

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

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
```


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.

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

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-constructor=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`.

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.

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

4.6 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`

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.


```

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/

```



```

      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)

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


```

        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

Arguments:

<i>size</i>	For normal coarrays, the byte size of the coarray to be allocated; for lock types and event types, the number of elements.
<i>type</i>	one of the <code>caf_register_t</code> types.
<i>token</i>	intent(out) An opaque pointer identifying the coarray.
<i>desc</i>	intent(inout) The (pseudo) array descriptor.
<i>stat</i>	intent(out) For allocatable coarrays, stores the <code>STAT=</code> ; may be <code>NULL</code>
<i>errmsg</i>	intent(out) When an error occurs, this is set to an error message; may be <code>NULL</code>
<i>errmsg_len</i>	the buffer size of <code>errmsg</code> .

Notes: Nonallocatable coarrays have to be registered prior use from remote images. In order to guarantee this, they have to be registered before the main program. This can be achieved by creating constructor functions. That is what GCC does such that also for nonallocatable coarrays the memory is allocated and no static memory is used. The token permits to identify the coarray; to the processor, the token is a nonaliasing pointer. The library can, for instance, store the base address of the coarray in the token, some handle or a more complicated struct. The library may also store the array descriptor *DESC* when its rank is nonzero. For lock types, the value shall only be used for checking the allocation status. Note that for critical blocks, the locking is only required on one image; in the locking statement, the processor shall always pass an image index of one for critical-block lock variables (`CAF_REGTYPE_CRITICAL`). For lock types and critical-block variables, the initial value shall be unlocked (or, respectively, not in critical section) such as the value `false`; for event types, the initial state should be no event, e.g. zero.

7.2.9 `_gfortran_caf_deregister` — Deregistering coarrays

Synopsis: `void _gfortran_caf_deregister (caf_token_t *token, caf_deregister_t type, int *stat, char *errmsg, size_t errmsg_len)`

Description:

Called to free or deregister the memory of a coarray; the processor calls this function for automatic and explicit deallocation. In case of an error, this function shall fail with an error message, unless the *STAT* variable is not null. The library is only expected to free memory it allocated itself during a call to `_gfortran_caf_register`.

Arguments:

<i>token</i>	the token to free.
<i>type</i>	the type of action to take for the coarray. A <code>CAF_DEREGTYPE_COARRAY_DEALLOCATE_ONLY</code> is allowed only for allocatable or pointer components of derived type coarrays. The action only deallocates the local memory without deleting the token.

<i>stat</i>	intent(out) Stores the STAT=; may be NULL
<i>errmsg</i>	intent(out) When an error occurs, this is set to an error message; may be NULL
<i>errmsg_len</i>	the buffer size of errmsg.

Notes: For nonallocatable coarrays this function is never called. If a cleanup is required, it has to be handled via the finish, stop and error stop functions, and via destructors.

7.2.10 `_gfortran_caf_register_accessor` — Register an accessor for remote access

Synopsis: `void _gfortran_caf_register_accessor (const int hash, void (*accessor)(void **, int32_t *, void *, void *, size_t *, size_t *))`

Description:

Identification of access functions across images is done using a unique hash. For each given hash an accessor has to be registered. This routine is expected to register an accessor function pointer for the given hash in nearly constant time. I.e. it is expected to add the hash and accessor to a buffer and return. Sorting shall be done in `_gfortran_caf_register_accessors_finish`.

Arguments:

<i>hash</i>	intent(in) The unique hash value this accessor is to be identified by.
<i>accessor</i>	intent(in) A pointer to the function on this image. The function has the signature <code>void accessor (void **dst_ptr, int32_t *free_dst, void *src_ptr, void *get_data, size_t *opt_src_charlen, size_t *opt_dst_charlen)</code> . GFortran ensures that functions provided to <code>_gfortran_caf_register_accessor</code> adhere to this interface.

Notes: This function is required to have a nearly constant runtime complexity, because it will be called to register multiple accessor in a sequence. GFortran ensures that before the first remote accesses commences `_gfortran_caf_register_accessors_finish` is called at least once. It is valid to register further accessors after a call to `_gfortran_caf_register_accessors_finish`. It is invalid to call `_gfortran_caf_register_accessor` after the first remote access has been done. See also Section 7.2.11 [`_gfortran_caf_register_accessors_finish`], page 96, and Section 7.2.12 [`_gfortran_caf_get_remote_function_index`], page 97,

7.2.11 `_gfortran_caf_register_accessors_finish` — Finish registering accessor functions

Synopsis: `void _gfortran_caf_register_accessors_finish ()`

Description:

Called to finalize registering of accessor functions. This function is expected to prepare a lookup table that has fast lookup time for the hash supplied to

`_gfortran_caf_get_remote_function_index` and constant access time for indexing operations.

Arguments:

No arguments.

Notes: This function may be called multiple times with and without new hash-accessors- pairs being added. The post-condition after each call has to be that hashes can be looked up quickly and indexing on the lookup table of hash-accessor-pairs is a constant time operation.

7.2.12 `_gfortran_caf_get_remote_function_index` — Get the index of an accessor

Synopsis: `int _gfortran_caf_get_remote_function_index (const int hash)`

Description:

Return the index of the accessor in the lookup table build by Section 7.2.10 [`_gfortran_caf_register_accessor`], page 96, and Section 7.2.11 [`_gfortran_caf_register_accessors_finish`], page 96. This function is expected to be fast, because it may be called often. A $\log(N)$ lookup time for a given hash is preferred. The reference implementation uses `bsearch()`, for example. The index returned shall be an array index to be used by Section 7.2.13 [`_gfortran_caf_get_from_remote`], page 97, i.e. a constant time operation is mandatory for quick access.

The GFortran compiler ensures that `_gfortran_caf_get_remote_function_index` is called once only for each hash and the result be stored in a static variable to prevent future redundant lookups.

Arguments:

`hash` `intent(in)` The hash of the accessor desired.

Result: The zero based index to access the accessor function in a lookup table. On error, -1 can be returned.

Notes: The function's complexity is expected to be significantly smaller than N , where N is the number of all accessors registered. Although returning -1 is valid, will this most likely crash the Fortran program when accessing the -1-th accessor function. It is therefore advised to terminate with an error message, when the hash could not be found.

7.2.13 `_gfortran_caf_get_from_remote` — Getting data from a remote image using a remote side accessor

Synopsis: `void _gfortran_caf_get_from_remote (caf_token_t token, const gfc_descriptor_t *opt_src_desc, const size_t *opt_src_charlen, const int image_index, const size_t dst_size, void **dst_data, size_t *opt_dst_charlen, gfc_descriptor_t *opt_dst_desc, const bool may_realloc_dst, const int getter_index, void *get_data, const size_t get_data_size, int *stat, caf_team_t *team, int *team_number)`

Description:

Called to get a scalar, an array section or a whole array from a remote image identified by the *image_index*.

Arguments:

<i>token</i>	intent(in) An opaque pointer identifying the coarray.
<i>opt_src_desc</i>	intent(in) A pointer to the descriptor when the object identified by <i>token</i> is an array with a descriptor. The parameter needs to be set to NULL, when <i>token</i> identifies a scalar.
<i>opt_src_charlen</i>	intent(in) When the object to get is a char array with deferred length, then this parameter needs to be set to point to its length. Else the parameter needs to be set to NULL.
<i>image_index</i>	intent(in) The ID of the remote image; must be a positive number. this_image () is valid.
<i>dst_size</i>	intent(in) The size of data expected to be transferred from the remote image. If the data type to get is a string or string array, then this needs to be set to the byte size of each character, i.e. 4 for a CHARACTER (KIND=4) string. The length of the string is then returned in opt_dst_charlen (also for string arrays).
<i>dst_data</i>	intent(inout) A pointer to the address the data is stored. To prevent copying of data into an output buffer the address to the live data is returned here. When a descriptor is provided also its data-member is set to that address. When <i>may_realloc_dst</i> is set, then the memory may be reallocated by the remote function, which needs to be replicated by this function.
<i>opt_dst_charlen</i>	intent(inout) When a char array is returned, this parameter is set to the length where applicable. The value can also be read to prevent reallocation in the accessor.
<i>opt_dst_desc</i>	intent(inout) When a descriptor array is returned, it is stored in the memory pointed to by this optional parameter. When <i>may_realloc_dst</i> is set, then the descriptor may be changed, i.e. its bounds, but upto now not its rank.
<i>may_realloc_dst</i>	intent(in) Set when the returned data may require reallocation of the output buffer in <i>dst_data</i> or <i>opt_dst_desc</i> .
<i>getter_index</i>	intent(in) The index of the accessor to execute as returned by _gfortran_caf_get_remote_function_index ().

<i>get_data</i>	intent(inout) Additional data needed in the accessor. I.e., when an array reference uses a local variable <i>v</i> , it is transported in this structure and all references in the accessor are rewritten to access the member. The data in the structure of <i>get_data</i> may be changed by the accessor, but these changes are lost to the calling Fortran program.
<i>get_data_size</i>	intent(in) The size of the <i>get_data</i> structure.
<i>stat</i>	intent(out) When non-NULL give the result of the operation, i.e., zero on success and non-zero on error. When NULL and an error occurs, then an error message is printed and the program is terminated.
<i>team</i>	intent(in) The opaque team handle as returned by <code>FORM TEAM</code> .
<i>team_number</i>	intent(in) The number of the team this access is to be part of.

Notes: It is permitted to have `image_index` equal the current image; the memory to get and the memory to store the data may (partially) overlap. The implementation has to take care that it handles this case, e.g. using `memmove` which handles (partially) overlapping memory.

7.2.14 `_gfortran_caf_is_present_on_remote` — Check that a coarray or a part of it is allocated on the remote image

Synopsis: `int32_t _gfortran_caf_is_present_on_remote (caf_token_t token, const int image_index, const int is_present_index, void *add_data, const size_t add_data_size)`

Description:

Check if an allocatable coarray or a component of a derived type coarray is allocated on the remote image identified by the *image_index*. The check is done by calling routine on the remote side.

Arguments:

<i>token</i>	intent(in) An opaque pointer identifying the coarray.
<i>image_index</i>	intent(in) The ID of the remote image; must be a positive number. <code>this_image ()</code> is valid.
<i>is_present_index</i>	intent(in) The index of the accessor to execute as returned by <code>_gfortran_caf_get_remote_function_index ()</code> .
<i>add_data</i>	intent(inout) Additional data needed in the accessor. I.e., when an array reference uses a local variable <i>v</i> , it is transported in this structure and all references in the accessor are rewritten to access the member. The data in the structure of <i>add_data</i> may be changed by the accessor, but these changes are lost to the calling Fortran program.
<i>add_data_size</i>	intent(in) The size of the <i>add_data</i> structure.

7.2.15 `_gfortran_caf_send_to_remote` — Send data to a remote image using a remote side accessor to store it

Synopsis: `void _gfortran_caf_send_to_remote (caf_token_t token, gfc_descriptor_t *opt_dst_desc, const size_t *opt_dst_charlen, const int image_index, const size_t src_size, const void *src_data, size_t *opt_src_charlen, const gfc_descriptor_t *opt_src_desc, const int setter_index, void *add_data, const size_t add_data_size, int *stat, caf_team_t *team, int *team_number)`

Description:

Called to send a scalar, an array section or a whole array to a remote image identified by the *image_index*. The call modifies the memory of the remote image.

Arguments:

<i>token</i>	intent(in) An opaque pointer identifying the coarray.
<i>opt_dst_desc</i>	intent(inout) A pointer to the descriptor when the object identified by <i>token</i> is an array with a descriptor. The parameter needs to be set to NULL, when <i>token</i> identifies a scalar or is an array without a descriptor.
<i>opt_dst_charlen</i>	intent(in) When the object to send is a char array with deferred length, then this parameter needs to be set to point to its length. Else the parameter needs to be set to NULL.
<i>image_index</i>	intent(in) The ID of the remote image; must be a positive number. <code>this_image ()</code> is valid.
<i>src_size</i>	intent(in) The size of data expected to be transferred to the remote image. If the data type to get is a string or string array, then this needs to be set to the byte size of each character, i.e. 4 for a CHARACTER (KIND=4) string. The length of the string is then given in <i>opt_src_charlen</i> (also for string arrays).
<i>src_data</i>	intent(in) A pointer the data to be send to the remote image. When a descriptor is provided in <i>opt_src_desc</i> then this parameter can be ignored by the library implementing the coarray functionality.
<i>opt_src_charlen</i>	intent(in) When a char array is send, this parameter is set to its length.
<i>opt_src_desc</i>	intent(in) When a descriptor array is send, then this parameter gives the handle.
<i>setter_index</i>	intent(in) The index of the accessor to execute as returned by <code>_gfortran_caf_get_remote_function_index ()</code> .

add_data intent(inout) Additional data needed in the accessor. I.e., when an array reference uses a local variable *v*, it is transported in this structure and all references in the accessor are rewritten to access the member. The data in the structure of *add_data* may be changed by the accessor, but these changes are lost to the calling Fortran program.

add_data_size intent(in) The size of the *add_data* structure.

stat intent(out) When non-NULL give the result of the operation, i.e., zero on success and non-zero on error. When NULL and an error occurs, then an error message is printed and the program is terminated.

team intent(in) The opaque team handle as returned by **FORM TEAM**.

team_number intent(in) The number of the team this access is to be part of.

Notes: It is permitted to have *image_index* equal the current image; the memory to send the data to and the memory to read for the data may (partially) overlap. The implementation has to take care that it handles this case, e.g. using **memmove** which handles (partially) overlapping memory.

7.2.16 **_gfortran_caf_transfer_between_remotes** — Initiate data transfer between to remote images

Synopsis: `void _gfortran_caf_transfer_between_remotes (caf_token_t dst_token, gfc_descriptor_t *opt_dst_desc, size_t *opt_dst_charlen, const int dst_image_index, const int dst_access_index, void *dst_add_data, const size_t dst_add_data_size, caf_token_t src_token, const gfc_descriptor_t *opt_src_desc, const size_t *opt_src_charlen, const int src_image_index, const int src_access_index, void *src_add_data, const size_t src_add_data_size, const size_t src_size, const bool scalar_transfer, int *dst_stat, int *src_stat, caf_team_t *dst_team, int *dst_team_number, caf_team_t *src_team, int *src_team_number)`

Description:

Initiates a transfer of data from one remote image to another remote image. The call modifies the memory of the receiving remote image given by *dst_image_index*. The *src_image_index*'s memory is not modified. The call returns when the transfer has commenced.

Arguments:

dst_token intent(in) An opaque pointer identifying the coarray on the receiving image.

- opt_dst_desc* intent(inout) A pointer to the descriptor when the object identified by *dst_token* is an array with a descriptor. The parameter needs to be set to NULL, when *dst_token* identifies a scalar or is an array without a descriptor.
- opt_dst_charlen* intent(in) When the object to modify on the receiving image is a char array with deferred length, then this parameter needs to be set to point to its length. Else the parameter needs to be set to NULL.
- dst_image_index* intent(in) The ID of the receiving/destination remote image; must be a positive number. `this_image ()` is valid.
- dst_access_index* intent(in) The index of the accessor to execute on the receiving image as returned by `_gfortran_caf_get_remote_function_index ()`.
- dst_add_data* intent(inout) Additional data needed in the accessor on the receiving side. I.e., when an array reference uses a local variable *v*, it is transported in this structure and all references in the accessor are rewritten to access the member. The data in the structure of *dst_add_data* may be changed by the accessor, but these changes are lost to the calling Fortran program.
- dst_add_data_size* intent(in) The size of the *dst_add_data* structure.
- src_token* intent(in) An opaque pointer identifying the coarray on the sending image.
- opt_src_desc* intent(inout) A pointer to the descriptor when the object identified by *src_token* is an array with a descriptor. The parameter needs to be set to NULL, when *src_token* identifies a scalar or is an array without a descriptor.
