# Bash 的 addpart 命令

```
addpart device partition start length
```

**功能**

通知内核添加硬盘分区（不会实际创建分区）。

**类型**

可执行文件（`/usr/bin/addpart`），属于 util-linux。

**参数**

- `device` - 硬盘设备
- `partition` - 分区号
- `start` - 分区的开始位置（以 512 字节扇区为单位）
- `length` - 分区的长度（以 512 字节扇区为单位）

## 示例

```
sudo addpart /dev/sda 3 2097152 1048576
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| delpart | 移除硬盘分区 |

## 推荐阅读

- addpart(8) — Linux manual page

## 手册

```
ADDPART(8)                  System Administration                  ADDPART(8)

NAME
```

```
       addpart - tell the kernel about the existence of a partition

SYNOPSIS
       addpart device partition start length

DESCRIPTION
       addpart tells the Linux kernel about the existence of the specified
       partition. The command is a simple wrapper around the "add partition"
       ioctl.

       This command doesn’t manipulate partitions on a block device.

PARAMETERS
       device
           The disk device.

       partition
           The partition number.

       start
           The beginning of the partition (in 512-byte sectors).

       length
           The length of the partition (in 512-byte sectors).

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

SEE ALSO
       delpart(8), fdisk(8), parted(8), partprobe(8), partx(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The addpart command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                        ADDPART(8)
```

# Bash 的 agetty 命令

```
agetty [OPTION]... PORT [BAUD_RATE]... [TERM]
```

**功能**

在终端设备上启动登录提示并等待用户登录。

> getty 的替代品。

**类型**

可执行文件（`/usr/sbin/agetty`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-8`，`--8bits` - 允许 8 位字符输入（禁用奇偶校验的第 8 位剥离）
  - `-a`，`--autologin username` - 自动以指定用户登录（无需输入用户名）
  - `-c`，`--noreset` - 不重置终端控制模式
  - `-E`，`--remote` - 表示连接来自远程终端（通常用于远程线路或终端服务器）
  - `-f`，`--issue-file path` - 指定显示登录提示信息的文件（默认通常为 /etc/issue）
  - `--show-issue` - 仅显示 issue 文件内容后退出
  - `-h`，`--flow-control` - 启用硬件流控制（RTS/CTS）
  - `-H`，`--host fakehost` - 指定伪造的远程主机名并传递给 login
  - `-i`，`--noissue` - 不显示 issue 文件内容
  - `-I`，`--init-string initstring` - 向终端或调制解调器发送初始化字符串
  - `-J`，`--noclear` - 不清屏直接显示登录提示
  - `-l`，`--login-program login_program` - 指定要执行的登录程序（默认是 /bin/login）
  - `-L`，`--local-line[=mode]` - 设置为本地线路（忽略 modem 控制信号）
  - `-m`，`--extract-baud` - 从终端连接字符串中提取波特率
  - `--list-speeds` - 显示支持的串口波特率列表
  - `-n`，`--skip-login` - 不提示用户名，直接执行登录程序
  - `-N`，`--nonewline` - 不在登录提示后输出换行符
  - `-o`，`--login-options login_options` - 向 login 程序传递额外参数
  - `-p`，`--login-pause` - 在执行 login 前等待用户按键
  - `-r`，`--chroot directory` - 在执行 login 前切换到指定的 chroot 目录
  - `-R`，`--hangup` - 在退出时对终端执行挂断操作

  - `-s`, `--keep-baud` - 保持当前串口波特率设置
  - `-t`, `--timeout timeout` - 指定等待用户名输入的超时时间（秒）
  - `-U`, `--detect-case` - 自动检测用户名输入的大小写并转换
  - `-w`, `--wait-cr` - 等待回车字符后再显示登录提示
  - `--nohints` - 不显示登录提示中的帮助信息
  - `--nohostname` - 不在登录提示中显示主机名
  - `--long-hostname` - 在登录提示中显示完整主机名（FQDN）
  - `--erase-chars string` - 设置退格（erase）字符
  - `--kill-chars string` - 设置删除整行（kill）字符
  - `--chdir directory` - 在执行 login 前切换到指定目录
  - `--delay number` - 在启动后延迟指定秒数再显示登录提示
  - `--nice number` - 设置进程的 nice 值（调整优先级）
  - `--reload` - 重新加载配置（通常用于 systemd 环境）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `PORT` - 端口；例如 `tty1`, `ttyS0`, `ttyUSB0` 等
- `BAUD_RATE` - 波特率；可以设置多个波特率，用逗号分隔
- `TERM` - 终端类型；例如 `vt100`, `linux` 等

## 示例

