

The C Preprocessor

For GCC version 16.0.0 (pre-release)

(GCC)

Richard M. Stallman, Zachary Weinberg

Copyright © 1987-2025 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. A copy of the license is included in the section entitled “GNU Free Documentation License”.

This manual contains no Invariant Sections. The Front-Cover Texts are (a) (see below), and the Back-Cover Texts are (b) (see below).

(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.

Table of Contents

1	Overview	1
1.1	Character sets	1
1.2	Initial processing	2
1.3	Tokenization	4
1.4	The preprocessing language	6
2	Header Files	7
2.1	Include Syntax	7
2.2	Include Operation	8
2.3	Search Path	9
2.4	Once-Only Headers	9
2.5	Alternatives to Wrapper <code>#ifndef</code>	10
2.6	Computed Includes	10
2.7	Wrapper Headers	11
2.8	System Headers	12
3	Macros	13
3.1	Object-like Macros	13
3.2	Function-like Macros	14
3.3	Macro Arguments	15
3.4	Stringizing	16
3.5	Concatenation	17
3.6	Variadic Macros	19
3.7	Predefined Macros	20
3.7.1	Standard Predefined Macros	20
3.7.2	Common Predefined Macros	22
3.7.3	System-specific Predefined Macros	34
3.7.4	C++ Named Operators	34
3.8	Undefining and Redefining Macros	35
3.9	Directives Within Macro Arguments	36
3.10	Macro Pitfalls	36
3.10.1	Misnesting	36
3.10.2	Operator Precedence Problems	36
3.10.3	Swallowing the Semicolon	37
3.10.4	Duplication of Side Effects	38
3.10.5	Self-Referential Macros	39
3.10.6	Argument Prescan	40
3.10.7	Newlines in Arguments	41
4	Conditionals	41
4.1	Conditional Uses	42
4.2	Conditional Syntax	42

GNU Free Documentation License	72
ADDENDUM: How to use this License for your documents	79
Index of Directives	80
Option Index	80
Concept Index	82

All preprocessing work (the subject of the rest of this manual) is carried out in the source character set. If you request textual output from the preprocessor with the `-E` option, it will be in UTF-8.

After preprocessing is complete, string and character constants are converted again, into the *execution* character set. This character set is under control of the user; the default is UTF-8, matching the source character set. Wide string and character constants have their own character set, which is not called out specifically in the standard. Again, it is under control of the user. The default is UTF-16 or UTF-32, whichever fits in the target’s `wchar_t` type, in the target machine’s byte order.¹ Octal and hexadecimal escape sequences do not undergo conversion; `‘\x12’` has the value 0x12 regardless of the currently selected execution character set. All other escapes are replaced by the character in the source character set that they represent, then converted to the execution character set, just like unescaped characters.

In identifiers, characters outside the ASCII range can be specified with the `‘\u’` and `‘\U’` escapes or used directly in the input encoding. If strict ISO C90 conformance is specified with an option such as `-std=c90`, or `-fno-extended-identifiers` is used, then those constructs are not permitted in identifiers.

1.2 Initial processing

The preprocessor performs a series of textual transformations on its input. These happen before all other processing. Conceptually, they happen in a rigid order, and the entire file is run through each transformation before the next one begins. CPP actually does them all at once, for performance reasons. These transformations correspond roughly to the first three “phases of translation” described in the C standard.

1. The input file is read into memory and broken into lines.

Different systems use different conventions to indicate the end of a line. GCC accepts the ASCII control sequences *LF*, *CR LF* and *CR* as end-of-line markers. These are the canonical sequences used by Unix, DOS and VMS, and the classic Mac OS (before OSX) respectively. You may therefore safely copy source code written on any of those systems to a different one and use it without conversion. (GCC may lose track of the current line number if a file doesn’t consistently use one convention, as sometimes happens when it is edited on computers with different conventions that share a network file system.)

If the last line of any input file lacks an end-of-line marker, the end of the file is considered to implicitly supply one. The C standard says that this condition provokes undefined behavior, so GCC will emit a warning message.

2. If trigraphs are enabled, they are replaced by their corresponding single characters. By default GCC ignores trigraphs, but if you request a strictly conforming mode with the `-std` option, or you specify the `-trigraphs` option, then it converts them.

