Using ld

The GNU linker

`ld` version 2
January 1994

Steve Chamberlain
Cygnus Support

Cygnus Support
steve@cygnus.com, doc@cygnus.com
Using LD, the GNU linker
Edited by Jeffrey Osier (jeffrey@cygnus.com)

Copyright © 1991, 92, 93, 94, 95, 96, 97, 1998 Free Software Foundation, Inc.

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions.

1 Overview

`ld` combines a number of object and archive files, relocates their data and ties up symbol references. Usually the last step in compiling a program is to run `ld`.

`ld` accepts Linker Command Language files written in a superset of AT&T's Link Editor Command Language syntax, to provide explicit and total control over the linking process.

This version of `ld` uses the general purpose BFD libraries to operate on object files. This allows `ld` to read, combine, and write object files in many different formats—for example, COFF or `a.out`. Different formats may be linked together to produce any available kind of object file. See Chapter 5 [BFD], page 47, for more information.

Aside from its flexibility, the GNU linker is more helpful than other linkers in providing diagnostic information. Many linkers abandon execution immediately upon encountering an error; whenever possible, `ld` continues executing, allowing you to identify other errors (or, in some cases, to get an output file in spite of the error).

2 Invocation

The GNU linker `ld` is meant to cover a broad range of situations, and to be as compatible as possible with other linkers. As a result, you have many choices to control its behavior.

2.1 Command Line Options

The linker supports a plethora of command-line options, but in actual practice few of them are used in any particular context. For instance, a frequent use of `ld` is to link standard Unix object files on a standard, supported Unix system. On such a system, to link a file `hello.o`:

```
ld -o output /lib/crt0.o hello.o -lc
```

This tells `ld` to produce a file called *output* as the result of linking the file `/lib/crt0.o` with `hello.o` and the library `libc.a`, which will come from the standard search directories. (See the discussion of the '`-l`' option below.)

The command-line options to `ld` may be specified in any order, and may be repeated at will. Repeating most options with a different argument will either have no further effect, or override prior occurrences (those further to the left on the command line) of that option. Options which may be meaningfully specified more than once are noted in the descriptions below.

Non-option arguments are objects files which are to be linked together. They may follow, precede, or be mixed in with command-line options, except that an object file argument may not be placed between an option and its argument.

Usually the linker is invoked with at least one object file, but you can specify other forms of binary input files using '`-l`', '`-R`', and the script command language. If *no* binary input files at all are specified, the linker does not produce any output, and issues the message '`No input files`'.

If the linker can not recognize the format of an object file, it will assume that it is a linker script. A script specified in this way augments the main linker script used for the link (either the default linker script or the one specified by using '`-T`'). This feature permits the linker to link against a file which appears to be an object or an archive, but actually merely defines some symbol values, or uses `INPUT` or `GROUP` to load other objects. Note that specifying a script in this way should only be used to augment the main linker script; if you want to use some command that logically can only appear once, such as the `SECTIONS` or `MEMORY` command, you must replace the default linker script using the '`-T`' option. See Chapter 3 [Commands], page 19.

For options whose names are a single letter, option arguments must either follow the option letter without intervening whitespace, or be given as separate arguments immediately following the option that requires them.

For options whose names are multiple letters, either one dash or two can precede the option name; for example, '`--oformat`' and '`--oformat`' are equivalent. Arguments to multiple-letter options must either be separated from the option name by an equals sign, or be given as separate arguments immediately following the option that requires them. For example, '`--oformat srec`' and '`--oformat=srec`' are equivalent. Unique abbreviations of the names of multiple-letter options are accepted.

`-a`*keyword*

This option is supported for HP/UX compatibility. The *keyword* argument must be one of the strings '`archive`', '`shared`', or '`default`'. '`-aarchive`' is functionally equivalent to '`-Bstatic`', and the other two keywords are functionally equivalent to '`-Bdynamic`'. This option may be used any number of times.

`-A`*architecture*
`--architecture=`*architecture*

In the current release of `ld`, this option is useful only for the Intel 960 family of architectures. In that `ld` configuration, the *architecture* argument identifies the particular architecture in the 960 family, enabling some safeguards and modifying the archive-library search path. See Section 4.2 [`ld` and the Intel 960 family], page 45, for details.

Future releases of `ld` may support similar functionality for other architecture families.

`-b` *input-format*
`--format=`*input-format*

`ld` may be configured to support more than one kind of object file. If your `ld` is configured this way, you can use the '`-b`' option to specify the binary format for input object files that follow this option on the command line. Even when `ld` is configured to support alternative object formats, you don't usually need to specify this, as `ld` should be configured to expect as a default input format the most usual format on each machine. *input-format* is a text string, the name of a particular format supported by the BFD libraries. (You can list the available binary formats with '`objdump -i`'.) See Chapter 5 [BFD], page 47.

You may want to use this option if you are linking files with an unusual binary format. You can also use '`-b`' to switch formats explicitly (when linking object files of different formats), by including '`-b` *input-format*' before each group of object files in a particular format.

The default format is taken from the environment variable `GNUTARGET`. See Section 2.2 [Environment], page 17. You can also define the input format from a script, using the command `TARGET`; see Section 3.8 [Option Commands], page 40.

`-c` *MRI-commandfile*
`--mri-script=`*MRI-commandfile*

For compatibility with linkers produced by MRI, `ld` accepts script files written in an alternate, restricted command language, described in Appendix A [MRI Compatible Script Files], page 55. Introduce MRI script files with the option '`-c`'; use the '`-T`' option to run linker scripts written in the general-purpose `ld` scripting language. If *MRI-cmdfile* does not exist, `ld` looks for it in the directories specified by any '`-L`' options.

`-d`
`-dc`
`-dp` These three options are equivalent; multiple forms are supported for compatibility with other linkers. They assign space to common symbols even if a relocatable output file is specified (with '`-r`'). The script command `FORCE_COMMON_ALLOCATION` has the same effect. See Section 3.8 [Option Commands], page 40.

`-e` *entry*
`--entry=`*entry*
Use *entry* as the explicit symbol for beginning execution of your program, rather than the default entry point. See Section 3.6 [Entry Point], page 37, for a discussion of defaults and other ways of specifying the entry point.

`-E`
`--export-dynamic`
When creating a dynamically linked executable, add all symbols to the dynamic symbol table. The dynamic symbol table is the set of symbols which are visible from dynamic objects at run time.

If you do not use this option, the dynamic symbol table will normally contain only those symbols which are referenced by some dynamic object mentioned in the link.

If you use `dlopen` to load a dynamic object which needs to refer back to the symbols defined by the program, rather than some other dynamic object, then you will probably need to use this option when linking the program itself.

`-f`
`--auxiliary` *name*
When creating an ELF shared object, set the internal DT_AUXILIARY field to the specified name. This tells the dynamic linker that the symbol table of the shared object should be used as an auxiliary filter on the symbol table of the shared object *name*.

If you later link a program against this filter object, then, when you run the program, the dynamic linker will see the DT_AUXILIARY field. If the dynamic linker resolves any symbols from the filter object, it will first check whether there is a definition in the shared object *name*. If there is one, it will be used instead of the definition in the filter object. The shared object *name* need not exist. Thus the shared object *name* may be used to provide an alternative implementation of certain functions, perhaps for debugging or for machine specific performance.

This option may be specified more than once. The DT_AUXILIARY entries will be created in the order in which they appear on the command line.

`-F` *name*
`--filter` *name*
When creating an ELF shared object, set the internal DT_FILTER field to the specified name. This tells the dynamic linker that the symbol table of the shared object which is being created should be used as a filter on the symbol table of the shared object *name*.

If you later link a program against this filter object, then, when you run the program, the dynamic linker will see the DT_FILTER field. The dynamic linker will resolve symbols according to the symbol table of the filter object as usual, but it will actually link to the definitions found in the shared object *name*. Thus the filter object can be used to select a subset of the symbols provided by the object *name*.

Some older linkers used the `-F` option throughout a compilation toolchain for specifying object-file format for both input and output object files. The GNU linker uses other mechanisms for this purpose: the `-b`, `--format`, `--oformat` options, the `TARGET` command in linker scripts, and the `GNUTARGET` environment variable. The GNU linker will ignore the `-F` option when not creating an ELF shared object.

`--force-exe-suffix`
Make sure that an output file has a .exe suffix.

If a successfully built fully linked output file does not have a `.exe` or `.dll` suffix, this option forces the linker to copy the output file to one of the same name with a `.exe` suffix. This option is useful when using unmodified Unix makefiles on a Microsoft Windows host, since some versions of Windows won't run an image unless it ends in a `.exe` suffix.

`-g` Ignored. Provided for compatibility with other tools.

`-G`*value*
`--gpsize=`*value*
Set the maximum size of objects to be optimized using the GP register to *size*. This is only meaningful for object file formats such as MIPS ECOFF which supports putting large and small objects into different sections. This is ignored for other object file formats.

`-h`*name*
`-soname=`*name*
When creating an ELF shared object, set the internal DT_SONAME field to the specified name. When an executable is linked with a shared object which has a DT_SONAME field, then when the executable is run the dynamic linker will attempt to load the shared object specified by the DT_SONAME field rather than the using the file name given to the linker.

`-i` Perform an incremental link (same as option '`-r`').

`-l`*archive*
`--library=`*archive*
Add archive file *archive* to the list of files to link. This option may be used any number of times. `ld` will search its path-list for occurrences of `lib`*archive*`.a` for every *archive* specified.

On systems which support shared libraries, `ld` may also search for libraries with extensions other than `.a`. Specifically, on ELF and SunOS systems, `ld` will search a directory for a library with an extension of `.so` before searching

for one with an extension of `.a`. By convention, a `.so` extension indicates a shared library.

The linker will search an archive only once, at the location where it is specified on the command line. If the archive defines a symbol which was undefined in some object which appeared before the archive on the command line, the linker will include the appropriate file(s) from the archive. However, an undefined symbol in an object appearing later on the command line will not cause the linker to search the archive again.

See the `-(` option for a way to force the linker to search archives multiple times.

You may list the same archive multiple times on the command line.

This type of archive searching is standard for Unix linkers. However, if you are using `ld` on AIX, note that it is different from the behaviour of the AIX linker.

`-L`*searchdir*
`--library-path=`*searchdir*
: Add path *searchdir* to the list of paths that `ld` will search for archive libraries and `ld` control scripts. You may use this option any number of times. The directories are searched in the order in which they are specified on the command line. Directories specified on the command line are searched before the default directories. All `-L` options apply to all `-l` options, regardless of the order in which the options appear.

  The default set of paths searched (without being specified with '`-L`') depends on which emulation mode `ld` is using, and in some cases also on how it was configured. See Section 2.2 [Environment], page 17.

  The paths can also be specified in a link script with the `SEARCH_DIR` command. Directories specified this way are searched at the point in which the linker script appears in the command line.

`-m`*emulation*
: Emulate the *emulation* linker. You can list the available emulations with the '`--verbose`' or '`-V`' options.

  If the '`-m`' option is not used, the emulation is taken from the `LDEMULATION` environment variable, if that is defined.

  Otherwise, the default emulation depends upon how the linker was configured.

`-M`
`--print-map`
: Print a link map to the standard output. A link map provides information about the link, including the following:

  - Where object files and symbols are mapped into memory.
  - How common symbols are allocated.
  - All archive members included in the link, with a mention of the symbol which caused the archive member to be brought in.

`-n`
`--nmagic` Set the text segment to be read only, and mark the output as `NMAGIC` if possible.

`-N`
`--omagic` Set the text and data sections to be readable and writable. Also, do not page-align the data segment. If the output format supports Unix style magic numbers, mark the output as `OMAGIC`.

`-o` *output*
`--output=`*output*
: Use *output* as the name for the program produced by `ld`; if this option is not specified, the name '`a.out`' is used by default. The script command `OUTPUT` can also specify the output file name.

`-r`
`--relocateable`
: Generate relocatable output—i.e., generate an output file that can in turn serve as input to `ld`. This is often called *partial linking*. As a side effect, in environments that support standard Unix magic numbers, this option also sets the output file's magic number to `OMAGIC`. If this option is not specified, an absolute file is produced. When linking C++ programs, this option *will not* resolve references to constructors; to do that, use '`-Ur`'.

  This option does the same thing as '`-i`'.

`-R` *filename*
`--just-symbols=`*filename*
: Read symbol names and their addresses from *filename*, but do not relocate it or include it in the output. This allows your output file to refer symbolically to absolute locations of memory defined in other programs. You may use this option more than once.

  For compatibility with other ELF linkers, if the `-R` option is followed by a directory name, rather than a file name, it is treated as the `-rpath` option.

`-s`
`--strip-all`
: Omit all symbol information from the output file.

`-S`
`--strip-debug`
: Omit debugger symbol information (but not all symbols) from the output file.

`-t`
`--trace` Print the names of the input files as `ld` processes them.

`-T` *commandfile*
`--script=`*commandfile*
: Read link commands from the file *commandfile*. These commands replace `ld`'s default link script (rather than adding to it), so *commandfile* must specify everything necessary to describe the target format. You must use this option if you want to use a command which can only appear once in a linker script, such as the `SECTIONS` or `MEMORY` command. See Chapter 3 [Commands], page 19. If *commandfile* does not exist, `ld` looks for it in the directories specified by any preceding '`-L`' options. Multiple '`-T`' options accumulate.