```
$ getty -t 60 tty1 38400     # 在 tty1 上以 38400 波特率启动，60秒内未操作则自动断开
```

## 推荐阅读

- agetty(8) — Linux manual page

## 手册

```
AGETTY(8)                   System Administration                   AGETTY(8)

NAME
       agetty - alternative Linux getty

SYNOPSIS
       agetty [options] port [baud_rate...] [term]

DESCRIPTION
       agetty opens a tty port, prompts for a login name and invokes the
       /bin/login command. It is normally invoked by init(8).
```

agetty has several non-standard features that are useful for hardwired and for dial-in lines:

- Adapts the tty settings to parity bits and to erase, kill, end-of-line and uppercase characters when it reads a login name. The program can handle 7-bit characters with even, odd, none or space parity, and 8-bit characters with no parity. The following special characters are recognized: Control-U (kill); DEL and backspace (erase); carriage return and line feed (end of line). See also the --erase-chars and --kill-chars options.

- Optionally deduces the baud rate from the CONNECT messages produced by Hayes(tm)-compatible modems.

- Optionally does not hang up when it is given an already opened line (useful for call-back applications).

- Optionally does not display the contents of the /etc/issue file.

- Optionally displays an alternative issue files or directories instead of /etc/issue or /etc/issue.d.

- Optionally does not ask for a login name.

- Optionally invokes a non-standard login program instead of /bin/login.

- Optionally turns on hardware flow control.

- Optionally forces the line to be local with no need for carrier detect.

This program does not use the /etc/gettydefs (System V) or /etc/gettytab (SunOS 4) files.

ARGUMENTS

port
: A path name relative to the /dev directory. If a "-" is specified, agetty assumes that its standard input is already connected to a tty port and that a connection to a remote user has already been established.

  Under System V, a "-" port argument should be preceded by a "--".

baud_rate,...
: A comma-separated list of one or more baud rates. Each time agetty receives a BREAK character it advances through the list, which is treated as if it were circular.

  Baud rates should be specified in descending order, so that the null character (Ctrl-@) can also be used for baud-rate switching.

```
            This argument is optional and unnecessary for virtual terminals.

            The default for serial terminals is keep the current baud rate (see
            --keep-baud) and if unsuccessful then default to '9600'.

        term
            The value to be used for the TERM environment variable. This
            overrides whatever init(1) may have set, and is inherited by login
            and the shell.

            The default is 'vt100', or 'linux' for Linux on a virtual terminal,
            or 'hurd' for GNU Hurd on a virtual terminal.

OPTIONS
        -8, --8bits
            Assume that the tty is 8-bit clean, hence disable parity detection.

        -a, --autologin username
            Automatically log in the specified user without asking for a
            username or password. Using this option causes an -f username
            option and argument to be added to the /bin/login command line. See
            --login-options, which can be used to modify this option’s
            behavior.

            Note that --autologin may affect the way in which getty initializes
            the serial line, because on auto-login agetty does not read from
            the line and it has no opportunity optimize the line setting.

        -c, --noreset
            Do not reset terminal cflags (control modes). See termios(3) for
            more details.

        -E, --remote
            Typically the login(1) command is given a remote hostname when
            called by something such as telnetd(8). This option allows agetty
            to pass what it is using for a hostname to login(1) for use in
            utmp(5). See --host, login(1), and utmp(5).

            If the --host fakehost option is given, then an -h fakehost option
            and argument are added to the /bin/login command line.

            If the --nohostname option is given, then an -H option is added to
            the /bin/login command line.

            See --login-options.