- opt_src_charlen* intent(in) When the object to get from the source image is a char array with deferred length, then this parameter needs to be set to point to its length. Else the parameter needs to be set to NULL.
- src_image_index* intent(in) The ID of the sending/source remote image; must be a positive number. `this_image ()` is valid.
- src_access_index* intent(in) The index of the accessor to execute on the sending image as returned by `_gfortran_caf_get_remote_function_index ()`.

<i>src_add_data</i>	intent(inout) Additional data needed in the accessor on the sending side. I.e., when an array reference uses a local variable <i>v</i> , it is transported in this structure and all references in the accessor are rewritten to access the member. The data in the structure of <i>src_add_data</i> may be changed by the accessor, but these changes are lost to the calling Fortran program.
<i>src_add_data_size</i>	intent(in) The size of the <i>src_add_data</i> structure.
<i>src_size</i>	intent(in) The size of data expected to be transferred between the images. If the data type to get is a string or string array, then this needs to be set to the byte size of each character, i.e. 4 for a <code>CHARACTER (KIND=4)</code> string. The length of the string is then given in <i>opt_src_charlen</i> and <i>opt_dst_charlen</i> (also for string arrays).
<i>scalar_transfer</i>	intent(in) Is set to true when the data to be transferred between the two images is not an array with a descriptor.
<i>dst_stat</i>	intent(out) When non-NULL give the result of the operation on the receiving side, i.e., zero on success and non-zero on error. When NULL and an error occurs, then an error message is printed and the program is terminated.
<i>src_stat</i>	intent(out) When non-NULL give the result of the operation on the sending side, i.e., zero on success and non-zero on error. When NULL and an error occurs, then an error message is printed and the program is terminated.
<i>dst_team</i>	intent(in) The opaque team handle as returned by <code>FORM TEAM</code> .
<i>dst_team_number</i>	intent(in) The number of the team this access is to be part of.
<i>src_team</i>	intent(in) The opaque team handle as returned by <code>FORM TEAM</code> .
<i>src_team_number</i>	intent(in) The number of the team this access is to be part of.

Notes: It is permitted to have both *dst_image_index* and *src_image_index* equal the current image; the memory to send the data to and the memory to read for the data may (partially) overlap. The implementation has to take care that it handles this case, e.g. using `memmove` which handles (partially) overlapping memory.

7.2.17 `_gfortran_caf_sendget_by_ref` — Sending data between remote images using enhanced references on both sides

Synopsis: `void _gfortran_caf_sendget_by_ref (caf_token_t dst_token, int dst_image_index, caf_reference_t *dst_refs, caf_token_t src_token, int src_image_index, caf_reference_t *src_refs, int dst_kind, int src_kind, bool may_require_tmp, int *dst_stat, int *src_stat, int dst_type, int src_type)`

Description:

Called to send a scalar, an array section or a whole array from a remote image identified by the `src_image_index` to a remote image identified by the `dst_image_index`.

Arguments:

<code>dst_token</code>	intent(in) An opaque pointer identifying the destination coarray.
<code>dst_image_index</code>	intent(in) The ID of the destination remote image; must be a positive number.
<code>dst_refs</code>	intent(in) The references on the remote array to store the data given by the source. Guaranteed to have at least one entry.
<code>src_token</code>	intent(in) An opaque pointer identifying the source coarray.
<code>src_image_index</code>	intent(in) The ID of the source remote image; must be a positive number.
<code>src_refs</code>	intent(in) The references to apply to the remote structure to get the data.
<code>dst_kind</code>	intent(in) Kind of the destination argument
<code>src_kind</code>	intent(in) Kind of the source argument
<code>may_require_tmp</code>	intent(in) The variable is false when it is known at compile time that the <code>dest</code> and <code>src</code> either cannot overlap or overlap (fully or partially) such that walking <code>src</code> and <code>dest</code> in elementwise order (honoring the stride value) does not lead to wrong results. Otherwise, the value is true .
<code>dst_stat</code>	intent(out) when non-NULL give the result of the send-operation, i.e., zero on success and nonzero on error. When NULL and an error occurs, then an error message is printed and the program is terminated.
<code>src_stat</code>	intent(out) When non-NULL give the result of the get-operation, i.e., zero on success and nonzero on error. When NULL and an error occurs, then an error message is printed and the program is terminated.
<code>dst_type</code>	intent(in) Give the type of the destination. When the destination is not an array, than the precise type, e.g. of a component in a derived type, is not known, but provided here.

src_type intent(in) Give the type of the source. When the source is not an array, than the precise type, e.g. of a component in a derived type, is not known, but provided here.

Notes: It is permitted to have the same image index for both *src_image_index* and *dst_image_index*; the memory of the send-to and the send-from might (partially) overlap in that case. The implementation has to take care that it handles this case, e.g. using `memmove` which handles (partially) overlapping memory. If *may_require_tmp* is true, the library might additionally create a temporary variable, unless additional checks show that this is not required (e.g. because walking backward is possible or because both arrays are contiguous and `memmove` takes care of overlap issues).

Note that the assignment of a scalar to an array is permitted. In addition, the library has to handle numeric-type conversion and for strings, padding and different character kinds.

Because of the more complicated references possible some operations may be unsupported by certain libraries. The library is expected to issue a precise error message why the operation is not permitted.

7.2.18 `_gfortran_caf_lock` — Locking a lock variable

Synopsis: `void _gfortran_caf_lock (caf_token_t token, size_t index, int image_index, int *acquired_lock, int *stat, char *errmsg, size_t errmsg_len)`

Description:

Acquire a lock on the given image on a scalar locking variable or for the given array element for an array-valued variable. If the *acquired_lock* is NULL, the function returns after having obtained the lock. If it is non-NULL, then *acquired_lock* is assigned the value true (one) when the lock could be obtained and false (zero) otherwise. Locking a lock variable that has already been locked by the same image is an error.

Arguments:

token intent(in) An opaque pointer identifying the coarray.
index intent(in) Array index; first array index is 0. For scalars, it is always 0.
image_index intent(in) The ID of the remote image; must be a positive number.
acquired_lock intent(out) If not NULL, it returns whether lock could be obtained.
stat intent(out) Stores the STAT=; may be NULL.
errmsg intent(out) When an error occurs, this is set to an error message; may be NULL.
errmsg_len intent(in) the buffer size of errmsg

Notes: This function is also called for critical blocks; for those, the array index is always zero and the image index is one. Libraries are permitted to use other images for critical-block locking variables.

Notes: The typical use is to check the local event variable to only call `event_wait` when the data is available. However, a coindexed variable is permitted; there is no ordering or synchronization implied. It acts like an atomic fetch of the value of the event variable.

7.2.23 `_gfortran_caf_sync_all` — All-image barrier

Synopsis: `void _gfortran_caf_sync_all (int *stat, char *errmsg, size_t errmsg_len)`

Description:

Synchronization of all images in the current team; the program only continues on a given image after this function has been called on all images of the current team. Additionally, it ensures that all pending data transfers of previous segment have completed.

Arguments:

<code>stat</code>	intent(out) Stores the status <code>STAT=</code> and may be NULL.
<code>errmsg</code>	intent(out) When an error occurs, this is set to an error message; may be NULL.
<code>errmsg_len</code>	intent(in) the buffer size of <code>errmsg</code>

7.2.24 `_gfortran_caf_sync_images` — Barrier for selected images

Synopsis: `void _gfortran_caf_sync_images (int count, int images[], int *stat, char *errmsg, size_t errmsg_len)`

Description:

Synchronization between the specified images; the program only continues on a given image after this function has been called on all images specified for that image. Note that one image can wait for all other images in the current team (e.g. via `sync images(*)`) while those only wait for that specific image. Additionally, `sync images` ensures that all pending data transfers of previous segments have completed.

Arguments:

<code>count</code>	intent(in) The number of images that are provided in the next argument. For a zero-sized array, the value is zero. For <code>sync images (*)</code> , the value is <code>-1</code> .
<code>images</code>	intent(in) An array with the images provided by the user. If <code>count</code> is zero, a NULL pointer is passed.
<code>stat</code>	intent(out) Stores the status <code>STAT=</code> and may be NULL.
<code>errmsg</code>	intent(out) When an error occurs, this is set to an error message; may be NULL.
<code>errmsg_len</code>	intent(in) the buffer size of <code>errmsg</code>

Description:

Atomic compare and swap of a kind-4 integer or logical variable. Assigns atomically the specified value to the atomic variable, if the latter has the value specified by the passed condition value.

Arguments:

<i>token</i>	intent(in) An opaque pointer identifying the coarray.
<i>offset</i>	intent(in) The number of bytes the actual data is shifted compared to the base address of the coarray.
<i>image_index</i>	intent(in) The ID of the remote image; must be a positive number; zero indicates the current image when used noncoindexed.
<i>old</i>	intent(out) The value the atomic variable had just before the cas operation.
<i>compare</i>	intent(in) The value used for comparison.
<i>new_val</i>	intent(in) The new value for the atomic variable, assigned to the atomic variable, if compare equals the value of the atomic variable.
<i>stat</i>	intent(out) Stores the status STAT= and may be NULL.
<i>type</i>	intent(in) the data type, i.e. BT_INTEGER (1) or BT_LOGICAL (2).
<i>kind</i>	intent(in) The kind value (only 4; always int)

7.2.32 _gfortran_caf_atomic_op — Atomic operation

Synopsis: `void _gfortran_caf_atomic_op (int op, caf_token_t token, size_t offset, int image_index, void *value, void *old, int *stat, int type, int kind)`

Description:

Apply an operation atomically to an atomic integer or logical variable. After the operation, *old* contains the value just before the operation, which, respectively, adds (GFC_CAF_ATOMIC_ADD) atomically the **value** to the atomic integer variable or does a bitwise AND, OR or exclusive OR between the atomic variable and *value*; the result is then stored in the atomic variable.

Arguments:

<i>op</i>	intent(in) the operation to be performed; possible values GFC_CAF_ATOMIC_ADD (1), GFC_CAF_ATOMIC_AND (2), GFC_CAF_ATOMIC_OR (3), GFC_CAF_ATOMIC_XOR (4).
<i>token</i>	intent(in) An opaque pointer identifying the coarray.
<i>offset</i>	intent(in) The number of bytes the actual data is shifted compared to the base address of the coarray.
<i>image_index</i>	intent(in) The ID of the remote image; must be a positive number; zero indicates the current image when used noncoindexed.

<i>old</i>	intent(out) The value the atomic variable had just before the atomic operation.
<i>val</i>	intent(in) The new value for the atomic variable, assigned to the atomic variable, if <code>compare</code> equals the value of the atomic variable.
<i>stat</i>	intent(out) Stores the status <code>STAT=</code> and may be NULL.
<i>type</i>	intent(in) the data type, i.e. <code>BT_INTEGER</code> (1) or <code>BT_LOGICAL</code> (2)
<i>kind</i>	intent(in) the kind value (only 4; always <code>int</code>)

7.2.33 `_gfortran_caf_co_broadcast` — Sending data to all images

Synopsis: `void _gfortran_caf_co_broadcast (gfc_descriptor_t *a, int source_image, int *stat, char *errmsg, size_t errmsg_len)`

Description:

Distribute a value from a given image to all other images in the team. Has to be called collectively.

Arguments:

<i>a</i>	intent(inout) An array descriptor with the data to be broadcasted (on <i>source_image</i>) or to be received (other images).
<i>source_image</i>	intent(in) The ID of the image from which the data should be broadcasted.
<i>stat</i>	intent(out) Stores the status <code>STAT=</code> and may be NULL.
<i>errmsg</i>	intent(out) When an error occurs, this is set to an error message; may be NULL.
<i>errmsg_len</i>	intent(in) the buffer size of <i>errmsg</i> .

7.2.34 `_gfortran_caf_co_max` — Collective maximum reduction

Synopsis: `void _gfortran_caf_co_max (gfc_descriptor_t *a, int result_image, int *stat, char *errmsg, int a_len, size_t errmsg_len)`

Description:

Calculates for each array element of the variable *a* the maximum value for that element in the current team; if *result_image* has the value 0, the result shall be stored on all images, otherwise, only on the specified image. This function operates on numeric values and character strings.

Arguments:

<i>a</i>	intent(inout) An array descriptor for the data to be processed. On the destination image(s) the result overwrites the old content.
<i>result_image</i>	intent(in) The ID of the image to which the reduced value should be copied to; if zero, it has to be copied to all images.

Arguments:

<i>a</i>	intent(inout) An array descriptor with the data to be processed. On the destination image(s) the result overwrites the old content.
<i>result_image</i>	intent(in) The ID of the image to which the reduced value should be copied to; if zero, it has to be copied to all images.
<i>stat</i>	intent(out) Stores the status <code>STAT=</code> and may be NULL.
<i>errmsg</i>	intent(out) When an error occurs, this is set to an error message; may be NULL.
<i>errmsg_len</i>	intent(in) the buffer size of <i>errmsg</i>

Notes: If *result_image* is nonzero, the data in the array descriptor *a* on all images except of the specified one become undefined; hence, the library may make use of this.

7.2.37 `_gfortran_caf_co_reduce` — Generic collective reduction

Synopsis: `void _gfortran_caf_co_reduce (gfc_descriptor_t *a, void * (*opr) (void *, void *), int opr_flags, int result_image, int *stat, char *errmsg, int a_len, size_t errmsg_len)`

Description:

Calculates for each array element of the variable *a* the reduction value for that element in the current team; if *result_image* has the value 0, the result shall be stored on all images, otherwise, only on the specified image. The *opr* is a pure function doing a mathematically commutative and associative operation.

The *opr_flags* denote the following; the values are bitwise ored. `GFC_CAF_BYREF` (1) if the result should be returned by reference; `GFC_CAF_HIDDENLEN` (2) whether the result and argument string lengths shall be specified as hidden arguments; `GFC_CAF_ARG_VALUE` (4) whether the arguments shall be passed by value, `GFC_CAF_ARG_DESC` (8) whether the arguments shall be passed by descriptor.

Arguments:

<i>a</i>	intent(inout) An array descriptor with the data to be processed. On the destination image(s) the result overwrites the old content.
<i>opr</i>	intent(in) Function pointer to the reduction function
<i>opr_flags</i>	intent(in) Flags regarding the reduction function
<i>result_image</i>	intent(in) The ID of the image to which the reduced value should be copied to; if zero, it has to be copied to all images.
<i>stat</i>	intent(out) Stores the status <code>STAT=</code> and may be NULL.
<i>errmsg</i>	intent(out) When an error occurs, this is set to an error message; may be NULL.

See also: Section 8.9 [ACOSPI], page 124,
 Section 8.28 [ATAN2PI], page 137,
 Section 8.31 [ATANPI], page 139,

8.24 ASSOCIATED — Status of a pointer or pointer/target pair

Synopsis: `RESULT = ASSOCIATED(POINTER [, TARGET])`

Description:

`ASSOCIATED(POINTER [, TARGET])` determines the status of the pointer *POINTER* or if *POINTER* is associated with the target *TARGET*.

Class: Inquiry function

Arguments:

POINTER *POINTER* shall have the `POINTER` attribute and it can be of any type.

TARGET (Optional) *TARGET* shall be a pointer or a target. It must have the same type, kind type parameter, and array rank as *POINTER*.

The association status of neither *POINTER* nor *TARGET* shall be undefined.

Return value:

`ASSOCIATED(POINTER)` returns a scalar value of type `LOGICAL(4)`. There are several cases:

- (A) When the optional *TARGET* is not present then
`ASSOCIATED(POINTER)` is true if *POINTER* is associated with a target; otherwise, it returns false.
- (B) If *TARGET* is present and a scalar target, the result is true if
TARGET is not a zero-sized storage sequence and the target associated with *POINTER* occupies the same storage units. If *POINTER* is disassociated, the result is false.
- (C) If *TARGET* is present and an array target, the result is true if
TARGET and *POINTER* have the same shape, are not zero-sized arrays, are arrays whose elements are not zero-sized storage sequences, and *TARGET* and *POINTER* occupy the same storage units in array element order. As in case(B), the result is false, if *POINTER* is disassociated.
- (D) If *TARGET* is present and an scalar pointer, the result is true
 if *TARGET* is associated with *POINTER*, the target associated with *TARGET* are not zero-sized storage sequences and occupy the same storage units. The result is false, if either *TARGET* or *POINTER* is disassociated.
- (E) If *TARGET* is present and an array pointer, the result is true if
 target associated with *POINTER* and the target associated with *TARGET* have the same shape, are not zero-sized arrays, are ar-

rays whose elements are not zero-sized storage sequences, and *TARGET* and *POINTER* occupy the same storage units in array element order. The result is false, if either *TARGET* or *POINTER* is disassociated.

Example:

```
program test_associated
  implicit none
  real, target :: tgt(2) = (/1., 2./)
  real, pointer :: ptr(:)
  ptr => tgt
  if (associated(ptr) .eqv. .false.) call abort
  if (associated(ptr,tgt) .eqv. .false.) call abort
end program test_associated
```

Standard: Fortran 90 and later

See also: Section 8.215 [NULL], page 259,

8.25 ATAN — Arctangent function

Synopsis:

```
RESULT = ATAN(X)
RESULT = ATAN(Y, X)
```

Description:

ATAN(X) computes the arctangent of X.

Class: Elemental function

Arguments:

X	The type shall be REAL or COMPLEX; if Y is present, X shall be REAL.
Y	The type and kind type parameter shall be the same as X.

Return value:

The return value is of the same type and kind as X. If Y is present, the result is identical to ATAN2(Y,X). Otherwise, it is the arctangent of X, where the real part of the result is in radians and lies in the range $-\pi/2 \leq \Re \operatorname{atan}(x) \leq \pi/2$.

Example:

```
program test_atan
  real(8) :: x = 2.866_8
  x = atan(x)
end program test_atan
```

Specific names:

Name	Argument	Return type	Standard
ATAN(X)	REAL(4) X	REAL(4)	Fortran 77 and later
DATAN(X)	REAL(8) X	REAL(8)	Fortran 77 and later

Standard: Fortran 77 and later, for a complex argument and for two arguments Fortran 2008 or later

See also: Inverse function:
 Section 8.276 [TAN], page 297,
 Degrees function:
 Section 8.29 [ATAND], page 138,

8.26 ATAN2 — Arctangent function

Synopsis: RESULT = ATAN2(Y, X)

Description:

ATAN2(Y, X) computes the principal value of the argument function of the complex number $X + iY$. This function can be used to transform from Cartesian into polar coordinates and allows to determine the angle in the correct quadrant.

Class: Elemental function

Arguments:

Y	The type shall be REAL.
X	The type and kind type parameter shall be the same as Y. If Y is zero, then X must be nonzero.

Return value:

The return value has the same type and kind type parameter as Y. It is the principal value of the complex number $X + iY$. If X is nonzero, then it lies in the range $-\pi \leq \text{atan}(x) \leq \pi$. The sign is positive if Y is positive. If Y is zero, then the return value is zero if X is strictly positive, π if X is negative and Y is positive zero (or the processor does not handle signed zeros), and $-\pi$ if X is negative and Y is negative zero. Finally, if X is zero, then the magnitude of the result is $\pi/2$.

Example:

```
program test_atan2
  real(4) :: x = 1.e0_4, y = 0.5e0_4
  x = atan2(y,x)
end program test_atan2
```

Specific names:

Name	Argument	Return type	Standard
ATAN2(X, Y)	REAL(4) X, Y	REAL(4)	Fortran 77 and later
DATAN2(X, Y)	REAL(8) X, Y	REAL(8)	Fortran 77 and later

Standard: Fortran 77 and later

See also: Alias:
 Section 8.25 [ATAN], page 135,
 Degrees function:
 Section 8.27 [ATAN2D], page 136,

8.27 ATAN2D — Arctangent function, degrees

Synopsis: RESULT = ATAN2D(Y, X)

Description:

ATAN2D(Y, X) computes the principal value of the argument function of the complex number $X + iY$ in degrees. This function can be used to transform from Cartesian into polar coordinates and allows to determine the angle in the correct quadrant.

Class: Elemental function

Arguments:

Y The type shall be **REAL**.
 X The type and kind type parameter shall be the same as Y. If Y is zero, then X must be nonzero.

Return value:

The return value has the same type and kind type parameter as Y. It is the principal value of the complex number $X + iY$. If X is nonzero, then it lies in the range $-180 \leq \text{atan}(x) \leq 180$. The sign is positive if Y is positive. If Y is zero, then the return value is zero if X is strictly positive, 180 if X is negative and Y is positive zero (or the processor does not handle signed zeros), and -180 if X is negative and Y is negative zero. Finally, if X is zero, then the magnitude of the result is 90.

Example:

```
program test_atan2d
  real(4) :: x = 1.e0_4, y = 0.5e0_4
  x = atan2d(y,x)
end program test_atan2d
```

Specific names:

Name	Argument	Return type	Standard
ATAN2D(X, Y)	REAL(4) X, Y	REAL(4)	Fortran 2023
DATAN2D(X, Y)	REAL(8) X, Y	REAL(8)	GNU extension

Standard: Fortran 2023

See also: Alias:

Section 8.29 [ATAND], page 138,
 Radians function:
 Section 8.26 [ATAN2], page 136,

8.28 ATAN2PI — Circular arc tangent function

Description:

ATAN2PI(Y, X) computes $\text{atan2}(y, x)/\pi$, and provides a measure of an angle in half-revolutions within the proper quadrant.

Standard: Fortran 2023

Class: Elemental function

Syntax: RESULT = ATAN2PI(Y, X)

Arguments:

Y The type shall be **REAL**.

X The type and kind type parameter shall be the same as *Y*. If *Y* is zero, then *X* shall be nonzero.

Return value:

The return value has the same type and kind type parameter as *Y* and satisfies $-1 \leq \text{atan2}(y, x)/\pi \leq 1$.

Example:

```
program test_atan2pi
  real(kind=4) :: x = 1.e0_4, y = 0.5e0_4
  x = atan2pi(y, x)
end program test_atan2pi
```

See also: Section 8.9 [ACOSPI], page 124,
 Section 8.23 [ASINPI], page 133,
 Section 8.31 [ATANPI], page 139,

8.29 ATAND — Arctangent function, degrees

Synopsis:

```
RESULT = ATAND(X)
RESULT = ATAND(Y, X)
```

Description:

ATAND(*X*) computes the arctangent of *X* in degrees (inverse of Section 8.277 [TAND], page 298).

Class: Elemental function

Arguments:

X The type shall be REAL.
Y The type and kind type parameter shall be the same as *X*.

Return value:

The return value is of the same type and kind as *X*. If *Y* is present, the result is identical to ATAN2D(*Y*, *X*). Otherwise, the result is in degrees and lies in the range $-90 \leq \text{atand}(x) \leq 90$.

Example:

```
program test_atand
  real(8) :: x = 2.866_8
  real(4) :: x1 = 1.e0_4, y1 = 0.5e0_4
  x = atand(x)
  x1 = atand(y1, x1)
end program test_atand
```

Specific names:

Name	Argument	Return type	Standard
ATAND(<i>X</i>)	REAL(4) <i>X</i>	REAL(4)	Fortran 2023
DATAND(<i>X</i>)	REAL(8) <i>X</i>	REAL(8)	GNU extension

Standard: Fortran 2023

See also: Inverse function:
 Section 8.277 [TAND], page 298,
 Radians function:
 Section 8.25 [ATAN], page 135,

8.30 ATANH — Inverse hyperbolic tangent function

Synopsis: RESULT = ATANH(X)

Description:

ATANH(X) computes the inverse hyperbolic tangent of X.

Class: Elemental function

Arguments:

X The type shall be REAL or COMPLEX.

Return value:

The return value has same type and kind as X. If X is complex, the imaginary part of the result is in radians and lies between $-\pi/2 \leq \Im \operatorname{atanh}(x) \leq \pi/2$.

Example:

```
PROGRAM test_atanh
  REAL, DIMENSION(3) :: x = (/ -1.0, 0.0, 1.0 /)
  WRITE (*,*) ATANH(x)
END PROGRAM
```

Specific names:

Name	Argument	Return type	Standard
DATANH(X)	REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008 and later

See also: Inverse function:
 Section 8.278 [TANH], page 298,

8.31 ATANPI — Circular arc tangent function

Description:

ATANPI(X) computes $\operatorname{atan}(x)/\pi$. ATANPI(Y, X) computes $\operatorname{atan2}(y, x)/\pi$. These provide a measure of an angle in half-revolutions.

Standard: Fortran 2023

Class: Elemental function

Syntax:

```
RESULT = ATANPI(X)
RESULT = ATANPI(Y, X)
```

Arguments:

Y The type shall be REAL.
 X If Y appears, X shall have the same type and kind as Y. If Y is zero, then X shall not be zero. If Y does not appear in a function reference, then X shall be REAL.

Return value:

The return value has the same type and kind as *X*. It is expressed in half-revolutions and satisfies $-0.5 \leq \text{atanpi}(x) \leq 0.5$.

Example:

```

program test_atanpi
  implicit none
  real, parameter :: x = 0.123, y(3) = [0.123, 0.45, 0.8]
  real, parameter :: a = atanpi(x), b(3) = atanpi(y)
  call foo(x, y)
contains
  subroutine foo(u, v)
    real, intent(in) :: u, v(:)
    real :: f, g(size(v))
    f = atanpi(u)
    g = atanpi(v)
    if (abs(a - f) > 8 * epsilon(f)) stop 1
    if (any(abs(g - b) > 8 * epsilon(f))) stop 2
  end subroutine foo
end program test_atanpi

```

See also: Section 8.9 [ACOSPI], page 124,
 Section 8.23 [ASINPI], page 133,
 Section 8.28 [ATAN2PI], page 137,

8.32 ATOMIC_ADD — Atomic ADD operation

Synopsis: CALL ATOMIC_ADD (ATOM, VALUE [, STAT])

Description:

ATOMIC_ADD(ATOM, VALUE) atomically adds the value of *VALUE* to the variable *ATOM*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's STAT_STOPPED_IMAGE and if the remote image has failed, the value STAT_FAILED_IMAGE.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with ATOMIC_INT_KIND kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*]
  call atomic_add (atom[1], this_image())
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.36 [ATOMIC_FETCH_ADD], page 143,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.33 [ATOMIC_AND], page 141,
 Section 8.40 [ATOMIC_OR], page 146,
 Section 8.42 [ATOMIC_XOR], page 148,

8.33 ATOMIC_AND — Atomic bitwise AND operation

Synopsis: CALL ATOMIC_AND (ATOM, VALUE [, STAT])

Description:

ATOMIC_AND(ATOM, VALUE) atomically defines *ATOM* with the bitwise AND between the values of *ATOM* and *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's STAT_STOPPED_IMAGE and if the remote image has failed, the value STAT_FAILED_IMAGE.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with ATOMIC_INT_KIND kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*]
  call atomic_and (atom[1], int(b'10100011101'))
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.37 [ATOMIC_FETCH_AND], page 144,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.32 [ATOMIC_ADD], page 140,
 Section 8.40 [ATOMIC_OR], page 146,
 Section 8.42 [ATOMIC_XOR], page 148,

8.34 ATOMIC_CAS — Atomic compare and swap

Synopsis: CALL ATOMIC_CAS (ATOM, OLD, COMPARE, NEW [, STAT])

Description:

ATOMIC_CAS compares the variable *ATOM* with the value of *COMPARE*; if the value is the same, *ATOM* is set to the value of *NEW*. Additionally, *OLD* is set

to the value of *ATOM* that was used for the comparison. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of *ISO_FORTRAN_ENV*'s *STAT_STOPPED_IMAGE* and if the remote image has failed, the value *STAT_FAILED_IMAGE*.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of either integer type with <i>ATOMIC_INT_KIND</i> kind or logical type with <i>ATOMIC_LOGICAL_KIND</i> kind.
<i>OLD</i>	Scalar of the same type and kind as <i>ATOM</i> .
<i>COMPARE</i>	Scalar variable of the same type and kind as <i>ATOM</i> .
<i>NEW</i>	Scalar variable of the same type as <i>ATOM</i> . If kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  logical(atomic_logical_kind) :: atom[*], prev
  call atomic_cas (atom[1], prev, .false., .true.)
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [*ATOMIC_DEFINE*], page 142,
 Section 8.41 [*ATOMIC_REF*], page 147,
 Section 9.1 [*ISO_FORTRAN_ENV*], page 311,

8.35 *ATOMIC_DEFINE* — Setting a variable atomically

Synopsis: CALL *ATOMIC_DEFINE* (*ATOM*, *VALUE* [, *STAT*])

Description:

ATOMIC_DEFINE(*ATOM*, *VALUE*) defines the variable *ATOM* with the value *VALUE* atomically. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of *ISO_FORTRAN_ENV*'s *STAT_STOPPED_IMAGE* and if the remote image has failed, the value *STAT_FAILED_IMAGE*.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of either integer type with <i>ATOMIC_INT_KIND</i> kind or logical type with <i>ATOMIC_LOGICAL_KIND</i> kind.
-------------	--

VALUE Scalar of the same type as *ATOM*. If the kind is different, the value is converted to the kind of *ATOM*.
STAT (optional) Scalar default-kind integer variable.

Example:

```
program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*]
  call atomic_define (atom[1], this_image())
end program atomic
```

Standard: Fortran 2008 and later; with *STAT*, TS 18508 or later

See also: Section 8.41 [ATOMIC_REF], page 147,
 Section 8.34 [ATOMIC_CAS], page 141,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.32 [ATOMIC_ADD], page 140,
 Section 8.33 [ATOMIC_AND], page 141,
 Section 8.40 [ATOMIC_OR], page 146,
 Section 8.42 [ATOMIC_XOR], page 148,

8.36 ATOMIC_FETCH_ADD — Atomic ADD operation with prior fetch

Synopsis: CALL ATOMIC_FETCH_ADD (ATOM, VALUE, OLD [, STAT])

Description:

ATOMIC_FETCH_ADD(ATOM, VALUE, OLD) atomically stores the value of *ATOM* in *OLD* and adds the value of *VALUE* to the variable *ATOM*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's STAT_STOPPED_IMAGE and if the remote image has failed, the value STAT_FAILED_IMAGE.

Class: Atomic subroutine

Arguments:

ATOM Scalar coarray or coindexed variable of integer type with ATOMIC_INT_KIND kind. ATOMIC_LOGICAL_KIND kind.
VALUE Scalar of the same type as *ATOM*. If the kind is different, the value is converted to the kind of *ATOM*.
OLD Scalar of the same type and kind as *ATOM*.
STAT (optional) Scalar default-kind integer variable.

Example:

```
program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*], old
  call atomic_add (atom[1], this_image(), old)
end program atomic
```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.32 [ATOMIC_ADD], page 140,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.37 [ATOMIC_FETCH_AND], page 144,
 Section 8.38 [ATOMIC_FETCH_OR], page 144,
 Section 8.39 [ATOMIC_FETCH_XOR], page 145,

8.37 ATOMIC_FETCH_AND — Atomic bitwise AND operation with prior fetch

Synopsis: CALL ATOMIC_FETCH_AND (ATOM, VALUE, OLD [, STAT])

Description:

ATOMIC_AND(ATOM, VALUE) atomically stores the value of *ATOM* in *OLD* and defines *ATOM* with the bitwise AND between the values of *ATOM* and *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's STAT_STOPPED_IMAGE and if the remote image has failed, the value STAT_FAILED_IMAGE.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with ATOMIC_INT_KIND kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>OLD</i>	Scalar of the same type and kind as <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*], old
  call atomic_fetch_and (atom[1], int(b'10100011101'), old)
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.33 [ATOMIC_AND], page 141,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.36 [ATOMIC_FETCH_ADD], page 143,
 Section 8.38 [ATOMIC_FETCH_OR], page 144,
 Section 8.39 [ATOMIC_FETCH_XOR], page 145,

8.38 ATOMIC_FETCH_OR — Atomic bitwise OR operation with prior fetch

Synopsis: CALL ATOMIC_FETCH_OR (ATOM, VALUE, OLD [, STAT])

Description:

`ATOMIC_OR(ATOM, VALUE)` atomically stores the value of *ATOM* in *OLD* and defines *ATOM* with the bitwise OR between the values of *ATOM* and *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of `ISO_FORTRAN_ENV`'s `STAT_STOPPED_IMAGE` and if the remote image has failed, the value `STAT_FAILED_IMAGE`.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with <code>ATOMIC_INT_KIND</code> kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>OLD</i>	Scalar of the same type and kind as <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*], old
  call atomic_fetch_or (atom[1], int(b'10100011101'), old)
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [`ATOMIC_DEFINE`], page 142,
 Section 8.40 [`ATOMIC_OR`], page 146,
 Section 9.1 [`ISO_FORTRAN_ENV`], page 311,
 Section 8.36 [`ATOMIC_FETCH_ADD`], page 143,
 Section 8.37 [`ATOMIC_FETCH_AND`], page 144,
 Section 8.39 [`ATOMIC_FETCH_XOR`], page 145,

8.39 `ATOMIC_FETCH_XOR` — Atomic bitwise XOR operation with prior fetch

Synopsis: `CALL ATOMIC_FETCH_XOR (ATOM, VALUE, OLD [, STAT])`

Description:

`ATOMIC_XOR(ATOM, VALUE)` atomically stores the value of *ATOM* in *OLD* and defines *ATOM* with the bitwise XOR between the values of *ATOM* and *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of `ISO_FORTRAN_ENV`'s `STAT_STOPPED_IMAGE` and if the remote image has failed, the value `STAT_FAILED_IMAGE`.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with <code>ATOMIC_INT_KIND</code> kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>OLD</i>	Scalar of the same type and kind as <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*], old
  call atomic_fetch_xor (atom[1], int(b'10100011101'), old)
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.42 [ATOMIC_XOR], page 148,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.36 [ATOMIC_FETCH_ADD], page 143,
 Section 8.37 [ATOMIC_FETCH_AND], page 144,
 Section 8.38 [ATOMIC_FETCH_OR], page 144,

8.40 ATOMIC_OR — Atomic bitwise OR operation

Synopsis: CALL ATOMIC_OR (ATOM, VALUE [, STAT])

Description:

ATOMIC_OR(ATOM, VALUE) atomically defines *ATOM* with the bitwise AND between the values of *ATOM* and *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's `STAT_STOPPED_IMAGE` and if the remote image has failed, the value `STAT_FAILED_IMAGE`.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with <code>ATOMIC_INT_KIND</code> kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*]
  call atomic_or (atom[1], int(b'10100011101'))
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.38 [ATOMIC_FETCH_OR], page 144,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.32 [ATOMIC_ADD], page 140,
 Section 8.40 [ATOMIC_OR], page 146,
 Section 8.42 [ATOMIC_XOR], page 148,

8.41 ATOMIC_REF — Obtaining the value of a variable atomically

Synopsis: CALL ATOMIC_REF(VALUE, ATOM [, STAT])

Description:

ATOMIC_DEFINE(ATOM, VALUE) atomically assigns the value of the variable *ATOM* to *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's STAT_STOPPED_IMAGE and if the remote image has failed, the value STAT_FAILED_IMAGE.

Class: Atomic subroutine

Arguments:

<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>ATOM</i>	Scalar coarray or coindexed variable of either integer type with ATOMIC_INT_KIND kind or logical type with ATOMIC_LOGICAL_KIND kind.
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  logical(atomic_logical_kind) :: atom[*]
  logical :: val
  call atomic_ref (atom, .false.)
  ! ...
  call atomic_ref (atom, val)
  if (val) then
    print *, "Obtained"
  end if
end program atomic

```

Standard: Fortran 2008 and later; with *STAT*, TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.34 [ATOMIC_CAS], page 141,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.36 [ATOMIC_FETCH_ADD], page 143,
 Section 8.37 [ATOMIC_FETCH_AND], page 144,

Section 8.38 [ATOMIC_FETCH_OR], page 144,
 Section 8.39 [ATOMIC_FETCH_XOR], page 145,

8.42 ATOMIC_XOR — Atomic bitwise OR operation

Synopsis: CALL ATOMIC_XOR (ATOM, VALUE [, STAT])

Description:

ATOMIC_AND(ATOM, VALUE) atomically defines *ATOM* with the bitwise XOR between the values of *ATOM* and *VALUE*. When *STAT* is present and the invocation was successful, it is assigned the value 0. If it is present and the invocation has failed, it is assigned a positive value; in particular, for a coindexed *ATOM*, if the remote image has stopped, it is assigned the value of ISO_FORTRAN_ENV's STAT_STOPPED_IMAGE and if the remote image has failed, the value STAT_FAILED_IMAGE.

Class: Atomic subroutine

Arguments:

<i>ATOM</i>	Scalar coarray or coindexed variable of integer type with ATOMIC_INT_KIND kind.
<i>VALUE</i>	Scalar of the same type as <i>ATOM</i> . If the kind is different, the value is converted to the kind of <i>ATOM</i> .
<i>STAT</i>	(optional) Scalar default-kind integer variable.

Example:

```

program atomic
  use iso_fortran_env
  integer(atomic_int_kind) :: atom[*]
  call atomic_xor (atom[1], int(b'10100011101'))
end program atomic

```

Standard: TS 18508 or later

See also: Section 8.35 [ATOMIC_DEFINE], page 142,
 Section 8.39 [ATOMIC_FETCH_XOR], page 145,
 Section 9.1 [ISO_FORTRAN_ENV], page 311,
 Section 8.32 [ATOMIC_ADD], page 140,
 Section 8.40 [ATOMIC_OR], page 146,
 Section 8.42 [ATOMIC_XOR], page 148,

8.43 BACKTRACE — Show a backtrace

Synopsis: CALL BACKTRACE

Description:

BACKTRACE shows a backtrace at an arbitrary place in user code. Program execution continues normally afterwards. The backtrace information is printed to the unit corresponding to ERROR_UNIT in ISO_FORTRAN_ENV.

Class: Subroutine

Arguments:

None

Standard: GNU extension

See also: Section 8.2 [ABORT], page 119,

8.44 BESSEL_J0 — Bessel function of the first kind of order 0

Synopsis: RESULT = BESSEL_J0(X)

Description:

BESSEL_J0(X) computes the Bessel function of the first kind of order 0 of X. This function is available under the name BESJ0 as a GNU extension.

Class: Elemental function

Arguments:

X The type shall be REAL.

Return value:

The return value is of type REAL and lies in the range $-0.4027... \leq Bessel(0, x) \leq 1$. It has the same kind as X.

Example:

```
program test_besj0
  real(8) :: x = 0.0_8
  x = bessej0(x)
end program test_besj0
```

Specific names:

Name	Argument	Return type	Standard
DBESJ0(X)	REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008 and later

8.45 BESSEL_J1 — Bessel function of the first kind of order 1

Synopsis: RESULT = BESSEL_J1(X)

Description:

BESSEL_J1(X) computes the Bessel function of the first kind of order 1 of X. This function is available under the name BESJ1 as a GNU extension.

Class: Elemental function

Arguments:

X The type shall be REAL.

Return value:

The return value is of type REAL and lies in the range $-0.5818... \leq Bessel(1, x) \leq 0.5818$. It has the same kind as X.

Example:

```
program test_besj1
  real(8) :: x = 1.0_8
  x = bessej1(x)
end program test_besj1
```

Specific names:

Name	Argument	Return type	Standard
DBESJ1(X)	REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008

8.46 BESSEL_JN — Bessel function of the first kind

Synopsis:

```
RESULT = BESSEL_JN(N, X)
RESULT = BESSEL_JN(N1, N2, X)
```

Description:

BESSEL_JN(N, X) computes the Bessel function of the first kind of order *N* of *X*. This function is available under the name BESJN as a GNU extension. If *N* and *X* are arrays, their ranks and shapes shall conform.

BESSEL_JN(N1, N2, X) returns an array with the Bessel functions of the first kind of the orders *N1* to *N2*.

Class: Elemental function, except for the transformational function BESSEL_JN(N1, N2, X)

Arguments:

<i>N</i>	Shall be a scalar or an array of type INTEGER.
<i>N1</i>	Shall be a non-negative scalar of type INTEGER.
<i>N2</i>	Shall be a non-negative scalar of type INTEGER.
<i>X</i>	Shall be a scalar or an array of type REAL; for BESSEL_JN(N1, N2, X) it shall be scalar.

Return value:

The return value is a scalar of type REAL. It has the same kind as *X*.

Notes: The transformational function uses a recurrence algorithm which might, for some values of *X*, lead to different results than calls to the elemental function.

Example:

```
program test_besjn
  real(8) :: x = 1.0_8
  x = besse_jn(5,x)
end program test_besjn
```

Specific names:

Name	Argument	Return type	Standard
DBESJN(N, X)	INTEGER N REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008 and later, negative *N* is allowed as GNU extension

8.47 BESSEL_Y0 — Bessel function of the second kind of order 0

Synopsis: RESULT = BESSEL_Y0(X)

Description:

BESSEL_Y0(X) computes the Bessel function of the second kind of order 0 of X. This function is available under the name BESY0 as a GNU extension.

Class: Elemental function

Arguments:

X The type shall be REAL.

Return value:

The return value is of type REAL. It has the same kind as X.

Example:

```
program test_besy0
  real(8) :: x = 0.0_8
  x = besseli_y0(x)
end program test_besy0
```

Specific names:

Name	Argument	Return type	Standard
DBESY0(X)	REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008 and later

8.48 BESSEL_Y1 — Bessel function of the second kind of order 1

Synopsis: RESULT = BESSEL_Y1(X)

Description:

BESSEL_Y1(X) computes the Bessel function of the second kind of order 1 of X. This function is available under the name BESY1 as a GNU extension.

Class: Elemental function

Arguments:

X The type shall be REAL.

Return value:

The return value is of type REAL. It has the same kind as X.

Example:

```
program test_besy1
  real(8) :: x = 1.0_8
  x = besseli_y1(x)
end program test_besy1
```

Specific names:

Name	Argument	Return type	Standard
DBESY1(X)	REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008 and later

Example:

```

program test_gamma
  real :: x = 1.0
  x = gamma(x) ! returns 1.0
end program test_gamma

```

Specific names:

Name	Argument	Return type	Standard
DGAMMA(X)	REAL(8) X	REAL(8)	GNU extension

Standard: Fortran 2008 and later

See also: Logarithm of the Gamma function:
Section 8.185 [LOG_GAMMA], page 241,

8.125 GERROR — Get last system error message

Synopsis: CALL GERROR(RESULT)

Description:

Returns the system error message corresponding to the last system error. This resembles the functionality of `strerror(3)` in C.

Class: Subroutine

Arguments:

RESULT Shall be of type CHARACTER and of default kind.

Example:

```

PROGRAM test_gerror
  CHARACTER(len=100) :: msg
  CALL gerror(msg)
  WRITE(*,*) msg
END PROGRAM

```

Standard: GNU extension

See also: Section 8.152 [IERRNO], page 222,
Section 8.221 [PERROR], page 264,

8.126 GETARG — Get command line arguments

Synopsis: CALL GETARG(POS, VALUE)

Description:

Retrieve the *POS*-th argument that was passed on the command line when the containing program was invoked.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.128 [GET_COMMAND_ARGUMENT], page 206, intrinsic defined by the Fortran 2003 standard.

Class: Subroutine

Arguments:

POS Shall be of type `INTEGER` and not wider than the default integer kind; $POS \geq 0$

VALUE Shall be of type `CHARACTER` and of default kind.

Return value:

After `GETARG` returns, the *VALUE* argument holds the *POS*th command line argument. If *VALUE* cannot hold the argument, it is truncated to fit the length of *VALUE*. If there are less than *POS* arguments specified at the command line, *VALUE* is filled with blanks. If $POS = 0$, *VALUE* is set to the name of the program (on systems that support this feature).

Example:

```
PROGRAM test_getarg
  INTEGER :: i
  CHARACTER(len=32) :: arg

  DO i = 1, iargc()
    CALL getarg(i, arg)
    WRITE (*,*) arg
  END DO
END PROGRAM
```

Standard: GNU extension

See also: GNU Fortran 77 compatibility function:
 Section 8.145 [IARGC], page 217,
 Fortran 2003 functions and subroutines:
 Section 8.127 [GET_COMMAND], page 205,
 Section 8.128 [GET_COMMAND_ARGUMENT], page 206,
 Section 8.72 [COMMAND_ARGUMENT_COUNT], page 168,

8.127 GET_COMMAND — Get the entire command line

Synopsis: `CALL GET_COMMAND([COMMAND, LENGTH, STATUS])`

Description:

Retrieve the entire command line that was used to invoke the program.

Class: Subroutine

Arguments:

COMMAND (Optional) shall be of type `CHARACTER` and of default kind.

LENGTH (Optional) Shall be of type `INTEGER` and of default kind.

STATUS (Optional) Shall be of type `INTEGER` and of default kind.

