# The GNU Binary Utilities

Version 2.2

May 1993

Roland H. Pesch
Jeffrey M. Osier
Cygnus Support

Cygnus Support
TEXinfo 2.137 (Cygnus)

Copyright © 1991, 1992, 1993, 1994 1995 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.

# Table of Contents

**Introduction** ..... **1**

**1 ar** ..... **3**
- 1.1 Controlling `ar` on the command line ..... 4
- 1.2 Controlling `ar` with a script ..... 6

**2 ld** ..... **11**

**3 nm** ..... **13**

**4 objcopy** ..... **17**

**5 objdump** ..... **21**

**6 ranlib** ..... **25**

**7 size** ..... **27**

**8 strings** ..... **29**

**9 strip** ..... **31**

**10 c++filt** ..... **33**

**11 nlmconv** ..... **35**

**12 Selecting the target system** ..... **37**
- 12.1 Target Selection ..... 37
- 12.2 Architecture selection ..... 39
- 12.3 Linker emulation selection ..... 39

**Index** ..... **41**

# Introduction

This brief manual contains preliminary documentation for the GNU binary utilities (collectively version 2.2):

| | |
|---|---|
| `ar` | Create, modify, and extract from archives |
| `nm` | List symbols from object files |
| `objcopy` | Copy and translate object files |
| `objdump` | Display information from object files |
| `ranlib` | Generate index to archive contents |
| `size` | List file section sizes and total size |
| `strings` | List printable strings from files |
| `strip` | Discard symbols |
| `c++filt` | Demangle encoded C++ symbols |
| `nlmconv` | Convert object code into a Netware Loadable Module |

# 1 ar

```
ar [-]p[mod [relpos]] archive [member...]
ar -M [ <mri-script ]
```

The GNU ar program creates, modifies, and extracts from archives. An *archive* is a single file holding a collection of other files in a structure that makes it possible to retrieve the original individual files (called *members* of the archive).

The original files' contents, mode (permissions), timestamp, owner, and group are preserved in the archive, and can be restored on extraction.

GNU ar can maintain archives whose members have names of any length; however, depending on how ar is configured on your system, a limit on member-name length may be imposed for compatibility with archive formats maintained with other tools. If it exists, the limit is often 15 characters (typical of formats related to a.out) or 16 characters (typical of formats related to coff).

ar is considered a binary utility because archives of this sort are most often used as *libraries* holding commonly needed subroutines.

ar creates an index to the symbols defined in relocatable object modules in the archive when you specify the modifier 's'. Once created, this index is updated in the archive whenever ar makes a change to its contents (save for the 'q' update operation). An archive with such an index speeds up linking to the library, and allows routines in the library to call each other without regard to their placement in the archive.

You may use 'nm -s' or 'nm --print-armap' to list this index table. If an archive lacks the table, another form of ar called ranlib can be used to add just the table.

GNU ar is designed to be compatible with two different facilities. You can control its activity using command-line options, like the different varieties of ar on Unix systems; or, if you specify the single command-line option '-M', you can control it with a script supplied via standard input, like the MRI "librarian" program.

## 1.1 Controlling `ar` on the command line

```
ar [-]p[mod [relpos]] archive [member...]
```

When you use `ar` in the Unix style, `ar` insists on at least two arguments to execute: one keyletter specifying the *operation* (optionally accompanied by other keyletters specifying *modifiers*), and the archive name to act on.

Most operations can also accept further *member* arguments, specifying particular files to operate on.

GNU `ar` allows you to mix the operation code *p* and modifier flags *mod* in any order, within the first command-line argument.

If you wish, you may begin the first command-line argument with a dash.

The *p* keyletter specifies what operation to execute; it may be any of the following, but you must specify only one of them:

`d` — *Delete* modules from the archive. Specify the names of modules to be deleted as *member*...; the archive is untouched if you specify no files to delete.

If you specify the 'v' modifier, `ar` lists each module as it is deleted.

`m` — Use this operation to *move* members in an archive.

The ordering of members in an archive can make a difference in how programs are linked using the library, if a symbol is defined in more than one member.

If no modifiers are used with `m`, any members you name in the *member* arguments are moved to the *end* of the archive; you can use the 'a', 'b', or 'i' modifiers to move them to a specified place instead.

`p` — *Print* the specified members of the archive, to the standard output file. If the 'v' modifier is specified, show the member name before copying its contents to standard output.

If you specify no *member* arguments, all the files in the archive are printed.

`q` — *Quick append*; add the files *member*... to the end of *archive*, without checking for replacement.

The modifiers 'a', 'b', and 'i' do *not* affect this operation; new members are always placed at the end of the archive.

The modifier 'v' makes `ar` list each file as it is appended.

Since the point of this operation is speed, the archive's symbol table index is not updated, even if it already existed; you can

use 'ar s' or ranlib explicitly to update the symbol table index.

r — Insert the files *member*... into *archive* (with *replacement*). This operation differs from 'q' in that any previously existing members are deleted if their names match those being added.

If one of the files named in *member*... does not exist, ar displays an error message, and leaves undisturbed any existing members of the archive matching that name.

By default, new members are added at the end of the file; but you may use one of the modifiers 'a', 'b', or 'i' to request placement relative to some existing member.

The modifier 'v' used with this operation elicits a line of output for each file inserted, along with one of the letters 'a' or 'r' to indicate whether the file was appended (no old member deleted) or replaced.

t — Display a *table* listing the contents of *archive*, or those of the files listed in *member*... that are present in the archive. Normally only the member name is shown; if you also want to see the modes (permissions), timestamp, owner, group, and size, you can request that by also specifying the 'v' modifier.