These are nine three-character sequences, all starting with `‘??’`, that are defined by ISO C to stand for single characters. They permit obsolete systems that lack some of

¹ UTF-16 does not meet the requirements of the C standard for a wide character set, but the choice of 16-bit `wchar_t` is enshrined in some system ABIs so we cannot fix this.

Line comments are not in the 1989 edition of the C standard, but they are recognized by GCC as an extension. In C++ and in the 1999 edition of the C standard, they are an official part of the language.

Since these transformations happen before all other processing, you can split a line mechanically with backslash-newline anywhere. You can comment out the end of a line. You can continue a line comment onto the next line with backslash-newline. You can even split `/*`, `*/`, and `/**` onto multiple lines with backslash-newline. For example:

```

/\
*
*/ # /*
*/ defi\
ne F0\
0 10\
20

```

is equivalent to `#define F00 1020`. All these tricks are extremely confusing and should not be used in code intended to be readable.

There is no way to prevent a backslash at the end of a line from being interpreted as a backslash-newline. This cannot affect any correct program, however.

1.3 Tokenization

After the textual transformations are finished, the input file is converted into a sequence of *preprocessing tokens*. These mostly correspond to the syntactic tokens used by the C compiler, but there are a few differences. White space separates tokens; it is not itself a token of any kind. Tokens do not have to be separated by white space, but it is often necessary to avoid ambiguities.

When faced with a sequence of characters that has more than one possible tokenization, the preprocessor is greedy. It always makes each token, starting from the left, as big as possible before moving on to the next token. For instance, `a+++++b` is interpreted as `a ++ ++ + b`, not as `a ++ + ++ b`, even though the latter tokenization could be part of a valid C program and the former could not.

Once the input file is broken into tokens, the token boundaries never change, except when the `##` preprocessing operator is used to paste tokens together. See Section 3.5 [Concatenation], page 17. For example,

```

#define foo() bar
foo()baz
    ↪ bar baz
not
    ↪ barbaz

```

The compiler does not re-tokenize the preprocessor's output. Each preprocessing token becomes one compiler token.

Preprocessing tokens fall into five broad classes: identifiers, preprocessing numbers, string literals, punctuators, and other. An *identifier* is the same as an identifier in C: any sequence of letters, digits, or underscores, which begins with a letter or underscore. Keywords of C have no significance to the preprocessor; they are ordinary identifiers. You can define a macro whose name is a keyword, for instance. The only identifier which can be considered a preprocessing keyword is `defined`. See Section 4.2.3 [Defined], page 44.

directive name. It specifies the operation to perform. Directives are commonly referred to as `#name` where *name* is the directive name. For example, `#define` is the directive that defines a macro.

The `#` which begins a directive cannot come from a macro expansion. Also, the directive name is not macro expanded. Thus, if `foo` is defined as a macro expanding to `define`, that does not make `#foo` a valid preprocessing directive.

The set of valid directive names is fixed. Programs cannot define new preprocessing directives.

Some directives require arguments; these make up the rest of the directive line and must be separated from the directive name by whitespace. For example, `#define` must be followed by a macro name and the intended expansion of the macro.

A preprocessing directive cannot cover more than one line. The line may, however, be continued with backslash-newline, or by a block comment which extends past the end of the line. In either case, when the directive is processed, the continuations have already been merged with the first line to make one long line.

2 Header Files

A header file is a file containing C declarations and macro definitions (see Chapter 3 [Macros], page 13) to be shared between several source files. You request the use of a header file in your program by *including* it, with the C preprocessing directive `#include`.

Header files serve two purposes.

- System header files declare the interfaces to parts of the operating system. You include them in your program to supply the definitions and declarations you need to invoke system calls and libraries.
- Your own header files contain declarations for interfaces between the source files of your program. Each time you have a group of related declarations and macro definitions all or most of which are needed in several different source files, it is a good idea to create a header file for them.

Including a header file produces the same results as copying the header file into each source file that needs it. Such copying would be time-consuming and error-prone. With a header file, the related declarations appear in only one place. If they need to be changed, they can be changed in one place, and programs that include the header file will automatically use the new version when next recompiled. The header file eliminates the labor of finding and changing all the copies as well as the risk that a failure to find one copy will result in inconsistencies within a program.