        -f, --issue-file path
            Specifies a ":" delimited list of files and directories to be
            displayed instead of /etc/issue (or other). All specified files and
            directories are displayed, missing or empty files are silently
            ignored. If the specified path is a directory then display all
            files with .issue file extension in version-sort order from the
            directory. This allows custom messages to be displayed on different
```

terminals. The --noissue option will override this option.

--show-issue
Display the current issue file (or other) on the current terminal and exit. Use this option to review the current setting, it is not designed for any other purpose. Note that output may use some default or incomplete information as proper output depends on terminal and agetty command line.

-h, --flow-control
Enable hardware (RTS/CTS) flow control. It is left up to the application to disable software (XON/XOFF) flow protocol where appropriate.

-H, --host fakehost
Write the specified fakehost into the utmp file. Normally, no login host is given, since agetty is used for local hardwired connections and consoles. However, this option can be useful for identifying terminal concentrators and the like.

-i, --noissue
Do not display the contents of /etc/issue (or other) before writing the login prompt. Terminals or communications hardware may become confused when receiving lots of text at the wrong baud rate; dial-up scripts may fail if the login prompt is preceded by too much text.

-I, --init-string initstring
Set an initial string to be sent to the tty or modem before sending anything else. This may be used to initialize a modem.
Non-printable characters may be sent by writing their octal code preceded by a backslash (\). For example, to send a linefeed character (ASCII 10, octal 012), write \12.

-J, --noclear
Do not clear the screen before prompting for the login name. By default the screen is cleared.

-l, --login-program login_program
Invoke the specified login_program instead of /bin/login. This allows the use of a non-standard login program. Such a program could, for example, ask for a dial-up password or use a different password file. See --login-options.

-L, --local-line[=mode]
Control the CLOCAL line flag. The optional mode argument is auto, always or never. If the mode argument is omitted, then the default is always. If the --local-line option is not given at all, then the default is auto.

always
Forces the line to be a local line with no need for carrier detect. This can be useful when you have a locally attached

terminal where the serial line does not set the carrier-detect signal.

never
Explicitly clears the CLOCAL flag from the line setting and the carrier-detect signal is expected on the line.

auto
The agetty default. Does not modify the CLOCAL setting and follows the setting enabled by the kernel.

-m, --extract-baud
Try to extract the baud rate from the CONNECT status message produced by Hayes(tm)-compatible modems. These status messages are of the form: "<junk><speed><junk>". agetty assumes that the modem emits its status message at the same speed as specified with (the first) baud_rate value on the command line.

Since the --extract-baud feature may fail on heavily-loaded systems, you still should enable BREAK processing by enumerating all expected baud rates on the command line.

--list-speeds
Display supported baud rates. These are determined at compilation time.

-n, --skip-login
Do not prompt the user for a login name. This can be used in connection with the --login-program option to invoke a non-standard login process such as a BBS system. Note that with the --skip-login option, agetty gets no input from the user who logs in and therefore will not be able to figure out parity, character size, and newline processing of the connection. It defaults to space parity, 7 bit characters, and ASCII CR (13) end-of-line character. Beware that the program that agetty starts (usually /bin/login) is run as root.

-N, --nonewline
Do not print a newline before writing out /etc/issue.

-o, --login-options login_options
Options and arguments that are passed to login(1). Where \u is replaced by the login name. For example:

--login-options '-h darkstar -- \u'

See --autologin, --login-program and --remote.

Please read the SECURITY NOTICE below before using this option.

-p, --login-pause
Wait for any key before dropping to the login prompt. Can be combined with --autologin to save memory by lazily spawning shells.

```
-r, --chroot directory
    Change root to the specified directory.

-R, --hangup
    Call vhangup(2) to do a virtual hangup of the specified terminal.

-s, --keep-baud
    Try to keep the existing baud rate. The baud rates from the command
    line are used when agetty receives a BREAK character. If another
    baud rates specified then the original baud rate is also saved to
    the end of the wanted baud rates list. This can be used to return
    to the original baud rate after unexpected BREAKs.

-t, --timeout timeout
    Terminate if no user name could be read within timeout seconds. Use
    of this option with hardwired terminal lines is not recommended.

-U, --detect-case
    Turn on support for detecting an uppercase-only terminal. This
    setting will detect a login name containing only capitals as
    indicating an uppercase-only terminal and turn on some
    upper-to-lower case conversions. Note that this has no support for
    any Unicode characters.

-w, --wait-cr
    Wait for the user or the modem to send a carriage-return or a
    linefeed character before sending the /etc/issue file (or others)
    and the login prompt. This is useful with the --init-string option.

