

# Package ‘unix’

January 31, 2020

**Title** POSIX System Utilities

**Version** 1.5.1

**Description** Bindings to system utilities found in most Unix systems such as POSIX functions which are not part of the Standard C Library.

**License** MIT + file LICENSE

**URL** <https://github.com/jeroen/unix>

**BugReports** <https://github.com/jeroen/unix/issues>

**OS\_type** unix

**SystemRequirements** POSIX.1-2001

**RoxygenNote** 6.1.1

**Suggests** testthat

**Language** en-US

**NeedsCompilation** yes

**Author** Jeroen Ooms [aut, cre] (<<https://orcid.org/0000-0002-4035-0289>>)

**Maintainer** Jeroen Ooms <[jeroen@berkeley.edu](mailto:jeroen@berkeley.edu)>

**Repository** CRAN

**Date/Publication** 2020-01-31 13:20:02 UTC

## R topics documented:

chroot . . . . .	2
eval_safe . . . . .	2
getuid . . . . .	4
rlimit . . . . .	6
sys_config . . . . .	8
userinfo . . . . .	8

<b>Index</b>	<b>9</b>
--------------	----------

---

chroot	<i>Change Root Dir</i>
--------	------------------------

---

### Description

Changes the root directory of the calling process to that specified in path. This directory will be used for pathnames beginning with /. **Only a privileged process (i.e. sudo) may call chroot().**

### Usage

```
chroot(path = getwd())
```

### Arguments

path	directory of the new root
------	---------------------------

### Details

This call changes an ingredient in the pathname resolution process and does nothing else. In particular, it is not intended to be used for any kind of security purpose, neither to fully sandbox a process nor to restrict filesystem system calls.

### References

[CHROOT\(2\)](#)

---

eval_safe	<i>Safe Evaluation</i>
-----------	------------------------

---

### Description

Evaluates an expression in a temporary fork and returns the value without any side effects on the main R session. For `eval_safe()` the expression is wrapped in additional R code to handle errors and graphics.

### Usage

```
eval_safe(expr, tmp = tempfile("fork"), std_out = stdout(),
  std_err = stderr(), timeout = 0, priority = NULL, uid = NULL,
  gid = NULL, rlimits = NULL, profile = NULL, device = pdf)
```

```
eval_fork(expr, tmp = tempfile("fork"), std_out = stdout(),
  std_err = stderr(), timeout = 0)
```

**Arguments**

expr	expression to evaluate
tmp	the value of <code>tempdir()</code> inside the forked process
std_out	if and where to direct child process STDOUT. Must be one of TRUE, FALSE, file-name, connection object or callback function. See section on <i>Output Streams</i> below for details.
std_err	if and where to direct child process STDERR. Must be one of TRUE, FALSE, file-name, connection object or callback function. See section on <i>Output Streams</i> below for details. Non root user may only raise this value (decrease priority)
timeout	maximum time in seconds to allow for call to return
priority	(integer) priority of the child process. High value is low priority.
uid	evaluate as given user (uid or name). See <code>unix::setuid()</code> , only for root.
gid	evaluate as given group (gid or name). See <code>unix::setgid()</code> only for root.
rlimits	named vector/list with rlimit values, for example: <code>c(cpu = 60, fsize = 1e6)</code> .
profile	AppArmor profile, see <code>RAppArmor::aa_change_profile()</code> . Requires the RAppArmor package (Debian/Ubuntu only)
device	graphics device to use in the fork, see <code>dev.new()</code>

**Details**

Some programs such as Java are not fork-safe and cannot be called from within a forked process if they have already been loaded in the main process. On MacOS any software calling CoreFoundation functionality might crash within the fork. This includes `libcurl` which has been built on OSX against native SecureTransport rather than OpenSSL for https connections. The same limitations hold for e.g. `parallel::mcpipeline()`.

**Examples**

```
# works like regular eval:
eval_safe(rnorm(5))

# Exceptions get propagated
test <- function() { doesnotexit() }
tryCatch(eval_safe(test()), error = function(e){
  cat("oh no!", e$message, "\n")
})

# Honor interrupt and timeout, even inside C evaluations
try(eval_safe(svd(matrix(rnorm(1e8), 1e4)), timeout = 2))

# Capture output
outcon <- rawConnection(raw(0), "r+")
eval_safe(print(sessionInfo()), std_out = outcon)
cat(rawToChar(rawConnectionValue(outcon)))
close(outcon)
```

---

`getuid`*Process Info*

---

**Description**

Get or set attributes of the current process.

**Usage**`getuid()``getgid()``geteuid()``getegid()``getpid()``getppid()``getpgid()``getpriority()``setuid(uid)``seteuid(uid)``setgid(gid)``setegid(gid)``setpgid(pgid = 0)``setpriority(prio)``kill(pid, signal = SIGTERM)`**Arguments**

<code>uid</code>	User ID from <code>/etc/passwd</code> .
<code>gid</code>	Group ID from <code>/etc/group</code> .
<code>pgid</code>	Process Group ID. Default <code>0</code> sets <code>pgid</code> to the current <code>pid</code> .
<code>prio</code>	Priority level
<code>pid</code>	process ID (integer)
<code>signal</code>	a signal number (integer), defaults to <a href="#">tools::SIGTERM</a> .

**Details**

Acronyms stand for:

- pid Process ID
- ppid Parent-Process ID
- pgid Process-Group ID
- uid User ID
- euid Effective User ID
- gid Group ID
- egid Effective Group ID
- prio Priority level

An unprivileged (non-root) process cannot change it's uid and only lower process priority (higher value).