In C, the usual convention is to give header files names that end with `.h`. It is most portable to use only letters, digits, dashes, and underscores in header file names, and at most one dot.

2.1 Include Syntax

Both user and system header files are included using the preprocessing directive `#include`. It has two variants:

file could even contain the beginning of a statement that is concluded in the containing file, or the end of a statement that was started in the including file. However, an included file must consist of complete tokens. Comments and string literals which have not been closed by the end of an included file are invalid. For error recovery, they are considered to end at the end of the file.

To avoid confusion, it is best if header files contain only complete syntactic units—function declarations or definitions, type declarations, etc.

The line following the `#include` directive is always treated as a separate line by the C preprocessor, even if the included file lacks a final newline.

2.3 Search Path

By default, the preprocessor looks for header files included by the quote form of the directive `#include "file"` first relative to the directory of the current file, and then in a preconfigured list of standard system directories. For example, if `/usr/include/sys/stat.h` contains `#include "types.h"`, GCC looks for `types.h` first in `/usr/include/sys`, then in its usual search path.

For the angle-bracket form `#include <file>`, the preprocessor's default behavior is to look only in the standard system directories. The exact search directory list depends on the target system, how GCC is configured, and where it is installed. You can find the default search directory list for your version of CPP by invoking it with the `-v` option. For example,

```
cpp -v /dev/null -o /dev/null
```

There are a number of command-line options you can use to add additional directories to the search path. The most commonly-used option is `-Idir`, which causes *dir* to be searched after the current directory (for the quote form of the directive) and ahead of the standard system directories. You can specify multiple `-I` options on the command line, in which case the directories are searched in left-to-right order.

If you need separate control over the search paths for the quote and angle-bracket forms of the `#include` directive, you can use the `-iquote` and/or `-isystem` options instead of `-I`. See Chapter 13 [Invocation], page 60, for a detailed description of these options, as well as others that are less generally useful.

If you specify other options on the command line, such as `-I`, that affect where the preprocessor searches for header files, the directory list printed by the `-v` option reflects the actual search path used by the preprocessor.

Note that you can also prevent the preprocessor from searching any of the default system header directories with the `-nostdinc` option. This is useful when you are compiling an operating system kernel or some other program that does not use the standard C library facilities, or the standard C library itself.

2.4 Once-Only Headers

If a header file happens to be included twice, the compiler will process its contents twice. This is very likely to cause an error, e.g. when the compiler sees the same structure definition twice. Even if it does not, it will certainly waste time.

The standard way to prevent this is to enclose the entire real contents of the file in a conditional, like this:

- There is also a directive, `#pragma GCC system_header`, which tells GCC to consider the rest of the current include file a system header, no matter where it was found. Code that comes before the `#pragma` in the file is not affected. `#pragma GCC system_header` has no effect in the primary source file.

On some targets, such as RS/6000 AIX, GCC implicitly surrounds all system headers with an `'extern "C"'` block when compiling as C++.

3 Macros

A *macro* is a fragment of code which has been given a name. Whenever the name is used, it is replaced by the contents of the macro. There are two kinds of macros. They differ mostly in what they look like when they are used. *Object-like* macros resemble data objects when used, *function-like* macros resemble function calls.

You may define any valid identifier as a macro, even if it is a C keyword. The preprocessor does not know anything about keywords. This can be useful if you wish to hide a keyword such as `const` from an older compiler that does not understand it. However, the preprocessor operator `defined` (see Section 4.2.3 [Defined], page 44) can never be defined as a macro, and C++'s named operators (see Section 3.7.4 [C++ Named Operators], page 34) cannot be macros when you are compiling C++.

3.1 Object-like Macros

An *object-like macro* is a simple identifier which will be replaced by a code fragment. It is called object-like because it looks like a data object in code that uses it. They are most commonly used to give symbolic names to numeric constants.