`-u` *symbol*
`--undefined=`*symbol*
: Force *symbol* to be entered in the output file as an undefined symbol. Doing this may, for example, trigger linking of additional modules from standard libraries. '`-u`' may be repeated with different option arguments to enter additional undefined symbols.

`-v`
`--version`
`-V`
: Display the version number for `ld`. The `-V` option also lists the supported emulations.

`-x`
`--discard-all`
: Delete all local symbols.

`-X`
`--discard-locals`
: Delete all temporary local symbols. For most targets, this is all local symbols whose names begin with 'L'.

`-y` *symbol*
`--trace-symbol=`*symbol*
: Print the name of each linked file in which *symbol* appears. This option may be given any number of times. On many systems it is necessary to prepend an underscore.

  This option is useful when you have an undefined symbol in your link but don't know where the reference is coming from.

`-Y` *path*
: Add *path* to the default library search path. This option exists for Solaris compatibility.

`-z` *keyword*
: This option is ignored for Solaris compatibility.

`-(` *archives* `-)`
`--start-group` *archives* `--end-group`
: The *archives* should be a list of archive files. They may be either explicit file names, or '`-l`' options.

  The specified archives are searched repeatedly until no new undefined references are created. Normally, an archive is searched only once in the order that it is specified on the command line. If a symbol in that archive is needed to resolve an undefined symbol referred to by an object in an archive that appears later on the command line, the linker would not be able to resolve that reference. By grouping the archives, they all be searched repeatedly until all possible references are resolved.

  Using this option has a significant performance cost. It is best to use it only when there are unavoidable circular references between two or more archives.

`-assert` *keyword*
: This option is ignored for SunOS compatibility.

`-Bdynamic`
`-dy`
`-call_shared`
: Link against dynamic libraries. This is only meaningful on platforms for which shared libraries are supported. This option is normally the default on such platforms. The different variants of this option are for compatibility with various systems. You may use this option multiple times on the command line: it affects library searching for `-l` options which follow it.

`-Bstatic`
`-dn`
`-non_shared`
`-static`
: Do not link against shared libraries. This is only meaningful on platforms for which shared libraries are supported. The different variants of this option are for compatibility with various systems. You may use this option multiple times on the command line: it affects library searching for `-l` options which follow it.

`-Bsymbolic`
: When creating a shared library, bind references to global symbols to the definition within the shared library, if any. Normally, it is possible for a program linked against a shared library to override the definition within the shared library. This option is only meaningful on ELF platforms which support shared libraries.

`--cref`
: Output a cross reference table. If a linker map file is being generated, the cross reference table is printed to the map file. Otherwise, it is printed on the standard output.

  The format of the table is intentionally simple, so that it may be easily processed by a script if necessary. The symbols are printed out, sorted by name. For each symbol, a list of file names is given. If the symbol is defined, the first file listed is the location of the definition. The remaining files contain references to the symbol.

`--defsym` *symbol*=*expression*
: Create a global symbol in the output file, containing the absolute address given by *expression*. You may use this option as many times as necessary to define multiple symbols in the command line. A limited form of arithmetic is supported for the *expression* in this context: you may give a hexadecimal constant or the name of an existing symbol, or use `+` and `-` to add or subtract hexadecimal constants or symbols. If you need more elaborate expressions, consider using the linker command language from a script (see Section 3.2.6 [Assignment: Symbol Definitions], page 22). *Note:* there should be no white space between *symbol*, the equals sign ("=="), and *expression*.

`--dynamic-linker` *file*
: Set the name of the dynamic linker. This is only meaningful when generating dynamically linked ELF executables. The default dynamic linker is normally correct; don't use this unless you know what you are doing.

`-EB`
: Link big-endian objects. This affects the default output format.

`-EL` Link little-endian objects. This affects the default output format.

`--embedded-relocs`
This option is only meaningful when linking MIPS embedded PIC code, generated by the -membedded-pic option to the GNU compiler and assembler. It causes the linker to create a table which may be used at runtime to relocate any data which was statically initialized to pointer values. See the code in testsuite/ld-empic for details.

`--help` Print a summary of the command-line options on the standard output and exit.

`-Map` *mapfile*
Print a link map to the file *mapfile*. See the description of the '`-M`' option, above.

`--no-keep-memory`
`ld` normally optimizes for speed over memory usage by caching the symbol tables of input files in memory. This option tells `ld` to instead optimize for memory usage, by rereading the symbol tables as necessary. This may be required if `ld` runs out of memory space while linking a large executable.

`--no-warn-mismatch`
Normally `ld` will give an error if you try to link together input files that are mismatched for some reason, perhaps because they have been compiled for different processors or for different endiannesses. This option tells `ld` that it should silently permit such possible errors. This option should only be used with care, in cases when you have taken some special action that ensures that the linker errors are inappropriate.

`--no-whole-archive`
Turn off the effect of the `--whole-archive` option for subsequent archive files.

`--noinhibit-exec`
Retain the executable output file whenever it is still usable. Normally, the linker will not produce an output file if it encounters errors during the link process; it exits without writing an output file when it issues any error whatsoever.

`--oformat` *output-format*
`ld` may be configured to support more than one kind of object file. If your `ld` is configured this way, you can use the '`--oformat`' option to specify the binary format for the output object file. Even when `ld` is configured to support alternative object formats, you don't usually need to specify this, as `ld` should be configured to produce as a default output format the most usual format on each machine. *output-format* is a text string, the name of a particular format supported by the BFD libraries. (You can list the available binary formats with '`objdump -i`'.) The script command `OUTPUT_FORMAT` can also specify the output format, but this option overrides it. See Chapter 5 [BFD], page 47.

`-qmagic` This option is ignored for Linux compatibility.

`-Qy` This option is ignored for SVR4 compatibility.

`--relax` An option with machine dependent effects. This option is only supported on a few targets. See Section 4.1 [`ld` and the H8/300], page 45. See Section 4.2 [`ld` and the Intel 960 family], page 45.

On some platforms, the '`--relax`' option performs global optimizations that become possible when the linker resolves addressing in the program, such as relaxing address modes and synthesizing new instructions in the output object file.

On platforms where this is not supported, '`--relax`' is accepted, but ignored.

`--retain-symbols-file` *filename*
Retain *only* the symbols listed in the file *filename*, discarding all others. *filename* is simply a flat file, with one symbol name per line. This option is especially useful in environments (such as VxWorks) where a large global symbol table is accumulated gradually, to conserve run-time memory.

'`--retain-symbols-file`' does *not* discard undefined symbols, or symbols needed for relocations.

You may only specify '`--retain-symbols-file`' once in the command line. It overrides '`-s`' and '`-S`'.

`-rpath` *dir*
Add a directory to the runtime library search path. This is used when linking an ELF executable with shared objects. All `-rpath` arguments are concatenated and passed to the runtime linker, which uses them to locate shared objects at runtime. The `-rpath` option is also used when locating shared objects which are needed by shared objects explicitly included in the link; see the description of the `-rpath-link` option. If `-rpath` is not used when linking an ELF executable, the contents of the environment variable `LD_RUN_PATH` will be used if it is defined.

The `-rpath` option may also be used on SunOS. By default, on SunOS, the linker will form a runtime search patch out of all the `-L` options it is given. If a `-rpath` option is used, the runtime search path will be formed exclusively using the `-rpath` options, ignoring the `-L` options. This can be useful when using gcc, which adds many `-L` options which may be on NFS mounted filesystems.

For compatibility with other ELF linkers, if the `-R` option is followed by a directory name, rather than a file name, it is treated as the `-rpath` option.

`-rpath-link` *DIR*
When using ELF or SunOS, one shared library may require another. This happens when an `ld -shared` link includes a shared library as one of the input files.

When the linker encounters such a dependency when doing a non-shared, non-relocateable link, it will automatically try to locate the required shared library and include it in the link, if it is not included explicitly. In such a case, the `-rpath-link` option specifies the first set of directories to search. The `-rpath-link` option may specify a sequence of directory names either by specifying a list of names separated by colons, or by appearing multiple times.

The linker uses the following search paths to locate required shared libraries.

1. Any directories specified by `-rpath-link` options.
2. Any directories specified by `-rpath` options. The difference between `-rpath` and `-rpath-link` is that directories specified by `-rpath` options are included in the executable and used at runtime, whereas the `-rpath-link` option is only effective at link time.
3. On an ELF system, if the `-rpath` and `rpath-link` options were not used, search the contents of the environment variable `LD_RUN_PATH`.
4. On SunOS, if the `-rpath` option was not used, search any directories specified using `-L` options.
5. For a native linker, the contents of the environment variable `LD_LIBRARY_PATH`.
6. The default directories, normally '`/lib`' and '`/usr/lib`'.

If the required shared library is not found, the linker will issue a warning and continue with the link.

`-shared`
`-Bshareable`
: Create a shared library. This is currently only supported on ELF, XCOFF and SunOS platforms. On SunOS, the linker will automatically create a shared library if the `-e` option is not used and there are undefined symbols in the link.

`--sort-common`
: This option tells `ld` to sort the common symbols by size when it places them in the appropriate output sections. First come all the one byte symbols, then all the two bytes, then all the four bytes, and then everything else. This is to prevent gaps between symbols due to alignment constraints.

`--split-by-file`
: Similar to `--split-by-reloc` but creates a new output section for each input file.

`--split-by-reloc` *count*
: Trys to creates extra sections in the output file so that no single output section in the file contains more than *count* relocations. This is useful when generating huge relocatable for downloading into certain real time kernels with the COFF object file format; since COFF cannot represent more than 65535 relocations in a single section. Note that this will fail to work with object file formats which do not support arbitrary sections. The linker will not split up individual input sections for redistribution, so if a single input section contains more than *count* relocations one output section will contain that many relocations.

`--stats`
: Compute and display statistics about the operation of the linker, such as execution time and memory usage.

`--traditional-format`
: For some targets, the output of `ld` is different in some ways from the output of some existing linker. This switch requests `ld` to use the traditional format instead.

For example, on SunOS, `ld` combines duplicate entries in the symbol string table. This can reduce the size of an output file with full debugging information by over 30 percent. Unfortunately, the SunOS `dbx` program can not read the resulting program (`gdb` has no trouble). The '`--traditional-format`' switch tells `ld` to not combine duplicate entries.

`-Tbss` *org*
`-Tdata` *org*
`-Ttext` *org*
: Use *org* as the starting address for—respectively—the `bss`, `data`, or the `text` segment of the output file. *org* must be a single hexadecimal integer; for compatibility with other linkers, you may omit the leading '`0x`' usually associated with hexadecimal values.

`-Ur`
: For anything other than C++ programs, this option is equivalent to '`-r`': it generates relocatable output—i.e., an output file that can in turn serve as input to `ld`. When linking C++ programs, '`-Ur`' *does* resolve references to constructors, unlike '`-r`'. It does not work to use '`-Ur`' on files that were themselves linked with '`-Ur`'; once the constructor table has been built, it cannot be added to. Use '`-Ur`' only for the last partial link, and '`-r`' for the others.

`--verbose`
: Display the version number for `ld` and list the linker emulations supported. Display which input files can and cannot be opened. Display the linker script if using a default builtin script.

`--version-script=`*version-scriptfile*
: Specify the name of a version script to the linker. This is typically used when creating shared libraries to specify additional information about the version heirarchy for the library being created. This option is only meaningful on ELF platforms which support shared libraries. See Section 3.7 [Version Script], page 38.

`--warn-common`
: Warn when a common symbol is combined with another common symbol or with a symbol definition. Unix linkers allow this somewhat sloppy practice, but linkers on some other operating systems do not. This option allows you to find potential problems from combining global symbols. Unfortunately, some C libraries use this practice, so you may get some warnings about symbols in the libraries as well as in your programs.

  There are three kinds of global symbols, illustrated here by C examples:

  '`int i = 1;`'
  : A definition, which goes in the initialized data section of the output file.

  '`extern int i;`'
  : An undefined reference, which does not allocate space. There must be either a definition or a common symbol for the variable somewhere.

'`int i;`' A common symbol. If there are only (one or more) common symbols for a variable, it goes in the uninitialized data area of the output file. The linker merges multiple common symbols for the same variable into a single symbol. If they are of different sizes, it picks the largest size. The linker turns a common symbol into a declaration, if there is a definition of the same variable.

The '`--warn-common`' option can produce five kinds of warnings. Each warning consists of a pair of lines: the first describes the symbol just encountered, and the second describes the previous symbol encountered with the same name. One or both of the two symbols will be a common symbol.

1. Turning a common symbol into a reference, because there is already a definition for the symbol.

   file(*section*): `warning: common of` '*symbol*'
   `overridden by definition`
   file(*section*): `warning: defined here`

2. Turning a common symbol into a reference, because a later definition for the symbol is encountered. This is the same as the previous case, except that the symbols are encountered in a different order.

   file(*section*): `warning: definition of` '*symbol*'
   `overriding common`
   file(*section*): `warning: common is here`

3. Merging a common symbol with a previous same-sized common symbol.

   file(*section*): `warning: multiple common`
   `of` '*symbol*'
   file(*section*): `warning: previous common is here`

4. Merging a common symbol with a previous larger common symbol.

   file(*section*): `warning: common of` '*symbol*'
   `overridden by larger common`
   file(*section*): `warning: larger common is here`

5. Merging a common symbol with a previous smaller common symbol. This is the same as the previous case, except that the symbols are encountered in a different order.

   file(*section*): `warning: common of` '*symbol*'
   `overriding smaller common`
   file(*section*): `warning: smaller common is here`

`--warn-constructors`
Warn if any global constructors are used. This is only useful for a few object file formats. For formats like COFF or ELF, the linker can not detect the use of global constructors.

