



Linux Ubuntu 26.04 Manual Pages on command 'printf.3'

\$ man printf.3

printf(3) Library Functions Manual printf(3)

NAME

printf, fprintf, dprintf, vprintf, vfprintf, vdprintf, - formatted output conversion

LIBRARY

Standard C library (libc, -lc)

SYNOPSIS

```
#include <stdio.h>
```

```
int printf(const char *restrict format, ...);
```

```
int fprintf(FILE *restrict stream,
            const char *restrict format, ...);
```

```
int dprintf(int fd,
            const char *restrict format, ...);
```

```
int vprintf(const char *restrict format, va_list ap);
```

```
int vfprintf(FILE *restrict stream,
            const char *restrict format, va_list ap);
```

```
int vdprintf(int fd,
            const char *restrict format, va_list ap);
```

Feature Test Macro Requirements for glibc (see feature_test_macros(7)):

dprintf(), vdprintf():

Since glibc 2.10:

```
_POSIX_C_SOURCE >= 200809L
```

Before glibc 2.10:

```
_GNU_SOURCE
```

DESCRIPTION

The functions in the `printf()` family produce output according to a format as described below. The functions `printf()` and `vprintf()` write output to `stdout`, the standard output stream; `fprintf()` and `vfprintf()` write output to the given output stream.

The function `dprintf()` is the same as `fprintf()` except that it outputs to a file descriptor, `fd`, instead of to a `stdio(3)` stream.

The functions `vprintf()`, `vfprintf()`, `vdprintf()` are equivalent to the functions `printf()`, `fprintf()`, `dprintf()`, respectively, except that they are called with a `va_list` instead of a variable number of arguments. These functions do not call the `va_end` macro. Because they invoke the `va_arg` macro, the value of `ap` is undefined after the call. See `stdarg(3)`.

All of these functions write the output under the control of a format string that specifies how subsequent arguments (or arguments accessed via the variable-length argument facilities of `stdarg(3)`) are converted for output.

Format of the format string

The format string is a character string, beginning and ending in its initial shift state, if any. The format string is composed of zero or more directives: ordinary characters (not `%`), which are copied unchanged to the output stream; and conversion specifications, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the character `%`, and ends with a conversion specifier. In between there may be (in this order) zero or more flags, an optional minimum field width, an optional precision and an optional length modifier.

The overall syntax of a conversion specification is:

```
%[argument$][flags][width][.precision][length modifier]conversion
```

The arguments must correspond properly (after type promotion) with the conversion specifier.

By default, the arguments are used in the order given, where each `*` (see Field width and Precision below) and each conversion specifier asks for the next argument (and it is an error if insufficiently many arguments are given). One can also specify explicitly which argument is taken, at each place where an argument is required, by writing `"%m$"` instead of `'%'` and `"*m$"` instead of `*`, where the decimal integer `m` denotes the position in the argument list of the desired argument, indexed starting from 1. Thus,

```
printf("%*d", width, num);
```

and

```
printf("%2$*1$d", width, num);
```

are equivalent. The second style allows repeated references to the same argument. The C99 standard does not include the style using '\$', which comes from the Single UNIX Specification. If the style using '\$' is used, it must be used throughout for all conversions taking an argument and all width and precision arguments, but it may be mixed with "%%" formats, which do not consume an argument. There may be no gaps in the numbers of arguments specified using '\$'; for example, if arguments 1 and 3 are specified, argument 2 must also be specified somewhere in the format string.

For some numeric conversions a radix character ("decimal point") or thousands' grouping character is used. The actual character used depends on the LC_NUMERIC part of the locale. (See `setlocale(3)`.) The POSIX locale uses '.' as radix character, and does not have a grouping character. Thus,

```
printf("%.2f", 1234567.89);
```

results in "1234567.89" in the POSIX locale, in "1234567,89" in the nl_NL locale, and in "1.234.567,89" in the da_DK locale.