If you do not specify a *member*, all files in the archive are listed.

If there is more than one file with the same name (say, 'fie') in an archive (say 'b.a'), 'ar t b.a fie' lists only the first instance; to see them all, you must ask for a complete listing—in our example, 'ar t b.a'.

x — *Extract* members (named *member*) from the archive. You can use the 'v' modifier with this operation, to request that ar list each name as it extracts it.

If you do not specify a *member*, all files in the archive are extracted.

A number of modifiers (*mod*) may immediately follow the *p* keyletter, to specify variations on an operation's behavior:

a — Add new files *after* an existing member of the archive. If you use the modifier 'a', the name of an existing archive member must be present as the *relpos* argument, before the *archive* specification.

b — Add new files *before* an existing member of the archive. If you use the modifier 'b', the name of an existing archive member must be present as the *relpos* argument, before the *archive* specification. (same as 'i').

`c` *Create* the archive. The specified *archive* is always created if it did not exist, when you request an update. But a warning is issued unless you specify in advance that you expect to create it, by using this modifier.

`f` Truncate names in the archive. GNU `ar` will normally permit file names of any length. This will cause it to create archives which are not compatible with the native `ar` program on some systems. If this is a concern, the 'f' modifier may be used to truncate file names when putting them in the archive.

`i` Insert new files *before* an existing member of the archive. If you use the modifier 'i', the name of an existing archive member must be present as the *relpos* argument, before the *archive* specification. (same as 'b').

`l` This modifier is accepted but not used.

`o` Preserve the *original* dates of members when extracting them. If you do not specify this modifier, files extracted from the archive are stamped with the time of extraction.

`s` Write an object-file index into the archive, or update an existing one, even if no other change is made to the archive. You may use this modifier flag either with any operation, or alone. Running 'ar s' on an archive is equivalent to running 'ranlib' on it.

`u` Normally, 'ar r'... inserts all files listed into the archive. If you would like to insert *only* those of the files you list that are newer than existing members of the same names, use this modifier. The 'u' modifier is allowed only for the operation 'r' (replace). In particular, the combination 'qu' is not allowed, since checking the timestamps would lose any speed advantage from the operation 'q'.

`v` This modifier requests the *verbose* version of an operation. Many operations display additional information, such as filenames processed, when the modifier 'v' is appended.

`V` This modifier shows the version number of `ar`.

## 1.2 Controlling `ar` with a script

```
ar -M [ <script ]
```

If you use the single command-line option '-M' with `ar`, you can control its operation with a rudimentary command language. This form of `ar` operates interactively if standard input is coming directly from a terminal. During interactive use, `ar` prompts for input (the prompt is 'AR

>'), and continues executing even after errors. If you redirect standard input to a script file, no prompts are issued, and `ar` abandons execution (with a nonzero exit code) on any error.

The `ar` command language is *not* designed to be equivalent to the command-line options; in fact, it provides somewhat less control over archives. The only purpose of the command language is to ease the transition to `GNU ar` for developers who already have scripts written for the MRI "librarian" program.

The syntax for the `ar` command language is straightforward:

- commands are recognized in upper or lower case; for example, `LIST` is the same as `list`. In the following descriptions, commands are shown in upper case for clarity.
- a single command may appear on each line; it is the first word on the line.
- empty lines are allowed, and have no effect.
- comments are allowed; text after either of the characters '`*`' or '`;`' is ignored.
- Whenever you use a list of names as part of the argument to an `ar` command, you can separate the individual names with either commas or blanks. Commas are shown in the explanations below, for clarity.
- '`+`' is used as a line continuation character; if '`+`' appears at the end of a line, the text on the following line is considered part of the current command.

Here are the commands you can use in `ar` scripts, or when using `ar` interactively. Three of them have special significance:

`OPEN` or `CREATE` specify a *current archive*, which is a temporary file required for most of the other commands.

`SAVE` commits the changes so far specified by the script. Prior to `SAVE`, commands affect only the temporary copy of the current archive.

`ADDLIB` *archive*
`ADDLIB` *archive* (*module*, *module*, ... *module*)
: Add all the contents of *archive* (or, if specified, each named *module* from *archive*) to the current archive.

  Requires prior use of `OPEN` or `CREATE`.

`ADDMOD` *member*, *member*, ... *member*
: Add each named *member* as a module in the current archive.

  Requires prior use of `OPEN` or `CREATE`.

CLEAR Discard the contents of the current archive, cancelling the effect of any operations since the last SAVE. May be executed (with no effect) even if no current archive is specified.

CREATE *archive*
Creates an archive, and makes it the current archive (required for many other commands). The new archive is created with a temporary name; it is not actually saved as *archive* until you use SAVE. You can overwrite existing archives; similarly, the contents of any existing file named *archive* will not be destroyed until SAVE.

DELETE *module*, *module*, ... *module*
Delete each listed *module* from the current archive; equivalent to 'ar -d *archive module* ... *module*'.

Requires prior use of OPEN or CREATE.

DIRECTORY *archive* (*module*, ... *module*)
DIRECTORY *archive* (*module*, ... *module*) *outputfile*
List each named *module* present in *archive*. The separate command VERBOSE specifies the form of the output: when verbose output is off, output is like that of 'ar -t *archive module*...'. When verbose output is on, the listing is like 'ar -tv *archive module*...'.

Output normally goes to the standard output stream; however, if you specify *outputfile* as a final argument, ar directs the output to that file.

END Exit from ar, with a 0 exit code to indicate successful completion. This command does not save the output file; if you have changed the current archive since the last SAVE command, those changes are lost.

EXTRACT *module*, *module*, ... *module*
Extract each named *module* from the current archive, writing them into the current directory as separate files. Equivalent to 'ar -x *archive module*...'.