`--warn-multiple-gp`
Warn if multiple global pointer values are required in the output file. This is only meaningful for certain processors, such as the Alpha. Specifically, some processors put large-valued constants in a special section. A special register

(the global pointer) points into the middle of this section, so that constants can be loaded efficiently via a base-register relative addressing mode. Since the offset in base-register relative mode is fixed and relatively small (e.g., 16 bits), this limits the maximum size of the constant pool. Thus, in large programs, it is often necessary to use multiple global pointer values in order to be able to address all possible constants. This option causes a warning to be issued whenever this case occurs.

`--warn-once`
Only warn once for each undefined symbol, rather than once per module which refers to it.

`--warn-section-align`
Warn if the address of an output section is changed because of alignment. Typically, the alignment will be set by an input section. The address will only be changed if it not explicitly specified; that is, if the `SECTIONS` command does not specify a start address for the section (see Section 3.4 [SECTIONS], page 26).

`--whole-archive`
For each archive mentioned on the command line after the `--whole-archive` option, include every object file in the archive in the link, rather than searching the archive for the required object files. This is normally used to turn an archive file into a shared library, forcing every object to be included in the resulting shared library. This option may be used more than once.

`--wrap` *symbol*
Use a wrapper function for *symbol*. Any undefined reference to *symbol* will be resolved to `__wrap_`*symbol*. Any undefined reference to `__real_`*symbol* will be resolved to *symbol*.

This can be used to provide a wrapper for a system function. The wrapper function should be called `__wrap_`*symbol*. If it wishes to call the system function, it should call `__real_`*symbol*.

Here is a trivial example:

```
void *
__wrap_malloc (int c)
{
  printf ("malloc called with %ld\n", c);
  return __real_malloc (c);
}
```

If you link other code with this file using `--wrap malloc`, then all calls to `malloc` will call the function `__wrap_malloc` instead. The call to `__real_malloc` in `__wrap_malloc` will call the real `malloc` function.

You may wish to provide a `__real_malloc` function as well, so that links without the `--wrap` option will succeed. If you do this, you should not put the definition of `__real_malloc` in the same file as `__wrap_malloc`; if you do, the assembler may resolve the call before the linker has a chance to wrap it to `malloc`.

2.2 Environment Variables

You can change the behavior of `ld` with the environment variables `GNUTARGET` and `LDEMULATION`.

`GNUTARGET` determines the input-file object format if you don't use '`-b`' (or its synonym '`--format`'). Its value should be one of the BFD names for an input format (see Chapter 5 [BFD], page 47). If there is no `GNUTARGET` in the environment, `ld` uses the natural format of the target. If `GNUTARGET` is set to `default` then BFD attempts to discover the input format by examining binary input files; this method often succeeds, but there are potential ambiguities, since there is no method of ensuring that the magic number used to specify object-file formats is unique. However, the configuration procedure for BFD on each system places the conventional format for that system first in the search-list, so ambiguities are resolved in favor of convention.

`LDEMULATION` determines the default emulation if you don't use the '`-m`' option. The emulation can affect various aspects of linker behaviour, particularly the default linker script. You can list the available emulations with the '`--verbose`' or '`-V`' options. If the '`-m`' option is not used, and the `LDEMULATION` environment variable is not defined, the default emulation depends upon how the linker was configured.

3 Command Language

The command language provides explicit control over the link process, allowing complete specification of the mapping between the linker's input files and its output. It controls:

- input files
- file formats
- output file layout
- addresses of sections
- placement of common blocks

You may supply a command file (also known as a linker script) to the linker either explicitly through the '`-T`' option, or implicitly as an ordinary file. Normally you should use the '`-T`' option. An implicit linker script should only be used when you want to augment, rather than replace, the default linker script; typically an implicit linker script would consist only of `INPUT` or `GROUP` commands.

If the linker opens a file which it cannot recognize as a supported object or archive format, nor as a linker script, it reports an error.

3.1 Linker Scripts

The `ld` command language is a collection of statements; some are simple keywords setting a particular option, some are used to select and group input files or name output files; and two statement types have a fundamental and pervasive impact on the linking process.

The most fundamental command of the `ld` command language is the `SECTIONS` command (see Section 3.4 [SECTIONS], page 26). Every meaningful command script must have a `SECTIONS` command: it specifies a "picture" of the output file's layout, in varying degrees of detail. No other command is required in all cases.

The `MEMORY` command complements `SECTIONS` by describing the available memory in the target architecture. This command is optional; if you don't use a `MEMORY` command, `ld` assumes sufficient memory is available in a contiguous block for all output. See Section 3.3 [MEMORY], page 25.

You may include comments in linker scripts just as in C: delimited by '`/*`' and '`*/`'. As in C, comments are syntactically equivalent to whitespace.

3.2 Expressions

Many useful commands involve arithmetic expressions. The syntax for expressions in the command language is identical to that of C expressions, with the following features:

- All expressions evaluated as integers and are of "long" or "unsigned long" type.
- All constants are integers.
- All of the C arithmetic operators are provided.
- You may reference, define, and create global variables.
- You may call special purpose built-in functions.

3.2.1 Integers

An octal integer is '`0`' followed by zero or more of the octal digits ('`01234567`').

```
_as_octal = 0157255;
```

A decimal integer starts with a non-zero digit followed by zero or more digits ('`0123456789`').

```
_as_decimal = 57005;
```

A hexadecimal integer is '`0x`' or '`0X`' followed by one or more hexadecimal digits chosen from '`0123456789abcdefABCDEF`'.

```
_as_hex = 0xdead;
```

To write a negative integer, use the prefix operator '`-`' (see Section 3.2.4 [Operators], page 21).

```
_as_neg = -57005;
```

Additionally the suffixes `K` and `M` may be used to scale a constant by 1024 or 1024^2 respectively. For example, the following all refer to the same quantity:

```
_fourk_1 = 4K;
_fourk_2 = 4096;
_fourk_3 = 0x1000;
```

3.2.2 Symbol Names

Unless quoted, symbol names start with a letter, underscore, or point and may include any letters, underscores, digits, points, and hyphens. Unquoted symbol names must not conflict with any keywords. You can specify a symbol which contains odd characters or has the same name as a keyword, by surrounding the symbol name in double quotes:

```
"SECTION" = 9;
"with a space" = "also with a space" + 10;
```

Since symbols can contain many non-alphabetic characters, it is safest to delimit symbols with spaces. For example, '`A-B`' is one symbol, whereas '`A - B`' is an expression involving subtraction.

3.2.3 The Location Counter

The special linker variable *dot* '`.`' always contains the current output location counter. Since the `.` always refers to a location in an output section, it must always appear in an expression within a `SECTIONS` command. The `.` symbol may appear anywhere that an ordinary symbol is allowed in an expression, but its assignments have a side effect. Assigning a value to the `.` symbol will cause the location counter to be moved. This may be used to create holes in the output section. The location counter may never be moved backwards.

```
SECTIONS
{
  output :
  {
  file1(.text)
  . = . + 1000;
  file2(.text)
```

```
    . += 1000;
    file3(.text)
    } = 0x1234;
  }
```

In the previous example, `file1` is located at the beginning of the output section, then there is a 1000 byte gap. Then `file2` appears, also with a 1000 byte gap following before `file3` is loaded. The notation '`= 0x1234`' specifies what data to write in the gaps (see Section 3.4.4 [Section Options], page 32).

3.2.4 Operators

The linker recognizes the standard C set of arithmetic operators, with the standard bindings and precedence levels:

Precedence	Associativity	Operators
highest		
1	left	`- ~ !` †
2	left	`* / %`
3	left	`+ -`
4	left	`>> <<`
5	left	`== != > < <= >=`
6	left	`&`
7	left	`|`
8	left	`&&`
9	left	`||`
10	right	`? :`
11	right	`&= += -= *= /=` ‡
lowest		

† Prefix operators.
‡ See Section 3.2.6 [Assignment], page 22.

3.2.5 Evaluation

The linker uses "lazy evaluation" for expressions; it only calculates an expression when absolutely necessary. The linker needs the value of the start address, and the lengths of memory regions, in order to do any linking at all; these values are computed as soon as possible when the linker reads in the command file. However, other values (such as symbol values) are not known or needed until after storage allocation. Such values are evaluated later, when other information (such as the sizes of output sections) is available for use in the symbol assignment expression.

3.2.6 Assignment: Defining Symbols

You may create global symbols, and assign values (addresses) to global symbols, using any of the C assignment operators:

symbol `=` *expression* `;`
symbol `&=` *expression* `;`
symbol `+=` *expression* `;`
symbol `-=` *expression* `;`
symbol `*=` *expression* `;`
symbol `/=` *expression* `;`

Two things distinguish assignment from other operators in `ld` expressions.

- Assignment may only be used at the root of an expression; '`a=b+3;`' is allowed, but '`a+b=3;`' is an error.
- You must place a trailing semicolon ("`;`") at the end of an assignment statement.

Assignment statements may appear:

- as commands in their own right in an `ld` script; or
- as independent statements within a `SECTIONS` command; or
- as part of the contents of a section definition in a `SECTIONS` command.

The first two cases are equivalent in effect—both define a symbol with an absolute address. The last case defines a symbol whose address is relative to a particular section (see Section 3.4 [SECTIONS], page 26).

When a linker expression is evaluated and assigned to a variable, it is given either an absolute or a relocatable type. An absolute expression type is one in which the symbol contains the value that it will have in the output file; a relocatable expression type is one in which the value is expressed as a fixed offset from the base of a section.

The type of the expression is controlled by its position in the script file. A symbol assigned within a section definition is created relative to the base of the section; a symbol assigned in any other place is created as an absolute symbol. Since a symbol created within a section definition is relative to the base of the section, it will remain relocatable if relocatable output is requested. A symbol may be created with an absolute value even when assigned to within a section definition by using the absolute assignment function `ABSOLUTE`. For example, to create an absolute symbol whose address is the last byte of an output section named `.data`:

```
  SECTIONS{ ...
    .data :
      {
        *(.data)
        _edata = ABSOLUTE(.) ;
      }
...
   }
```

The linker tries to put off the evaluation of an assignment until all the terms in the source expression are known (see Section 3.2.5 [Evaluation], page 21). For instance, the sizes of sections cannot be known until after allocation, so assignments dependent upon these are

not performed until after allocation. Some expressions, such as those depending upon the location counter *dot*, '`.`' must be evaluated during allocation. If the result of an expression is required, but the value is not available, then an error results. For example, a script like the following

```
SECTIONS { ...
  text 9+this_isnt_constant :
    { ...
    }
...
 }
```

will cause the error message "`Non constant expression for initial address`".

In some cases, it is desirable for a linker script to define a symbol only if it is referenced, and only if it is not defined by any object included in the link. For example, traditional linkers defined the symbol '`etext`'. However, ANSI C requires that the user be able to use '`etext`' as a function name without encountering an error. The `PROVIDE` keyword may be used to define a symbol, such as '`etext`', only if it is referenced but not defined. The syntax is `PROVIDE(`*symbol* `=` *expression*`)`.

3.2.7 Arithmetic Functions

The command language includes a number of built-in functions for use in link script expressions.

`ABSOLUTE(`*exp*`)`
: Return the absolute (non-relocatable, as opposed to non-negative) value of the expression *exp*. Primarily useful to assign an absolute value to a symbol within a section definition, where symbol values are normally section-relative.

`ADDR(`*section*`)`
: Return the absolute address of the named *section*. Your script must previously have defined the location of that section. In the following example, `symbol_1` and `symbol_2` are assigned identical values:

```
SECTIONS{ ...
  .output1 :
    {
    start_of_output_1 = ABSOLUTE(.);
    ...
    }
  .output :
    {
    symbol_1 = ADDR(.output1);
    symbol_2 = start_of_output_1;
    }
...
 }
```

`LOADADDR(`*section*`)`
: Return the absolute load address of the named *section*. This is normally the same as `ADDR`, but it may be different if the `AT` keyword is used in the section definition (see Section 3.4.4 [Section Options], page 32).

`ALIGN(`*exp*`)`
: Return the result of the current location counter (`.`) aligned to the next *exp* boundary. *exp* must be an expression whose value is a power of two. This is equivalent to

```
(. + exp - 1) & ~(exp - 1)
```

`ALIGN` doesn't change the value of the location counter—it just does arithmetic on it. As an example, to align the output `.data` section to the next `0x2000` byte boundary after the preceding section and to set a variable within the section to the next `0x8000` boundary after the input sections:

```
SECTIONS{ ...
  .data ALIGN(0x2000): {
    *(.data)
    variable = ALIGN(0x8000);
  }
...
 }
```

The first use of `ALIGN` in this example specifies the location of a section because it is used as the optional *start* attribute of a section definition (see Section 3.4.4 [Section Options], page 32). The second use simply defines the value of a variable.

The built-in `NEXT` is closely related to `ALIGN`.