Flag characters

The character % is followed by zero or more of the following flags:

- # The value should be converted to an "alternate form". For o conversions, the first character of the output string is made zero (by prefixing a 0 if it was not zero already). For x and X conversions, a nonzero result has the string "0x" (or "0X" for X conversions) prepended to it. For a, A, e, E, f, F, g, and G conversions, the result will always contain a decimal point, even if no digits follow it (normally, a decimal point appears in the results of those conversions only if a digit follows). For g and G conversions, trailing zeros are not removed from the result as they would otherwise be. For m, if `errno` contains a valid error code, the output of `strerror_r(name_np(errno))` is printed; otherwise, the value stored in `errno` is printed as a decimal number. For other conversions, the result is undefined.
- 0 The value should be zero padded. For d, i, o, u, x, X, a, A, e, E, f, F, g, and G conversions, the converted value is padded on the left with zeros rather than blanks. If the 0 and - flags both appear, the 0 flag is ignored. If a precision is given with an integer conversion (d, i, o, u, x, and X), the 0 flag is ignored. For other conversions, the behavior is undefined.
- The converted value is to be left adjusted on the field boundary. (The default is right justification.) The converted value is padded on the right with blanks, rather

than on the left with blanks or zeros. A - overrides a 0 if both are given.

- ' ' (a space) A blank should be left before a positive number (or empty string) produced by a signed conversion.
- + A sign (+ or -) should always be placed before a number produced by a signed conversion. By default, a sign is used only for negative numbers. A + overrides a space if both are used.

The five flag characters above are defined in the C99 standard. POSIX specifies one further flag character.

- ' For decimal conversion (i, d, u, f, F, g, G) the output is to be grouped with thousands' grouping characters as a non-monetary quantity. Misleadingly, this isn't necessarily every thousand: for example Karbi ("mjw_IN"), groups its digits into 3 once, then 2 repeatedly. Compare locale(7) grouping and thousands_sep, contrast with mon_grouping/mon_thousands_sep and strfmon(3). This is a no-op in the default "C" locale.

glibc 2.2 adds one further flag character.

- l For decimal integer conversion (i, d, u) the output uses the locale's alternative output digits, if any. For example, since glibc 2.2.3 this will give Arabic-Indic digits in the Persian ("fa_IR") locale.

Field width

An optional decimal digit string (with nonzero first digit) specifying a minimum field width. If the converted value has fewer characters than the field width, it will be padded with spaces on the left (or right, if the left-adjustment flag has been given). Instead of a decimal digit string one may write "*" or "*m\$" (for some decimal integer m) to specify that the field width is given in the next argument, or in the m-th argument, respectively, which must be of type int. A negative field width is taken as a '-' flag followed by a positive field width. In no case does a nonexistent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is expanded to contain the conversion result.

Precision

An optional precision, in the form of a period ('.') followed by an optional decimal digit string. Instead of a decimal digit string one may write "*" or "*m\$" (for some decimal integer m) to specify that the precision is given in the next argument, or in the m-th argument, respectively, which must be of type int. If the precision is given as just '.', the

precision is taken to be zero. A negative precision is taken as if the precision were omitted. This gives the minimum number of digits to appear for d, i, o, u, x, and X conversions, the number of digits to appear after the radix character for a, A, e, E, f, and F conversions, the maximum number of significant digits for g and G conversions, or the maximum number of characters to be printed from a string for s and S conversions.

Length modifier

Here, "integer conversion" stands for d, i, o, u, x, or X conversion.

hh A following integer conversion corresponds to a signed char or unsigned char argument, or a following n conversion corresponds to a pointer to a signed char argument.

h A following integer conversion corresponds to a short or unsigned short argument, or a following n conversion corresponds to a pointer to a short argument.

l (ell) A following integer conversion corresponds to a long or unsigned long argument, or a following n conversion corresponds to a pointer to a long argument, or a following c conversion corresponds to a `wint_t` argument, or a following s conversion corresponds to a pointer to `wchar_t` argument. On a following a, A, e, E, f, F, g, or G conversion, this length modifier is ignored (C99; not in SUSv2).

ll (ell-ell). A following integer conversion corresponds to a long long or unsigned long long argument, or a following n conversion corresponds to a pointer to a long long argument.

q A synonym for ll. This is a nonstandard extension, derived from BSD; avoid its use in new code.