--nohints
    Do not print hints about Num, Caps and Scroll Locks.

--nohostname
    By default the hostname will be printed. With this option enabled,
    no hostname at all will be shown. This setting is also possible to
    able by LOGIN_PLAIN_PROMPT option in the /etc/login.defs
    configuration file (see below for more details).

--long-hostname
    By default the hostname is only printed until the first dot. With
    this option enabled, the fully qualified hostname by
    gethostname(3P) or (if not found) by getaddrinfo(3) is shown.

--erase-chars string
    This option specifies additional characters that should be
    interpreted as a backspace ("ignore the previous character") when
    the user types the login name. The default additional 'erase' has
    been '#', but since util-linux 2.23 no additional erase characters
    are enabled by default.

--kill-chars string
    This option specifies additional characters that should be
```

```
            interpreted as a kill ("ignore all previous characters") when the
            user types the login name. The default additional 'kill' has been
            '@', but since util-linux 2.23 no additional kill characters are
            enabled by default.

        --chdir directory
            Change directory before the login.

        --delay number
            Sleep seconds before open tty.

        --nice number
            Run login with this priority.

        --reload
            Ask all running agetty instances to reload and update their
            displayed prompts, if the user has not yet commenced logging in.
            After doing so the command will exit. This feature might be
            unsupported on systems without Linux inotify(7).

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

CONFIG FILE ITEMS
        agetty reads the /etc/login.defs configuration file (see
        login.defs(5)). Note that the configuration file could be distributed
        with another package (usually shadow-utils). The following
        configuration items are relevant for agetty:

        LOGIN_PLAIN_PROMPT (boolean)
            Tell agetty that printing the hostname should be suppressed in the
            login: prompt. This is an alternative to the --nohostname command
            line option. The default value is no.

EXAMPLE
        This section shows examples for the process field of an entry in the
        /etc/inittab file. You'll have to prepend appropriate values for the
        other fields. See inittab(5) for more details.

        For a hardwired line or a console tty:

           /sbin/agetty 9600 ttyS1

        For a directly connected terminal without proper carrier-detect wiring
        (try this if your terminal just sleeps instead of giving you a
        password: prompt):

           /sbin/agetty --local-line 9600 ttyS1 vt100

        For an old-style dial-in line with a 9600/2400/1200 baud modem:
```

```
   /sbin/agetty --extract-baud --timeout 60 ttyS1 9600,2400,1200
```

For a Hayes modem with a fixed 115200 bps interface to the machine (the example init string turns off modem echo and result codes, makes modem/computer DCD track modem/modem DCD, makes a DTR drop cause a disconnection, and turns on auto-answer after 1 ring):