`DEFINED(`*symbol*`)`
: Return 1 if *symbol* is in the linker global symbol table and is defined, otherwise return 0. You can use this function to provide default values for symbols. For example, the following command-file fragment shows how to set a global symbol `begin` to the first location in the `.text` section—but if a symbol called `begin` already existed, its value is preserved:

```
SECTIONS{ ...
  .text : {
    begin = DEFINED(begin) ? begin : . ;
    ...
  }
...
 }
```

`NEXT(`*exp*`)`
: Return the next unallocated address that is a multiple of *exp*. This function is closely related to `ALIGN(`*exp*`)`; unless you use the `MEMORY` command to define discontinuous memory for the output file, the two functions are equivalent.

`SIZEOF(`*section*`)`
: Return the size in bytes of the named *section*, if that section has been allocated. In the following example, `symbol_1` and `symbol_2` are assigned identical values:

```
SECTIONS{ ...
  .output {
    .start = . ;
    ...
    .end = . ;
    }
  symbol_1 = .end - .start ;
  symbol_2 = SIZEOF(.output);
...  }
```

`SIZEOF_HEADERS`
`sizeof_headers`
: Return the size in bytes of the output file's headers. You can use this number as the start address of the first section, if you choose, to facilitate paging.

`MAX(`*exp1*`,` *exp2*`)`
: Returns the maximum of *exp1* and *exp2*.

`MIN(`*exp1*`,` *exp2*`)`
: Returns the minimum of *exp1* and *exp2*.

3.2.8 Semicolons

Semicolons ("`;`") are required in the following places. In all other places they can appear for aesthetic reasons but are otherwise ignored.

`Assignment`
: Semicolons must appear at the end of assignment expressions. See Section 3.2.6 [Assignment], page 22

`PHDRS`
: Semicolons must appear at the end of a `PHDRS` statement. See Section 3.5 [PHDRS], page 35

3.3 Memory Layout

The linker's default configuration permits allocation of all available memory. You can override this configuration by using the `MEMORY` command. The `MEMORY` command describes the location and size of blocks of memory in the target. By using it carefully, you can describe which memory regions may be used by the linker, and which memory regions it must avoid. The linker does not shuffle sections to fit into the available regions, but does move the requested sections into the correct regions and issue errors when the regions become too full.

A command file may contain at most one use of the `MEMORY` command; however, you can define as many blocks of memory within it as you wish. The syntax is:

```
MEMORY
  {
    name (attr) : ORIGIN = origin, LENGTH = len
    ...
  }
```


name
: is a name used internally by the linker to refer to the region. Any symbol name may be used. The region names are stored in a separate name space, and will not conflict with symbols, file names or section names. Use distinct names to specify multiple regions.

(*attr*)
: is an optional list of attributes that specify whether to use a particular memory to place sections that are not listed in the linker script. Valid attribute lists must be made up of the characters "`ALIRWX`" that match section attributes. If you omit the attribute list, you may omit the parentheses around it as well. The attributes currently supported are:

'Letter'	Section Attribute
'R'	Read-only sections.
'W'	Read/write sections.
'X'	Sections containing executable code.
'A'	Allocated sections.
'I'	Initialized sections.
'L'	Same as I.
'!'	Invert the sense of any of the following attributes.

origin
: is the start address of the region in physical memory. It is an expression that must evaluate to a constant before memory allocation is performed. The keyword `ORIGIN` may be abbreviated to `org` or `o` (but not, for example, '`ORG`').

len
: is the size in bytes of the region (an expression). The keyword `LENGTH` may be abbreviated to `len` or `l`.

For example, to specify that memory has two regions available for allocation—one starting at 0 for 256 kilobytes, and the other starting at `0x40000000` for four megabytes. The `rom` memory region will get all sections without an explicit memory register that are either read-only or contain code, while the `ram` memory region will get the sections.

```
MEMORY
  {
  rom (rx)  : ORIGIN = 0, LENGTH = 256K
  ram (!rx) : org = 0x40000000, l = 4M
  }
```

Once you have defined a region of memory named *mem*, you can direct specific output sections there by using a command ending in '`>`*mem*' within the `SECTIONS` command (see Section 3.4.4 [Section Options], page 32). If the combined output sections directed to a region are too big for the region, the linker will issue an error message.

3.4 Specifying Output Sections

The `SECTIONS` command controls exactly where input sections are placed into output sections, their order in the output file, and to which output sections they are allocated.

You may use at most one `SECTIONS` command in a script file, but you can have as many statements within it as you wish. Statements within the `SECTIONS` command can do one of three things:

- define the entry point;
- assign a value to a symbol;
- describe the placement of a named output section, and which input sections go into it.

You can also use the first two operations—defining the entry point and defining symbols—outside the `SECTIONS` command: see Section 3.6 [Entry Point], page 37, and Section 3.2.6 [Assignment], page 22. They are permitted here as well for your convenience in reading the script, so that symbols and the entry point can be defined at meaningful points in your output-file layout.

If you do not use a `SECTIONS` command, the linker places each input section into an identically named output section in the order that the sections are first encountered in the input files. If all input sections are present in the first file, for example, the order of sections in the output file will match the order in the first input file.

3.4.1 Section Definitions

The most frequently used statement in the `SECTIONS` command is the *section definition*, which specifies the properties of an output section: its location, alignment, contents, fill pattern, and target memory region. Most of these specifications are optional; the simplest form of a section definition is

```
SECTIONS { ...
  secname : {
    contents
  }
...
}
```

secname is the name of the output section, and *contents* a specification of what goes there—for example, a list of input files or sections of input files (see Section 3.4.2 [Section Placement], page 28). The whitespace around *secname* is required, so that the section name is unambiguous. The other whitespace shown is optional. You do need the colon '`:`' and the braces '`{}`', however.

secname must meet the constraints of your output format. In formats which only support a limited number of sections, such as `a.out`, the name must be one of the names supported by the format (`a.out`, for example, allows only `.text`, `.data` or `.bss`). If the output format supports any number of sections, but with numbers and not names (as is the case for Oasys), the name should be supplied as a quoted numeric string. A section name may consist of any sequence of characters, but any name which does not conform to the standard `ld` symbol name syntax must be quoted. See Section 3.2.2 [Symbol Names], page 20.

The special *secname* '`/DISCARD/`' may be used to discard input sections. Any sections which are assigned to an output section named '`/DISCARD/`' are not included in the final link output.

The linker will not create output sections which do not have any contents. This is for convenience when referring to input sections that may or may not exist. For example,


```
.foo { *(.foo) }
```

will only create a '`.foo`' section in the output file if there is a '`.foo`' section in at least one input file.

3.4.2 Section Placement

In a section definition, you can specify the contents of an output section by listing particular input files, by listing particular input-file sections, or by a combination of the two. You can also place arbitrary data in the section, and define symbols relative to the beginning of the section.

The *contents* of a section definition may include any of the following kinds of statement. You can include as many of these as you like in a single section definition, separated from one another by whitespace.

filename You may simply name a particular input file to be placed in the current output section; *all* sections from that file are placed in the current section definition. If the file name has already been mentioned in another section definition, with an explicit section name list, then only those sections which have not yet been allocated are used.

To specify a list of particular files by name:

```
.data : { afile.o bfile.o cfile.o }
```

The example also illustrates that multiple statements can be included in the contents of a section definition, since each file name is a separate statement.

filename`(` *section* `)`
filename`(` *section* `,` *section*`, ... )`
filename`(` *section section* `... )`

You can name one or more sections from your input files, for insertion in the current output section. If you wish to specify a list of input-file sections inside the parentheses, separate the section names with whitespace.

`*` `(`*section*`)`
`*` `(`*section*`,` *section*`, ...)`
`*` `(`*section section* `...)`

Instead of explicitly naming particular input files in a link control script, you can refer to *all* files from the `ld` command line: use '`*`' instead of a particular file name before the parenthesized input-file section list.

If you have already explicitly included some files by name, '`*`' refers to all *remaining* files—those whose places in the output file have not yet been defined.

For example, to copy sections `1` through `4` from an Oasys file into the `.text` section of an `a.out` file, and sections `13` and `14` into the `.data` section:

```
SECTIONS {
  .text :{
    *("1" "2" "3" "4")
  }

  .data :{
    *("13" "14")
  }
}
```

'`[` *section* `...` `]`' used to be accepted as an alternate way to specify named sections from all unallocated input files. Because some operating systems (VMS) allow brackets in file names, that notation is no longer supported.

filename`( COMMON )`
`*( COMMON )`
Specify where in your output file to place uninitialized data with this notation. `*(COMMON)` by itself refers to all uninitialized data from all input files (so far as it is not yet allocated); *filename*`(COMMON)` refers to uninitialized data from a particular file. Both are special cases of the general mechanisms for specifying where to place input-file sections: `ld` permits you to refer to uninitialized data as if it were in an input-file section named `COMMON`, regardless of the input file's format.

In any place where you may use a specific file or section name, you may also use a wildcard pattern. The linker handles wildcards much as the Unix shell does. A '`*`' character matches any number of characters. A '`?`' character matches any single character. The sequence '`[`*chars*`]`' will match a single instance of any of the *chars*; the '`-`' character may be used to specify a range of characters, as in '`[a-z]`' to match any lower case letter. A '`\`' character may be used to quote the following character.

When a file name is matched with a wildcard, the wildcard characters will not match a '`/`' character (used to separate directory names on Unix). A pattern consisting of a single '`*`' character is an exception; it will always match any file name. In a section name, the wildcard characters will match a '`/`' character.

Wildcards only match files which are explicitly specified on the command line. The linker does not search directories to expand wildcards. However, if you specify a simple file name—a name with no wildcard characters—in a linker script, and the file name is not also specified on the command line, the linker will attempt to open the file as though it appeared on the command line.

In the following example, the command script arranges the output file into three consecutive sections, named `.text`, `.data`, and `.bss`, taking the input for each from the correspondingly named sections of all the input files:

```
SECTIONS {
  .text : { *(.text) }
  .data : { *(.data) }
  .bss :  { *(.bss)  *(COMMON) }
}
```


The following example reads all of the sections from file `all.o` and places them at the start of output section `outputa` which starts at location `0x10000`. All of section `.input1` from file `foo.o` follows immediately, in the same output section. All of section `.input2` from `foo.o` goes into output section `outputb`, followed by section `.input1` from `foo1.o`. All of the remaining `.input1` and `.input2` sections from any files are written to output section `outputc`.

```
SECTIONS {
  outputa 0x10000 :
    {
    all.o
    foo.o (.input1)
    }
  outputb :
    {
    foo.o (.input2)
    foo1.o (.input1)
    }
  outputc :
    {
    *(.input1)
    *(.input2)
    }
}
```

This example shows how wildcard patterns might be used to partition files. All `.text` sections are placed in `.text`, and all `.bss` sections are placed in `.bss`. For all files beginning with an upper case character, the `.data` section is placed into `.DATA`; for all other files, the `.data` section is placed into `.data`.

```
SECTIONS {
  .text : { *(.text) }
  .DATA : { [A-Z]*(.data) }
  .data : { *(.data) }
  .bss : { *(.bss) }
}
```

3.4.3 Section Data Expressions

The foregoing statements arrange, in your output file, data originating from your input files. You can also place data directly in an output section from the link command script. Most of these additional statements involve expressions (see Section 3.2 [Expressions], page 19). Although these statements are shown separately here for ease of presentation, no such segregation is needed within a section definition in the `SECTIONS` command; you can intermix them freely with any of the statements we've just described.

`CREATE_OBJECT_SYMBOLS`
Create a symbol for each input file in the current section, set to the address of the first byte of data written from that input file. For instance, with `a.out` files it is conventional to have a symbol for each input file. You can accomplish this by defining the output `.text` section as follows:

```
SECTIONS {
  .text 0x2020 :
     {
     CREATE_OBJECT_SYMBOLS
     *(.text)
     _etext = ALIGN(0x2000);
     }
  ...
}
```

If `sample.ld` is a file containing this script, and `a.o`, `b.o`, `c.o`, and `d.o` are four input files with contents like the following—

```
/* a.c */

afunction() { }
int adata=1;
int abss;
```

'`ld -M -T sample.ld a.o b.o c.o d.o`' would create a map like this, containing symbols matching the object file names:

```
00000000 A __DYNAMIC
00004020 B _abss
00004000 D _adata
00002020 T _afunction
00004024 B _bbss
00004008 D _bdata
00002038 T _bfunction
00004028 B _cbss
00004010 D _cdata
00002050 T _cfunction
0000402c B _dbss
00004018 D _ddata
00002068 T _dfunction
00004020 D _edata
00004030 B _end
00004000 T _etext
00002020 t a.o
00002038 t b.o
00002050 t c.o
00002068 t d.o
```

symbol `=` *expression* `;`
symbol *f*`=` *expression* `;`

symbol is any symbol name (see Section 3.2.2 [Symbols], page 20). "*f*=" refers to any of the operators `&= += -= *= /=` which combine arithmetic and assignment.

When you assign a value to a symbol within a particular section definition, the value is relative to the beginning of the section (see Section 3.2.6 [Assignment], page 22). If you write


```
SECTIONS {
  abs = 14 ;
  ...
  .data : { ... rel = 14 ; ... }
  abs2 = 14 + ADDR(.data);
  ...
}
```

`abs` and `rel` do not have the same value; `rel` has the same value as `abs2`.