Return value:

If *COMMAND* is present, stores the entire command line that was used to invoke the program in *COMMAND*. If *LENGTH* is present, it is assigned the length of the command line. If *STATUS* is present, it is assigned 0 upon success

of the command, -1 if *COMMAND* is too short to store the command line, or a positive value in case of an error.

Example:

```
PROGRAM test_get_command
  CHARACTER(len=255) :: cmd
  CALL get_command(cmd)
  WRITE (*,*) TRIM(cmd)
END PROGRAM
```

Standard: Fortran 2003 and later

See also: Section 8.128 [GET_COMMAND_ARGUMENT], page 206,
Section 8.72 [COMMAND_ARGUMENT_COUNT], page 168,

8.128 GET_COMMAND_ARGUMENT — Get command line arguments

Synopsis: CALL GET_COMMAND_ARGUMENT(NUMBER [, VALUE, LENGTH, STATUS])

Description:

Retrieve the *NUMBER*-th argument that was passed on the command line when the containing program was invoked.

Class: Subroutine

Arguments:

<i>NUMBER</i>	Shall be a scalar of type INTEGER and of default kind, $NUMBER \geq 0$
<i>VALUE</i>	(Optional) Shall be a scalar of type CHARACTER and of default kind.
<i>LENGTH</i>	(Optional) Shall be a scalar of type INTEGER and of default kind.
<i>STATUS</i>	(Optional) Shall be a scalar of type INTEGER and of default kind.

Return value:

After GET_COMMAND_ARGUMENT returns, the *VALUE* argument holds the *NUMBER*-th command line argument. If *VALUE* cannot hold the argument, it is truncated to fit the length of *VALUE*. If there are less than *NUMBER* arguments specified at the command line, *VALUE* is filled with blanks. If *NUMBER* = 0, *VALUE* is set to the name of the program (on systems that support this feature). The *LENGTH* argument contains the length of the *NUMBER*-th command line argument. If the argument retrieval fails, *STATUS* is a positive number; if *VALUE* contains a truncated command line argument, *STATUS* is -1; and otherwise the *STATUS* is zero.

Example:

```
PROGRAM test_get_command_argument
  INTEGER :: i
  CHARACTER(len=32) :: arg

  i = 0
  DO
```

```

      CALL get_command_argument(i, arg)
      IF (LEN_TRIM(arg) == 0) EXIT

      WRITE (*,*) TRIM(arg)
      i = i+1
    END DO
  END PROGRAM

```

Standard: Fortran 2003 and later

See also: Section 8.127 [GET_COMMAND], page 205,
Section 8.72 [COMMAND_ARGUMENT_COUNT], page 168,

8.129 GETCWD — Get current working directory

Synopsis:

```

CALL GETCWD(C [, STATUS])
STATUS = GETCWD(C)

```

Description:

Get current working directory.

This intrinsic is provided in both subroutine and function forms; however, only one form can be used in any given program unit.

Class: Subroutine, function

Arguments:

<i>C</i>	The type shall be <code>CHARACTER</code> and of default kind.
<i>STATUS</i>	(Optional) status flag. Returns 0 on success, a system specific and nonzero error code otherwise.

Example:

```

PROGRAM test_getcwd
  CHARACTER(len=255) :: cwd
  CALL getcwd(cwd)
  WRITE(*,*) TRIM(cwd)
END PROGRAM

```

Standard: GNU extension

See also: Section 8.64 [CHDIR], page 161,

8.130 GETENV — Get an environmental variable

Synopsis: CALL GETENV(NAME, VALUE)

Description:

Get the *VALUE* of the environmental variable *NAME*.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.131 [GET_ENVIRONMENT_VARIABLE], page 208, intrinsic defined by the Fortran 2003 standard.

Note that `GETENV` need not be thread-safe. It is the responsibility of the user to ensure that the environment is not being updated concurrently with a call to the `GETENV` intrinsic.

Class: Subroutine

Arguments:

NAME Shall be of type `CHARACTER` and of default kind.
VALUE Shall be of type `CHARACTER` and of default kind.

Return value:

Stores the value of *NAME* in *VALUE*. If *VALUE* is not large enough to hold the data, it is truncated. If *NAME* is not set, *VALUE* is filled with blanks.

Example:

```
PROGRAM test_getenv
  CHARACTER(len=255) :: homedir
  CALL getenv("HOME", homedir)
  WRITE (*,*) TRIM(homedir)
END PROGRAM
```

Standard: GNU extension

See also: Section 8.131 [GET_ENVIRONMENT_VARIABLE], page 208,

8.131 GET_ENVIRONMENT_VARIABLE — Get an environmental variable

Synopsis: CALL GET_ENVIRONMENT_VARIABLE(*NAME*[, *VALUE*, *LENGTH*, *STATUS*, *TRIM_NAME*])

Description:

Get the *VALUE* of the environmental variable *NAME*.

Note that `GET_ENVIRONMENT_VARIABLE` need not be thread-safe. It is the responsibility of the user to ensure that the environment is not being updated concurrently with a call to the `GET_ENVIRONMENT_VARIABLE` intrinsic.

Class: Subroutine

Arguments:

NAME Shall be a scalar of type `CHARACTER` and of default kind.
VALUE (Optional) Shall be a scalar of type `CHARACTER` and of default kind.
LENGTH (Optional) Shall be a scalar of type `INTEGER` and of default kind.
STATUS (Optional) Shall be a scalar of type `INTEGER` and of default kind.
TRIM_NAME(Optional) Shall be a scalar of type `LOGICAL` and of default kind.

Return value:

Stores the value of *NAME* in *VALUE*. If *VALUE* is not large enough to hold the data, it is truncated. If *NAME* is not set, *VALUE* is filled with blanks. Argument *LENGTH* contains the length needed for storing the environment variable *NAME* or zero if it is not present. *STATUS* is -1 if *VALUE* is present

but too short for the environment variable; it is 1 if the environment variable does not exist and 2 if the processor does not support environment variables; in all other cases *STATUS* is zero. If *TRIM_NAME* is present with the value *.FALSE.*, the trailing blanks in *NAME* are significant; otherwise they are not part of the environment variable name.

Example:

```
PROGRAM test_getenv
  CHARACTER(len=255) :: homedir
  CALL get_environment_variable("HOME", homedir)
  WRITE (*,*) TRIM(homedir)
END PROGRAM
```

Standard: Fortran 2003 and later

8.132 GETGID — Group ID function

Synopsis: RESULT = GETGID()

Description:

Returns the numerical group ID of the current process.

Class: Function

Return value:

The return value of GETGID is an INTEGER of the default kind.

Example: See GETPID for an example.

Standard: GNU extension

See also: Section 8.134 [GETPID], page 210,
Section 8.136 [GETUID], page 211,

8.133 GETLOG — Get login name

Synopsis: CALL GETLOG(C)

Description:

Gets the username under which the program is running.

Class: Subroutine

Arguments:

C Shall be of type CHARACTER and of default kind.

Return value:

Stores the current user name in *C*. (On systems where POSIX functions *geteuid* and *getpwuid* are not available, and the *getlogin* function is not implemented either, this returns a blank string.)

Example:

```
PROGRAM TEST_GETLOG
  CHARACTER(32) :: login
  CALL GETLOG(login)
  WRITE(*,*) login
END PROGRAM
```

Standard: GNU extension

See also: Section 8.136 [GETUID], page 211,

8.134 GETPID — Process ID function

Synopsis: `RESULT = GETPID()`

Description:

Returns the numerical process identifier of the current process.

Class: Function

Return value:

The return value of `GETPID` is an `INTEGER` of the default kind.

Example:

```
program info
  print *, "The current process ID is ", getpid()
  print *, "Your numerical user ID is ", getuid()
  print *, "Your numerical group ID is ", getgid()
end program info
```

Standard: GNU extension

See also: Section 8.132 [GETGID], page 209,
Section 8.136 [GETUID], page 211,

8.135 GET_TEAM — Get the handle of a team

Synopsis: `RESULT = GET_TEAM([LEVEL])`

Description:

Returns the handle of the current team, if *LEVEL* is not given. Or the team specified by *LEVEL*, where *LEVEL* is one of the constants `INITIAL_TEAM`, `PARENT_TEAM` or `CURRENT_TEAM` from the intrinsic module `ISO_FORTRAN_ENV`. Calling the function with `PARENT_TEAM` while being on the initial team, returns a handle to the initial team. This ensures that always a valid team is returned, given that team handles can neither be checked for validity nor compared with each other or null.

Class: Transformational function

Return value:

An opaque handle of `TEAM_TYPE` from the intrinsic module `ISO_FORTRAN_ENV`.

Example:

```
program info
  use, intrinsic :: iso_fortran_env
  type(team_type) :: init, curr, par, nt

  init = get_team()
  curr = get_team(current_team) ! init equals curr here
  form team(1, nt)
  change team(nt)
  curr = get_team() ! or get_team(current_team)
```

```

        par = get_team(parent_team) ! par equals init here
    end team
end program info

```

Standard: Fortran 2018 or later

See also: Section 8.281 [THIS_IMAGE], page 300,
Section 9.1 [ISO_FORTRAN_ENV], page 311,

8.136 GETUID — User ID function

Synopsis: RESULT = GETUID()

Description:

Returns the numerical user ID of the current process.

Class: Function

Return value:

The return value of GETUID is an INTEGER of the default kind.

Example: See GETPID for an example.

Standard: GNU extension

See also: Section 8.134 [GETPID], page 210,
Section 8.133 [GETLOG], page 209,

8.137 GMTIME — Convert time to GMT info

Synopsis: CALL GMTIME(TIME, VALUES)

Description:

Given a system time value *TIME* (as provided by the Section 8.282 [TIME], page 301, intrinsic), fills *VALUES* with values extracted from it appropriate to the UTC time zone (Universal Coordinated Time, also known in some countries as GMT, Greenwich Mean Time), using `gmtime(3)`.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.87 [DATE_AND_TIME], page 178, intrinsic defined by the Fortran 95 standard.

Class: Subroutine

Arguments:

<i>TIME</i>	An INTEGER scalar expression corresponding to a system time, with INTENT(IN).
<i>VALUES</i>	A default INTEGER array with 9 elements, with INTENT(OUT).

Return value:

The elements of *VALUES* are assigned as follows:

1. Seconds after the minute, range 0–59 or 0–61 to allow for leap seconds
2. Minutes after the hour, range 0–59

3. Hours past midnight, range 0–23
4. Day of month, range 1–31
5. Number of months since January, range 0–11
6. Years since 1900
7. Number of days since Sunday, range 0–6
8. Days since January 1, range 0–365
9. Daylight savings indicator: positive if daylight savings is in effect, zero if not, and negative if the information is not available.

Standard: GNU extension

See also: Section 8.87 [DATE_AND_TIME], page 178,
 Section 8.86 [CTIME], page 177,
 Section 8.189 [LTIME], page 243,
 Section 8.282 [TIME], page 301,
 Section 8.283 [TIME8], page 301,

8.138 HOSTNM — Get system host name

Synopsis:

```
CALL HOSTNM(C [, STATUS])
STATUS = HOSTNM(NAME)
```

Description:

Retrieves the host name of the system on which the program is running. This intrinsic is provided in both subroutine and function forms; however, only one form can be used in any given program unit.

Class: Subroutine, function

Arguments:

<i>C</i>	Shall of type CHARACTER and of default kind.
<i>STATUS</i>	(Optional) status flag of type INTEGER . Returns 0 on success, or a system specific error code otherwise.

Return value:

In either syntax, *NAME* is set to the current hostname if it can be obtained, or to a blank string otherwise.

Standard: GNU extension

8.139 HUGE — Largest number of a kind

Synopsis: `RESULT = HUGE(X)`

Description:

HUGE(X) returns the largest number that is not an infinity in the model of the type of **X**.

Class: Inquiry function

Arguments:

X Shall be of type **REAL**, **INTEGER** or **UNSIGNED**.

Return value:

The return value is of the same type and kind as *X*

Example:

```
program test_huge_tiny
  print *, huge(0), huge(0.0), huge(0.0d0)
  print *, tiny(0.0), tiny(0.0d0)
end program test_huge_tiny
```

Standard: Fortran 90 and later, extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67)

8.140 HYPOT — Euclidean distance function

Synopsis: **RESULT = HYPOT(X, Y)**

Description:

HYPOT(X,Y) is the Euclidean distance function. It is equal to $\sqrt{X^2 + Y^2}$, without undue underflow or overflow.

Class: Elemental function

Arguments:

X The type shall be **REAL**.
Y The type and kind type parameter shall be the same as *X*.

Return value:

The return value has the same type and kind type parameter as *X*.

Example:

```
program test_hypot
  real(4) :: x = 1.e0_4, y = 0.5e0_4
  x = hypot(x,y)
end program test_hypot
```

Standard: Fortran 2008 and later

8.141 IACHAR — Code in ASCII collating sequence

Synopsis: **RESULT = IACHAR(C [, KIND])**

Description:

IACHAR(C) returns the code for the ASCII character in the first character position of *C*.

Class: Elemental function

Arguments:

C Shall be a scalar **CHARACTER**, with **INTENT(IN)**
KIND (Optional) A scalar **INTEGER** constant expression indicating the kind parameter of the result.

Return value:

The return value is of type `INTEGER` and of kind *KIND*. If *KIND* is absent, the return value is of default integer kind.

Example:

```
program test_iachar
  integer i
  i = iachar(' ')
end program test_iachar
```

Notes: See Section 8.149 [ICHAR], page 219, for a discussion of converting between numerical values and formatted string representations.

Standard: Fortran 95 and later, with *KIND* argument Fortran 2003 and later

See also: Section 8.5 [ACHAR], page 121,
 Section 8.63 [CHAR], page 160,
 Section 8.149 [ICHAR], page 219,

8.142 IALL — Bitwise AND of array elements

Synopsis:

```
RESULT = IALL(ARRAY[, MASK])
RESULT = IALL(ARRAY, DIM[, MASK])
```

Description:

Reduces with bitwise AND the elements of *ARRAY* along dimension *DIM* if the corresponding element in *MASK* is `TRUE`.

Class: Transformational function

Arguments:

<i>ARRAY</i>	Shall be an array of type <code>INTEGER</code> or <code>UNSIGNED</code>
<i>DIM</i>	(Optional) shall be a scalar of type <code>INTEGER</code> with a value in the range from 1 to n, where n equals the rank of <i>ARRAY</i> .
<i>MASK</i>	(Optional) shall be of type <code>LOGICAL</code> and either be a scalar or an array of the same shape as <i>ARRAY</i> .

Return value:

The result is of the same type as *ARRAY*.

If *DIM* is absent, a scalar with the bitwise AND of all elements in *ARRAY* is returned. Otherwise, an array of rank n-1, where n equals the rank of *ARRAY*, and a shape similar to that of *ARRAY* with dimension *DIM* dropped is returned.

Example:

```
PROGRAM test_iall
  INTEGER(1) :: a(2)

  a(1) = b'00100100'
  a(2) = b'01101010'
```

```

      ! prints 00100000
      PRINT '(b8.8)', IALL(a)
END PROGRAM

```

Standard: Fortran 2008 and later, extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.144 [IANY], page 216,
 Section 8.159 [IPARITY], page 225,
 Section 8.143 [IAND], page 215,

8.143 IAND — Bitwise logical and

Synopsis: `RESULT = IAND(I, J)`

Description:

Bitwise logical AND.