**References**

[GETUID\(2\)](#) [GETPID\(2\)](#) [GETPGID\(2\)](#) [GETPRIORITY\(2\)](#)

**Examples**

```
# Current User:
getuid()
# Current UserGroup:
getgid()
# Current UserGroup:
geteuid()
# Current UserGroup:
getegid()
# Process ID
getpid()
# parent PID:
getppid()
# Process group id:
getpgid()

# Detach process group
setpgid(0)
getpgid()
# Process priority:
getpriority()
# Decrease priority
setpriority(getpriority() + 1)
```

---

`rlimit`*Resource Limits*

---

### Description

Get and set process resource limits. Each function returns the current limits, and can optionally update the limit by passing argument values. The `rlimit_all()` function is a convenience wrapper which prints all current hard and soft limits.

### Usage

```
rlimit_all()

rlimit_as(cur = NULL, max = NULL)

rlimit_core(cur = NULL, max = NULL)

rlimit_cpu(cur = NULL, max = NULL)

rlimit_data(cur = NULL, max = NULL)

rlimit_fsize(cur = NULL, max = NULL)

rlimit_memlock(cur = NULL, max = NULL)

rlimit_nofile(cur = NULL, max = NULL)

rlimit_nproc(cur = NULL, max = NULL)

rlimit_stack(cur = NULL, max = NULL)
```

### Arguments

<code>cur</code>	set the current (soft) limit for this resource. See details.
<code>max</code>	set the max (hard) limit for this resource. See details.

### Details

Each resource has an associated soft and hard limit. The soft limit is the value that the kernel enforces for the corresponding resource. The hard limit acts as a ceiling for the soft limit: an unprivileged process may set only its soft limit to a value in the range from 0 up to the hard limit, and (irreversibly) lower its hard limit.

Definitons from the [Linux manual page](#) are as follows:

- `RLIMIT_AS` : the maximum size of the process's virtual memory (address space) in bytes.
- `RLIMIT_CORE` : the maximum size of a core file that the process may dump.

- RLIMIT\_CPU : a limit in seconds on the amount of CPU time (**not** elapsed time) that the process may consume. When the process reaches the soft limit, it is sent a SIGXCPU signal.
- RLIMIT\_DATA : the maximum size of the process's data segment (initialized data, uninitialized data, and heap).
- RLIMIT\_FSIZE : the maximum size of files that the process may create. Attempts to extend a file beyond this limit result in delivery of a SIGXFSZ signal.
- RLIMIT\_MEMLOCK : the maximum number of bytes of memory that may be locked into RAM.
- RLIMIT\_NOFILE : a value one greater than the maximum file descriptor number that can be opened by this process.
- RLIMIT\_NPROC : the maximum number of processes that can be created for the real user ID of the calling process. Upon encountering this limit, fork fails with the error EAGAIN. Not enforced for root user.
- RLIMIT\_STACK : the maximum size of the process stack, in bytes.

Note that the support for enforcing limits vary widely by system. In particular RLIMIT\_AS has a different meaning depending on how memory allocation is managed by the operating system (and doesn't work at all on MacOS).

## References

[GETRLIMIT\(2\)](#)

## Examples

```
# Print all limits
rlimit_all()

# Get one limit
rlimit_as()

# Set a soft limit
lim <- rlimit_as(1e9)
print(lim)

# Reset the limit to max
rlimit_as(cur = lim$max)

## Not run:
# Set a hard limit (irreversible)
rlimit_as(max = 1e10)

## End(Not run)
```

sys\_config

*Package config*

---

**Description**

Shows which features are enabled in the package configuration.

**Usage**

```
sys_config()
```

```
aa_config()
```

**Examples**

```
sys_config()
```

---

userinfo

*User / Group Info*

---

**Description**

Lookup a user or group info via user uid/name or group gid/name.

**Usage**

```
user_info(uid = getuid())
```

```
group_info(gid = getgid())
```

**Arguments**

```
uid          user ID (integer) or name (string)
```

```
gid          group ID (integer) or name (string)
```

**References**

[GETPWNAM\(3\)](#) [GETGRNAM\(3\)](#)

**Examples**

```
# Get info current user
user_info()
group_info()
```



# Index

aa\_config (sys\_config), 8

chroot, 2

dev.new(), 3

eval\_fork (eval\_safe), 2

eval\_safe, 2

eval\_safe(), 2

getegid (getuid), 4

geteuid (getuid), 4

getgid (getuid), 4

getpgid (getuid), 4

getpid (getuid), 4

getppid (getuid), 4

getpriority (getuid), 4

getuid, 4

group\_info (userinfo), 8

kill (getuid), 4

rlimit, 6

rlimit\_all (rlimit), 6

rlimit\_as (rlimit), 6

rlimit\_core (rlimit), 6

rlimit\_cpu (rlimit), 6

rlimit\_data (rlimit), 6

rlimit\_fsize (rlimit), 6

rlimit\_memlock (rlimit), 6

rlimit\_nofile (rlimit), 6

rlimit\_nproc (rlimit), 6

rlimit\_stack (rlimit), 6

setegid (getuid), 4

seteuid (getuid), 4

setgid (getuid), 4

setpgid (getuid), 4

setpriority (getuid), 4

setuid (getuid), 4

sys\_config, 8

tempdir(), 3

tools::SIGTERM, 4

unix::setgid(), 3

unix::setuid(), 3

user\_info (userinfo), 8

userinfo, 8