



Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'getcwd.3p' command

\$ man getcwd.3p

GETCWD(3P) POSIX Programmer's Manual GETCWD(3P)

PROLOG

This manual page is part of the POSIX Programmer's Manual. The Linux implementation of this interface may differ (consult the corresponding Linux manual page for details of Linux behavior), or the interface may not be implemented on Linux.

NAME

getcwd ? get the pathname of the current working directory

SYNOPSIS

```
#include <unistd.h>

char *getcwd(char *buf, size_t size);
```

DESCRIPTION

The `getcwd()` function shall place an absolute pathname of the current working directory in the array pointed to by `buf`, and return `buf`. The pathname shall contain no components that are dot or dot-dot, or are symbolic links.

If there are multiple pathnames that `getcwd()` could place in the array pointed to by `buf`, one beginning with a single `<slash>` character and one or more beginning with two `<slash>` characters, then `getcwd()` shall place the pathname beginning with a single `<slash>` character in the array. The pathname shall not contain any unnecessary `<slash>` characters after the leading one or two `<slash>` characters.

The size argument is the size in bytes of the character array pointed

to by the `buf` argument. If `buf` is a null pointer, the behavior of `getcwd()` is unspecified.

RETURN VALUE

Upon successful completion, `getcwd()` shall return the `buf` argument.

Otherwise, `getcwd()` shall return a null pointer and set `errno` to indicate the error. The contents of the array pointed to by `buf` are then undefined.

ERRORS

The `getcwd()` function shall fail if:

`EINVAL` The size argument is 0.

`ERANGE` The size argument is greater than 0, but is smaller than the length of the string +1.

The `getcwd()` function may fail if:

`EACCES` Search permission was denied for the current directory, or read or search permission was denied for a directory above the current directory in the file hierarchy.

`ENOMEM` Insufficient storage space is available.

The following sections are informative.

EXAMPLES

The following example uses `{PATH_MAX}` as the initial buffer size (unless it is indeterminate or very large), and calls `getcwd()` with progressively larger buffers until it does not give an `[ERANGE]` error.

```
#include <stdlib.h>
#include <errno.h>
#include <unistd.h>
...
long path_max;
size_t size;
char *buf;
char *ptr;
path_max = pathconf(".", _PC_PATH_MAX);
if (path_max == -1)
    size = 1024;
```

```

else if (path_max > 10240)
    size = 10240;
else
    size = path_max;
for (buf = ptr = NULL; ptr == NULL; size *= 2)
{
    if ((buf = realloc(buf, size)) == NULL)
    {
        ... handle error ...
    }
    ptr = getcwd(buf, size);
    if (ptr == NULL && errno != ERANGE)
    {
        ... handle error ...
    }
}
...
free (buf);

```

APPLICATION USAGE

If the pathname obtained from `getcwd()` is longer than `{PATH_MAX}` bytes, it could produce an `[ENAMETOOLONG]` error if passed to `chdir()`. Therefore, in order to return to that directory it may be necessary to break the pathname into sections shorter than `{PATH_MAX}` bytes and call `chdir()` on each section in turn (the first section being an absolute pathname and subsequent sections being relative pathnames). A simpler way to handle saving and restoring the working directory when it may be deeper than `{PATH_MAX}` bytes in the file hierarchy is to use a file descriptor and `fchdir()`, rather than `getcwd()` and `chdir()`. However, the two methods do have some differences. The `fchdir()` approach causes the program to restore a working directory even if it has been renamed in the meantime, whereas the `chdir()` approach restores to a directory with the same name as the original, even if the directories were renamed in the meantime. Since the `fchdir()` approach does not access parent direc?

ories, it can succeed when `getcwd()` would fail due to permissions problems. In applications conforming to earlier versions of this standard, it was not possible to use the `fchdir()` approach when the working directory is searchable but not readable, as the only way to open a directory was with `O_RDONLY`, whereas the `getcwd()` approach can succeed in this case.

RATIONALE

Having `getcwd()` take no arguments and instead use the `malloc()` function to produce space for the returned argument was considered. The advantage is that `getcwd()` knows how big the working directory pathname is and can allocate an appropriate amount of space. But the programmer would have to use the `free()` function to free the resulting object, or each use of `getcwd()` would further reduce the available memory. Finally, `getcwd()` is taken from the SVID where it has the two arguments used in this volume of POSIX.1-2017.

The older function `getwd()` was rejected for use in this context because it had only a buffer argument and no size argument, and thus had no way to prevent overwriting the buffer, except to depend on the programmer to provide a large enough buffer.

On some implementations, if `buf` is a null pointer, `getcwd()` may obtain size bytes of memory using `malloc()`. In this case, the pointer returned by `getcwd()` may be used as the argument in a subsequent call to `free()`. Invoking `getcwd()` with `buf` as a null pointer is not recommended in conforming applications.

Earlier implementations of `getcwd()` sometimes generated pathnames like `"../..../subdirname"` internally, using them to explore the path of ancestor directories back to the root. If one of these internal pathnames exceeded `{PATH_MAX}` in length, the implementation could fail with `errno` set to `[ENAMETOOLONG]`. This is no longer allowed.

If a program is operating in a directory where some (grand)parent directory does not permit reading, `getcwd()` may fail, as in most implementations it must read the directory to determine the name of the file. This can occur if search, but not read, permission is granted in

an intermediate directory, or if the program is placed in that directory by some more privileged process (for example, login). Including the [EACCES] error condition makes the reporting of the error consistent and warns the application developer that `getcwd()` can fail for reasons beyond the control of the application developer or user. Some implementations can avoid this occurrence (for example, by implementing `getcwd()` using `pwd`, where `pwd` is a set-user-root process), thus the error was made optional. Since this volume of POSIX.1-2017 permits the addition of other errors, this would be a common addition and yet one that applications could not be expected to deal with without this addition.

FUTURE DIRECTIONS

None.

SEE ALSO

`malloc()`

The Base Definitions volume of POSIX.1-2017, `<unistd.h>`

COPYRIGHT

Portions of this text are reprinted and reproduced in electronic form from IEEE Std 1003.1-2017, Standard for Information Technology -- Portable Operating System Interface (POSIX), The Open Group Base Specifications Issue 7, 2018 Edition, Copyright (C) 2018 by the Institute of Electrical and Electronics Engineers, Inc and The Open Group. In the event of any discrepancy between this version and the original IEEE and The Open Group Standard, the original IEEE and The Open Group Standard is the referee document. The original Standard can be obtained online at <http://www.opengroup.org/unix/online.html>.

Any typographical or formatting errors that appear in this page are most likely to have been introduced during the conversion of the source files to man page format. To report such errors, see https://www.kernel.org/doc/man-pages/reporting_bugs.html.