Class: Elemental function

Arguments:

I The type shall be **INTEGER**, **UNSIGNED** or a boz-literal-constant.

J The type shall be the same type as *I* with the same kind type parameter or a boz-literal-constant. *I* and *J* shall not both be boz-literal-constants.

Return value:

The return type is with the kind type parameter of the arguments. A boz-literal-constant is converted to an **INTEGER** or **UNSIGNED** with the kind type parameter of the other argument as-if a call to Section 8.155 [INT], page 223, or Section 8.292 [UINT], page 306, respectively, occurred.

Example:

```

PROGRAM test_iand
  INTEGER :: a, b
  DATA a / Z'F' /, b / Z'3' /
  WRITE (*,*) IAND(a, b)
END PROGRAM

```

Specific names:

Name	Argument	Return type	Standard
IAND(A)	INTEGER A	INTEGER	Fortran 90 and later
BIAND(A)	INTEGER(1) A	INTEGER(1)	GNU extension
IIAND(A)	INTEGER(2) A	INTEGER(2)	GNU extension
JIAND(A)	INTEGER(4) A	INTEGER(4)	GNU extension
KIAND(A)	INTEGER(8) A	INTEGER(8)	GNU extension

Standard: Fortran 90 and later, with boz-literal-constant Fortran 2008 and later, has overloads that are GNU extensions. Extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.158 [IOR], page 225,
 Section 8.151 [IEOR], page 221,

Section 8.147 [IBITS], page 218,
 Section 8.148 [IBSET], page 218,
 Section 8.146 [IBCLR], page 217,
 Section 8.214 [NOT], page 259,

8.144 IANY — Bitwise OR of array elements

Synopsis:

```
RESULT = IANY(ARRAY[, MASK])
RESULT = IANY(ARRAY, DIM[, MASK])
```

Description:

Reduces with bitwise OR (inclusive or) the elements of *ARRAY* along dimension *DIM* if the corresponding element in *MASK* is **TRUE**.

Class: Transformational function

Arguments:

ARRAY Shall be an array of type **INTEGER** or **UNSIGNED**
DIM (Optional) shall be a scalar of type **INTEGER** with a value in the range from 1 to n, where n equals the rank of *ARRAY*.
MASK (Optional) shall be of type **LOGICAL** and either be a scalar or an array of the same shape as *ARRAY*.

Return value:

The result is of the same type as *ARRAY*.

If *DIM* is absent, a scalar with the bitwise OR of all elements in *ARRAY* is returned. Otherwise, an array of rank n-1, where n equals the rank of *ARRAY*, and a shape similar to that of *ARRAY* with dimension *DIM* dropped is returned.

Example:

```
PROGRAM test_iany
  INTEGER(1) :: a(2)

  a(1) = b'00100100'
  a(2) = b'01101010'

  ! prints 01101110
  PRINT '(b8.8)', IANY(a)
END PROGRAM
```

Standard: Fortran 2008 and later, extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.159 [IPARITY], page 225,
 Section 8.142 [IALL], page 214,
 Section 8.158 [IOR], page 225,

8.145 IARGC — Get the number of command line arguments

Synopsis: `RESULT = IARGC()`

Description:

IARGC returns the number of arguments passed on the command line when the containing program was invoked.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.72 [COMMAND_ARGUMENT_COUNT], page 168, intrinsic defined by the Fortran 2003 standard.

Class: Function

Arguments:

None

Return value:

The number of command line arguments, type `INTEGER(4)`.

Example: See Section 8.126 [GETARG], page 204,

Standard: GNU extension

See also: GNU Fortran 77 compatibility subroutine:

Section 8.126 [GETARG], page 204,

Fortran 2003 functions and subroutines:

Section 8.127 [GET_COMMAND], page 205,

Section 8.128 [GET_COMMAND_ARGUMENT], page 206,

Section 8.72 [COMMAND_ARGUMENT_COUNT], page 168,

8.146 IBCLR — Clear bit

Synopsis: `RESULT = IBCLR(I, POS)`

Description:

IBCLR returns the value of *I* with the bit at position *POS* set to zero.

Class: Elemental function

Arguments:

I The type shall be `INTEGER` or `UNSIGNED`.

POS The type shall be `INTEGER`.

Return value:

The return value is of the same type as *I*.

Specific names:

Name	Argument	Return type	Standard
IBCLR(A)	INTEGER A	INTEGER	Fortran 90 and later
BBCLR(A)	INTEGER(1) A	INTEGER(1)	GNU extension
IIBCLR(A)	INTEGER(2) A	INTEGER(2)	GNU extension
JIBCLR(A)	INTEGER(4) A	INTEGER(4)	GNU extension
KIBCLR(A)	INTEGER(8) A	INTEGER(8)	GNU extension

Standard: Fortran 90 and later, has overloads that are GNU extensions. Extension for UNSIGNED (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.147 [IBITS], page 218,
 Section 8.148 [IBSET], page 218,
 Section 8.143 [IAND], page 215,
 Section 8.158 [IOR], page 225,
 Section 8.151 [IEOR], page 221,
 Section 8.209 [MVBITS], page 256,

8.147 IBITS — Bit extraction

Synopsis: RESULT = IBITS(I, POS, LEN)

Description:

IBITS extracts a field of length *LEN* from *I*, starting from bit position *POS* and extending left for *LEN* bits. The result is right-justified and the remaining bits are zeroed. The value of *POS*+*LEN* must be less than or equal to the value BIT_SIZE(*I*).

Class: Elemental function

Arguments:

<i>I</i>	The type shall be INTEGER or UNSIGNED.
<i>POS</i>	The type shall be INTEGER.
<i>LEN</i>	The type shall be INTEGER.

Return value:

The return value is of type as *I*.

Specific names:

Name	Argument	Return type	Standard
IBITS(A)	INTEGER A	INTEGER	Fortran 90 and later
BBITS(A)	INTEGER(1) A	INTEGER(1)	GNU extension
IIBITS(A)	INTEGER(2) A	INTEGER(2)	GNU extension
JIBITS(A)	INTEGER(4) A	INTEGER(4)	GNU extension
KIBITS(A)	INTEGER(8) A	INTEGER(8)	GNU extension

Standard: Fortran 90 and later, has overloads that are GNU extensions. Extension for UNSIGNED (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.52 [BIT_SIZE], page 153,
 Section 8.146 [IBCLR], page 217,
 Section 8.148 [IBSET], page 218,
 Section 8.143 [IAND], page 215,
 Section 8.158 [IOR], page 225,
 Section 8.151 [IEOR], page 221,

8.148 IBSET — Set bit

Synopsis: RESULT = IBSET(I, POS)

Description:

IBSET returns the value of *I* with the bit at position *POS* set to one.

Class: Elemental function

Arguments:

I The type shall be `INTEGER` or `UNSIGNED`.
POS The type shall be `INTEGER`.

Return value:

The return value is of the same type as *I*.

Specific names:

Name	Argument	Return type	Standard
IBSET(A)	INTEGER A	INTEGER	Fortran 90 and later
BBSET(A)	INTEGER(1) A	INTEGER(1)	GNU extension
IIBSET(A)	INTEGER(2) A	INTEGER(2)	GNU extension
JIBSET(A)	INTEGER(4) A	INTEGER(4)	GNU extension
KIBSET(A)	INTEGER(8) A	INTEGER(8)	GNU extension

Standard: Fortran 90 and later, has overloads that are GNU extensions. Extension for `UNSIGNED` (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.146 [IBCLR], page 217,
 Section 8.147 [IBITS], page 218,
 Section 8.143 [IAND], page 215,
 Section 8.158 [IOR], page 225,
 Section 8.151 [IEOR], page 221,
 Section 8.209 [MVBITS], page 256,

8.149 ICHAR — Character-to-integer conversion function

Synopsis: `RESULT = ICHAR(C [, KIND])`

Description:

ICHAR(*C*) returns the code for the character in the first character position of *C* in the system's native character set. The correspondence between characters and their codes is not necessarily the same across different GNU Fortran implementations.

Class: Elemental function

Arguments:

C Shall be a scalar `CHARACTER`, with `INTENT(IN)`
KIND (Optional) A scalar `INTEGER` constant expression indicating the kind parameter of the result.

Return value:

The return value is of type `INTEGER` and of kind *KIND*. If *KIND* is absent, the return value is of default integer kind.

Example:

```
program test_ichar
```

```

      integer i
      i = ichar(' ')
end program test_ichar

```

Specific names:

Name	Argument	Return type	Standard
ICHAR(C)	CHARACTER C	INTEGER(4)	Fortran 77 and later

Notes: No intrinsic exists to convert between a numeric value and a formatted character string representation – for instance, given the **CHARACTER** value '154', obtaining an **INTEGER** or **REAL** value with the value 154, or vice versa. Instead, this functionality is provided by internal-file I/O, as in the following example:

```

program read_val
  integer value
  character(len=10) string, string2
  string = '154'

  ! Convert a string to a numeric value
  read (string,'(I10)') value
  print *, value

  ! Convert a value to a formatted string
  write (string2,'(I10)') value
  print *, string2
end program read_val

```

Standard: Fortran 77 and later, with *KIND* argument Fortran 2003 and later

See also: Section 8.5 [ACHAR], page 121,
 Section 8.63 [CHAR], page 160,
 Section 8.141 [IACHAR], page 213,

8.150 IDATE — Get current local time subroutine (day/month/year)

Synopsis: CALL IDATE(VALUE)

Description:

IDATE(VALUE) Fills *VALUES* with the numerical values at the current local time. The day (in the range 1-31), month (in the range 1-12), and year appear in elements 1, 2, and 3 of *VALUES*, respectively. The year has four significant digits.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.87 [DATE-AND-TIME], page 178, intrinsic defined by the Fortran 95 standard.

Class: Subroutine

Arguments:

<i>VALUES</i>	The type shall be INTEGER , DIMENSION(3) and the kind shall be the default integer kind.
---------------	--

Return value:

Does not return anything.

Example:

```

program test_idate
  integer, dimension(3) :: tarray
  call idate(tarray)
  print *, tarray(1)
  print *, tarray(2)
  print *, tarray(3)
end program test_idate

```

Standard: GNU extension

See also: Section 8.87 [DATE_AND_TIME], page 178,

8.151 IEOR — Bitwise logical exclusive or

Synopsis: `RESULT = IEOR(I, J)`

Description:

IEOR returns the bitwise Boolean exclusive-OR of *I* and *J*.

Class: Elemental function

Arguments:

I The type shall be `INTEGER`, `UNSIGNED` or a boz-literal-constant.

J The type shall be the same type as *I* with the same kind type parameter or a boz-literal-constant. *I* and *J* shall not both be boz-literal-constants.

Return value:

The return type is with the kind type parameter of the arguments. A boz-literal-constant is converted to an `INTEGER` or `UNSIGNED` with the kind type parameter of the other argument as-if a call to Section 8.155 [INT], page 223, or Section 8.292 [UINT], page 306, respectively, occurred.

Specific names:

Name	Argument	Return type	Standard
IEOR(A)	INTEGER A	INTEGER	Fortran 90 and later
BIEOR(A)	INTEGER(1) A	INTEGER(1)	GNU extension
IIIEOR(A)	INTEGER(2) A	INTEGER(2)	GNU extension
JIEOR(A)	INTEGER(4) A	INTEGER(4)	GNU extension
KIEOR(A)	INTEGER(8) A	INTEGER(8)	GNU extension

Standard: Fortran 90 and later, with boz-literal-constant Fortran 2008 and later, has overloads that are GNU extensions. Extension for `UNSIGNED` (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.158 [IOR], page 225,
 Section 8.143 [IAND], page 215,
 Section 8.147 [IBITS], page 218,
 Section 8.148 [IBSET], page 218,
 Section 8.146 [IBCLR], page 217,
 Section 8.214 [NOT], page 259,

8.152 IERRNO — Get the last system error number

Synopsis: `RESULT = IERRNO()`

Description:

Returns the last system error number, as given by the C `errno` variable.

Class: Function

Arguments:

None

Return value:

The return value is of type `INTEGER` and of the default integer kind.

Standard: GNU extension

See also: Section 8.221 [`PERERROR`], page 264,

8.153 IMAGE_INDEX — Function that converts a cosubscript to an image index

Synopsis: `RESULT = IMAGE_INDEX(COARRAY, SUB)`

Description:

Returns the image index belonging to a cosubscript.

Class: Inquiry function.

Arguments:

`COARRAY` Coarray of any type.

`SUB` default integer rank-1 array of a size equal to the corank of `COARRAY`.

Return value:

Scalar default integer with the value of the image index that corresponds to the cosubscripts. For invalid cosubscripts the result is zero.

Example:

```
INTEGER :: array[2,-1:4,8,*]
! Writes 28 (or 0 if there are fewer than 28 images)
WRITE (*,*) IMAGE_INDEX (array, [2,0,3,1])
```

Standard: Fortran 2008 and later

See also: Section 8.281 [`THIS_IMAGE`], page 300,
Section 8.216 [`NUM_IMAGES`], page 260,

8.154 INDEX — Position of a substring within a string

Synopsis: `RESULT = INDEX(STRING, SUBSTRING [, BACK [, KIND]])`

Description:

Returns the position of the start of the first occurrence of string *SUBSTRING* as a substring in *STRING*, counting from one. If *SUBSTRING* is not present in *STRING*, zero is returned. If the *BACK* argument is present and true, the return value is the start of the last occurrence rather than the first.

Class: Elemental function

Arguments:

STRING Shall be a scalar **CHARACTER**, with **INTENT(IN)**
SUBSTRING Shall be a scalar **CHARACTER**, with **INTENT(IN)**
BACK (Optional) Shall be a scalar **LOGICAL**, with **INTENT(IN)**
KIND (Optional) A scalar **INTEGER** constant expression indicating the kind parameter of the result.

Return value:

The return value is of type **INTEGER** and of kind *KIND*. If *KIND* is absent, the return value is of default integer kind.

Specific names:

Name	Argument	Return type	Standard
INDEX(<i>STRING</i> , <i>SUBSTRING</i>)	CHARACTER	INTEGER(4)	Fortran 77 and later

Standard: Fortran 77 and later, with *KIND* argument Fortran 2003 and later

See also: Section 8.243 [SCAN], page 276,
 Section 8.298 [VERIFY], page 309,

8.155 INT — Convert to integer type

Synopsis: **RESULT = INT(A [, KIND])**

Description:

Convert to integer type

Class: Elemental function, extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67).

Arguments:

A Shall be of type **INTEGER**, **REAL**, **COMPLEX** or **UNSIGNED** or a *boz*-literal-constant.
KIND (Optional) A scalar **INTEGER** constant expression indicating the kind parameter of the result.

Return value:

These functions return a **INTEGER** variable or array under the following rules:

- (A) If *A* is of type **INTEGER**, **INT(A)** = *A*
- (B) If *A* is of type **REAL** and $|A| < 1$, **INT(A)** equals 0. If $|A| \geq 1$, then **INT(A)** is the integer whose magnitude is the largest integer that does not exceed the magnitude of *A* and whose sign is the same as the sign of *A*.
- (C) If *A* is of type **COMPLEX**, rule B is applied to the real part of *A*.
- (D) If *A* is of type **UNSIGNED** and $0 \leq A \leq \text{HUGE}(A)$, **INT(A)** = *A*. Outside that range, the result is interpreted using two's complement.