Requires prior use of OPEN or CREATE.

LIST Display full contents of the current archive, in "verbose" style regardless of the state of VERBOSE. The effect is like 'ar tv *archive*'). (This single command is a GNU ld enhancement, rather than present for MRI compatibility.)

Requires prior use of OPEN or CREATE.

OPEN *archive*
Opens an existing archive for use as the current archive (required for many other commands). Any changes as the result

of subsequent commands will not actually affect *archive* until you next use SAVE.

REPLACE *module*, *module*, ... *module*
: In the current archive, replace each existing *module* (named in the REPLACE arguments) from files in the current working directory. To execute this command without errors, both the file, and the module in the current archive, must exist.

  Requires prior use of OPEN or CREATE.

VERBOSE
: Toggle an internal flag governing the output from DIRECTORY. When the flag is on, DIRECTORY output matches output from 'ar -tv'...

SAVE
: Commit your changes to the current archive, and actually save it as a file with the name specified in the last CREATE or OPEN command.

  Requires prior use of OPEN or CREATE.

# 2 ld

The GNU linker `ld` is now described in a separate manual. See section "Overview" in *Using LD: the GNU linker*.

# 3 nm

```
nm [ -a | --debug-syms ]  [ -g | --extern-only ]
   [ -B ]  [ -C | --demangle ] [ -D | --dynamic ]
   [ -s | --print-armap ]  [ -A | -o | --print-file-name ]
   [ -n | -v | --numeric-sort ]  [ -p | --no-sort ]
   [ -r | --reverse-sort ]  [ --size-sort ] [ -u | --undefined-
only ]
   [ -t radix | --radix=radix ] [ -P | --portability ]
   [ --target=bfdname ] [ -f format | --format=format ]
   [ --no-demangle ] [ -V | --version ]  [ --help ]
                              [ objfile... ]
```

GNU nm lists the symbols from object files *objfile*.... If no object files are listed as arguments, nm assumes 'a.out'.

For each symbol, nm shows:

- The symbol value, in the radix selected by options (see below), or hexadecimal by default.
- The symbol type. At least the following types are used; others are, as well, depending on the object file format. If lowercase, the symbol is local; if uppercase, the symbol is global (external).

  A — Absolute.

  B — BSS (uninitialized data).

  C — Common.

  D — Initialized data.

  I — Indirect reference.

  T — Text (program code).

  U — Undefined.
- The symbol name.

The long and short forms of options, shown here as alternatives, are equivalent.

`-A`
`-o`
`--print-file-name`
: Precede each symbol by the name of the input file (or archive element) in which it was found, rather than identifying the input file once only, before all of its symbols.

`-a`
`--debug-syms`
: Display all symbols, even debugger-only symbols; normally these are not listed.

`-B` The same as '`--format=bsd`' (for compatibility with the MIPS `nm`).

`-C`
`--demangle`
Decode (*demangle*) low-level symbol names into user-level names. Besides removing any initial underscore prepended by the system, this makes C++ function names readable. See Chapter 10 "c++filt," page 33, for more information on demangling.

`--no-demangle`
Do not demangle low-level symbol names. This is the default.

`-D`
`--dynamic`
Display the dynamic symbols rather than the normal symbols. This is only meaningful for dynamic objects, such as certain types of shared libraries.

`-f` *format*
`--format=`*format*
Use the output format *format*, which can be `bsd`, `sysv`, or `posix`. The default is `bsd`. Only the first character of *format* is significant; it can be either upper or lower case.

`-g`
`--extern-only`
Display only external symbols.

`-n`
`-v`
`--numeric-sort`
Sort symbols numerically by their addresses, rather than alphabetically by their names.

`-p`
`--no-sort`
Do not bother to sort the symbols in any order; print them in the order encountered.

`-P`
`--portability`
Use the POSIX.2 standard output format instead of the default format. Equivalent to '`-f posix`'.

`-s`
`--print-armap`
When listing symbols from archive members, include the index: a mapping (stored in the archive by `ar` or `ranlib`) of which modules contain definitions for which names.

`-r`
`--reverse-sort`
: Reverse the order of the sort (whether numeric or alphabetic); let the last come first.

`--size-sort`
: Sort symbols by size. The size is computed as the difference between the value of the symbol and the value of the symbol with the next higher value. The size of the symbol is printed, rather than the value.

`-t` *`radix`*
`--radix=`*`radix`*
: Use *`radix`* as the radix for printing the symbol values. It must be 'd' for decimal, 'o' for octal, or 'x' for hexadecimal.

`--target=`*`bfdname`*
: Specify an object code format other than your system's default format. See Section 12.1 "Target Selection," page 37, for more information.

`-u`
`--undefined-only`
: Display only undefined symbols (those external to each object file).

`-V`
`--version`
: Show the version number of `nm` and exit.

`--help`
: Show a summary of the options to `nm` and exit.

# 4 objcopy

```
objcopy [ -F bfdname | --target=bfdname ]
        [ -I bfdname | --input-target=bfdname ]
        [ -O bfdname | --output-target=bfdname ]
        [ -S | --strip-all ]  [ -g | --strip-debug ]
        [ -K symbolname | --keep-symbol=symbolname ]
        [ -N symbolname | --strip-symbol=symbolname ]
        [ -x | --discard-all ]  [ -X | --discard-locals ]
        [ -b byte | --byte=byte ]
        [ -i interleave | --interleave=interleave ]
        [ -R sectionname | --remove-section=sectionname ]
        [ --gap-fill=val ] [ --pad-to=address ]
        [ --set-start=val ] [ --adjust-start=incr ]
        [ --adjust-vma=incr ]
        [ --adjust-section-vma=section{=,+,-}val ]
        [ --adjust-warnings ] [ --no-adjust-warnings ]
        [ --set-section-flags=section=flags ]
        [ --add-section=sectionname=filename ]
        [ -v | --verbose ] [ -V | --version ]  [ --help ]
        infile [outfile]
```

The GNU `objcopy` utility copies the contents of an object file to another. `objcopy` uses the GNU BFD Library to read and write the object files. It can write the destination object file in a format different from that of the source object file. The exact behavior of `objcopy` is controlled by command-line options.