L A following a, A, e, E, f, F, g, or G conversion corresponds to a long double argument. (C99 allows %LF, but SUSv2 does not.)

j A following integer conversion corresponds to an `intmax_t` or `uintmax_t` argument, or a following n conversion corresponds to a pointer to an `intmax_t` argument.

z A following integer conversion corresponds to a `size_t` or `ssize_t` argument, or a following n conversion corresponds to a pointer to a `size_t` argument.

Z A nonstandard synonym for z that predates the appearance of z. Do not use in new code.

t A following integer conversion corresponds to a `ptrdiff_t` argument, or a following n conversion corresponds to a pointer to a `ptrdiff_t` argument.

SUSv3 specifies all of the above, except for those modifiers explicitly noted as being nonstandard extensions. SUSv2 specified only the length modifiers h (in hd, hi, ho, hx, hX,

hn) and l (in ld, li, lo, lx, lX, ln, lc, ls) and L (in Le, LE, Lf, Lg, LG).

As a nonstandard extension, the GNU implementations treat ll and L as synonyms, so that one can, for example, write llg (as a synonym for the standards-compliant Lg) and Ld (as a synonym for the standards-compliant ld). Such usage is nonportable.

Conversion specifiers

A character that specifies the type of conversion to be applied. The conversion specifiers and their meanings are:

d, i The `int` argument is converted to signed decimal notation. The precision, if any, gives the minimum number of digits that must appear; if the converted value requires fewer digits, it is padded on the left with zeros. The default precision is 1. When 0 is printed with an explicit precision 0, the output is empty.

o, u, x, X

The `unsigned int` argument is converted to unsigned octal (`o`), unsigned decimal (`u`), or unsigned hexadecimal (`x` and `X`) notation. The letters `abcdef` are used for `x` conversions; the letters `ABCDEF` are used for `X` conversions. The precision, if any, gives the minimum number of digits that must appear; if the converted value requires fewer digits, it is padded on the left with zeros. The default precision is 1. When 0 is printed with an explicit precision 0, the output is empty.

e, E The `double` argument is rounded and converted in the style `[-]d.ddde?dd` where there is one digit (which is nonzero if the argument is nonzero) before the decimal-point character and the number of digits after it is equal to the precision; if the precision is missing, it is taken as 6; if the precision is zero, no decimal-point character appears. An `E` conversion uses the letter `E` (rather than `e`) to introduce the exponent. The exponent always contains at least two digits; if the value is zero, the exponent is `00`.

f, F The `double` argument is rounded and converted to decimal notation in the style `[-]ddd.ddd`, where the number of digits after the decimal-point character is equal to the precision specification. If the precision is missing, it is taken as 6; if the precision is explicitly zero, no decimal-point character appears. If a decimal point appears, at least one digit appears before it.

(SUSv2 does not know about `F` and says that character string representations for infinity and NaN may be made available. SUSv3 adds a specification for `F`. The C99 standard specifies `"[-]inf"` or `"[-]infinity"` for infinity, and a string starting with

"nan" for NaN, in the case of f conversion, and "[-]INF" or "[-]INFINITY" or "NAN" in the case of F conversion.)

g, G The double argument is converted in style f or e (or F or E for G conversions). The precision specifies the number of significant digits. If the precision is missing, 6 digits are given; if the precision is zero, it is treated as 1. Style e is used if the exponent from its conversion is less than -4 or greater than or equal to the precision. Trailing zeros are removed from the fractional part of the result; a decimal point appears only if it is followed by at least one digit.

a, A (C99; not in SUSv2, but added in SUSv3) For a conversion, the double argument is converted to hexadecimal notation (using the letters abcdef) in the style [-]0xh.hhp?d; for A conversion the prefix 0X, the letters ABCDEF, and the exponent separator P is used. There is one hexadecimal digit before the radix point, and the number of digits after it is equal to the precision. The default precision suffices for an exact representation of the value if an exact representation in base 2 exists and otherwise is sufficiently large to distinguish values of type double. The digit before the radix point is unspecified for nonnormalized numbers, and nonzero but otherwise unspecified for normalized numbers. The exponent, d, is the appropriate exponent of 2 expressed as a decimal integer; it always contains at least one digit; if the value is zero, the exponent is 0.