`BYTE(`*expression*`)`
`SHORT(`*expression*`)`
`LONG(`*expression*`)`
`QUAD(`*expression*`)`
`SQUAD(`*expression*`)`

By including one of these four statements in a section definition, you can explicitly place one, two, four, eight unsigned, or eight signed bytes (respectively) at the current address of that section. When using a 64 bit host or target, `QUAD` and `SQUAD` are the same. When both host and target are 32 bits, `QUAD` uses an unsigned 32 bit value, and `SQUAD` sign extends the value. Both will use the correct endianness when writing out the value.

Multiple-byte quantities are represented in whatever byte order is appropriate for the output file format (see Chapter 5 [BFD], page 47).

`FILL(`*expression*`)`

Specify the "fill pattern" for the current section. Any otherwise unspecified regions of memory within the section (for example, regions you skip over by assigning a new value to the location counter '`.`') are filled with the two least significant bytes from the *expression* argument. A `FILL` statement covers memory locations *after* the point it occurs in the section definition; by including more than one `FILL` statement, you can have different fill patterns in different parts of an output section.

3.4.4 Optional Section Attributes

Here is the full syntax of a section definition, including all the optional portions:

```
SECTIONS {
...
  secname start BLOCK(align) (NOLOAD) : AT ( ldadr )
    { contents } >region :phdr =fill
...
}
```

secname and *contents* are required. See Section 3.4.1 [Section Definition], page 27, and Section 3.4.2 [Section Placement], page 28, for details on *contents*. The remaining elements—*start*, `BLOCK(`*align*`)`, `(NOLOAD)`, `AT ( `*ldadr*` )`, `>`*region*, `:`*phdr*, and `=`*fill*—are all optional.

start

You can force the output section to be loaded at a specified address by specifying *start* immediately following the section name. *start* can be represented

as any expression. The following example generates section *output* at location `0x40000000`:

```
SECTIONS {
  ...
  output 0x40000000: {
    ...
    }
  ...
}
```

`BLOCK(`*align*`)`

You can include `BLOCK()` specification to advance the location counter `.` prior to the beginning of the section, so that the section will begin at the specified alignment. *align* is an expression.

`(NOLOAD)` The '`(NOLOAD)`' directive will mark a section to not be loaded at run time. The linker will process the section normally, but will mark it so that a program loader will not load it into memory. For example, in the script sample below, the `ROM` section is addressed at memory location '`0`' and does not need to be loaded when the program is run. The contents of the `ROM` section will appear in the linker output file as usual.

```
SECTIONS {
  ROM  0  (NOLOAD)  : { ... }
  ...
}
```

`AT ( `*ldadr*` )`

The expression *ldadr* that follows the `AT` keyword specifies the load address of the section. The default (if you do not use the `AT` keyword) is to make the load address the same as the relocation address. This feature is designed to make it easy to build a ROM image. For example, this `SECTIONS` definition creates two output sections: one called '`.text`', which starts at `0x1000`, and one called '`.mdata`', which is loaded at the end of the '`.text`' section even though its relocation address is `0x2000`. The symbol `_data` is defined with the value `0x2000`:

```
SECTIONS
  {
  .text 0x1000 : { *(.text) _etext = . ; }
  .mdata 0x2000 :
    AT ( ADDR(.text) + SIZEOF ( .text ) )
    { _data = . ; *(.data); _edata = . ;  }
  .bss 0x3000 :
    { _bstart = . ;  *(.bss) *(COMMON) ; _bend = . ;}
}
```

The run-time initialization code (for C programs, usually `crt0`) for use with a ROM generated this way has to include something like the following, to copy the initialized data from the ROM image to its runtime address:


```
char *src = _etext;
char *dst = _data;

/* ROM has data at end of text; copy it. */
while (dst < _edata) {
  *dst++ = *src++;
}

/* Zero bss */
for (dst = _bstart; dst< _bend; dst++)
  *dst = 0;
```

`>`*region* Assign this section to a previously defined region of memory. See Section 3.3 [MEMORY], page 25.

`:`*phdr* Assign this section to a segment described by a program header. See Section 3.5 [PHDRS], page 35. If a section is assigned to one or more segments, then all subsequent allocated sections will be assigned to those segments as well, unless they use an explicitly `:`*phdr* modifier. To prevent a section from being assigned to a segment when it would normally default to one, use `:NONE`.

`=`*fill* Including `=`*fill* in a section definition specifies the initial fill value for that section. You may use any expression to specify *fill*. Any unallocated holes in the current output section when written to the output file will be filled with the two least significant bytes of the value, repeated as necessary. You can also change the fill value with a `FILL` statement in the *contents* of a section definition.

3.4.5 Overlays

The `OVERLAY` command provides an easy way to describe sections which are to be loaded as part of a single memory image but are to be run at the same memory address. At run time, some sort of overlay manager will copy the overlaid sections in and out of the runtime memory address as required, perhaps by simply manipulating addressing bits. This approach can be useful, for example, when a certain region of memory is faster than another.

The `OVERLAY` command is used within a `SECTIONS` command. It appears as follows:

```
OVERLAY start : [ NOCROSSREFS ] AT ( ldaddr )
 {
   secname1 { contents } :phdr =fill
   secname2 { contents } :phdr =fill
   ...
 } >region :phdr =fill
```

Everything is optional except `OVERLAY` (a keyword), and each section must have a name (*secname1* and *secname2* above). The section definitions within the `OVERLAY` construct are identical to those within the general `SECTIONS` contruct (see Section 3.4 [SECTIONS], page 26), except that no addresses and no memory regions may be defined for sections within an `OVERLAY`.

The sections are all defined with the same starting address. The load addresses of the sections are arranged such that they are consecutive in memory starting at the load address

used for the `OVERLAY` as a whole (as with normal section definitions, the load address is optional, and defaults to the start address; the start address is also optional, and defaults to .).

If the `NOCROSSREFS` keyword is used, and there any references among the sections, the linker will report an error. Since the sections all run at the same address, it normally does not make sense for one section to refer directly to another. See Section 3.8 [Option Commands], page 40.

For each section within the `OVERLAY`, the linker automatically defines two symbols. The symbol `__load_start_`*secname* is defined as the starting load address of the section. The symbol `__load_stop_`*secname* is defined as the final load address of the section. Any characters within *secname* which are not legal within C identifiers are removed. C (or assembler) code may use these symbols to move the overlaid sections around as necessary.

At the end of the overlay, the value of `.` is set to the start address of the overlay plus the size of the largest section.

Here is an example. Remember that this would appear inside a `SECTIONS` construct.

```
OVERLAY 0x1000 : AT (0x4000)
 {
   .text0 { o1/*.o(.text) }
   .text1 { o2/*.o(.text) }
 }
```

This will define both `.text0` and `.text1` to start at address 0x1000. `.text0` will be loaded at address 0x4000, and `.text1` will be loaded immediately after `.text0`. The following symbols will be defined: `__load_start_text0`, `__load_stop_text0`, `__load_start_text1`, `__load_stop_text1`.

C code to copy overlay `.text1` into the overlay area might look like the following.

```
extern char __load_start_text1, __load_stop_text1;
memcpy ((char *) 0x1000, &__load_start_text1,
        &__load_stop_text1 - &__load_start_text1);
```

Note that the `OVERLAY` command is just syntactic sugar, since everything it does can be done using the more basic commands. The above example could have been written identically as follows.

```
.text0 0x1000 : AT (0x4000) { o1/*.o(.text) }
__load_start_text0 = LOADADDR (.text0);
__load_stop_text0 = LOADADDR (.text0) + SIZEOF (.text0);
.text1 0x1000 : AT (0x4000 + SIZEOF (.text0)) { o2/*.o(.text) }
__load_start_text1 = LOADADDR (.text1);
__load_stop_text1 = LOADADDR (.text1) + SIZEOF (.text1);
. = 0x1000 + MAX (SIZEOF (.text0), SIZEOF (.text1));
```

3.5 ELF Program Headers

The ELF object file format uses *program headers*, which are read by the system loader and describe how the program should be loaded into memory. These program headers must be set correctly in order to run the program on a native ELF system. The linker will create reasonable program headers by default. However, in some cases, it is desirable to specify

the program headers more precisely; the `PHDRS` command may be used for this purpose. When the `PHDRS` command is used, the linker will not generate any program headers itself.

The `PHDRS` command is only meaningful when generating an ELF output file. It is ignored in other cases. This manual does not describe the details of how the system loader interprets program headers; for more information, see the ELF ABI. The program headers of an ELF file may be displayed using the '`-p`' option of the `objdump` command.

This is the syntax of the `PHDRS` command. The words `PHDRS`, `FILEHDR`, `AT`, and `FLAGS` are keywords.

```
PHDRS
{
  name type [ FILEHDR ] [ PHDRS ] [ AT ( address ) ]
        [ FLAGS ( flags ) ] ;
}
```

The *name* is used only for reference in the `SECTIONS` command of the linker script. It does not get put into the output file.

Certain program header types describe segments of memory which are loaded from the file by the system loader. In the linker script, the contents of these segments are specified by directing allocated output sections to be placed in the segment. To do this, the command describing the output section in the `SECTIONS` command should use '`:`*name*', where *name* is the name of the program header as it appears in the `PHDRS` command. See Section 3.4.4 [Section Options], page 32.

It is normal for certain sections to appear in more than one segment. This merely implies that one segment of memory contains another. This is specified by repeating '`:`*name*', using it once for each program header in which the section is to appear.

If a section is placed in one or more segments using '`:`*name*', then all subsequent allocated sections which do not specify '`:`*name*' are placed in the same segments. This is for convenience, since generally a whole set of contiguous sections will be placed in a single segment. To prevent a section from being assigned to a segment when it would normally default to one, use `:NONE`.

The `FILEHDR` and `PHDRS` keywords which may appear after the program header type also indicate contents of the segment of memory. The `FILEHDR` keyword means that the segment should include the ELF file header. The `PHDRS` keyword means that the segment should include the ELF program headers themselves.

The *type* may be one of the following. The numbers indicate the value of the keyword.

`PT_NULL` (0)
: Indicates an unused program header.

`PT_LOAD` (1)
: Indicates that this program header describes a segment to be loaded from the file.

`PT_DYNAMIC` (2)
: Indicates a segment where dynamic linking information can be found.

`PT_INTERP` (3)
: Indicates a segment where the name of the program interpreter may be found.

`PT_NOTE` (4)
: Indicates a segment holding note information.

`PT_SHLIB` (5)
: A reserved program header type, defined but not specified by the ELF ABI.

`PT_PHDR` (6)
: Indicates a segment where the program headers may be found.

expression
: An expression giving the numeric type of the program header. This may be used for types not defined above.

It is possible to specify that a segment should be loaded at a particular address in memory. This is done using an `AT` expression. This is identical to the `AT` command used in the `SECTIONS` command (see Section 3.4.4 [Section Options], page 32). Using the `AT` command for a program header overrides any information in the `SECTIONS` command.

Normally the segment flags are set based on the sections. The `FLAGS` keyword may be used to explicitly specify the segment flags. The value of *flags* must be an integer. It is used to set the `p_flags` field of the program header.

Here is an example of the use of `PHDRS`. This shows a typical set of program headers used on a native ELF system.

```
PHDRS
{
  headers PT_PHDR PHDRS ;
  interp PT_INTERP ;
  text PT_LOAD FILEHDR PHDRS ;
  data PT_LOAD ;
  dynamic PT_DYNAMIC ;
}

SECTIONS
{
  . = SIZEOF_HEADERS;
  .interp : { *(.interp) } :text :interp
  .text : { *(.text) } :text
  .rodata : { *(.rodata) } /* defaults to :text */
  ...
  . = . + 0x1000; /* move to a new page in memory */
  .data : { *(.data) } :data
  .dynamic : { *(.dynamic) } :data :dynamic
  ...
}
```

3.6 The Entry Point

The linker command language includes a command specifically for defining the first executable instruction in an output file (its *entry point*). Its argument is a symbol name:

```
ENTRY(symbol)
```


Like symbol assignments, the `ENTRY` command may be placed either as an independent command in the command file, or among the section definitions within the `SECTIONS` command—whatever makes the most sense for your layout.

`ENTRY` is only one of several ways of choosing the entry point. You may indicate it in any of the following ways (shown in descending order of priority: methods higher in the list override methods lower down).

- the '`-e`' *entry* command-line option;
- the `ENTRY(`*symbol*`)` command in a linker control script;
- the value of the symbol `start`, if present;
- the address of the first byte of the `.text` section, if present;
- The address `0`.

For example, you can use these rules to generate an entry point with an assignment statement: if no symbol `start` is defined within your input files, you can simply define it, assigning it an appropriate value—

```
start = 0x2020;
```

The example shows an absolute address, but you can use any expression. For example, if your input object files use some other symbol-name convention for the entry point, you can just assign the value of whatever symbol contains the start address to `start`:

```
start = other_symbol ;
```

3.7 Version Script

The linker command script includes a command specifically for specifying a version script, and is only meaningful for ELF platforms that support shared libraries. A version script can be build directly into the linker script that you are using, or you can supply the version script as just another input file to the linker at the time that you link. The command script syntax is:

```
VERSION { version script contents }
```

The version script can also be specified to the linker by means of the '`--version-script`' linker command line option. Version scripts are only meaningful when creating shared libraries.