`objcopy` creates temporary files to do its translations and deletes them afterward. `objcopy` uses BFD to do all its translation work; it has access to all the formats described in BFD and thus is able to recognize most formats without being told explicitly. See section "BFD" in *Using LD*.

`objcopy` can be used to generate S-records by using an output target of '`srec`' (e.g., use '`-O srec`').

`objcopy` can be used to generate a raw binary file by using an output target of '`binary`' (e.g., use '`-O binary`'). When `objcopy` generates a raw binary file, it will essentially produce a memory dump of the contents of the input object file. All symbols and relocation information will be discarded. The memory dump will start at the virtual address of the lowest section copied into the output file.

When generating an S-record or a raw binary file, it may be helpful to use '`-S`' to remove sections containing debugging information. In some cases '`-R`' will be useful to remove sections which contain information which is not needed by the binary file.

*infile*
*outfile* The source and output files, respectively. If you do not specify *outfile*, objcopy creates a temporary file and destructively renames the result with the name of *infile*.

-I *bfdname*
--input-target=*bfdname*
Consider the source file's object format to be *bfdname*, rather than attempting to deduce it. See Section 12.1 "Target Selection," page 37, for more information.

-O *bfdname*
--output-target=*bfdname*
Write the output file using the object format *bfdname*. See Section 12.1 "Target Selection," page 37, for more information.

-F *bfdname*
--target=*bfdname*
Use *bfdname* as the object format for both the input and the output file; i.e., simply transfer data from source to destination with no translation. See Section 12.1 "Target Selection," page 37, for more information.

-R *sectionname*
--remove-section=*sectionname*
Remove any section named *sectionname* from the output file. This option may be given more than once. Using this option inappropriately may make the output file unusable.

-S
--strip-all
Do not copy relocation and symbol information from the source file.

-g
--strip-debug
Do not copy debugging symbols from the source file.

--strip-unneeded
Strip all symbols that are not needed for relocation processing.

-K *symbolname*
--keep-symbol=*symbolname*
Copy only symbol *symbolname* from the source file. This option may be given more than once.

`-N` *symbolname*
`--strip-symbol=`*symbolname*
: Do not copy symbol *symbolname* from the source file. This option may be given more than once, and may be combined with strip options other than `-K`.

`-x`
`--discard-all`
: Do not copy non-global symbols from the source file.

`-X`
`--discard-locals`
: Do not copy compiler-generated local symbols. (These usually start with 'L' or '.'.)

`-b` *byte*
`--byte=`*byte*
: Keep only every *byte*th byte of the input file (header data is not affected). *byte* can be in the range from 0 to *interleave*-1, where *interleave* is given by the '`-i`' or '`--interleave`' option, or the default of 4. This option is useful for creating files to program ROM. It is typically used with an `srec` output target.

`-i` *interleave*
`--interleave=`*interleave*
: Only copy one out of every *interleave* bytes. Select which byte to copy with the *-b* or '`--byte`' option. The default is 4. `objcopy` ignores this option if you do not specify either '`-b`' or '`--byte`'.

`--gap-fill` *val*
: Fill gaps between sections with *val*. This is done by increasing the size of the section with the lower address, and filling in the extra space created with *val*.

`--pad-to` *address*
: Pad the output file up to the virtual address *address*. This is done by increasing the size of the last section. The extra space is filled in with the value specified by '`--gap-fill`' (default zero).

`--set-start` *val*
: Set the address of the new file to *val*. Not all object file formats support setting the start address.

`--adjust-start` *incr*
: Adjust the start address by adding *incr*. Not all object file formats support setting the start address.

`--adjust-vma` *incr*
: Adjust the address of all sections, as well as the start address, by adding *incr*. Some object file formats do not permit section addresses to be changed arbitrarily. Note that this does not relocate the sections; if the program expects sections to be loaded at a certain address, and this option is used to change the sections such that they are loaded at a different address, the program may fail.

`--adjust-section-vma` *section*`{=,+,-}`*val*
: Set or adjust the address of the named *section*. If '`=`' is used, the section address is set to *val*. Otherwise, *val* is added to or subtracted from the section address. See the comments under '`--adjust-vma`', above. If *section* does not exist in the input file, a warning will be issued, unless '`--no-adjust-warnings`' is used.

`--adjust-warnings`
: If '`--adjust-section-vma`' is used, and the named section does not exist, issue a warning. This is the default.

`--no-adjust-warnings`
: Do not issue a warning if '`--adjust-section-vma`' is used, even if the named section does not exist.

`--set-section-flags` *section*`=`*flags*
: Set the flags for the named section. The *flags* argument is a comma separated string of flag names. The recognized names are '`alloc`', '`load`', '`readonly`', '`code`', '`data`', and '`rom`'. Not all flags are meaningful for all object file formats.

`--add-section` *sectionname*`=`*filename*
: Add a new section named *sectionname* while copying the file. The contents of the new section are taken from the file *filename*. The size of the section will be the size of the file. This option only works on file formats which can support sections with arbitrary names.

`-V`
`--version`
: Show the version number of `objcopy`.

