



## ***Linux Ubuntu 22.4.5 Manual Pages on command 'clock\_settime.3'***

***\$ man clock\_settime.3***

CLOCK\_GETRES(2)                      Linux Programmer's Manual                      CLOCK\_GETRES(2)

NAME

clock\_getres, clock\_gettime, clock\_settime - clock and time functions

SYNOPSIS

```
#include <time.h>
```

```
int clock_getres(clockid_t clk_id, struct timespec *res);
```

```
int clock_gettime(clockid_t clk_id, struct timespec *tp);
```

```
int clock_settime(clockid_t clk_id, const struct timespec *tp);
```

Link with -lrt (only for glibc versions before 2.17).

Feature Test Macro Requirements for glibc (see feature\_test\_macros(7)):

```
clock_getres(), clock_gettime(), clock_settime():
```

```
    _POSIX_C_SOURCE >= 199309L
```

DESCRIPTION

The function `clock_getres()` finds the resolution (precision) of the specified clock `clk_id`, and, if `res` is non-NULL, stores it in the struct `timespec` pointed to by `res`. The resolution of clocks depends on the implementation and cannot be configured by a particular process. If the time value pointed to by the argument `tp` of `clock_settime()` is not a multiple of `res`, then it is truncated to a multiple of `res`.

The functions `clock_gettime()` and `clock_settime()` retrieve and set the time of the specified clock `clk_id`.

The `res` and `tp` arguments are `timespec` structures, as specified in `<time.h>`:

```

struct timespec {
    time_t  tv_sec;    /* seconds */
    long   tv_nsec;   /* nanoseconds */
};

```

The `clk_id` argument is the identifier of the particular clock on which to act. A clock may be system-wide and hence visible for all processes, or per-process if it measures time only within a single process.

All implementations support the system-wide real-time clock, which is identified by `CLOCK_REALTIME`. Its time represents seconds and nanoseconds since the Epoch. When its time is changed, timers for a relative interval are unaffected, but timers for an absolute point in time are affected.

More clocks may be implemented. The interpretation of the corresponding time values and the effect on timers is unspecified.

Sufficiently recent versions of glibc and the Linux kernel support the following clocks:

#### `CLOCK_REALTIME`

System-wide clock that measures real (i.e., wall-clock) time. Setting this clock requires appropriate privileges. This clock is affected by discontinuous jumps in the system time (e.g., if the system administrator manually changes the clock), and by the incremental adjustments performed by `adjtime(3)` and NTP.

#### `CLOCK_REALTIME_COARSE` (since Linux 2.6.32; Linux-specific)

A faster but less precise version of `CLOCK_REALTIME`. Use when you need very fast, but not fine-grained timestamps. Requires per-architecture support, and probably also architecture support for this flag in the `vdso(7)`.

#### `CLOCK_MONOTONIC`

Clock that cannot be set and represents monotonic time since some unspecified point in the past. On Linux, that point corresponds to the number of seconds that the system has been running since it was booted.

The `CLOCK_MONOTONIC` clock is not affected by discontinuous jumps in the system time (e.g., if the system administrator manually changes the clock), but is affected by the incremental adjustments performed by `adjtime(3)` and NTP.

This clock does not count time that the system is suspended.

`CLOCK_MONOTONIC_COARSE` (since Linux 2.6.32; Linux-specific)

A faster but less precise version of `CLOCK_MONOTONIC`. Use when you need very fast, but not fine-grained timestamps. Requires per-architecture support, and probably also architecture support for this flag in the `vdso(7)`.

`CLOCK_MONOTONIC_RAW` (since Linux 2.6.28; Linux-specific)

Similar to `CLOCK_MONOTONIC`, but provides access to a raw hardware-based time that is not subject to NTP adjustments or the incremental adjustments performed by `adjtime(3)`. This clock does not count time that the system is suspended.

`CLOCK_BOOTTIME` (since Linux 2.6.39; Linux-specific)

Identical to `CLOCK_MONOTONIC`, except it also includes any time that the system is suspended. This allows applications to get a suspend-aware monotonic clock without having to deal with the complications of `CLOCK_REALTIME`, which may have discontinuities if the time is changed using `settimeofday(2)` or similar.

`CLOCK_PROCESS_CPUTIME_ID` (since Linux 2.6.12)

Per-process CPU-time clock (measures CPU time consumed by all threads in the process).

`CLOCK_THREAD_CPUTIME_ID` (since Linux 2.6.12)

Thread-specific CPU-time clock.