You create macros with the `'#define'` directive. `'#define'` is followed by the name of the macro and then the token sequence it should be an abbreviation for, which is variously referred to as the macro's *body*, *expansion* or *replacement list*. For example,

```
#define BUFFER_SIZE 1024
```

defines a macro named `BUFFER_SIZE` as an abbreviation for the token `1024`. If somewhere after this `'#define'` directive there comes a C statement of the form

```
foo = (char *) malloc (BUFFER_SIZE);
```

then the C preprocessor will recognize and *expand* the macro `BUFFER_SIZE`. The C compiler will see the same tokens as it would if you had written

```
foo = (char *) malloc (1024);
```

By convention, macro names are written in uppercase. Programs are easier to read when it is possible to tell at a glance which names are macros.

The macro's body ends at the end of the `'#define'` line. You may continue the definition onto multiple lines, if necessary, using backslash-newline. When the macro is expanded, however, it will all come out on one line. For example,

```
#define NUMBERS 1, \
                2, \
                3
int x[] = { NUMBERS };
↪ int x[] = { 1, 2, 3 };
```


preprocessing number. When pasted, they make a longer identifier. This isn't the only valid case. It is also possible to concatenate two numbers (or a number and a name, such as 1.5 and e3) into a number. Also, multi-character operators such as += can be formed by token pasting.

However, two tokens that don't together form a valid token cannot be pasted together. For example, you cannot concatenate x with + in either order. If you try, the preprocessor issues a warning and emits the two tokens. Whether it puts white space between the tokens is undefined. It is common to find unnecessary uses of '##' in complex macros. If you get this warning, it is likely that you can simply remove the '##'.

Both the tokens combined by '##' could come from the macro body, but you could just as well write them as one token in the first place. Token pasting is most useful when one or both of the tokens comes from a macro argument. If either of the tokens next to an '##' is a parameter name, it is replaced by its actual argument before '##' executes. As with stringizing, the actual argument is not macro-expanded first. If the argument is empty, that '##' has no effect.

Keep in mind that the C preprocessor converts comments to whitespace before macros are even considered. Therefore, you cannot create a comment by concatenating '/' and '*'. You can put as much whitespace between '##' and its operands as you like, including comments, and you can put comments in arguments that will be concatenated. However, it is an error if '##' appears at either end of a macro body.

Consider a C program that interprets named commands. There probably needs to be a table of commands, perhaps an array of structures declared as follows:

```
struct command
{
    char *name;
    void (*function) (void);
};

struct command commands[] =
{
    { "quit", quit_command },
    { "help", help_command },
    ...
};
```

It would be cleaner not to have to give each command name twice, once in the string constant and once in the function name. A macro which takes the name of a command as an argument can make this unnecessary. The string constant can be created with stringizing, and the function name by concatenating the argument with '_command'. Here is how it is done:

```
#define COMMAND(NAME) { #NAME, NAME ## _command }

struct command commands[] =
{
    COMMAND (quit),
    COMMAND (help),
    ...
};
```


__REGISTER_PREFIX__

This macro expands to a single token (not a string constant) which is the prefix applied to CPU register names in assembly language for this target. You can use it to write assembly that is usable in multiple environments. For example, in the **m68k-aout** environment it expands to nothing, but in the **m68k-coff** environment it expands to a single **'%'**.

__USER_LABEL_PREFIX__

This macro expands to a single token which is the prefix applied to user labels (symbols visible to C code) in assembly. For example, in the **m68k-aout** environment it expands to an **'_'**, but in the **m68k-coff** environment it expands to nothing.

This macro will have the correct definition even if **-f(no-)underscores** is in use, but it will not be correct if target-specific options that adjust this prefix are used (e.g. the OSF/rose **-mno-underscores** option).

```

__SIZE_TYPE__
__PTRDIFF_TYPE__
__WCHAR_TYPE__
__WINT_TYPE__
__INTMAX_TYPE__
__UINTMAX_TYPE__
__SIG_ATOMIC_TYPE__
__INT8_TYPE__
__INT16_TYPE__
__INT32_TYPE__
__INT64_TYPE__
__UINT8_TYPE__
__UINT16_TYPE__
__UINT32_TYPE__
__UINT64_TYPE__
__INT_LEAST8_TYPE__
__INT_LEAST16_TYPE__
__INT_LEAST32_TYPE__
__INT_LEAST64_TYPE__
__UINT_LEAST8_TYPE__
__UINT_LEAST16_TYPE__
__UINT_LEAST32_TYPE__
__UINT_LEAST64_TYPE__
__INT_FAST8_TYPE__
__INT_FAST16_TYPE__
__INT_FAST32_TYPE__
__INT_FAST64_TYPE__
__UINT_FAST8_TYPE__
__UINT_FAST16_TYPE__
__UINT_FAST32_TYPE__
__UINT_FAST64_TYPE__
__INTPTR_TYPE__
__UINTPTR_TYPE__