`-v`
`--verbose`
: Verbose output: list all object files modified. In the case of archives, '`objcopy -V`' lists all members of the archive.

`--help` Show a summary of the options to `objcopy`.

# 5 objdump

```
objdump [ -a | --archive-headers ]
        [ -b bfdname | --target=bfdname ]
        [ -d | --disassemble ]  [ -D | --disassemble-all ]
        [ -f | --file-headers ]
        [ -h | --section-headers | --headers ]  [ -i | --info ]
        [ -j section | --section=section ]
        [ -l | --line-numbers ] [ -S | --source ]
        [ -m machine | --architecture=machine ]
        [ -r | --reloc ] [ -R | --dynamic-reloc ]
        [ -s | --full-contents ]  [ --stabs ]
        [ -t | --syms ] [ -T | --dynamic-syms ]
        [ -x | --all-headers ] [ -w | --wide ]
        [ --version ]  [ --help ] objfile...
```

`objdump` displays information about one or more object files. The options control what particular information to display. This information is mostly useful to programmers who are working on the compilation tools, as opposed to programmers who just want their program to compile and work.

*`objfile`*... are the object files to be examined. When you specify archives, `objdump` shows information on each of the member object files.

The long and short forms of options, shown here as alternatives, are equivalent. At least one option besides '`-l`' must be given.

`-a`
`--archive-header`
: If any of the *`objfile`* files are archives, display the archive header information (in a format similar to '`ls -l`'). Besides the information you could list with '`ar tv`', '`objdump -a`' shows the object file format of each archive member.

`-b` *`bfdname`*
`--target=`*`bfdname`*
: Specify that the object-code format for the object files is *`bfdname`*. This option may not be necessary; *`objdump`* can automatically recognize many formats.

  For example,

  ```
  objdump -b oasys -m vax -h fu.o
  ```

  displays summary information from the section headers ('`-h`') of '`fu.o`', which is explicitly identified ('`-m`') as a VAX object file in the format produced by Oasys compilers. You can list the formats available with the '`-i`' option. See Section 12.1 "Target Selection," page 37, for more information.

`-d`
`--disassemble`
: Display the assembler mnemonics for the machine instructions from *objfile*. This option only disassembles those sections which are expected to contain instructions.

`-D`
`--disassemble-all`
: Like '`-d`', but disassemble the contents of all sections, not just those expected to contain instructions.

`-f`
`--file-header`
: Display summary information from the overall header of each of the *objfile* files.

`-h`
`--section-header`
`--header`
: Display summary information from the section headers of the object file.

  File segments may be relocated to nonstandard addresses, for example by using the '`-Ttext`', '`-Tdata`', or '`-Tbss`' options to `ld`. However, some object file formats, such as a.out, do not store the starting address of the file segments. In those situations, although `ld` relocates the sections correctly, using '`objdump -h`' to list the file section headers cannot show the correct addresses. Instead, it shows the usual addresses, which are implicit for the target.

`--help`
: Print a summary of the options to `objdump` and exit.

`-i`
`--info`
: Display a list showing all architectures and object formats available for specification with '`-b`' or '`-m`'.

`-j` *name*
`--section=`*name*
: Display information only for section *name*.

`-l`
`--line-numbers`
: Label the display (using debugging information) with the filename and source line numbers corresponding to the object code shown. Only useful with '`-d`' or '`-D`'.

`-m` *machine*
`--architecture=`*machine*
: Specify that the object files *objfile* are for architecture *machine*. List available architectures using the '`-i`' option.

`-r`
`--reloc` Print the relocation entries of the file. If used with '`-d`' or '`-D`', the relocations are printed interspersed with the disassembly.

`-R`
`--dynamic-reloc`
Print the dynamic relocation entries of the file. This is only meaningful for dynamic objects, such as certain types of shared libraries.

`-s`
`--full-contents`
Display the full contents of any sections requested.

`-S`
`--source` Display source code intermixed with disassembly, if possible. Implies '`-d`'.

`--stabs` Display the full contents of any sections requested. Display the contents of the .stab and .stab.index and .stab.excl sections from an ELF file. This is only useful on systems (such as Solaris 2.0) in which `.stab` debugging symbol-table entries are carried in an ELF section. In most other file formats, debugging symbol-table entries are interleaved with linkage symbols, and are visible in the '`--syms`' output.

`-t`
`--syms` Print the symbol table entries of the file. This is similar to the information provided by the '`nm`' program.

`-T`
`--dynamic-syms`
Print the dynamic symbol table entries of the file. This is only meaningful for dynamic objects, such as certain types of shared libraries. This is similar to the information provided by the '`nm`' program when given the '`-D`' ('`--dynamic`') option.

`--version`
Print the version number of `objdump` and exit.

`-x`
`--all-header`
Display all available header information, including the symbol table and relocation entries. Using '`-x`' is equivalent to specifying all of '`-a -f -h -r -t`'.

`-w`
`--wide` Format some lines for output devices that have more than 80 columns.

# 6 ranlib

```
ranlib [-vV] archive
```

`ranlib` generates an index to the contents of an archive and stores it in the archive. The index lists each symbol defined by a member of an archive that is a relocatable object file.

You may use '`nm -s`' or '`nm --print-armap`' to list this index.

An archive with such an index speeds up linking to the library and allows routines in the library to call each other without regard to their placement in the archive.

The `GNU ranlib` program is another form of `GNU ar`; running `ranlib` is completely equivalent to executing '`ar -s`'. See Chapter 1 "ar," page 3.

`-v`
`-V` Show the version number of `ranlib`.