c If no l modifier is present, the int argument is converted to an unsigned char, and the resulting character is written. If an l modifier is present, the wint_t (wide character) argument is converted to a multibyte sequence by a call to the wctomb(3) function, with a conversion state starting in the initial state, and the resulting multibyte string is written.

s If no l modifier is present: the const char * argument is expected to be a pointer to an array of character type (pointer to a string). Characters from the array are written up to (but not including) a terminating null byte ('\0'); if a precision is specified, no more than the number specified are written. If a precision is given, no null byte need be present; if the precision is not specified, or is greater than the size of the array, the array must contain a terminating null byte.
If an l modifier is present: the const wchar_t * argument is expected to be a pointer to an array of wide characters. Wide characters from the array are converted to multibyte characters (each by a call to the wctomb(3) function, with a conversion

state starting in the initial state before the first wide character), up to and including a terminating null wide character. The resulting multibyte characters are written up to (but not including) the terminating null byte. If a precision is specified, no more bytes than the number specified are written, but no partial multibyte characters are written. Note that the precision determines the number of bytes written, not the number of wide characters or screen positions. The array must contain a terminating null wide character, unless a precision is given and it is so small that the number of bytes written exceeds it before the end of the array is reached.

C (Not in C99 or C11, but in SUSv2, SUSv3, and SUSv4.) Synonym for **lc**. Don't use.

S (Not in C99 or C11, but in SUSv2, SUSv3, and SUSv4.) Synonym for **ls**. Don't use.

p The void * pointer argument is printed in hexadecimal (as if by `%#x` or `%#lx`).

n The number of characters written so far is stored into the integer pointed to by the corresponding argument. That argument shall be an `int *`, or variant whose size matches the (optionally) supplied integer length modifier. No argument is converted. (This specifier is not supported by the bionic C library.) The behavior is undefined if the conversion specification includes any flags, a field width, or a precision.

m (glibc extension; supported by uClibc and musl, and on Android from API level 29.)

Print output of `strerror(errno)` (or `strerrorname_np(errno)` in the alternate form).

No argument is required.

% A `'%'` is written. No argument is converted. The complete conversion specification is `'%%'`.

RETURN VALUE

Upon successful return, these functions return the number of bytes printed (excluding the null byte used to end output to strings).

On error, a negative value is returned, and `errno` is set to indicate the error.

ERRORS

See `write(2)` and `putwc(3)`. In addition, the following error may occur:

EOVERFLOW

The value to be returned is greater than `INT_MAX`.

The `dprintf()` function may fail additionally if:

EBADF The `fd` argument is not a valid file descriptor.

ATTRIBUTES

For an explanation of the terms used in this section, see `attributes(7)`.

To print Pi to five decimal places:

```
#include <math.h>

#include <stdio.h>

fprintf(stdout, "pi = %.5f\n", 4 * atan(1.0));
```

To print a date and time in the form "Sunday, July 3, 10:02", where weekday and month are pointers to strings:

```
#include <stdio.h>

fprintf(stdout, "%s, %s %d, %.2d:%.2d\n",

        weekday, month, day, hour, min);
```

Many countries use the day-month-year order. Hence, an internationalized version must be able to print the arguments in an order specified by the format:

```
#include <stdio.h>

fprintf(stdout, format,

        weekday, month, day, hour, min);
```

where format depends on locale, and may permute the arguments. With the value:

```
"%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

one might obtain "Sonntag, 3. Juli, 10:02".

SEE ALSO

printf(1), asprintf(3), puts(3), scanf(3), setlocale(3), snprintf(3), strfromd(3), wcr?
tomb(3), wprintf(3), locale(5)

Linux man-pages 6.17

2026-02-08

printf(3)