

---

# The GCC Compilers

## Preface

If all you really want to know is how to compile your C or C++ program using GCC, and you don't have the time or interest in understanding what you're doing or what GCC is, you can skip most of this lesson and cut to the chase by jumping to the [Examples](#). I think you will be better off if you read the whole thing, since I believe that when you understand *what* something is, you are better able to figure out *how* to use it.

If you have never used GCC, or if you have used it without really knowing what you did, (because you were pretty much using it by rote), then you should read this. If you think you *do* understand GCC and do not use it by rote, you may still benefit from reading this; you might learn something anyway. Because I believe in the importance of historical context, I begin with a brief history of GCC.

## Brief History

Richard Stallman started the GNU Project in 1984 with the purpose of creating a free, Unix-like operating system. His motivation was to promote freedom and cooperation among users and programmers. Since Unix requires a C compiler and there were no free C compilers at the time, the GNU Project had to build a C compiler from the ground up. The Free Software Foundation was a non-profit organization created to support the work of the GNU Project.

GCC was first released in 1987. This was a significant breakthrough, being the first portable ANSI C optimizing compiler released as free software. Since that time GCC has become one of the most important tools in the development of free software.

In 1992, it was revised and released as GCC 2.0, with the added feature of a C++ compiler. It was revised again in 1997, with improved optimization and C++ support. These features became widely available in the 3.0 release of GCC in 2001.

## Languages Supported by GCC

GCC stands for “GNU Compiler Collection”. GCC is an integrated collection of compilers for several major programming languages, which as of this writing are C, C++, Objective-C, Java, Fortran, and Ada. The GNU compilers all generate machine code, not higher-level language code which is then translated via another compiler.

## Language Standards Supported by GCC

For the most part, GCC supports the major standards and provides the ability to turn them on or off. GCC itself provides features for languages like C and C++ that deviate from certain standards, but turning on the appropriate compiler options will make it check programs against the standards. Consult the manual for details.

## Command Options

### Options and File Extensions Controlling the Kind of Output

At this point, I use lowercase *gcc* instead of *GCC* because the name of the executable program is *gcc*. When *gcc* looks at the file extension for guidance as to which compiler to use and what kind of output to generate. For example, a file ending in *.c* is taken as C source code, and files ending in either *.cc*, *.cpp*, *.c++*, *.C*, and *.cxx* are taken to be C++ source code. (There are other extensions that are also taken to imply C++ source code.) Consult the manual for other extensions and languages.

**-c** Compile or assemble the source files, but do not link. The linking stage simply is not done. The ultimate output is in the form of an object file for each source file. By default, the object file name for a source file is made by replacing the suffix *.c*, *.i*, *.s*, etc., with *.o*.

E.g. `gcc -c myprog.c` produces `myprog.o`

**-S** Stop after the stage of compilation proper; do not assemble. The output is in the form of an assembler code file for each non-assembler input file specified. By default, the assembler file name for a source file is made by replacing the suffix *.c*, *.i*, etc., with *.s*.

**-E** Stop after the preprocessing stage; do not run the compiler proper. The output is in the form of preprocessed source code, which is sent to the standard output.

E.g., `gcc -E myprog.c > myprog.i`

**-o *file*** Place output in file *file*. This applies regardless to whatever sort of output is being produced, whether it is an executable file, an object file, an assembler file or preprocessed C code.

E.g., `gcc -o myprog myprog.c`

**-v** Print (on standard error output) the commands executed to run the stages of compilation. Also print the version number of the compiler driver program and of the preprocessor and the compiler proper.

**--help** Print (on the standard output) a description of the command line options understood by *gcc*.

### Compiling C++ Programs

*gcc* comes with a compiler named *g++* that specifically compiles C++ programs, regardless of the file extension. Sometimes you need to use the C++ compiler even though the file extension is not a C++ extension; in this case you need to use *g++*.