# 7 size

```
size [ -A | -B | --format=compatibility ]
     [ --help ]  [ -d | -o | -x | --radix=number ]
     [ --target=bfdname ]  [ -V | --version ]
     objfile...
```

The GNU `size` utility lists the section sizes—and the total size—for each of the object or archive files *objfile* in its argument list. By default, one line of output is generated for each object file or each module in an archive.

*objfile*... are the object files to be examined.

The command line options have the following meanings:

`-A`
`-B`
`--format=`*compatibility*
: Using one of these options, you can choose whether the output from GNU `size` resembles output from System V `size` (using '`-A`', or '`--format=sysv`'), or Berkeley `size` (using '`-B`', or '`--format=berkeley`'). The default is the one-line format similar to Berkeley's.

  Here is an example of the Berkeley (default) format of output from `size`:

```
size --format=Berkeley ranlib size
text    data    bss     dec     hex     filename
294880  81920   11592   388392  5ed28   ranlib
294880  81920   11888   388688  5ee50   size
```

  This is the same data, but displayed closer to System V conventions:

```
size --format=SysV ranlib size
ranlib  :
section         size         addr
.text         294880         8192
.data          81920       303104
.bss           11592       385024
Total         388392


size  :
section         size         addr
.text         294880         8192
.data          81920       303104
.bss           11888       385024
Total         388688
```

`--help`
: Show a summary of acceptable arguments and options.

`-d`
`-o`
`-x`
`--radix=`*`number`*

Using one of these options, you can control whether the size of each section is given in decimal ('`-d`', or '`--radix=10`'); octal ('`-o`', or '`--radix=8`'); or hexadecimal ('`-x`', or '`--radix=16`'). In '`--radix=`*`number`*', only the three values (8, 10, 16) are supported. The total size is always given in two radices; decimal and hexadecimal for '`-d`' or '`-x`' output, or octal and hexadecimal if you're using '`-o`'.

`--target=`*`bfdname`*

Specify that the object-code format for *`objfile`* is *`bfdname`*. This option may not be necessary; `size` can automatically recognize many formats. See Section 12.1 "Target Selection," page 37, for more information.

`-V`
`--version`

Display the version number of `size`.

# 8 strings

```
strings [-afov] [-min-len] [-n min-len] [-t radix] [-]
        [--all] [--print-file-name] [--bytes=min-len]
        [--radix=radix] [--target=bfdname]
        [--help] [--version] file...
```

For each *file* given, GNU `strings` prints the printable character sequences that are at least 4 characters long (or the number given with the options below) and are followed by an unprintable character. By default, it only prints the strings from the initialized and loaded sections of object files; for other types of files, it prints the strings from the whole file.

`strings` is mainly useful for determining the contents of non-text files.

`-a`
`--all`
`-` Do not scan only the initialized and loaded sections of object files; scan the whole files.

`-f`
`--print-file-name`
Print the name of the file before each string.

`--help` Print a summary of the program usage on the standard output and exit.

`-n` *min-len*
*-min-len*
`--bytes=`*min-len*
Print sequences of characters that are at least *min-len* characters long, instead of the default 4.

`-o` Like '`-t o`'. Some other versions of `strings` have '`-o`' act like '`-t d`' instead. Since we can not be compatible with both ways, we simply chose one.

`-t` *radix*
`--radix=`*radix*
Print the offset within the file before each string. The single character argument specifies the radix of the offset—'`o`' for octal, '`x`' for hexadecimal, or '`d`' for decimal.

`--target=`*bfdname*
Specify an object code format other than your system's default format. See Section 12.1 "Target Selection," page 37, for more information.

`-v`
`--version`
: Print the program version number on the standard output and exit.

# 9 strip

```
strip [ -F bfdname | --target=bfdname | --target=bfdname ]
      [ -I bfdname | --input-target=bfdname ]
      [ -O bfdname | --output-target=bfdname ]
      [ -s | --strip-all ] [ -S | -g | --strip-debug ]
      [ -K symbolname | --keep-symbol=symbolname ]
      [ -N symbolname | --strip-symbol=symbolname ]
      [ -x | --discard-all ] [ -X | --discard-locals ]
      [ -R sectionname | --remove-section=sectionname ]
      [ -v | --verbose ]  [ -V | --version ]  [ --help ]
      objfile...
```

GNU `strip` discards all symbols from object files *objfile*. The list of object files may include archives. At least one object file must be given.

`strip` modifies the files named in its argument, rather than writing modified copies under different names.

`-F` *bfdname*
`--target=`*bfdname*
: Treat the original *objfile* as a file with the object code format *bfdname*, and rewrite it in the same format. See Section 12.1 "Target Selection," page 37, for more information.

`--help`
: Show a summary of the options to `strip` and exit.

`-I` *bfdname*
`--input-target=`*bfdname*
: Treat the original *objfile* as a file with the object code format *bfdname*. See Section 12.1 "Target Selection," page 37, for more information.

`-O` *bfdname*
`--output-target=`*bfdname*
: Replace *objfile* with a file in the output format *bfdname*. See Section 12.1 "Target Selection," page 37, for more information.

`-R` *sectionname*
`--remove-section=`*sectionname*
: Remove any section named *sectionname* from the output file. This option may be given more than once. Note that using this option inappropriately may make the output file unusable.

`-s`
`--strip-all`
: Remove all symbols.

`-g`
`-S`
`--strip-debug`
: Remove debugging symbols only.

`--strip-unneeded`
: Remove all symbols that are not needed for relocation processing.

`-K` *`symbolname`*
`--keep-symbol=`*`symbolname`*
: Keep only symbol *`symbolname`* from the source file. This option may be given more than once.