```

These macros are defined to the correct underlying types for the `size_t`, `ptrdiff_t`, `wchar_t`, `wint_t`, `intmax_t`, `uintmax_t`, `sig_atomic_t`, `int8_t`, `int16_t`, `int32_t`, `int64_t`, `uint8_t`, `uint16_t`, `uint32_t`, `uint64_t`, `int_least8_t`, `int_least16_t`, `int_least32_t`, `int_least64_t`, `uint_least8_t`, `uint_least16_t`, `uint_least32_t`, `uint_least64_t`, `int_fast8_t`, `int_fast16_t`, `int_fast32_t`, `int_fast64_t`, `uint_fast8_t`, `uint_fast16_t`, `uint_fast32_t`, `uint_fast64_t`, `intptr_t`, and `uintptr_t` typedefs, respectively. They exist to make the standard header files `stddef.h`, `stdint.h`, and `wchar.h` work correctly. You should not use these macros directly; instead, include the appropriate headers and use the typedefs. Some of these macros may not be defined on particular systems if GCC does not provide a `stdint.h` header on those systems.

__CHAR_BIT__

Defined to the number of bits used in the representation of the `char` data type. It exists to make the standard header given numerical limits work correctly. You should not use this macro directly; instead, include the appropriate headers.

```

__SCHAR_MAX__
__WCHAR_MAX__
__SHRT_MAX__
__INT_MAX__
__LONG_MAX__
__LONG_LONG_MAX__
__WINT_MAX__
__SIZE_MAX__
__PTRDIFF_MAX__
__INTMAX_MAX__
__UINTMAX_MAX__
__SIG_ATOMIC_MAX__
__INT8_MAX__
__INT16_MAX__
__INT32_MAX__
__INT64_MAX__
__UINT8_MAX__
__UINT16_MAX__
__UINT32_MAX__
__UINT64_MAX__
__INT_LEAST8_MAX__
__INT_LEAST16_MAX__
__INT_LEAST32_MAX__
__INT_LEAST64_MAX__
__UINT_LEAST8_MAX__
__UINT_LEAST16_MAX__
__UINT_LEAST32_MAX__
__UINT_LEAST64_MAX__
__INT_FAST8_MAX__
__INT_FAST16_MAX__
__INT_FAST32_MAX__
__INT_FAST64_MAX__
__UINT_FAST8_MAX__
__UINT_FAST16_MAX__
__UINT_FAST32_MAX__
__UINT_FAST64_MAX__
__INTPTR_MAX__
__UINTPTR_MAX__
__WCHAR_MIN__
__WINT_MIN__
__SIG_ATOMIC_MIN__

```

Defined to the maximum value of the signed char, wchar_t, signed short, signed int, signed long, signed long long, wint_t, size_t, ptrdiff_t, intmax_t, uintmax_t, sig_atomic_t, int8_t, int16_t, int32_t, int64_t, uint8_t, uint16_t, uint32_t, uint64_t, int_least8_t, int_least16_t, int_least32_t, int_least64_t, uint_least8_t, uint_least16_t, uint_least32_t, uint_least64_t, int_fast8_t, int_fast16_t, int_

`fast32_t`, `int_fast64_t`, `uint_fast8_t`, `uint_fast16_t`, `uint_fast32_t`, `uint_fast64_t`, `intptr_t`, and `uintptr_t` types and to the minimum value of the `wchar_t`, `wint_t`, and `sig_atomic_t` types respectively. They exist to make the standard header given numerical limits work correctly. You should not use these macros directly; instead, include the appropriate headers. Some of these macros may not be defined on particular systems if GCC does not provide a `stdint.h` header on those systems.

```
__INT8_C
__INT16_C
__INT32_C
__INT64_C
__UINT8_C
__UINT16_C
__UINT32_C
__UINT64_C
__INTMAX_C
__UINTMAX_C
```