---

## **Options that Control the C Dialect**

By “dialect” is meant a specific collection of features of C. For example, the ANSI standard known as ISO90 C is a dialect of C. The full set of features supported by *gcc* is much larger than the ANSI standard, and this is also a dialect. Another dialect is obtained by adding GNU extensions to the ANSI ISO C90 standard. You can selectively remove features from the full GNU set of extensions. The basic options, however, are:

- ansi In C mode, support all ISO C90 programs. In C++ mode, remove GNU extensions that conflict with ISO C++. This turns off certain features of *gcc* that are incompatible with ISO C90 (when compiling C code), or of standard C++ (when compiling C++ code).
- std= When followed by a specific dialect designating string such as ‘c90’ or ‘gnu9x’, it specifies that dialect.
- fno-... There are many options that begin with *-fno-* and are followed by a string that represents some feature to disable. Consult the manual.

## **Options that Control Warnings**

Warnings are diagnostic messages about constructions that are not errors, but are often associated with errors. For example, when a variable is declared but never used in a program, it is possible that the programmer overlooked something, so the compiler could issue a warning when it finds such a construction. *gcc* allows you to suppress certain warnings and request others. Some of the more common options are listed below.

- w Inhibit all warning messages.
- Wall Enable all warning messages.
- Wextra Enable warnings that are not checked by *-Wall*. For example, comparing an unsigned `int` against `-1` suggests that you forgot the `int` was unsigned.

## **Options For Debugging**

- g This produces debugging information in the operating system’s native format. GDB can work with this debugging information. On most systems that use stabs format, ‘-g’ enables use of extra debugging information that only GDB can use; this extra information makes debugging work better in GDB but will probably make other debuggers crash or refuse to read the program.

E.g., `gcc -g -o myprog myprog.c`

Once you have compiled a program with debugging information, you can run it under the control of the debugger. Of course you have to learn how to use the debugger, *gdb*, from the command line, or from within a GUI-based SDK that uses it as the underlying debugger.

### **Options Controlling the Preprocessor**

Of course you are aware of the fact that when your program is compiled, the very first step that the compiler takes is to run the preprocessor, which processes all of the preprocessor directives, those lines that begin with the pound sign '#'. There are options that control how the preprocessor behaves, and these are important to know and understand. The most important are (1) how to define symbols on the command line, and (2) how to tell the preprocessor where to look for include-files.

`-D name` This predefines *name* as a macro symbol, with the value 1, or true if you want to think of it that way.

`-D name=definition` The contents of *definition* are tokenized and processed as if they appeared during translation phase three in a '#define' directive.

E.g., `gcc -D testvalue=6 -o myprog myprog.c`

is the same as if you had placed

```
#define testvalue 6
in myprog.c
```

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

```
gcc -D 'introduction=AArgH$#^*!!' -o myprog myprog.c
```

would be how to protect the metacharacters in my string from the shell.

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

`-undef` Do not predefine any system-specific or *gcc*-specific macros. The standard predefined macros remain defined.

`-I dir` Add the directory *dir* to the list of directories to be searched for header files. Directories named by '-I' are searched **before** the standard system include-directories. If the directory *dir* is a standard system include directory, the option is ignored to ensure that the default search order for system directories and the special treatment of system headers are not defeated.

Any definitions or undefinitions in the program files override definitions or undefinitions you make on the command line. Thus, if you define a symbol on the command line but within the program you undef it, it will be undef-ed from that point forward.

The `-I` option is important. It is the way to tell `gcc` that the included files are not in any standard places. It does not matter whether there is space between the `I` and the directory name.

### **Options for Linking**

When `gcc` finds unresolved symbols in your program, it has to resolve them by searching in library files. You specify nonstandard libraries using the following option.

`-llibrary`

`-l library` Search the library named `library` when linking. (The second alternative with the `library` as a separate argument is only for POSIX compliance and is not recommended.) It makes a difference where in the command you write this option; the linker searches and processes libraries and object files in the order they are specified.

**IMPORTANT:** If your program references the symbol `supersort` and that symbol is defined in the library `libgreatstuff.a`, then your command line would have to be

```
gcc -o myprog myprog.c -lgreatstuff
```

because `gcc` does not know what it has to look for until it reads `myprog.c`'s unresolved symbol list, and so it is only after seeing `myprog.c` that it will search for the symbol `supersort`. If you reverse the two on the command line, you will get a linker error.

Note that the name you supply to `-l` is not the full library name, but the name with the `'lib'` and the `'.a'` removed.