`-N` *`symbolname`*
`--strip-symbol=`*`symbolname`*
: Remove symbol *`symbolname`* from the source file. This option may be given more than once, and may be combined with strip options other than `-K`.

`-x`
`--discard-all`
: Remove non-global symbols.

`-X`
`--discard-locals`
: Remove compiler-generated local symbols. (These usually start with 'L' or '.'.)

`-V`
`--version`
: Show the version number for `strip`.

`-v`
`--verbose`
: Verbose output: list all object files modified. In the case of archives, '`strip -v`' lists all members of the archive.

# 10 c++filt

```
c++filt [ -_ | --strip-underscores ]
[ -n | --no-strip-underscores ]
        [ -s format | --format=format ]
        [ --help ]  [ --version ]  [ symbol... ]
```

The C++ language provides function overloading, which means that you can write many functions with the same name (providing each takes parameters of different types). All C++ function names are encoded into a low-level assembly label (this process is known as *mangling*). The `c++filt` program does the inverse mapping: it decodes (*demangles*) low-level names into user-level names so that the linker can keep these overloaded functions from clashing.

Every alphanumeric word (consisting of letters, digits, underscores, dollars, or periods) seen in the input is a potential label. If the label decodes into a C++ name, the C++ name replaces the low-level name in the output.

You can use `c++filt` to decipher individual symbols:

```
c++filt symbol
```

If no *symbol* arguments are given, `c++filt` reads symbol names from the standard input and writes the demangled names to the standard output. All results are printed on the standard output.

`-_`
`--strip-underscores`
: On some systems, both the C and C++ compilers put an underscore in front of every name. For example, the C name `foo` gets the low-level name `_foo`. This option removes the initial underscore. Whether `c++filt` removes the underscore by default is target dependent.

`-n`
`--no-strip-underscores`
: Do not remove the initial underscore.

`-s` *format*
`--format=`*format*
: GNU `nm` can decode three different methods of mangling, used by different C++ compilers. The argument to this option selects which method it uses:

  `gnu` — the one used by the GNU compiler (the default method)

  `lucid` — the one used by the Lucid compiler

  `arm` — the one specified by the C++ Annotated Reference Manual

`--help` Print a summary of the options to `c++filt` and exit.

`--version`
Print the version number of `c++filt` and exit.

*Warning:* `c++filt` is a new utility, and the details of its user interface are subject to change in future releases. In particular, a command-line option may be required in the the future to decode a name passed as an argument on the command line; in other words,

```
c++filt symbol
```

may in a future release become

```
c++filt option symbol
```

# 11 nlmconv

`nlmconv` converts a relocatable object file into a NetWare Loadable Module.

*Warning:* `nlmconv` is not always built as part of the binary utilities, since it is only useful for NLM targets.

```
nlmconv [ -I bfdname | --input-target=bfdname ]
        [ -O bfdname | --output-target=bfdname ]
        [ -T headerfile | --header-file=headerfile ]
        [ -d | --debug]  [ -l linker | --linker=linker ]
        [ -h | --help ]  [ -V | --version ]
        infile outfile
```

`nlmconv` converts the relocatable 'i386' object file *infile* into the NetWare Loadable Module *outfile*, optionally reading *headerfile* for NLM header information. For instructions on writing the NLM command file language used in header files, see the 'linkers' section, 'NLMLINK' in particular, of the *NLM Development and Tools Overview*, which is part of the NLM Software Developer's Kit ("NLM SDK"), available from Novell, Inc. `nlmconv` uses the GNU Binary File Descriptor library to read *infile*; see section "BFD" in *Using LD*, for more information.

`nlmconv` can perform a link step. In other words, you can list more than one object file for input if you list them in the definitions file (rather than simply specifying one input file on the command line). In this case, `nlmconv` calls the linker for you.

`-I` *bfdname*
`--input-target=`*bfdname*
: Object format of the input file. `nlmconv` can usually determine the format of a given file (so no default is necessary). See Section 12.1 "Target Selection," page 37, for more information.

`-O` *bfdname*
`--output-target=`*bfdname*
: Object format of the output file. `nlmconv` infers the output format based on the input format, e.g. for a 'i386' input file the output format is 'nlm32-i386'. See Section 12.1 "Target Selection," page 37, for more information.

`-T` *headerfile*
`--header-file=`*headerfile*
: Reads *headerfile* for NLM header information. For instructions on writing the NLM command file language used in header files, see see the 'linkers' section, of the *NLM Development and Tools Overview*, which is part of the NLM Software Developer's Kit, available from Novell, Inc.

`-d`
`--debug` Displays (on standard error) the linker command line used by `nlmconv`.

`-l` *`linker`*
`--linker=`*`linker`*
Use *`linker`* for any linking. *`linker`* can be an abosolute or a relative pathname.

`-h`
`--help` Prints a usage summary.

`-V`
`--version`
Prints the version number for `nlmconv`.

# 12 Selecting the target system

You can specify three aspects of the target system to the GNU binary file utilities, each in several ways:

- the target
- the architecture
- the linker emulation (which applies to the linker only)

In the following summaries, the lists of ways to specify values are in order of decreasing precedence. The ways listed first override those listed later.

The commands to list valid values only list the values for which the programs you are running were configured. If they were configured with '`--with-targets=all`', the commands list most of the available values, but a few are left out; not all targets can be configured in at once because some of them can only be configured *native* (on hosts with the same type as the target system).

## 12.1 Target Selection

A *target* is an object file format. A given target may be supported for multiple architectures (see Section 12.2 "Architecture Selection," page 39). A target selection may also have variations for different operating systems or architectures.

The command to list valid target values is '`objdump -i`' (the first column of output contains the relevant information).