```
   /sbin/agetty --wait-cr --init-string 'ATE0Q1&D2&C1S0=1\015'
   115200 ttyS1
```

## SECURITY NOTICE

If you use the --login-program and --login-options options, be aware that a malicious user may try to enter lognames with embedded options, which then get passed to the used login program. agetty does check for a leading "-" and makes sure the logname gets passed as one parameter (so embedded spaces will not create yet another parameter), but depending on how the login binary parses the command line that might not be sufficient. Check that the used login program cannot be abused this way.

Some programs use "--" to indicate that the rest of the command line should not be interpreted as options. Use this feature if available by passing "--" before the username gets passed by \u.

## ISSUE FILES

The default issue file is /etc/issue. If the file exists, then agetty also checks for /etc/issue.d directory. The directory is optional extension to the default issue file and content of the directory is printed after /etc/issue content. If the /etc/issue does not exist, then the directory is ignored. All files with .issue extension from the directory are printed in version-sort order. The directory can be used to maintain 3rd-party messages independently on the primary system /etc/issue file.

Since version 2.35 additional locations for issue file and directory are supported. If the default /etc/issue does not exist, then agetty checks for /run/issue and /run/issue.d, thereafter for /usr/lib/issue and /usr/lib/issue.d. The directory /etc is expected for host specific configuration, /run is expected for generated stuff and /usr/lib for static distribution maintained configuration.

The default path maybe overridden by --issue-file option. In this case specified path has to be file or directory and all the default issue file and directory locations are ignored.

The issue file feature can be completely disabled by --noissue option.

It is possible to review the current issue file by agetty --show-issue on the current terminal.

The issue files may contain certain escape codes to display the system name, date, time et cetera. All escape codes consist of a backslash (\)

immediately followed by one of the characters listed below.

4 or 4{interface}
    Insert the IPv4 address of the specified network interface (for example: \4{eth0}). If the interface argument is not specified, then select the first fully configured (UP, non-LOCALBACK, RUNNING) interface. If no configured interface is found, fall back to the IP address of the machine's hostname.

6 or 6{interface}
    The same as \4 but for IPv6.

b
    Insert the baudrate of the current line.

d
    Insert the current date.

e or e{name}
    Translate the human-readable name to an escape sequence and insert it (for example: \e{red}Alert text.\e{reset}). If the name argument is not specified, then insert \033. The currently supported names are: black, blink, blue, bold, brown, cyan, darkgray, gray, green, halfbright, lightblue, lightcyan, lightgray, lightgreen, lightmagenta, lightred, magenta, red, reset, reverse, yellow and white. All unknown names are silently ignored.

s
    Insert the system name (the name of the operating system). Same as 'uname -s'. See also the \S escape code.

S or S{VARIABLE}
    Insert the VARIABLE data from /etc/os-release. If this file does not exist then fall back to /usr/lib/os-release. If the VARIABLE argument is not specified, then use PRETTY_NAME from the file or the system name (see \s). This escape code can be used to keep /etc/issue distribution and release independent. Note that \S{ANSI_COLOR} is converted to the real terminal escape sequence.

l
    Insert the name of the current tty line.

m
    Insert the architecture identifier of the machine. Same as uname -m.

n
    Insert the nodename of the machine, also known as the hostname. Same as uname -n.

o
    Insert the NIS domainname of the machine. Same as hostname -d.

O
    Insert the DNS domainname of the machine.

r
    Insert the release number of the OS. Same as uname -r.

t
    Insert the current time.

u
    Insert the number of current users logged in.

U
    Insert the string "1 user" or "<n> users" where <n> is the number
    of current users logged in.

v
    Insert the version of the OS, that is, the build-date and such.

An example. On my system, the following /etc/issue file:

    This is \n.\o (\s \m \r) \t

displays as:

    This is thingol.orcan.dk (Linux i386 1.1.9) 18:29:30

FILES
/var/run/utmp
    the system status file.

/etc/issue
    printed before the login prompt.

/etc/os-release /usr/lib/os-release
    operating system identification data.

/dev/console
    problem reports (if syslog(3) is not used).

/etc/inittab
    init(8) configuration file for SysV-style init daemon.

BUGS
The baud-rate detection feature (the --extract-baud option) requires
that agetty be scheduled soon enough after completion of a dial-in call
(within 30 ms with modems that talk at 2400 baud). For robustness,
always use the --extract-baud option in combination with a multiple
baud rate command-line argument, so that BREAK processing is enabled.

The text in the /etc/issue file (or other) and the login prompt are
always output with 7-bit characters and space parity.

The baud-rate detection feature (the --extract-baud option) requires that the modem emits its status message after raising the DCD line.

## DIAGNOSTICS

Depending on how the program was configured, all diagnostics are written to the console device or reported via the syslog(3) facility. Error messages are produced if the port argument does not specify a terminal device; if there is no utmp entry for the current process (System V only); and so on.

## AUTHORS

Werner Fink <werner@suse.de>, Karel Zak <kzak@redhat.com>

The original agetty for serial terminals was written by W.Z. Venema <wietse@wzv.win.tue.nl> and ported to Linux by Peter Orbaek <poe@daimi.aau.dk>.

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The agetty command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 arch 命令

```
arch [OPTION]...
```

**功能**

查看硬件架构。

**类型**

可执行文件（`/usr/bin/arch`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 推荐阅读

- Print machine hardware name - GNU Coreutils
- arch(1) — Linux manual page

## 示例

```
$ arch
x86_64
```

## 手册

```
ARCH(1)                          User Commands                          ARCH(1)

NAME
       arch - print machine hardware name (same as uname -m)

SYNOPSIS
       arch [OPTION]...