### **Options to Control Directory Search**

If you are thinking that the `-I` option controls directory search, you are correct, but it is not the only directory searching that `gcc` must do. When `gcc` links your program, it has to find all library modules that your program uses. If you are using libraries that are not in `gcc`'s library search path, then you must tell it on the command line to include the containing directories. The `-L` option is what you need:

`-L dir` This adds `dir` to `gcc`'s search path for libraries.

## **Environment Variables Affecting GCC**

Certain environment variables change the way `gcc` behaves. The ones of some importance initially are the variables that tell it where to search for included files, where to search for libraries, and where to search for dynamically loaded libraries. If you are writing programs that

are internationalized and you need to make sure that locale information is specified, you will also need to set some environment variables. The variables of interest are therefore

<code>C_INCLUDE_PATH</code>	A colon-separated list of directories in which to look for include files for C programs.
<code>CPLUS_INCLUDE_PATH</code>	A colon-separated list of directories in which to look for include files for C++ programs.
<code>LIBRARY_PATH</code>	A colon separated list of directories that the linker uses to look for static libraries.
<code>LD_LIBRARY_PATH</code>	A colon separated list of directories that the linking loader uses to look for dynamic libraries.
<code>LANG</code>	A variable that controls the locale information used when the compiler is parsing strings and comments in the program.

## Examples

Here is a collection of examples to demonstrate how to do the typical tasks. If you did not read about [Languages Supported by GCC](#) above, then make a note to yourself that GCC will accept C++ programs with any of the extensions `.cc`, `.cpp`, `.cxx`, `.c++`, and `.C`. I will use `.cpp` extensions to denote C++ programs.

### Single Source File Programs

1. To compile a program entirely contained in a single file named `fudge.c`, putting the executable code into a file named `fudge`:

```
gcc -o fudge fudge.c
```

2. To compile the C++ program `fudge.cpp` putting the executable code into `fudge`:

```
g++ -o fudge fudge.cpp
```

3. If you are really in a hurry to compile `fudge.cpp` and you don't even have the time to give the executable a name, just type:

```
g++ fudge.cpp
```

In this case, `g++` will place the executable in a file named `a.out`, overwriting any other `a.out` that existed in your current working directory. Other than sheer laziness, I cannot think of a reason anyone would want to do this.

---

## Multi-Source File Programs

4. If you have a program distributed in the three files *nuts.c*, *nuts.h* (the header file for *nuts.c*), *fudge.c*, *fudge.h* (the header file for *fudge.c*), and *ice\_cream.c*, and the main program, *sundae.c*, includes *nuts.h*, *fudge.h*, and *ice\_cream.h*, then enter the following to create the *sundae* executable:

```
gcc -o sundae nuts.c fudge.c ice_cream.c sundae.c
```

Whereas *you* may care whether the nuts precede the fudge or the ice cream in your sundae, GCC's linker does not. GCC will be just as happy if you write

```
gcc -o sundae sundae.c fudge.c ice_cream.c nuts.c
```

or any other rearrangement of the source code file names.

5. Suppose that you modified the *nuts.c* file (maybe you're using walnuts instead of pecans now). If you type the entire line written above, you will be recompiling every other file needlessly. It is much faster and more efficient to compile each of the source code files separately. This is a multi-step procedure:

```
gcc -c nuts.c fudge.c ice_cream.c sundae.c
```

```
gcc -o sundae nuts.o fudge.o ice_cream.o sundae.o
```

The first line produces **object files**, i.e. *.o* files, for each *.c* file. Thus, GCC will create *nuts.o*, *fudge.o*, *ice\_cream.o*, and *sundae.o*. The second line invokes the linker to link all of these object files together and produce the executable, *sundae*. It does not matter which order you write the object files on the line.

The linker also links any libraries that the program requires into the executable (provided that it can find them; please read [Options for Linking](#) and [Options to Control Directory Search](#) above to better understand.) Now, if you change *nuts.c* all you have to do is re-link it using the second line:

```
gcc -o sundae nuts.o fudge.o ice_cream.o sundae.o
```

Linking is usually faster than compiling, so this will save you time.

6. If you do this often enough you will discover that it is a drag to have to keep re-typing the same commands over and over. This is why Make-files were invented. But that is another lesson.