Some sample values are: '`a.out-hp300bsd`', '`ecoff-littlemips`', '`a.out-sunos-big`'.

### `objdump` Target

Ways to specify:

1. command line option: '`-b`' or '`--target`'
2. environment variable `GNUTARGET`
3. deduced from the input file

### `objcopy` and `strip` Input Target

Ways to specify:

1. command line options: '`-I`' or '`--input-target`', or '`-F`' or '`--target`'
2. environment variable `GNUTARGET`
3. deduced from the input file

### `objcopy` and `strip` Output Target

Ways to specify:

1. command line options: '`-O`' or '`--output-target`', or '`-F`' or '`--target`'
2. the input target (see "`objcopy` and `strip` Input Target" above)
3. environment variable `GNUTARGET`
4. deduced from the input file

### `nm`, `size`, and `strings` Target

Ways to specify:

1. command line option: '`--target`'
2. environment variable `GNUTARGET`
3. deduced from the input file

### Linker Input Target

Ways to specify:

1. command line option: '`-b`' or '`--format`' (see section "Options" in *Using LD*)
2. script command `TARGET` (see section "Option Commands" in *Using LD*)
3. environment variable `GNUTARGET` (see section "Environment" in *Using LD*)
4. the default target of the selected linker emulation (see Section 12.3 "Linker Emulation Selection," page 39)

### Linker Output Target

Ways to specify:

1. command line option: '`-oformat`' (see section "Options" in *Using LD*)
2. script command `OUTPUT_FORMAT` (see section "Option Commands" in *Using LD*)
3. the linker input target (see "Linker Input Target" above)

## 12.2 Architecture selection

An *architecture* is a type of CPU on which an object file is to run. Its name may contain a colon, separating the name of the processor family from the name of the particular CPU.

The command to list valid architecture values is 'objdump -i' (the second column contains the relevant information).

Sample values: 'm68k:68020', 'mips:3000', 'sparc'.

### objdump Architecture

Ways to specify:

1. command line option: '-m' or '--architecture'
2. deduced from the input file

### objcopy, nm, size, strings Architecture

Ways to specify:

1. deduced from the input file

### Linker Input Architecture

Ways to specify:

1. deduced from the input file

### Linker Output Architecture

Ways to specify:

1. script command OUTPUT_ARCH (see section "Option Commands" in *Using LD*)
2. the default architecture from the linker output target (see Section 12.1 "Target Selection," page 37)

## 12.3 Linker emulation selection

A linker *emulation* is a "personality" of the linker, which gives the linker default values for the other aspects of the target system. In particular, it consists of

- the linker script
- the target

- several "hook" functions that are run at certain stages of the linking process to do special things that some targets require

The command to list valid linker emulation values is '`ld -V`'.

Sample values: '`hp300bsd`', '`mipslit`', '`sun4`'.

Ways to specify:

1. command line option: '`-m`' (see section "Options" in *Using LD*)
2. environment variable `LDEMULATION`
3. compiled-in `DEFAULT_EMULATION` from '`Makefile`', which comes from `EMUL` in '`config/`*`target`*`.mt`'

# Index


## .

.stab ..... 23

## A

all header information, object file ..... 23
ar ..... 3
`ar` compatibility ..... 3
architecture ..... 22
architectures available ..... 22
archive contents ..... 25
archive headers ..... 21
archives ..... 3

## C

c++filt ..... 33
collections of files ..... 3
compatibility, `ar` ..... 3
contents of archive ..... 5
creating archives ..... 5

## D

dates in archive ..... 6
debug symbols ..... 23
debugging symbols ..... 13
deleting from archive ..... 4
demangling C++ symbols ..... 14, 33
disassembling object code ..... 22
disassembly, with source ..... 23
discarding symbols ..... 31
dynamic relocation entries, in object file ..... 23
dynamic symbol table entries, printing ..... 23
dynamic symbols ..... 14

## E

ELF object file format ..... 23
external symbols ..... 14, 15
extract from archive ..... 5

## F

file name ..... 13

## H

header information, all ..... 23

## I

input file name ..... 13

## L

ld ..... 11
libraries ..... 3
linker ..... 11
listings strings ..... 29

## M

machine instructions ..... 22
moving in archive ..... 4
MRI compatibility, `ar` ..... 6

## N

name duplication in archive ..... 5
name length ..... 3
nm ..... 13
`nm` compatibility ..... 13, 14
`nm` format ..... 13, 14

## O

objdump ..... 21
object code format ..... 15, 21, 28, 29
object file header ..... 22
object file information ..... 21
object file sections ..... 23
object formats available ..... 22
operations on archive ..... 4

## P

printing from archive ..... 4
printing strings ..... 29

## Q

quick append to archive ..... 4

## R

radix for section sizes ..... 28
ranlib ..... 25
relative placement in archive ..... 5
relocation entries, in object file ..... 22
removing symbols ..... 31
repeated names in archive ..... 5
replacement in archive ..... 5

## S

scripts, `ar` ..... 6
section headers ..... 22
section information ..... 22
section sizes ..... 27
sections, full contents ..... 23
size ..... 27
`size` display format ..... 27
`size` number format ..... 28
sorting symbols ..... 14
source disassembly ..... 23
source file name ..... 13
source filenames for object files ..... 22
stab ..... 23
strings ..... 29
strings, printing ..... 29
strip ..... 31
symbol index ..... 3, 25
symbol index, listing ..... 14
symbol table entries, printing ..... 23
symbols ..... 13
symbols, discarding ..... 31

## U

undefined symbols ..... 15
Unix compatibility, `ar` ..... 4
updating an archive ..... 6

## V

version ..... 1

## W

wide output, printing ..... 23
writing archive index ..... 6