Example:

```

program test_int
  integer :: i = 42
  complex :: z = (-3.7, 1.0)
  print *, int(i)
  print *, int(z), int(z,8)
end program

```

Specific names:

Name	Argument	Return type	Standard
INT(A)	REAL(4) A	INTEGER	Fortran 77 and later
IFIX(A)	REAL(4) A	INTEGER	Fortran 77 and later
IDINT(A)	REAL(8) A	INTEGER	Fortran 77 and later

Standard: Fortran 77 and later, with `boz-literal-constant` Fortran 2008 and later.

8.156 INT2 — Convert to 16-bit integer type

Synopsis: `RESULT = INT2(A)`

Description:

Convert to a `KIND=2` integer type. This is equivalent to the standard `INT` intrinsic with an optional argument of `KIND=2`, and is only included for backwards compatibility.

Class: Elemental function

Arguments:

A Shall be of type `INTEGER`, `REAL`, or `COMPLEX`.

Return value:

The return value is a `INTEGER(2)` variable.

Standard: GNU extension

See also: Section 8.155 [INT], page 223,
Section 8.157 [INT8], page 224,

8.157 INT8 — Convert to 64-bit integer type

Synopsis: `RESULT = INT8(A)`

Description:

Convert to a `KIND=8` integer type. This is equivalent to the standard `INT` intrinsic with an optional argument of `KIND=8`, and is only included for backwards compatibility.

Class: Elemental function

Arguments:

A Shall be of type `INTEGER`, `REAL`, or `COMPLEX`.

Return value:

The return value is a `INTEGER(8)` variable.

Standard: GNU extension

See also: Section 8.155 [INT], page 223,
Section 8.156 [INT2], page 224,

8.158 IOR — Bitwise logical or

Synopsis: `RESULT = IOR(I, J)`

Description:

IOR returns the bitwise Boolean inclusive-OR of *I* and *J*.

Class: Elemental function

Arguments:

I The type shall be `INTEGER`, `UNSIGNED` or a boz-literal-constant.
J The type shall be the same type as *I* with the same kind type parameter or a boz-literal-constant. *I* and *J* shall not both be boz-literal-constants.

Return value:

The return type is `INTEGER` with the kind type parameter of the arguments. A boz-literal-constant is converted to an `INTEGER` or `UNSIGNED` with the kind type parameter of the other argument as-if a call to Section 8.155 [INT], page 223, or Section 8.292 [UINT], page 306, respectively, occurred.

Specific names:

Name	Argument	Return type	Standard
<code>IOR(A)</code>	<code>INTEGER A</code>	<code>INTEGER</code>	Fortran 90 and later
<code>BIOR(A)</code>	<code>INTEGER(1) A</code>	<code>INTEGER(1)</code>	GNU extension
<code>IIOR(A)</code>	<code>INTEGER(2) A</code>	<code>INTEGER(2)</code>	GNU extension
<code>JIOR(A)</code>	<code>INTEGER(4) A</code>	<code>INTEGER(4)</code>	GNU extension
<code>KIOR(A)</code>	<code>INTEGER(8) A</code>	<code>INTEGER(8)</code>	GNU extension

Standard: Fortran 90 and later, with boz-literal-constant Fortran 2008 and later, has overloads that are GNU extensions. Extension for `UNSIGNED` (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.151 [IEOR], page 221,
Section 8.143 [IAND], page 215,
Section 8.147 [IBITS], page 218,
Section 8.148 [IBSET], page 218,
Section 8.146 [IBCLR], page 217,
Section 8.214 [NOT], page 259,

8.159 IPARITY — Bitwise XOR of array elements

Synopsis:

`RESULT = IPARITY(ARRAY[, MASK])`
`RESULT = IPARITY(ARRAY, DIM[, MASK])`

Description:

Reduces with bitwise XOR (exclusive or) the elements of *ARRAY* along dimension *DIM* if the corresponding element in *MASK* is `TRUE`.

Class: Transformational function

Arguments:

ARRAY Shall be an array of type `INTEGER` or `UNSIGNED`.
DIM (Optional) shall be a scalar of type `INTEGER` with a value in the range from 1 to n, where n equals the rank of *ARRAY*.
MASK (Optional) shall be of type `LOGICAL` and either be a scalar or an array of the same shape as *ARRAY*.

Return value:

The result is of the same type as *ARRAY*.

If *DIM* is absent, a scalar with the bitwise XOR of all elements in *ARRAY* is returned. Otherwise, an array of rank n-1, where n equals the rank of *ARRAY*, and a shape similar to that of *ARRAY* with dimension *DIM* dropped is returned.

Example:

```
PROGRAM test_iparity
  INTEGER(1) :: a(2)

  a(1) = int(b'00100100', 1)
  a(2) = int(b'01101010', 1)

  ! prints 01001110
  PRINT '(b8.8)', IPARITY(a)
END PROGRAM
```

Standard: Fortran 2008 and later. Extension for `UNSIGNED` (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.144 [IANY], page 216,
 Section 8.142 [IALL], page 214,
 Section 8.151 [IEOR], page 221,
 Section 8.220 [PARITY], page 263,

8.160 IRAND — Integer pseudo-random number

Synopsis: `RESULT = IRAND(I)`

Description:

`IRAND(FLAG)` returns a pseudo-random number from a uniform distribution between 0 and a system-dependent limit (which is in most cases 2147483647). If *FLAG* is 0, the next number in the current sequence is returned; if *FLAG* is 1, the generator is restarted by `CALL SRAND(0)`; if *FLAG* has any other value, it is used as a new seed with `SRAND`.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. It implements a simple modulo generator as provided by g77. For new code, one should consider the use of Section 8.231 [RANDOM_NUMBER], page 269, as it implements a superior algorithm.

Class: Function

Return value:

Returns a default-kind LOGICAL. The returned value is TRUE if *X* is a NaN and FALSE otherwise.

Example:

```

program test_nan
  implicit none
  real :: x
  x = -1.0
  x = sqrt(x)
  if (isnan(x)) stop '"x" is a NaN'
end program test_nan

```

Standard: GNU extension

8.168 ITIME — Get current local time subroutine (hour/minutes/seconds)

Synopsis: CALL ITIME(VALUES)

Description:

ITIME(VALUES) Fills *VALUES* with the numerical values at the current local time. The hour (in the range 1-24), minute (in the range 1-60), and seconds (in the range 1-60) appear in elements 1, 2, and 3 of *VALUES*, respectively.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.87 [DATE_AND_TIME], page 178, intrinsic defined by the Fortran 95 standard.

Class: Subroutine

Arguments:

VALUES The type shall be INTEGER, DIMENSION(3) and the kind shall be the default integer kind.

Return value:

Does not return anything.

Example:

```

program test_itime
  integer, dimension(3) :: tarray
  call itime(tarray)
  print *, tarray(1)
  print *, tarray(2)
  print *, tarray(3)
end program test_itime

```

Standard: GNU extension

See also: Section 8.87 [DATE_AND_TIME], page 178,

8.169 KILL — Send a signal to a process

Synopsis:

```

CALL KILL(PID, SIG [, STATUS])
STATUS = KILL(PID, SIG)

```

Description:

Sends the signal specified by *SIG* to the process *PID*. See `kill(2)`.

This intrinsic is provided in both subroutine and function forms; however, only one form can be used in any given program unit.

Class: Subroutine, function

Arguments:

<i>PID</i>	Shall be a scalar <code>INTEGER</code> with <code>INTENT(IN)</code> .
<i>SIG</i>	Shall be a scalar <code>INTEGER</code> with <code>INTENT(IN)</code> .
<i>STATUS</i>	[Subroutine](Optional) Shall be a scalar <code>INTEGER</code> . Returns 0 on success; otherwise a system-specific error code is returned.
<i>STATUS</i>	[Function] The kind type parameter is that of <code>pid</code> . Returns 0 on success; otherwise a system-specific error code is returned.

Standard: GNU extension

See also: Section 8.2 [ABORT], page 119,
Section 8.106 [EXIT], page 191,

8.170 KIND — Kind of an entity

Synopsis: `K = KIND(X)`

Description:

`KIND(X)` returns the kind value of the entity *X*.

Class: Inquiry function

Arguments:

<i>X</i>	Shall be of type <code>LOGICAL</code> , <code>INTEGER</code> , <code>REAL</code> , <code>COMPLEX</code> or <code>CHARACTER</code> . It may be scalar or array valued.
----------	---

Return value:

The return value is a scalar of type `INTEGER` and of the default integer kind.

Example:

```

program test_kind
  integer,parameter :: kc = kind(' ')
  integer,parameter :: kl = kind(.true.)

  print *, "The default character kind is ", kc
  print *, "The default logical kind is ", kl
end program test_kind

```

Standard: Fortran 95 and later

8.171 LBOUND — Lower dimension bounds of an array

Synopsis: `RESULT = LBOUND(ARRAY [, DIM [, KIND]])`

Description:

Returns the lower bounds of an array, or a single lower bound along the *DIM* dimension.

Class: Inquiry function

Arguments:

ARRAY Shall be an array, of any type.
DIM (Optional) Shall be a scalar **INTEGER**.
KIND (Optional) A scalar **INTEGER** constant expression indicating the kind parameter of the result.

Return value:

The return value is of type **INTEGER** and of kind *KIND*. If *KIND* is absent, the return value is of default integer kind. If *DIM* is absent, the result is an array of the lower bounds of *ARRAY*. If *DIM* is present, the result is a scalar corresponding to the lower bound of the array along that dimension. If *ARRAY* is an expression rather than a whole array or array structure component, or if it has a zero extent along the relevant dimension, the lower bound is taken to be 1.

Standard: Fortran 90 and later, with *KIND* argument Fortran 2003 and later

See also: Section 8.290 [UBOUND], page 305,
 Section 8.172 [LCOBOUND], page 233,

8.172 LCOBOUND — Lower codimension bounds of an array

Synopsis: **RESULT = LCOBOUND(COARRAY [, DIM [, KIND]])**

Description:

Returns the lower bounds of a coarray, or a single lower cobound along the *DIM* codimension.

Class: Inquiry function

Arguments:

ARRAY Shall be an coarray, of any type.
DIM (Optional) Shall be a scalar **INTEGER**.
KIND (Optional) A scalar **INTEGER** constant expression indicating the kind parameter of the result.

Return value:

The return value is of type **INTEGER** and of kind *KIND*. If *KIND* is absent, the return value is of default integer kind. If *DIM* is absent, the result is an array of the lower cobounds of *COARRAY*. If *DIM* is present, the result is a scalar corresponding to the lower cobound of the array along that codimension.

Standard: Fortran 2008 and later

See also: Section 8.291 [UCOBUND], page 306,
 Section 8.171 [LBOUND], page 232,

Return value:

Returns `.TRUE.` if `STRING_A >= STRING_B`, and `.FALSE.` otherwise, based on the ASCII ordering.

Specific names:

Name	Argument	Return type	Standard
<code>LGE(STRING_A,STRING_B)</code>	CHARACTER	LOGICAL	Fortran 77 and later

Standard: Fortran 77 and later

See also: Section 8.177 [LGT], page 236,
 Section 8.179 [LLE], page 237,
 Section 8.180 [LLT], page 238,

8.177 LGT — Lexical greater than

Synopsis: `RESULT = LGT(STRING_A, STRING_B)`

Description:

Determines whether one string is lexically greater than another string, where the two strings are interpreted as containing ASCII character codes. If the String A and String B are not the same length, the shorter is compared as if spaces were appended to it to form a value that has the same length as the longer.

In general, the lexical comparison intrinsics `LGE`, `LGT`, `LLE`, and `LLT` differ from the corresponding intrinsic operators `.GE.`, `.GT.`, `.LE.`, and `.LT.`, in that the latter use the processor's character ordering (which is not ASCII on some targets), whereas the former always use the ASCII ordering.

Class: Elemental function

Arguments:

`STRING_A` Shall be of default CHARACTER type.
`STRING_B` Shall be of default CHARACTER type.

Return value:

Returns `.TRUE.` if `STRING_A > STRING_B`, and `.FALSE.` otherwise, based on the ASCII ordering.

Specific names:

Name	Argument	Return type	Standard
<code>LGT(STRING_A,STRING_B)</code>	CHARACTER	LOGICAL	Fortran 77 and later

Standard: Fortran 77 and later

See also: Section 8.176 [LGE], page 235,
 Section 8.179 [LLE], page 237,
 Section 8.180 [LLT], page 238,

8.178 LINK — Create a hard link

Synopsis:

```
CALL LINK(PATH1, PATH2 [, STATUS])
STATUS = LINK(PATH1, PATH2)
```

Description:

Makes a (hard) link from file *PATH1* to *PATH2*. A null character (CHAR(0)) can be used to mark the end of the names in *PATH1* and *PATH2*; otherwise, trailing blanks in the file names are ignored. If the *STATUS* argument is supplied, it contains 0 on success or a nonzero error code upon return; see `link(2)`.

This intrinsic is provided in both subroutine and function forms; however, only one form can be used in any given program unit.

Class: Subroutine, function

Arguments:

PATH1 Shall be of default CHARACTER type.
PATH2 Shall be of default CHARACTER type.
STATUS (Optional) Shall be of default INTEGER type.

Standard: GNU extension

See also: Section 8.273 [SYMLNK], page 295,
 Section 8.296 [UNLINK], page 308,

8.179 LLE — Lexical less than or equal

Synopsis: RESULT = LLE(STRING_A, STRING_B)

Description:

Determines whether one string is lexically less than or equal to another string, where the two strings are interpreted as containing ASCII character codes. If the String A and String B are not the same length, the shorter is compared as if spaces were appended to it to form a value that has the same length as the longer.

In general, the lexical comparison intrinsics LGE, LGT, LLE, and LLT differ from the corresponding intrinsic operators .GE., .GT., .LE., and .LT., in that the latter use the processor's character ordering (which is not ASCII on some targets), whereas the former always use the ASCII ordering.

Class: Elemental function

Arguments:

STRING_A Shall be of default CHARACTER type.
STRING_B Shall be of default CHARACTER type.

Return value:

Returns .TRUE. if *STRING_A* <= *STRING_B*, and .FALSE. otherwise, based on the ASCII ordering.

Specific names:

Name	Argument	Return type	Standard
LLE(String_A,String_B)	CHARACTER	LOGICAL	Fortran 77 and later

Standard: Fortran 77 and later

See also: Section 8.176 [LGE], page 235,
Section 8.177 [LGT], page 236,
Section 8.180 [LLT], page 238,

8.180 LLT — Lexical less than

Synopsis: RESULT = LLT(String_A, String_B)

Description:

Determines whether one string is lexically less than another string, where the two strings are interpreted as containing ASCII character codes. If the String A and String B are not the same length, the shorter is compared as if spaces were appended to it to form a value that has the same length as the longer.

In general, the lexical comparison intrinsics LGE, LGT, LLE, and LLT differ from the corresponding intrinsic operators .GE., .GT., .LE., and .LT., in that the latter use the processor's character ordering (which is not ASCII on some targets), whereas the former always use the ASCII ordering.

Class: Elemental function

Arguments:

String_A Shall be of default CHARACTER type.
String_B Shall be of default CHARACTER type.

