code and/or ideas and significant help to the GNU Fortran project (in alphabetical order):

- Janne Blomqvist
- Steven Bosscher
- Paul Brook
- Tobias Burnus
- François-Xavier Coudert
- Bud Davis
- Jerry DeLisle
- Erik Edelman
- Bernhard Fischer
- Daniel Franke
- Richard Guenther
- Richard Henderson
- Katherine Holcomb
- Jakub Jelinek
- Niels Kristian Bech Jensen
- Steven Johnson
- Steven G. Kargl
- Thomas Koenig
- Asher Langton
- H. J. Lu
- Toon Moene
- Brooks Moses
- Andrew Pinski

- Tim Prince
- Christopher D. Rickett
- Richard Sandiford
- Tobias Schlüter
- Roger Sayle
- Paul Thomas
- Andy Vaught
- Feng Wang
- Janus Weil
- Daniel Kraft

The following people have contributed bug reports, smaller or larger patches, and much needed feedback and encouragement for the GNU Fortran project:

- Bill Clodius
- Dominique d’Humières
- Kate Hedstrom
- Erik Schnetter
- Gerhard Steinmetz
- Joost VandeVondele

Many other individuals have helped debug, test and improve the GNU Fortran compiler over the past few years, and we welcome you to do the same! If you already have done so, and you would like to see your name listed in the list above, please contact us.

Projects

Help build the test suite

Solicit more code for donation to the test suite: the more extensive the testsuite, the smaller the risk of breaking things in the future! We can keep code private on request.

Bug hunting/squishing

Find bugs and write more test cases! Test cases are especially very welcome, because it allows us to concentrate on fixing bugs instead of isolating them. Going through the bugzilla database at <https://gcc.gnu.org/bugzilla/> to reduce testcases posted there and add more information (for example, for which version does the testcase work, for which versions does it fail?) is also very helpful.

Missing features

For a larger project, consider working on the missing features required for Fortran language standards compliance (see Section 1.3 [Standards], page 3), or contributing to the implementation of extensions such as OpenMP (see Section 5.1.18 [OpenMP], page 56) or OpenACC (see Section 5.1.19 [OpenACC], page 57) that are under active development. Again, contributing test cases for these features is useful too!

GNU General Public License

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <https://www.fsf.org>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

The “System Libraries” of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A “Major Component”, in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The “Corresponding Source” for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work’s System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users’ Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a. The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b. The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c. You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d. If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b. Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.
- c. Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d. Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e. Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source.

The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms.

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a. Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b. Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c. Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or

- d. Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e. Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f. Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance.

available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient’s use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others’ Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their

titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.3
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover
Texts.  A copy of the license is included in the section entitled ``GNU
Free Documentation License''.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Funding Free Software

If you want to have more free software a few years from now, it makes sense for you to help encourage people to contribute funds for its development. The most effective approach known is to encourage commercial redistributors to donate.

Users of free software systems can boost the pace of development by encouraging for-a-fee distributors to donate part of their selling price to free software developers—the Free Software Foundation, and others.

The way to convince distributors to do this is to demand it and expect it from them. So when you compare distributors, judge them partly by how much they give to free software development. Show distributors they must compete to be the one who gives the most.

To make this approach work, you must insist on numbers that you can compare, such as, “We will donate ten dollars to the Frobnitz project for each disk sold.” Don’t be satisfied with a vague promise, such as “A portion of the profits are donated,” since it doesn’t give a basis for comparison.

Even a precise fraction “of the profits from this disk” is not very meaningful, since creative accounting and unrelated business decisions can greatly alter what fraction of the sales price counts as profit. If the price you pay is \$50, ten percent of the profit is probably less than a dollar; it might be a few cents, or nothing at all.

Some redistributors do development work themselves. This is useful too; but to keep everyone honest, you need to inquire how much they do, and what kind. Some kinds of development make much more long-term difference than others. For example, maintaining a separate version of a program contributes very little; maintaining the standard version of a program for the whole community contributes much. Easy new ports contribute little, since someone else would surely do them; difficult ports such as adding a new CPU to the GNU Compiler Collection contribute more; major new features or packages contribute the most.

By establishing the idea that supporting further development is “the proper thing to do” when distributing free software for a fee, we can assure a steady flow of resources into making more free software.

Copyright © 1994 Free Software Foundation, Inc.

Verbatim copying and redistribution of this section is permitted without royalty; alteration is not permitted.