DESCRIPTION
       Print machine architecture.
```

```
       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie and Karel Zak.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       uname(1), uname(2)

       Full documentation <https://www.gnu.org/software/coreutils/arch>
       or available locally via: info '(coreutils) arch invocation'

GNU coreutils 9.4                April 2024                          ARCH(1)
```

# Bash 的 b2sum 命令

```
b2sum [OPTION]... [FILE]...
```

**功能**

计算或校验 BLAKE2 值。

**类型**

可执行文件（`/usr/bin/b2sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 BLAKE2 值进行校验
  - `-l`，`--length=BITS` - 结果保留的长度（位），不能超过 512 且必须是 8 的倍数
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 BLAKE2 值**

```
$ b2sum 1.txt                                          # 单个文件
364dce569668afdf7620e6a0a118c7ef718f4df1b419333cd41373d1af91ccc4f619b1a245c543e715bc6c3d95
871c0936f94b7875687bd09154e4af70cfd987  1.txt
$ b2sum -l 64 1.txt                                    # 结果保留 64 位
31e4af026ab94028  1.txt
$ b2sum 1.txt 2.txt 3.txt                              # 多个文件
```

```
364dce569668afdf7620e6a0a118c7ef718f4df1b419333cd41373d1af91ccc4f619b1a245c543e715bc6c3d95
871c0936f94b7875687bd09154e4af70cfd987  1.txt
cdf8234a9831d36804b2ccf54288127cbe6c876a9167d9f863334c9714ccfa8d73e9b644819e2ec37b17de1c5d
7381d5c7d1806f6bd021dcfda5def7e958e336  2.txt
b9fd7b32f6e9ecaa4f48b1bb22a160fe716c67835e8d2e67f0a3fb2f3b6cdd266499d2e1926dfae79c436202e8
136ff930e8bf384899dd7be7ae4ee7b12afa4e  3.txt
$ b2sum --tag 1.txt 2.txt 3.txt                    # BSD 风格
BLAKE2b (1.txt) =
364dce569668afdf7620e6a0a118c7ef718f4df1b419333cd41373d1af91ccc4f619b1a245c543e715bc6c3d95
871c0936f94b7875687bd09154e4af70cfd987
BLAKE2b (2.txt) =
cdf8234a9831d36804b2ccf54288127cbe6c876a9167d9f863334c9714ccfa8d73e9b644819e2ec37b17de1c5d
7381d5c7d1806f6bd021dcfda5def7e958e336
BLAKE2b (3.txt) =
b9fd7b32f6e9ecaa4f48b1bb22a160fe716c67835e8d2e67f0a3fb2f3b6cdd266499d2e1926dfae79c436202e8
136ff930e8bf384899dd7be7ae4ee7b12afa4e
```

**校验 BLAKE2 值**

```
$ b2sum 1.txt 2.txt 3.txt > sum.txt
$ b2sum -c sum.txt # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

# 推荐阅读

- Print or check BLAKE2 digests - GNU Coreutils
- b2sum(1) — Linux manual page

# 手册

```
B2SUM(1)                          User Commands                          B2SUM(1)

NAME
       b2sum - compute and check BLAKE2 message digest

SYNOPSIS
       b2sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check BLAKE2b (512-bit) checksums.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.
```

```
       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       -l, --length=BITS
              digest length in bits; must not exceed the max  for  the  blake2
              algorithm and must be a multiple of 8

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in RFC 7693.  When checking, the in-
       put should be a former output of this program.  The default mode is  to
       print a line with: checksum, a space, a character indicating input mode
       ('*'  for  binary,  ' ' for text or where binary is insignificant), and
       name for each FILE.

       Note: There is no difference between binary mode and text mode  on  GNU
       systems.

AUTHOR
       Written by Padraig Brady and Samuel Neves.

REPORTING BUGS
```

```
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       cksum(1)

       Full documentation <https://www.gnu.org/software/coreutils/b2sum>
       or available locally via: info '(coreutils) b2sum invocation'

GNU coreutils 9.4                April 2024                          B2SUM(1)
```