Defined to implementations of the standard `stdint.h` macros with the same names without the leading `__`. They exist to make the implementation of that header work correctly. You should not use these macros directly; instead, include the appropriate headers. Some of these macros may not be defined on particular systems if GCC does not provide a `stdint.h` header on those systems.

```
__SCHAR_WIDTH__
__SHRT_WIDTH__
__INT_WIDTH__
__LONG_WIDTH__
__LONG_LONG_WIDTH__
__PTRDIFF_WIDTH__
__SIG_ATOMIC_WIDTH__
__SIZE_WIDTH__
__WCHAR_WIDTH__
__WINT_WIDTH__
__INT_LEAST8_WIDTH__
__INT_LEAST16_WIDTH__
__INT_LEAST32_WIDTH__
__INT_LEAST64_WIDTH__
__INT_FAST8_WIDTH__
__INT_FAST16_WIDTH__
__INT_FAST32_WIDTH__
__INT_FAST64_WIDTH__
__INTPTR_WIDTH__
__INTMAX_WIDTH__
```

Defined to the bit widths of the corresponding types. They exist to make the implementations of `limits.h` and `stdint.h` behave correctly. You should not use these macros directly; instead, include the appropriate headers. Some of

these macros may not be defined on particular systems if GCC does not provide a `stdint.h` header on those systems.

```
__SIZEOF_INT__
__SIZEOF_LONG__
__SIZEOF_LONG_LONG__
__SIZEOF_SHORT__
__SIZEOF_POINTER__
__SIZEOF_FLOAT__
__SIZEOF_DOUBLE__
__SIZEOF_LONG_DOUBLE__
__SIZEOF_SIZE_T__
__SIZEOF_WCHAR_T__
__SIZEOF_WINT_T__
__SIZEOF_PTRDIFF_T__
```

Defined to the number of bytes of the C standard data types: `int`, `long`, `long long`, `short`, `void *`, `float`, `double`, `long double`, `size_t`, `wchar_t`, `wint_t` and `ptrdiff_t`.

```
__BYTE_ORDER__
__ORDER_LITTLE_ENDIAN__
__ORDER_BIG_ENDIAN__
__ORDER_PDP_ENDIAN__
```

`__BYTE_ORDER__` is defined to one of the values `__ORDER_LITTLE_ENDIAN__`, `__ORDER_BIG_ENDIAN__`, or `__ORDER_PDP_ENDIAN__` to reflect the layout of multi-byte and multi-word quantities in memory. If `__BYTE_ORDER__` is equal to `__ORDER_LITTLE_ENDIAN__` or `__ORDER_BIG_ENDIAN__`, then multi-byte and multi-word quantities are laid out identically: the byte (word) at the lowest address is the least significant or most significant byte (word) of the quantity, respectively. If `__BYTE_ORDER__` is equal to `__ORDER_PDP_ENDIAN__`, then bytes in 16-bit words are laid out in a little-endian fashion, whereas the 16-bit subwords of a 32-bit quantity are laid out in big-endian fashion.

You should use these macros for testing like this:

```
/* Test for a little-endian machine */
#if __BYTE_ORDER__ == __ORDER_LITTLE_ENDIAN__
```

```
__FLOAT_WORD_ORDER__
```

`__FLOAT_WORD_ORDER__` is defined to one of the values `__ORDER_LITTLE_ENDIAN__` or `__ORDER_BIG_ENDIAN__` to reflect the layout of the words of multi-word floating-point quantities.

```
__DEPRECATED
```

This macro is defined, with value 1, when compiling a C++ source file with warnings about deprecated constructs enabled. These warnings are enabled by default, but can be disabled with `-Wno-deprecated`.

```
__EXCEPTIONS
```

This macro is defined, with value 1, when compiling a C++ source file with exceptions enabled. If `-fno-exceptions` is used when compiling the file, then this macro is not defined.