The format of the version script itself is identical to that used by Sun's linker in Solaris 2.5. Versioning is done by defining a tree of version nodes with the names and interdependencies specified in the version script. The version script can specify which symbols are bound to which version nodes, and it can reduce a specified set of symbols to local scope so that they are not globally visible outside of the shared library.

The easiest way to demonstrate the version script language is with a few examples.

```
VERS_1.1 {
 global:
 foo1;
 local:
 old*;
 original*;
```

```
 new*;
};

VERS_1.2 {
 foo2;
} VERS_1.1;

VERS_2.0 {
 bar1; bar2;
} VERS_1.2;
```

In this example, three version nodes are defined. '`VERS_1.1`' is the first version node defined, and has no other dependencies. The symbol '`foo1`' is bound to this version node, and a number of symbols that have appeared within various object files are reduced in scope to local so that they are not visible outside of the shared library.

Next, the node '`VERS_1.2`' is defined. It depends upon '`VERS_1.1`'. The symbol '`foo2`' is bound to this version node.

Finally, the node '`VERS_2.0`' is defined. It depends upon '`VERS_1.2`'. The symbols '`bar1`' and '`bar2`' are bound to this version node.

Symbols defined in the library which aren't specifically bound to a version node are effectively bound to an unspecified base version of the library. It is possible to bind all otherwise unspecified symbols to a given version node using '`global: *`' somewhere in the version script.

Lexically the names of the version nodes have no specific meaning other than what they might suggest to the person reading them. The '`2.0`' version could just as well have appeared in between '`1.1`' and '`1.2`'. However, this would be a confusing way to write a version script.

When you link an application against a shared library that has versioned symbols, the application itself knows which version of each symbol it requires, and it also knows which version nodes it needs from each shared library it is linked against. Thus at runtime, the dynamic loader can make a quick check to make sure that the libraries you have linked against do in fact supply all of the version nodes that the application will need to resolve all of the dynamic symbols. In this way it is possible for the dynamic linker to know with certainty that all external symbols that it needs will be resolvable without having to search for each symbol reference.

The symbol versioning is in effect a much more sophisticated way of doing minor version checking that SunOS does. The fundamental problem that is being addressed here is that typically references to external functions are bound on an as-needed basis, and are not all bound when the application starts up. If a shared library is out of date, a required interface may be missing; when the application tries to use that interface, it may suddenly and unexpectedly fail. With symbol versioning, the user will get a warning when they start their program if the libraries being used with the application are too old.

There are several GNU extensions to Sun's versioning approach. The first of these is the ability to bind a symbol to a version node in the source file where the symbol is defined instead of in the versioning script. This was done mainly to reduce the burden on the library maintainer. This can be done by putting something like:

```
__asm__(".symver original_foo,foo@VERS_1.1");
```


in the C source file. This renamed the function '`original_foo`' to be an alias for '`foo`' bound to the version node '`VERS_1.1`'. The '`local:`' directive can be used to prevent the symbol '`original_foo`' from being exported.

The second GNU extension is to allow multiple versions of the same function to appear in a given shared library. In this way an incompatible change to an interface can take place without increasing the major version number of the shared library, while still allowing applications linked against the old interface to continue to function.

This can only be accomplished by using multiple '`.symver`' directives in the assembler. An example of this would be:

```
__asm__(".symver original_foo,foo@");
__asm__(".symver old_foo,foo@VERS_1.1");
__asm__(".symver old_foo1,foo@VERS_1.2");
__asm__(".symver new_foo,foo@@VERS_2.0");
```

In this example, '`foo@`' represents the symbol '`foo`' bound to the unspecified base version of the symbol. The source file that contains this example would define 4 C functions: '`original_foo`', '`old_foo`', '`old_foo1`', and '`new_foo`'.

When you have multiple definitions of a given symbol, there needs to be some way to specify a default version to which external references to this symbol will be bound. This can be accomplished with the '`foo@@VERS_2.0`' type of '`.symver`' directive. Only one version of a symbol can be declared 'default' in this manner - otherwise you would effectively have multiple definitions of the same symbol.

If you wish to bind a reference to a specific version of the symbol within the shared library, you can use the aliases of convenience (i.e. '`old_foo`'), or you can use the '`.symver`' directive to specifically bind to an external version of the function in question.

3.8 Option Commands

The command language includes a number of other commands that you can use for specialized purposes. They are similar in purpose to command-line options.

`CONSTRUCTORS`

When linking using the `a.out` object file format, the linker uses an unusual set construct to support C++ global constructors and destructors. When linking object file formats which do not support arbitrary sections, such as `ECOFF` and `XCOFF`, the linker will automatically recognize C++ global constructors and destructors by name. For these object file formats, the `CONSTRUCTORS` command tells the linker where this information should be placed. The `CONSTRUCTORS` command is ignored for other object file formats.

The symbol `__CTOR_LIST__` marks the start of the global constructors, and the symbol `__DTOR_LIST` marks the end. The first word in the list is the number of entries, followed by the address of each constructor or destructor, followed by a zero word. The compiler must arrange to actually run the code. For these object file formats GNU C++ calls constructors from a subroutine `__main`; a call to `__main` is automatically inserted into the startup code for `main`. GNU C++ runs destructors either by using `atexit`, or directly from the function `exit`.

For object file formats such as `COFF` or `ELF` which support multiple sections, GNU C++ will normally arrange to put the addresses of global constructors and destructors into the `.ctors` and `.dtors` sections. Placing the following sequence into your linker script will build the sort of table which the GNU C++ runtime code expects to see.

```
__CTOR_LIST__ = .;
LONG((__CTOR_END__ - __CTOR_LIST__) / 4 - 2)
*(.ctors)
LONG(0)
__CTOR_END__ = .;
__DTOR_LIST__ = .;
LONG((__DTOR_END__ - __DTOR_LIST__) / 4 - 2)
*(.dtors)
LONG(0)
__DTOR_END__ = .;
```

Normally the compiler and linker will handle these issues automatically, and you will not need to concern yourself with them. However, you may need to consider this if you are using C++ and writing your own linker scripts.

`FLOAT`
`NOFLOAT` These keywords were used in some older linkers to request a particular math subroutine library. `ld` doesn't use the keywords, assuming instead that any necessary subroutines are in libraries specified using the general mechanisms for linking to archives; but to permit the use of scripts that were written for the older linkers, the keywords `FLOAT` and `NOFLOAT` are accepted and ignored.

`FORCE_COMMON_ALLOCATION`
This command has the same effect as the '`-d`' command-line option: to make `ld` assign space to common symbols even if a relocatable output file is specified ('`-r`').

`INCLUDE` *filename*
Include the linker script *filename* at this point. The file will be searched for in the current directory, and in any directory specified with the `-L` option. You can nest calls to `INCLUDE` up to 10 levels deep.

`INPUT (` *file*, *file*, ... `)`
`INPUT (` *file file* ... `)`
Use this command to include binary input files in the link, without including them in a particular section definition. Specify the full name for each *file*, including '`.a`' if required.

`ld` searches for each *file* through the archive-library search path, just as for files you specify on the command line. See the description of '`-L`' in Section 2.1 [Command Line Options], page 3.

If you use '`-l`*file*', `ld` will transform the name to `lib`*file*`.a` as with the command line argument '`-l`'.

`GROUP (` *file*, *file*, ... `)`
`GROUP (` *file file* ... `)`
This command is like `INPUT`, except that the named files should all be archives, and they are searched repeatedly until no new undefined references are created. See the description of '`-(`' in Section 2.1 [Command Line Options], page 3.

`OUTPUT (` *filename* `)`
Use this command to name the link output file *filename*. The effect of `OUTPUT(`*filename*`)` is identical to the effect of '`-o` *filename*', which overrides it. You can use this command to supply a default output-file name other than `a.out`.

`OUTPUT_ARCH (` *bfdname* `)`
Specify a particular output machine architecture, with one of the names used by the BFD back-end routines (see Chapter 5 [BFD], page 47). This command is often unnecessary; the architecture is most often set implicitly by either the system BFD configuration or as a side effect of the `OUTPUT_FORMAT` command.

`OUTPUT_FORMAT (` *bfdname* `)`
When `ld` is configured to support multiple object code formats, you can use this command to specify a particular output format. *bfdname* is one of the names used by the BFD back-end routines (see Chapter 5 [BFD], page 47). The effect is identical to the effect of the '`--oformat`' command-line option. This selection affects only the output file; the related command `TARGET` affects primarily input files.

`SEARCH_DIR (` *path* `)`
Add *path* to the list of paths where `ld` looks for archive libraries. `SEARCH_DIR(`*path*`)` has the same effect as '`-L`*path*' on the command line.

`STARTUP (` *filename* `)`
Ensure that *filename* is the first input file used in the link process.

`TARGET (` *format* `)`
When `ld` is configured to support multiple object code formats, you can use this command to change the input-file object code format (like the command-line option '`-b`' or its synonym '`--format`'). The argument *format* is one of the strings used by BFD to name binary formats. If `TARGET` is specified but `OUTPUT_FORMAT` is not, the last `TARGET` argument is also used as the default format for the `ld` output file. See Chapter 5 [BFD], page 47.

If you don't use the `TARGET` command, `ld` uses the value of the environment variable `GNUTARGET`, if available, to select the output file format. If that variable is also absent, `ld` uses the default format configured for your machine in the BFD libraries.

`NOCROSSREFS (` *section section* ... `)`
This command may be used to tell `ld` to issue an error about any references among certain sections.

In certain types of programs, particularly on embedded systems, when one section is loaded into memory, another section will not be. Any direct references

between the two sections would be errors. For example, it would be an error if code in one section called a function defined in the other section.

The `NOCROSSREFS` command takes a list of section names. If `ld` detects any cross references between the sections, it reports an error and returns a non-zero exit status. The `NOCROSSREFS` command uses output section names, defined in the `SECTIONS` command. It does not use the names of input sections.

4 Machine Dependent Features

`ld` has additional features on some platforms; the following sections describe them. Machines where `ld` has no additional functionality are not listed.

4.1 `ld` and the H8/300

For the H8/300, `ld` can perform these global optimizations when you specify the '`--relax`' command-line option.

relaxing address modes
: `ld` finds all `jsr` and `jmp` instructions whose targets are within eight bits, and turns them into eight-bit program-counter relative `bsr` and `bra` instructions, respectively.

synthesizing instructions
: `ld` finds all `mov.b` instructions which use the sixteen-bit absolute address form, but refer to the top page of memory, and changes them to use the eight-bit address form. (That is: the linker turns '`mov.b @`*aa*`:16`' into '`mov.b @`*aa*`:8`' whenever the address *aa* is in the top page of memory).

4.2 `ld` and the Intel 960 family

You can use the '`-A`*architecture*' command line option to specify one of the two-letter names identifying members of the 960 family; the option specifies the desired output target, and warns of any incompatible instructions in the input files. It also modifies the linker's search strategy for archive libraries, to support the use of libraries specific to each particular architecture, by including in the search loop names suffixed with the string identifying the architecture.

For example, if your `ld` command line included '`-ACA`' as well as '`-ltry`', the linker would look (in its built-in search paths, and in any paths you specify with '`-L`') for a library with the names

```
try
libtry.a
tryca
libtryca.a
```

The first two possibilities would be considered in any event; the last two are due to the use of '`-ACA`'.

You can meaningfully use '`-A`' more than once on a command line, since the 960 architecture family allows combination of target architectures; each use will add another pair of name variants to search for when '`-l`' specifies a library.

`ld` supports the '`--relax`' option for the i960 family. If you specify '`--relax`', `ld` finds all `balx` and `calx` instructions whose targets are within 24 bits, and turns them into 24-bit program-counter relative `bal` and `cal` instructions, respectively. `ld` also turns `cal` instructions into `bal` instructions when it determines that the target subroutine is a leaf routine (that is, the target subroutine does not itself call any subroutines).

5 BFD

The linker accesses object and archive files using the BFD libraries. These libraries allow the linker to use the same routines to operate on object files whatever the object file format. A different object file format can be supported simply by creating a new BFD back end and adding it to the library. To conserve runtime memory, however, the linker and associated tools are usually configured to support only a subset of the object file formats available. You can use `objdump -i` (see section "objdump" in *The GNU Binary Utilities*) to list all the formats available for your configuration.

As with most implementations, BFD is a compromise between several conflicting requirements. The major factor influencing BFD design was efficiency: any time used converting between formats is time which would not have been spent had BFD not been involved. This is partly offset by abstraction payback; since BFD simplifies applications and back ends, more time and care may be spent optimizing algorithms for a greater speed.

One minor artifact of the BFD solution which you should bear in mind is the potential for information loss. There are two places where useful information can be lost using the BFD mechanism: during conversion and during output. See Section 5.1.1 [BFD information loss], page 47.

5.1 How it works: an outline of BFD

When an object file is opened, BFD subroutines automatically determine the format of the input object file. They then build a descriptor in memory with pointers to routines that will be used to access elements of the object file's data structures.

As different information from the the object files is required, BFD reads from different sections of the file and processes them. For example, a very common operation for the linker is processing symbol tables. Each BFD back end provides a routine for converting between the object file's representation of symbols and an internal canonical format. When the linker asks for the symbol table of an object file, it calls through a memory pointer to the routine from the relevant BFD back end which reads and converts the table into a canonical form. The linker then operates upon the canonical form. When the link is finished and the linker writes the output file's symbol table, another BFD back end routine is called to take the newly created symbol table and convert it into the chosen output format.