# Bash 的 base32 命令

```
base32 [OPTION]... [FILE]...
```

**功能**

进行 BASE32 编码或解码。

**类型**

可执行文件（`/usr/bin/base32`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-d`，`--decode` - 解码；不带此选项则为编码
  - `-i`，`--ignore-garbage` - 解码时忽略无效字符
  - `-w`，`--wrap=COLS` - 编码输出时一行的字符数；默认为 76，设为 0 则不换行
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

```
$ man base32 > 1.txt
$ cat 1.txt
BASE32(1)                       User Commands                       BASE32(1)

NAME
       base32 - base32 encode/decode data and print to standard output

SYNOPSIS
       base32 [OPTION]... [FILE]
...
$ base32 1.txt > 2.txt      # 编码
$ cat 2.txt
IJAVGRJTGIUDCKJAEAQCAIBAEAQCAIBAEAQCAIBAEAQCAIBAEAQCAVLTMVZCAQ3PNVWWC3TEOMQC
AIBAEAQCAIBAEAQCAIBAEAQCAIBAEAQCAICCIFJUKMZSFAYSSCQKJZAU2RIKEAQCAIBAEAQGEYLT
MUZTEIBNEBRGC43FGMZCAZLOMNXWIZJPMRSWG33EMUQGIYLUMEQGC3TEEBYHE2LOOQQHI3ZAON2G
C3TEMFZGIIDPOV2HA5LUBIFFGWKOJ5IFGSKTBIQCAIBAEAQCAYTBONSTGMRALNHVAVCJJ5HF2LRO
...
$ base32 -d 2.txt           # 解码
BASE32(1)                       User Commands                       BASE32(1)
```

```
NAME
       base32 - base32 encode/decode data and print to standard output

SYNOPSIS
       base32 [OPTION]... [FILE]
...
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| base64 | 进行 BASE64 编码或解码 |
| basenc | 进行 BASE2 BASE16 BASE32 BASE64 编码或解码 |

## 推荐阅读

- Transform data into printable data - GNU Coreutils
- base32(1) — Linux manual page

## 手册

```
BASE32(1)                          User Commands                          BASE32(1)

NAME
       base32 - base32 encode/decode data and print to standard output

SYNOPSIS
       base32 [OPTION]... [FILE]

DESCRIPTION
       Base32 encode or decode FILE, or standard input, to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -d, --decode
              decode data

       -i, --ignore-garbage
              when decoding, ignore non-alphabet characters
```

```
       -w, --wrap=COLS
              wrap encoded lines after COLS character (default 76).  Use 0  to
              disable line wrapping

       --help display this help and exit

       --version
              output version information and exit

       The  data are encoded as described for the base32 alphabet in RFC 4648.
       When decoding, the input may contain newlines in addition to the  bytes
       of  the formal base32 alphabet.  Use --ignore-garbage to attempt to re-
       cover from any other non-alphabet bytes in the encoded stream.

AUTHOR
       Written by Simon Josefsson.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       basenc(1)

       Full documentation <https://www.gnu.org/software/coreutils/base32>
       or available locally via: info '(coreutils) base32 invocation'

GNU coreutils 9.4                April 2024                        BASE32(1)
```

# Bash 的 base64 命令

```
base64 [OPTION]... [FILE]...
```

**功能**

进行 BASE64 编码或解码。

**类型**

可执行文件（`/usr/bin/base64`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-d`，`--decode` - 解码；不带此选项则为编码
  - `-i`，`--ignore-garbage` - 解码时忽略无效字符
  - `-w`，`--wrap=COLS` - 编码输出时一行的字符数；默认为 76，设为 0 则不换行
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

```
$ man base64 > 1.txt
$ cat 1.txt
BASE64(1)                       User Commands                       BASE64(1)

NAME
       base64 - base64 encode/decode data and print to standard output

SYNOPSIS
       base64 [OPTION]... [FILE]
...
$ base64 1.txt > 2.txt      # 编码
$ cat 2.txt
QkFTRTY0KDEpICAgICAgICAgICAgICAgICAgICAgICAgVXNlciBDb21tYW5kcyAgICAgICAgICAg
ICAgICAgICAgICAgICAgQkFTRTY0KDEpCgpOQU1FCiAgICAgICBiYXNlNjQgLSBiYXNlNjQgZW5jb2Rl
L2RlY29kZSBkYXRhIGFuZCBwcmludCB0byBzdGFuZGFyZCBvdXRwdXQKClNZTk9QU0lTCiAgICAgICAg
ICBiYXNlNjQgW09QVElPTl0uLi4gW0ZJTEVdCgpERVNDUklQVElPTgogICAgICAgQmFzZTY0IGVu
...
$ base64 -d 2.txt           # 解码
BASE64(1)                       User Commands                       BASE64(1)
```