## RETURN VALUE

`clock_gettime()`, `clock_settime()`, and `clock_getres()` return 0 for success, or -1 for failure (in which case `errno` is set appropriately).

## ERRORS

`EFAULT` `tp` points outside the accessible address space.

`EINVAL` The `clk_id` specified is not supported on this system.

`EINVAL` (`clock_settime()`): `tp.tv_sec` is negative or `tp.tv_nsec` is outside the range `[0..999,999,999]`.

`EINVAL` (since Linux 4.3)

A call to `clock_settime()` with a `clk_id` of `CLOCK_REALTIME` attempted to set the time to a value less than the current value of the `CLOCK_MONOTONIC` clock.

EPERM clock\_settime() does not have permission to set the clock indicated.

## VERSIONS

These system calls first appeared in Linux 2.6.

## ATTRIBUTES

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

??

?Interface                   ? Attribute   ? Value   ?

??

?clock\_getres(), clock\_gettime(), ? Thread safety ? MT-Safe ?

?clock\_settime()           ?           ?           ?

??

## CONFORMING TO

POSIX.1-2001, POSIX.1-2008, SUSv2.

## AVAILABILITY

On POSIX systems on which these functions are available, the symbol \_POSIX\_TIMERS is defined in <unistd.h> to a value greater than 0. The symbols \_POSIX\_MONOTONIC\_CLOCK, \_POSIX\_CPUTIME, \_POSIX\_THREAD\_CPUTIME indicate that CLOCK\_MONOTONIC, CLOCK\_PROCESS\_CPUTIME\_ID, CLOCK\_THREAD\_CPUTIME\_ID are available. (See also sysconf(3).)

## NOTES

POSIX.1 specifies the following:

Setting the value of the CLOCK\_REALTIME clock via clock\_settime() shall have no effect on threads that are blocked waiting for a relative time service based upon this clock, including the nanosleep() function; nor on the expiration of relative timers based upon this clock. Consequently, these time services shall expire when the requested relative interval elapses, independently of the new or old value of the clock.

### C library/kernel differences

On some architectures, an implementation of clock\_gettime() is provided in the vdso(7).

### Historical note for SMP systems

Before Linux added kernel support for CLOCK\_PROCESS\_CPUTIME\_ID and CLOCK\_THREAD\_CPUTIME\_ID, glibc implemented these clocks on many platforms using

timer registers from the CPUs (TSC on i386, AR.ITC on Itanium). These registers may differ between CPUs and as a consequence these clocks may return bogus results if a process is migrated to another CPU.

If the CPUs in an SMP system have different clock sources, then there is no way to maintain a correlation between the timer registers since each CPU will run at a slightly different frequency. If that is the case, then `clock_getcpuclockid(0)` will return `ENOENT` to signify this condition. The two clocks will then be useful only if it can be ensured that a process stays on a certain CPU.

The processors in an SMP system do not start all at exactly the same time and therefore the timer registers are typically running at an offset. Some architectures include code that attempts to limit these offsets on bootup. However, the code cannot guarantee to accurately tune the offsets. Glibc contains no provisions to deal with these offsets (unlike the Linux Kernel). Typically these offsets are small and therefore the effects may be negligible in most cases.

Since glibc 2.4, the wrapper functions for the system calls described in this page avoid the abovementioned problems by employing the kernel implementation of `CLOCK_PROCESS_CPUTIME_ID` and `CLOCK_THREAD_CPUTIME_ID`, on systems that provide such an implementation (i.e., Linux 2.6.12 and later).

## BUGS

According to POSIX.1-2001, a process with "appropriate privileges" may set the `CLOCK_PROCESS_CPUTIME_ID` and `CLOCK_THREAD_CPUTIME_ID` clocks using `clock_settime()`. On Linux, these clocks are not settable (i.e., no process has "appropriate privileges").

## SEE ALSO

`date(1)`, `gettimeofday(2)`, `settimeofday(2)`, `time(2)`, `adjtime(3)`, `clock_getcpu?`  
`clockid(3)`, `ctime(3)`, `ftime(3)`, `pthread_getcpuclockid(3)`, `sysconf(3)`, `time(7)`,  
`vdso(7)`, `hwclock(8)`

## COLOPHON

This page is part of release 5.05 of the Linux man-pages project. A description of the project, information about reporting bugs, and the latest version of this page, can be found at <https://www.kernel.org/doc/man-pages/>.