5.1.1 Information Loss

Information can be lost during output. The output formats supported by BFD do not provide identical facilities, and information which can be described in one form has nowhere to go in another format. One example of this is alignment information in `b.out`. There is nowhere in an `a.out` format file to store alignment information on the contained data, so when a file is linked from `b.out` and an `a.out` image is produced, alignment information will not propagate to the output file. (The linker will still use the alignment information internally, so the link is performed correctly).

Another example is COFF section names. COFF files may contain an unlimited number of sections, each one with a textual section name. If the target of the link is a format which

does not have many sections (e.g., `a.out`) or has sections without names (e.g., the Oasys format), the link cannot be done simply. You can circumvent this problem by describing the desired input-to-output section mapping with the linker command language.

Information can be lost during canonicalization. The BFD internal canonical form of the external formats is not exhaustive; there are structures in input formats for which there is no direct representation internally. This means that the BFD back ends cannot maintain all possible data richness through the transformation between external to internal and back to external formats.

This limitation is only a problem when an application reads one format and writes another. Each BFD back end is responsible for maintaining as much data as possible, and the internal BFD canonical form has structures which are opaque to the BFD core, and exported only to the back ends. When a file is read in one format, the canonical form is generated for BFD and the application. At the same time, the back end saves away any information which may otherwise be lost. If the data is then written back in the same format, the back end routine will be able to use the canonical form provided by the BFD core as well as the information it prepared earlier. Since there is a great deal of commonality between back ends, there is no information lost when linking or copying big endian COFF to little endian COFF, or `a.out` to `b.out`. When a mixture of formats is linked, the information is only lost from the files whose format differs from the destination.

5.1.2 The BFD canonical object-file format

The greatest potential for loss of information occurs when there is the least overlap between the information provided by the source format, that stored by the canonical format, and that needed by the destination format. A brief description of the canonical form may help you understand which kinds of data you can count on preserving across conversions.

files Information stored on a per-file basis includes target machine architecture, particular implementation format type, a demand pageable bit, and a write protected bit. Information like Unix magic numbers is not stored here—only the magic numbers' meaning, so a `ZMAGIC` file would have both the demand pageable bit and the write protected text bit set. The byte order of the target is stored on a per-file basis, so that big- and little-endian object files may be used with one another.

sections Each section in the input file contains the name of the section, the section's original address in the object file, size and alignment information, various flags, and pointers into other BFD data structures.

symbols Each symbol contains a pointer to the information for the object file which originally defined it, its name, its value, and various flag bits. When a BFD back end reads in a symbol table, it relocates all symbols to make them relative to the base of the section where they were defined. Doing this ensures that each symbol points to its containing section. Each symbol also has a varying amount of hidden private data for the BFD back end. Since the symbol points to the original file, the private data format for that symbol is accessible. `ld` can operate on a collection of symbols of wildly different formats without problems.

Normal global and simple local symbols are maintained on output, so an output file (no matter its format) will retain symbols pointing to functions and to global, static, and common variables. Some symbol information is not worth retaining; in `a.out`, type information is stored in the symbol table as long symbol names. This information would be useless to most COFF debuggers; the linker has command line switches to allow users to throw it away.

There is one word of type information within the symbol, so if the format supports symbol type information within symbols (for example, COFF, IEEE, Oasys) and the type is simple enough to fit within one word (nearly everything but aggregates), the information will be preserved.

relocation level

Each canonical BFD relocation record contains a pointer to the symbol to relocate to, the offset of the data to relocate, the section the data is in, and a pointer to a relocation type descriptor. Relocation is performed by passing messages through the relocation type descriptor and the symbol pointer. Therefore, relocations can be performed on output data using a relocation method that is only available in one of the input formats. For instance, Oasys provides a byte relocation format. A relocation record requesting this relocation type would point indirectly to a routine to perform this, so the relocation may be performed on a byte being written to a 68k COFF file, even though 68k COFF has no such relocation type.

line numbers

Object formats can contain, for debugging purposes, some form of mapping between symbols, source line numbers, and addresses in the output file. These addresses have to be relocated along with the symbol information. Each symbol with an associated list of line number records points to the first record of the list. The head of a line number list consists of a pointer to the symbol, which allows finding out the address of the function whose line number is being described. The rest of the list is made up of pairs: offsets into the section and line numbers. Any format which can simply derive this information can pass it successfully between formats (COFF, IEEE and Oasys).

6 Reporting Bugs

Your bug reports play an essential role in making `ld` reliable.

Reporting a bug may help you by bringing a solution to your problem, or it may not. But in any case the principal function of a bug report is to help the entire community by making the next version of `ld` work better. Bug reports are your contribution to the maintenance of `ld`.

In order for a bug report to serve its purpose, you must include the information that enables us to fix the bug.

6.1 Have you found a bug?

If you are not sure whether you have found a bug, here are some guidelines:

- If the linker gets a fatal signal, for any input whatever, that is a `ld` bug. Reliable linkers never crash.
- If `ld` produces an error message for valid input, that is a bug.
- If `ld` does not produce an error message for invalid input, that may be a bug. In the general case, the linker can not verify that object files are correct.
- If you are an experienced user of linkers, your suggestions for improvement of `ld` are welcome in any case.

6.2 How to report bugs

A number of companies and individuals offer support for GNU products. If you obtained `ld` from a support organization, we recommend you contact that organization first.

You can find contact information for many support companies and individuals in the file '`etc/SERVICE`' in the GNU Emacs distribution.

In any event, we also recommend that you send bug reports for `ld` to '`bug-gnu-utils@gnu.org`'.

The fundamental principle of reporting bugs usefully is this: **report all the facts**. If you are not sure whether to state a fact or leave it out, state it!

Often people omit facts because they think they know what causes the problem and assume that some details do not matter. Thus, you might assume that the name of a symbol you use in an example does not matter. Well, probably it does not, but one cannot be sure. Perhaps the bug is a stray memory reference which happens to fetch from the location where that name is stored in memory; perhaps, if the name were different, the contents of that location would fool the linker into doing the right thing despite the bug. Play it safe and give a specific, complete example. That is the easiest thing for you to do, and the most helpful.

Keep in mind that the purpose of a bug report is to enable us to fix the bug if it is new to us. Therefore, always write your bug reports on the assumption that the bug has not been reported previously.

Sometimes people give a few sketchy facts and ask, "Does this ring a bell?" Those bug reports are useless, and we urge everyone to *refuse to respond to them* except to chide the sender to report bugs properly.

To enable us to fix the bug, you should include all these things:

- The version of `ld`. `ld` announces it if you start it with the '`--version`' argument.

  Without this, we will not know whether there is any point in looking for the bug in the current version of `ld`.
- Any patches you may have applied to the `ld` source, including any patches made to the `BFD` library.
- The type of machine you are using, and the operating system name and version number.
- What compiler (and its version) was used to compile `ld`—e.g. "`gcc-2.7`".
- The command arguments you gave the linker to link your example and observe the bug. To guarantee you will not omit something important, list them all. A copy of the Makefile (or the output from make) is sufficient.

  If we were to try to guess the arguments, we would probably guess wrong and then we might not encounter the bug.
- A complete input file, or set of input files, that will reproduce the bug. It is generally most helpful to send the actual object files, uuencoded if necessary to get them through the mail system. Making them available for anonymous FTP is not as good, but may be the only reasonable choice for large object files.

  If the source files were assembled using `gas` or compiled using `gcc`, then it may be OK to send the source files rather than the object files. In this case, be sure to say exactly what version of `gas` or `gcc` was used to produce the object files. Also say how `gas` or `gcc` were configured.
- A description of what behavior you observe that you believe is incorrect. For example, "It gets a fatal signal."

  Of course, if the bug is that `ld` gets a fatal signal, then we will certainly notice it. But if the bug is incorrect output, we might not notice unless it is glaringly wrong. You might as well not give us a chance to make a mistake.

  Even if the problem you experience is a fatal signal, you should still say so explicitly. Suppose something strange is going on, such as, your copy of `ld` is out of synch, or you have encountered a bug in the C library on your system. (This has happened!) Your copy might crash and ours would not. If you told us to expect a crash, then when ours fails to crash, we would know that the bug was not happening for us. If you had not told us to expect a crash, then we would not be able to draw any conclusion from our observations.
- If you wish to suggest changes to the `ld` source, send us context diffs, as generated by `diff` with the '`-u`', '`-c`', or '`-p`' option. Always send diffs from the old file to the new file. If you even discuss something in the `ld` source, refer to it by context, not by line number.

  The line numbers in our development sources will not match those in your sources. Your line numbers would convey no useful information to us.

Here are some things that are not necessary:

- A description of the envelope of the bug.

  Often people who encounter a bug spend a lot of time investigating which changes to the input file will make the bug go away and which changes will not affect it.

This is often time consuming and not very useful, because the way we will find the bug is by running a single example under the debugger with breakpoints, not by pure deduction from a series of examples. We recommend that you save your time for something else.

Of course, if you can find a simpler example to report *instead* of the original one, that is a convenience for us. Errors in the output will be easier to spot, running under the debugger will take less time, and so on.

However, simplification is not vital; if you do not want to do this, report the bug anyway and send us the entire test case you used.

- A patch for the bug.

  A patch for the bug does help us if it is a good one. But do not omit the necessary information, such as the test case, on the assumption that a patch is all we need. We might see problems with your patch and decide to fix the problem another way, or we might not understand it at all.

  Sometimes with a program as complicated as `ld` it is very hard to construct an example that will make the program follow a certain path through the code. If you do not send us the example, we will not be able to construct one, so we will not be able to verify that the bug is fixed.

  And if we cannot understand what bug you are trying to fix, or why your patch should be an improvement, we will not install it. A test case will help us to understand.

- A guess about what the bug is or what it depends on.

  Such guesses are usually wrong. Even we cannot guess right about such things without first using the debugger to find the facts.

Appendix A MRI Compatible Script Files

To aid users making the transition to GNU `ld` from the MRI linker, `ld` can use MRI compatible linker scripts as an alternative to the more general-purpose linker scripting language described in Chapter 3 [Command Language], page 19. MRI compatible linker scripts have a much simpler command set than the scripting language otherwise used with `ld`. GNU `ld` supports the most commonly used MRI linker commands; these commands are described here.

In general, MRI scripts aren't of much use with the `a.out` object file format, since it only has three sections and MRI scripts lack some features to make use of them.

You can specify a file containing an MRI-compatible script using the '`-c`' command-line option.

Each command in an MRI-compatible script occupies its own line; each command line starts with the keyword that identifies the command (though blank lines are also allowed for punctuation). If a line of an MRI-compatible script begins with an unrecognized keyword, `ld` issues a warning message, but continues processing the script.

Lines beginning with '`*`' are comments.

You can write these commands using all upper-case letters, or all lower case; for example, '`chip`' is the same as '`CHIP`'. The following list shows only the upper-case form of each command.

`ABSOLUTE` *secname*
`ABSOLUTE` *secname*`,` *secname*`, ...` *secname*
: Normally, `ld` includes in the output file all sections from all the input files. However, in an MRI-compatible script, you can use the `ABSOLUTE` command to restrict the sections that will be present in your output program. If the `ABSOLUTE` command is used at all in a script, then only the sections named explicitly in `ABSOLUTE` commands will appear in the linker output. You can still use other input sections (whatever you select on the command line, or using `LOAD`) to resolve addresses in the output file.

`ALIAS` *out-secname*`,` *in-secname*
: Use this command to place the data from input section *in-secname* in a section called *out-secname* in the linker output file.

  in-secname may be an integer.

`ALIGN` *secname* `=` *expression*
: Align the section called *secname* to *expression*. The *expression* should be a power of two.

`BASE` *expression*
: Use the value of *expression* as the lowest address (other than absolute addresses) in the output file.

`CHIP` *expression*
`CHIP` *expression*`,` *expression*
: This command does nothing; it is accepted only for compatibility.

`END`
: This command does nothing whatever; it's only accepted for compatibility.

`FORMAT` *output-format*
: Similar to the `OUTPUT_FORMAT` command in the more general linker language, but restricted to one of these output formats:

  1. S-records, if *output-format* is '`S`'
  2. IEEE, if *output-format* is '`IEEE`'
  3. COFF (the '`coff-m68k`' variant in BFD), if *output-format* is '`COFF`'

`LIST` *anything ...*
: Print (to the standard output file) a link map, as produced by the `ld` command-line option '`-M`'.

  The keyword `LIST` may be followed by anything on the same line, with no change in its effect.

`LOAD` *filename*
`LOAD` *filename*`,` *filename*`, ...` *filename*
: Include one or more object file *filename* in the link; this has the same effect as specifying *filename* directly on the `ld` command line.

`NAME` *output-name*
: *output-name* is the name for the program produced by `ld`; the MRI-compatible command `NAME` is equivalent to the command-line option '`-o`' or the general script language command `OUTPUT`.

`ORDER` *secname*`,` *secname*`, ...` *secname*
`ORDER` *secname secname secname*
: Normally, `ld` orders the sections in its output file in the order in which they first appear in the input files. In an MRI-compatible script, you can override this ordering with the `ORDER` command. The sections you list with `ORDER` will appear first in your output file, in the order specified.