Return value:

Returns .TRUE. if String_A < String_B, and .FALSE. otherwise, based on the ASCII ordering.

Specific names:

Name	Argument	Return type	Standard
LLT(String_A,String_B)	CHARACTER	LOGICAL	Fortran 77 and later

Standard: Fortran 77 and later

See also: Section 8.176 [LGE], page 235,
Section 8.177 [LGT], page 236,
Section 8.179 [LLE], page 237,

8.181 LNBLNK — Index of the last non-blank character in a string

Synopsis: RESULT = LNBLNK(String)

VALUES The type shall be `INTEGER, DIMENSION(13)` of either kind 4 or kind 8.

STATUS (Optional) status flag of type `INTEGER` of kind 2 or larger. Returns 0 on success and a system specific error code otherwise.

Example: See Section 8.270 [STAT], page 292, for an example.

Standard: GNU extension

See also: To stat an open file:
 Section 8.122 [FSTAT], page 202,
 To stat a file:
 Section 8.270 [STAT], page 292,

8.189 LTIME — Convert time to local time info

Synopsis: `CALL LTIME(TIME, VALUES)`

Description:

Given a system time value *TIME* (as provided by the Section 8.282 [TIME], page 301, intrinsic), fills *VALUES* with values extracted from it appropriate to the local time zone using `localtime(3)`.

This intrinsic routine is provided for backwards compatibility with GNU Fortran 77. In new code, programmers should consider the use of the Section 8.87 [DATE_AND_TIME], page 178, intrinsic defined by the Fortran 95 standard.

Class: Subroutine

Arguments:

TIME An `INTEGER` scalar expression corresponding to a system time, with `INTENT(IN)`.

VALUES A default `INTEGER` array with 9 elements, with `INTENT(OUT)`.

Return value:

The elements of *VALUES* are assigned as follows:

1. Seconds after the minute, range 0–59 or 0–61 to allow for leap seconds
2. Minutes after the hour, range 0–59
3. Hours past midnight, range 0–23
4. Day of month, range 1–31
5. Number of months since January, range 0–11
6. Years since 1900
7. Number of days since Sunday, range 0–6
8. Days since January 1, range 0–365
9. Daylight savings indicator: positive if daylight savings is in effect, zero if not, and negative if the information is not available.

Standard: GNU extension

See also: Section 8.87 [DATE_AND_TIME], page 178,
 Section 8.86 [CTIME], page 177,
 Section 8.137 [GMTIME], page 211,
 Section 8.282 [TIME], page 301,
 Section 8.283 [TIME8], page 301,

8.190 MALLOC — Allocate dynamic memory

Synopsis: PTR = MALLOC(SIZE)

Description:

MALLOC(SIZE) allocates SIZE bytes of dynamic memory and returns the address of the allocated memory. The MALLOC intrinsic is an extension intended to be used with Cray pointers, and is provided in GNU Fortran to allow the user to compile legacy code. For new code using Fortran 95 pointers, the memory allocation intrinsic is ALLOCATE.

Class: Function

Arguments:

SIZE The type shall be INTEGER.

Return value:

The return value is of type INTEGER(K), with K such that variables of type INTEGER(K) have the same size as C pointers (sizeof(void *)).

Example: The following example demonstrates the use of MALLOC and FREE with Cray pointers.

```

program test_malloc
  implicit none
  integer i
  real*8 x(*), z
  pointer(ptr_x,x)

  ptr_x = malloc(20*8)
  do i = 1, 20
    x(i) = sqrt(1.0d0 / i)
  end do
  z = 0
  do i = 1, 20
    z = z + x(i)
  end do
  print *, z
  call free(ptr_x)
end program test_malloc

```

Standard: GNU extension

See also: Section 8.120 [FREE], page 200,

8.191 MASKL — Left justified mask

Synopsis: RESULT = MASKL(I[, KIND])

8.195 MAXEXPONENT — Maximum exponent of a real kind

Synopsis: `RESULT = MAXEXPONENT(X)`

Description:

`MAXEXPONENT(X)` returns the maximum exponent in the model of the type of `X`.

Class: Inquiry function

Arguments:

`X` Shall be of type `REAL`.

Return value:

The return value is of type `INTEGER` and of the default integer kind.

Example:

```
program exponents
  real(kind=4) :: x
  real(kind=8) :: y

  print *, minexponent(x), maxexponent(x)
  print *, minexponent(y), maxexponent(y)
end program exponents
```

Standard: Fortran 90 and later

8.196 MAXLOC — Location of the maximum value within an array

Synopsis:

```
RESULT = MAXLOC(ARRAY, DIM [, MASK] [,KIND] [,BACK])
RESULT = MAXLOC(ARRAY [, MASK] [,KIND] [,BACK])
```

Description:

Determines the location of the element in the array with the maximum value, or, if the *DIM* argument is supplied, determines the locations of the maximum element along each row of the array in the *DIM* direction. If *MASK* is present, only the elements for which *MASK* is `.TRUE.` are considered. If more than one element in the array has the maximum value, the location returned is that of the first such element in array element order if the *BACK* is not present, or is false; if *BACK* is true, the location returned is that of the last such element. If the array has zero size, or all of the elements of *MASK* are `.FALSE.`, then the result is an array of zeroes. Similarly, if *DIM* is supplied and all of the elements of *MASK* along a given row are zero, the result value for that row is zero.

Class: Transformational function

Arguments:

ARRAY Shall be an array of type `INTEGER`, `REAL`, `UNSIGNED` or `CHARACTER`.

DIM (Optional) Shall be a scalar of type `INTEGER`, with a value between one and the rank of *ARRAY*, inclusive. It may not be an optional dummy argument.

<i>MASK</i>	Shall be of type LOGICAL , and conformable with <i>ARRAY</i> .
<i>KIND</i>	(Optional) A scalar INTEGER constant expression indicating the kind parameter of the result.
<i>BACK</i>	(Optional) A scalar of type LOGICAL .

Return value:

If *DIM* is absent, the result is a rank-one array with a length equal to the rank of *ARRAY*. If *DIM* is present, the result is an array with a rank one less than the rank of *ARRAY*, and a size corresponding to the size of *ARRAY* with the *DIM* dimension removed. If *DIM* is present and *ARRAY* has a rank of one, the result is a scalar. If the optional argument *KIND* is present, the result is an integer of kind *KIND*, otherwise it is of default kind.

Standard: Fortran 95 and later; *ARRAY* of **CHARACTER** and the *KIND* argument are available in Fortran 2003 and later. The *BACK* argument is available in Fortran 2008 and later. Extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67).

See also: Section 8.113 [FINDLOC], page 195,
 Section 8.194 [MAX], page 246,
 Section 8.197 [MAXVAL], page 248,

8.197 MAXVAL — Maximum value of an array

Synopsis:

```
RESULT = MAXVAL(ARRAY, DIM [, MASK])
RESULT = MAXVAL(ARRAY [, MASK])
```

Description:

Determines the maximum value of the elements in an array value, or, if the *DIM* argument is supplied, determines the maximum value along each row of the array in the *DIM* direction. If *MASK* is present, only the elements for which *MASK* is **.TRUE.** are considered. If the array has zero size, or all of the elements of *MASK* are **.FALSE.**, then the result is **-HUGE(ARRAY)** if *ARRAY* is of type **INTEGER** or **REAL**, 0 if it is type **UNSIGNED**. or a string of nulls if *ARRAY* is of character type.

Class: Transformational function

Arguments:

<i>ARRAY</i>	Shall be an array of type INTEGER , REAL , UNSIGNED or CHARACTER .
<i>DIM</i>	(Optional) Shall be a scalar of type INTEGER , with a value between one and the rank of <i>ARRAY</i> , inclusive. It may not be an optional dummy argument.
<i>MASK</i>	(Optional) Shall be of type LOGICAL , and conformable with <i>ARRAY</i> .

Description:

Returns a disassociated pointer.

If *MOLD* is present, a disassociated pointer of the same type is returned, otherwise the type is determined by context.

In Fortran 95, *MOLD* is optional. Please note that Fortran 2003 includes cases where it is required.

Class: Transformational function

Arguments:

MOLD (Optional) shall be a pointer of any association status and of any type.

Return value:

A disassociated pointer.

Example:

```
REAL, POINTER, DIMENSION(:) :: VEC => NULL ()
```

Standard: Fortran 95 and later

See also: Section 8.24 [ASSOCIATED], page 134,

8.216 NUM_IMAGES — Function that returns the number of images

Synopsis:

```
RESULT = NUM_IMAGES([TEAM])
RESULT = NUM_IMAGES(TEAM_NUMBER)
```

Description:

Returns the number of images in the current team or the given team.

Class: Transformational function

Arguments:

TEAM (optional, intent(in)) If present, return the number of images in the given team; if absent, return the number of images in the current team.

TEAM_NUMBER (intent(in)) The number as given in the **FORM TEAM** statement.

Return value:

Scalar default-kind integer. Can be called without any arguments or a team type argument or a team_number argument.

Example:

```
use, intrinsic :: iso_fortran_env
INTEGER :: value[*]
INTEGER :: i
type(team_type) :: t

! When running with 4 images
```


Class: Transformational function

Arguments:

ARRAY Shall be an array of any type.
MASK Shall be an array of type **LOGICAL** and of the same size as *ARRAY*. Alternatively, it may be a **LOGICAL** scalar.
VECTOR (Optional) shall be an array of the same type as *ARRAY* and of rank one. If present, the number of elements in *VECTOR* shall be equal to or greater than the number of true elements in *MASK*. If *MASK* is scalar, the number of elements in *VECTOR* shall be equal to or greater than the number of elements in *ARRAY*.

Return value:

The result is an array of rank one and the same type as that of *ARRAY*. If *VECTOR* is present, the result size is that of *VECTOR*, the number of **TRUE** values in *MASK* otherwise.

Example: Gathering nonzero elements from an array:

```
PROGRAM test_pack_1
  INTEGER :: m(6)
  m = (/ 1, 0, 0, 0, 5, 0 /)
  WRITE(*, FMT="(6(I0, ' '))") pack(m, m /= 0) ! "1 5"
END PROGRAM
```

Gathering nonzero elements from an array and appending elements from *VECTOR*:

```
PROGRAM test_pack_2
  INTEGER :: m(4)
  m = (/ 1, 0, 0, 2 /)
  ! The following results in "1 2 3 4"
  WRITE(*, FMT="(4(I0, ' '))") pack(m, m /= 0, (/ 0, 0, 3, 4 /))
END PROGRAM
```

Standard: Fortran 90 and later

See also: Section 8.297 [UNPACK], page 308,

8.220 PARITY — Reduction with exclusive OR

Synopsis:

```
RESULT = PARITY(MASK[, DIM])
```

Description:

Calculates the parity, i.e. the reduction using **.XOR.**, of *MASK* along dimension *DIM*.

Class: Transformational function

Arguments:

MASK Shall be an array of type **LOGICAL**

PUT (Optional) Shall be an array of type default **INTEGER** and rank one. It is **INTENT(IN)** and the size of the array must be larger than or equal to the number returned by the *SIZE* argument.

GET (Optional) Shall be an array of type default **INTEGER** and rank one. It is **INTENT(OUT)** and the size of the array must be larger than or equal to the number returned by the *SIZE* argument.

Example:

```

program test_random_seed
  implicit none
  integer, allocatable :: seed(:)
  integer :: n

  call random_seed(size = n)
  allocate(seed(n))
  call random_seed(get=seed)
  write (*, *) seed
end program test_random_seed

```

Standard: Fortran 90 and later

See also: Section 8.231 [RANDOM_NUMBER], page 269,
Section 8.230 [RANDOM_INIT], page 268,

8.233 RANGE — Decimal exponent range

Synopsis: **RESULT = RANGE(X)**

Description:

RANGE(X) returns the decimal exponent range in the model of the type of **X**.

Class: Inquiry function

Arguments:

X Shall be of type **INTEGER**, **REAL**, **COMPLEX** or **UNSIGNED**.

Return value:

The return value is of type **INTEGER** and of the default integer kind.

Example: See **PRECISION** for an example.

Standard: Fortran 90 and later, extension for **UNSIGNED** (see Section 5.1.34 [Unsigned integers], page 67)

See also: Section 8.249 [SELECTED_REAL_KIND], page 280,
Section 8.224 [PRECISION], page 265,

8.234 RANK — Rank of a data object

Synopsis: **RESULT = RANK(A)**

Description:

RANK(A) returns the rank of a scalar or array data object.

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.

- 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.

However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An “entity transaction” is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party’s predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents.

A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor’s “contributor version”.

A contributor’s “essential patent claims” are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, “control” includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor’s essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a “patent license” is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To “grant” such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so

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.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

```
one line to give the program's name and a brief idea of what it does.
Copyright (C) year name of author
```

```
This program is free software: you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation, either version 3 of the License, or (at
your option) any later version.
```

```
This program is distributed in the hope that it will be useful, but
WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program. If not, see https://www.gnu.org/licenses/.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
program Copyright (C) year name of author
This program comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type 'show c' for details.
```

The hypothetical commands ‘show w’ and ‘show c’ should show the appropriate parts of the General Public License. Of course, your program’s commands might be different; for a GUI interface, you would use an “about box”.

You should also get your employer (if you work as a programmer) or school, if any, to sign a “copyright disclaimer” for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <https://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <https://www.gnu.org/licenses/why-not-lgpl.html>.

GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 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.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any,

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.

type alias print 64
 type cast 303

U

UBOUND 305
 UCOBOUND 306
 UINT 306
 UMASK 306
 UMASKL 307
 UMASKR 307
 underflow 21
 underscore 27, 28
 unformatted sequential 46
 UNION 61
 UNLINK 308
 UNPACK 308
 UNROLL directive 80
 Unsigned integers 67
 unsigned kind 281
 unused dummy argument 21
 unused parameter 21
 user id 211

V

variable attributes 64
 variable interoperability with C 74
 Varying length strings 3
 vector product 181
 VECTOR directive 80
 VERIFY 309
 version of the compiler 169
 VOLATILE 71

W

warning, C binding type 19
 warnings, aliasing 19
 warnings, alignment of COMMON blocks 21
 warnings, all 18
 warnings, ampersand 19
 warnings, argument mismatch 19
 warnings, array temporaries 19
 warnings, character truncation 19
 warnings, conversion 19

warnings, division of integers 20
 warnings, extra 20
 warnings, function elimination 21
 warnings, implicit interface 20
 warnings, implicit procedure 20
 warnings, integer division 20
 warnings, intrinsic 21
 warnings, intrinsics of other standards 20
 warnings, line truncation 19
 warnings, loop interchange 20
 warnings, nonstandard intrinsics 20
 warnings, overwrite recursive 20
 warnings, q exponent-letter 20
 warnings, suppressing 18
 warnings, suspicious code 20
 warnings, tabs 21
 warnings, to errors 22
 warnings, undefined do loop 21
 warnings, underflow 21
 warnings, unused dummy argument 21
 warnings, unused parameter 21
 warnings, use statements 21
 write character, stream mode 198, 199

X

XOR 310
 XOR reduction 263

Z

ZABS 120
 ZCOS 171
 ZCOSD 171
 zero bits 234, 302
 ZEXP 191
 ZLOG 239
 ZSIN 285
 ZSIND 286
 ZSQRT 291