__GCC_IEC_559

This macro is defined to indicate the intended level of support for IEEE 754 (IEC 60559) floating-point arithmetic. It expands to a nonnegative integer value. If 0, it indicates that the combination of the compiler configuration and the command-line options is not intended to support IEEE 754 arithmetic for `float` and `double` as defined in C99 and C11 Annex F (for example, that the standard rounding modes and exceptions are not supported, or that optimizations are enabled that conflict with IEEE 754 semantics). If 1, it indicates that IEEE 754 arithmetic is intended to be supported; this does not mean that all relevant language features are supported by GCC. If 2 or more, it additionally indicates support for IEEE 754-2008 (in particular, that the binary encodings for quiet and signaling NaNs are as specified in IEEE 754-2008).

This macro does not indicate the default state of command-line options that control optimizations that C99 and C11 permit to be controlled by standard pragmas, where those standards do not require a particular default state. It does not indicate whether optimizations respect signaling NaN semantics (the macro for that is `__SUPPORT_SNAN__`). It does not indicate support for decimal floating point or the IEEE 754 binary16 and binary128 types.

__GCC_IEC_559_COMPLEX

This macro is defined to indicate the intended level of support for IEEE 754 (IEC 60559) floating-point arithmetic for complex numbers, as defined in C99 and C11 Annex G. It expands to a nonnegative integer value. If 0, it indicates that the combination of the compiler configuration and the command-line options is not intended to support Annex G requirements (for example, because `-fcx-limited-range` was used). If 1 or more, it indicates that it is intended to support those requirements; this does not mean that all relevant language features are supported by GCC.

__NO_MATH_ERRNO__

This macro is defined if `-fno-math-errno` is used, or enabled by another option such as `-ffast-math` or by default.

__RECIPROCAL_MATH__

This macro is defined if `-freciprocal-math` is used, or enabled by another option such as `-ffast-math` or by default.

__NO_SIGNED_ZEROS__

This macro is defined if `-fno-signed-zeros` is used, or enabled by another option such as `-ffast-math` or by default.

__NO_TRAPPING_MATH__

This macro is defined if `-fno-trapping-math` is used.

__ASSOCIATIVE_MATH__

This macro is defined if `-fassociative-math` is used, or enabled by another option such as `-ffast-math` or by default.

__ROUNDING_MATH__

This macro is defined if `-frounding-math` is used.


```
compl      ~
not        !
not_eq     !=
or         ||
or_eq      |=
xor        ^
xor_eq     ^=
```

3.8 Undefined and Redefining Macros

If a macro ceases to be useful, it may be *undefined* with the ‘`#undef`’ directive. ‘`#undef`’ takes a single argument, the name of the macro to undefine. You use the bare macro name, even if the macro is function-like. It is an error if anything appears on the line after the macro name. ‘`#undef`’ has no effect if the name is not a macro.

```
#define FOO 4
x = FOO;      ↦ x = 4;
#undef FOO
x = FOO;      ↦ x = FOO;
```

Once a macro has been undefined, that identifier may be *redefined* as a macro by a subsequent ‘`#define`’ directive. The new definition need not have any resemblance to the old definition.

However, if an identifier which is currently a macro is redefined, then the new definition must be *effectively the same* as the old one. Two macro definitions are effectively the same if:

- Both are the same type of macro (object- or function-like).
- All the tokens of the replacement list are the same.
- If there are any parameters, they are the same.
- Whitespace appears in the same places in both. It need not be exactly the same amount of whitespace, though. Remember that comments count as whitespace.

These definitions are effectively the same:

```
#define FOUR (2 + 2)
#define FOUR      (2    +    2)
#define FOUR (2 /* two */ + 2)
```

but these are not:

```
#define FOUR (2 + 2)
#define FOUR ( 2+2 )
#define FOUR (2 * 2)
#define FOUR(score,and,seven,years,ago) (2 + 2)
```

If a macro is redefined with a definition that is not effectively the same as the old one, the preprocessor issues a warning and changes the macro to use the new definition. If the new definition is effectively the same, the redefinition is silently ignored. This allows, for instance, two different headers to define a common macro. The preprocessor will only complain if the definitions do not match.

- 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

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.