`PUBLIC` *name*`=`*expression*
`PUBLIC` *name*`,`*expression*
`PUBLIC` *name expression*
: Supply a value (*expression*) for external symbol *name* used in the linker input files.

`SECT` *secname*`,` *expression*
`SECT` *secname*`=`*expression*
`SECT` *secname expression*
: You can use any of these three forms of the `SECT` command to specify the start address (*expression*) for section *secname*. If you have more than one `SECT` statement for the same *secname*, only the *first* sets the start address.

Index

*

*(COMMON) ... 29
*(*section*) ... 28

-

-(... 9
--architecture=*arch* ... 4
--auxiliary ... 5
--cref ... 10
--defsym *symbol=exp* ... 10
--discard-all ... 9
--discard-locals ... 9
--dynamic-linker *file* ... 10
--embedded-relocs ... 11
--entry=*entry* ... 5
--export-dynamic ... 5
--filter ... 5
--force-exe-suffix ... 6
--format=*format* ... 4
--gpsize ... 6
--help ... 11
--just-symbols=*file* ... 8
--library-path=*dir* ... 7
--library=*archive* ... 6
--mri-script=*MRI-cmdfile* ... 4
--nmagic ... 7
--no-keep-memory ... 11
--no-warn-mismatch ... 11
--no-whole-archive ... 11
--noinhibit-exec ... 11
--oformat ... 11
--omagic ... 7
--output=*output* ... 8
--print-map ... 7
--relax ... 11
--relax on i960 ... 45
--relocateable ... 8
--script=*script* ... 8
--sort-common ... 13
--split-by-file ... 13
--split-by-reloc ... 13
--stats ... 13
--strip-all ... 8
--strip-debug ... 8
--trace ... 8
--trace-symbol=*symbol* ... 9
--traditional-format ... 13
--undefined=*symbol* ... 8
--verbose ... 14
--version ... 9
--version-script=*version-scriptfile* ... 14
--warn-comon ... 14
--warn-constructors ... 15
--warn-multiple-gp ... 15
--warn-once ... 16
--warn-section-align ... 16
--whole-archive ... 16
--wrap ... 16
-A*arch* ... 4
-a*keyword* ... 4
-assert *keyword* ... 9
-b *format* ... 4
-Bdynamic ... 9
-Bshareable ... 13
-Bstatic ... 10
-Bsymbolic ... 10
-c *MRI-cmdfile* ... 4
-call_shared ... 9
-d ... 4
-dc ... 4
-dn ... 10
-dp ... 4
-dy ... 9
-E ... 5
-e *entry* ... 5
-EB ... 10
-EL ... 10
-f ... 5
-F ... 5
-g ... 6
-G ... 6
-h*name* ... 6
-i ... 6
-l*archive* ... 6
-L*dir* ... 7
-M ... 7
-m *emulation* ... 7
-Map ... 11
-n ... 7
-N ... 7
-non_shared ... 10
-o *output* ... 8
-qmagic ... 11
-Qy ... 11
-r ... 8
-R *file* ... 8
-rpath ... 12
-rpath-link ... 12
-s ... 8
-S ... 8
-shared ... 13
-soname=*name* ... 6
-static ... 10
-t ... 8
-T *script* ... 8
-Tbss *org* ... 14
-Tdata *org* ... 14
-Ttext *org* ... 14
-u *symbol* ... 8
-Ur ... 14
-v ... 9
-V ... 9
-x ... 9
-X ... 9
-Y *path* ... 9
-y *symbol* ... 9
-z *keyword* ... 9

.

. ... 20

:

:*phdr* ... 34

;

; ... 22

=

=*fill* ... 34

[

[*section*...], not supported ... 29

"

" ... 20

>

>*region* ... 34

0

0x ... 20

A

ABSOLUTE (MRI) ... 55
absolute and relocatable symbols ... 22
ABSOLUTE(*exp*) ... 23
ADDR(*section*) ... 23
ALIAS (MRI) ... 55
ALIGN (MRI) ... 55
ALIGN(*exp*) ... 23
aligning sections ... 33
allocating memory ... 25
architectures ... 4
archive files, from cmd line ... 6
arithmetic ... 19
arithmetic operators ... 21
assignment in scripts ... 22
assignment, in section defn ... 31
AT (*ldadr*) ... 33

B

back end ... 47
BASE (MRI) ... 55
BFD canonical format ... 48
BFD requirements ... 47
big-endian objects ... 10
binary input files ... 41
binary input format ... 4
BLOCK(*align*) ... 33
bug criteria ... 51
bug reports ... 51
bugs in ld ... 51
BYTE(*expression*) ... 32

C

C++ constructors, arranging in link ... 40
CHIP (MRI) ... 55
combining symbols, warnings on ... 14
command files ... 19
command line ... 3
commands, fundamental ... 19
comments ... 19
common allocation ... 4, 41
commons in output ... 29
compatibility, MRI ... 4
constructors ... 14
CONSTRUCTORS ... 40
constructors, arranging in link ... 40
contents of a section ... 28
crash of linker ... 51
CREATE_OBJECT_SYMBOLS ... 30
cross reference table ... 10

cross references ... 42
current output location ... 20

D

dbx ... 13
decimal integers ... 20
default emulation ... 17
default input format ... 17
DEFINED(*symbol*) ... 24
deleting local symbols ... 9
direct output ... 32
discontinuous memory ... 25
dot ... 20
dynamic linker, from command line ... 10
dynamic symbol table ... 5

E

ELF program headers ... 35
emulation ... 7
emulation, default ... 17
END (MRI) ... 55
endianness ... 10
entry point, defaults ... 38
entry point, from command line ... 5
ENTRY(*symbol*) ... 37
error on valid input ... 51
expression evaluation order ... 21
expression syntax ... 19
expression, absolute ... 23
expressions in a section ... 30

F

fatal signal ... 51
filename ... 28
filename symbols ... 30
filename(*section*) ... 28
files and sections, section defn ... 28
files, including in output sections ... 28
fill pattern, entire section ... 34
FILL(*expression*) ... 32
first input file ... 42
first instruction ... 37
FLOAT ... 41
FORCE_COMMON_ALLOCATION ... 41
FORMAT (MRI) ... 56
format, output file ... 42
functions in expression language ... 23
fundamental script commands ... 19

G

GNU linker ... 1
GNUTARGET ... 17, 42
GROUP (*files*) ... 41
grouping input files ... 41
groups of archives ... 9

H

H8/300 support ... 45
header size ... 25
help ... 11
hexadecimal integers ... 20
holes ... 20
holes, filling ... 32

I

i960 support ... 45
INCLUDE *filename* ... 41
including a linker script ... 41
including an entire archive ... 16
incremental link ... 6
INPUT (*files*) ... 41
input file format ... 42
input filename symbols ... 30
input files, displaying ... 8
input files, section defn ... 28
input format ... 4
input sections to output section ... 28
integer notation ... 20
integer suffixes ... 20
internal object-file format ... 48
invalid input ... 51

K

K and M integer suffixes ... 20

L

l = ... 26
L, deleting symbols beginning ... 9
layout of output file ... 19
lazy evaluation ... 21
ld bugs, reporting ... 51
LDEMULATION ... 17
len = ... 26
LENGTH = ... 26
link map ... 7
link-time runtime library search path ... 12
linker crash ... 51
LIST (MRI) ... 56
little-endian objects ... 10

LOAD (MRI) ... 56
load address, specifying ... 33
LOADADDR(*section*) ... 23
loading, preventing ... 33
local symbols, deleting ... 9
location counter ... 20
LONG(*expression*) ... 32

M

M and K integer suffixes ... 20
machine architecture, output ... 42
machine dependencies ... 45
MAX ... 25
MEMORY ... 25
memory region attributes ... 26
memory regions and sections ... 34
memory usage ... 11
MIN ... 25
MIPS embedded PIC code ... 11
MRI compatibility ... 55

N

NAME (MRI) ... 56
names ... 20
naming memory regions ... 25
naming output sections ... 27
naming the output file ... 8, 42
negative integers ... 20
NEXT(*exp*) ... 24
NMAGIC ... 7
NOCROSSREFS (*sections*) ... 42
NOFLOAT ... 41
NOLOAD ... 33
Non constant expression ... 23

O

o = ... 26
objdump -i ... 47
object file management ... 47
object files ... 3
object formats available ... 47
object size ... 6
octal integers ... 20
OMAGIC ... 7
opening object files ... 47
Operators for arithmetic ... 21
options ... 3
ORDER (MRI) ... 56
org = ... 26
ORIGIN = ... 26
OUTPUT (*filename*) ... 42
output file after errors ... 11
output file layout ... 19
OUTPUT_ARCH (*bfdname*) ... 42
OUTPUT_FORMAT (*bfdname*) ... 42
OVERLAY ... 34
overlays ... 34

P

partial link ... 8
path for libraries ... 42
PHDRS ... 35
precedence in expressions ... 21
prevent unnecessary loading ... 33
program headers ... 35
program headers and sections ... 34
provide ... 23
PUBLIC (MRI) ... 56

Q

QUAD(*expression*) ... 32
quoted symbol names ... 20

R

read-only text ... 7
read/write from cmd line ... 7
regions of memory ... 25
relaxing addressing modes ... 11
relaxing on H8/300 ... 45
relaxing on i960 ... 45
relocatable and absolute symbols ... 22
relocatable output ... 8
reporting bugs in ld ... 51
requirements for BFD ... 47
retaining specified symbols ... 12
rounding up location counter ... 23
runtime library name ... 6
runtime library search path ... 12

S

scaled integers ... 20
script files ... 8
search directory, from cmd line ... 7
search path, libraries ... 42
SEARCH_DIR (*path*) ... 42
SECT (MRI) ... 56
section address ... 23, 32
section alignment ... 33
section alignment, warnings on ... 16
section definition ... 27

section defn, full syntax ... 32
section fill pattern ... 34
section load address ... 23
section size ... 24
section start ... 32
section, assigning to memory region ... 34
section, assigning to program header ... 34
SECTIONS ... 26
segment origins, cmd line ... 14
semicolon ... 22
shared libraries ... 13
SHORT(*expression*) ... 32
SIZEOF(*section*) ... 24
sizeof_headers ... 25
SIZEOF_HEADERS ... 25
specify load address ... 33
SQUAD(*expression*) ... 32
standard Unix system ... 3
start address, section ... 32
start of execution ... 37
STARTUP (*filename*) ... 42
strip all symbols ... 8
strip debugger symbols ... 8
stripping all but some symbols ... 12
suffixes for integers ... 20
symbol = *expression* ; ... 31
symbol defaults ... 24
symbol definition, scripts ... 22
symbol f= *expression* ; ... 31
symbol names ... 20
symbol tracing ... 9
symbol versions ... 38
symbol-only input ... 8
symbols, from command line ... 10
symbols, relocatable and absolute ... 22
symbols, retaining selectively ... 12
synthesizing linker ... 11
synthesizing on H8/300 ... 45

T

TARGET (*format*) ... 42
traditional format ... 13

U

unallocated address, next ... 24
undefined symbol ... 8
undefined symbols, warnings on ... 16
uninitialized data ... 29
unspecified memory ... 32
usage ... 11

V

variables, defining ... 22
verbose ... 14
version ... 9
VERSION {script text} ... 38
version script ... 38
version script, symbol versions ... 14
versions of symbols ... 38

W

warnings, on combining symbols ... 14
warnings, on section alignment ... 16
warnings, on undefined symbols ... 16
what is this? ... 1

The body of this manual is set in
cmr10 at 10.95pt,
with headings in **cmb10 at 10.95pt**
and examples in `cmtt10 at 10.95pt`.
cmti10 at 10.95pt and
cmsl10 at 10.95pt
are used for emphasis.

Table of Contents

1 Overview ... 1

2 Invocation ... 3

2.1 Command Line Options ... 3
2.2 Environment Variables ... 17

3 Command Language ... 19

3.1 Linker Scripts ... 19
3.2 Expressions ... 19
3.2.1 Integers ... 20
3.2.2 Symbol Names ... 20
3.2.3 The Location Counter ... 20
3.2.4 Operators ... 21
3.2.5 Evaluation ... 21
3.2.6 Assignment: Defining Symbols ... 22
3.2.7 Arithmetic Functions ... 23
3.2.8 Semicolons ... 25
3.3 Memory Layout ... 25
3.4 Specifying Output Sections ... 26
3.4.1 Section Definitions ... 27
3.4.2 Section Placement ... 28
3.4.3 Section Data Expressions ... 30
3.4.4 Optional Section Attributes ... 32
3.4.5 Overlays ... 34
3.5 ELF Program Headers ... 35
3.6 The Entry Point ... 37
3.7 Version Script ... 38
3.8 Option Commands ... 40

4 Machine Dependent Features ... 45

4.1 `ld` and the H8/300 ... 45
4.2 `ld` and the Intel 960 family ... 45

5 BFD ... 47

5.1 How it works: an outline of BFD ... 47
5.1.1 Information Loss ... 47
5.1.2 The BFD canonical object-file format ... 48

6 Reporting Bugs ... 51

6.1 Have you found a bug? ... 51
6.2 How to report bugs ... 51

Appendix A MRI Compatible Script Files ... 55

Index ... 57