```
NAME
       base64 - base64 encode/decode data and print to standard output

SYNOPSIS
       base64 [OPTION]... [FILE]
...
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| base32 | 进行 BASE32 编码或解码 |
| basenc | 进行 BASE2 BASE16 BASE32 BASE64 编码或解码 |

## 推荐阅读

- Transform data into printable data - GNU Coreutils
- base64(1) — Linux manual page

## 手册

```
BASE64(1)                        User Commands                        BASE64(1)

NAME
       base64 - base64 encode/decode data and print to standard output

SYNOPSIS
       base64 [OPTION]... [FILE]

DESCRIPTION
       Base64 encode or decode FILE, or standard input, to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -d, --decode
              decode data

       -i, --ignore-garbage
              when decoding, ignore non-alphabet characters
```

```
       -w, --wrap=COLS
              wrap encoded lines after COLS character (default 76).  Use 0  to
              disable line wrapping

       --help display this help and exit

       --version
              output version information and exit

       The  data are encoded as described for the base64 alphabet in RFC 4648.
       When decoding, the input may contain newlines in addition to the  bytes
       of  the formal base64 alphabet.  Use --ignore-garbage to attempt to re-
       cover from any other non-alphabet bytes in the encoded stream.

AUTHOR
       Written by Simon Josefsson.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       basenc(1)

       Full documentation <https://www.gnu.org/software/coreutils/base64>
       or available locally via: info '(coreutils) base64 invocation'

GNU coreutils 9.4                April 2024                        BASE64(1)
```

# Bash 的 basename 命令

```
basename OPTION... NAME [SUFFIX]
```

### 功能

从文件路径中去除目录和后缀。

### 类型

可执行文件（`/usr/bin/basename`），属于 coreutils。

### 参数

- `OPTION` 选项：
  - `-a`，`--multiple` - 输入多个文件名参数，而不是要删除的后缀
  - `-s`，`--suffix=SUFFIX` - 删除后缀的 `SUFFIX`
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NAME` - 文件路径；不要求文件存在
- `SUFFIX` - 要删除的后缀；同 `-s` 选项

## 示例

### 单个文件

```
$ basename /path/to/file.txt        # 去除目录
rgb.txt
$ basename /path/to/file.txt .txt   # 去除目录和后缀
```

### 多个文件

```
$ basename -a /path/to/file1.txt /path/to/file2.txt             # 去除目录
file1.txt
file2.txt
$ basename -s .txt -a /path/to/file1.txt /path/to/file2.txt     # 去除目录和后缀
file1
file2
```

- `-s` 选项参数必须在 `-a` 选项之前，否则 `-a` 会将 `-s` 和 `.txt` 也视为两个文件路径参数

## 相关命令

| 命令 | 说明 |
| --- | --- |
| dirname | 提取文件路径中的目录 |
| realpath | 解析路径 |

## 推荐阅读

- Strip directory and suffix from a file name - GNU Coreutils
- basename(1) — Linux manual page

## 手册

```
BASENAME(1)                      User Commands                     BASENAME(1)

NAME
       basename - strip directory and suffix from filenames

SYNOPSIS
       basename NAME [SUFFIX]
       basename OPTION... NAME...

DESCRIPTION
       Print  NAME  with  any leading directory components removed.  If speci-
       fied, also remove a trailing SUFFIX.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -a, --multiple
              support multiple arguments and treat each as a NAME

       -s, --suffix=SUFFIX
              remove a trailing SUFFIX; implies -a

       -z, --zero
              end each output line with NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

EXAMPLES
       basename /usr/bin/sort
              -> "sort"
```

```
       basename include/stdio.h .h
              -> "stdio"

       basename -s .h include/stdio.h
              -> "stdio"

       basename -a any/str1 any/str2
              -> "str1" followed by "str2"

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       dirname(1), readlink(1)

       Full documentation <https://www.gnu.org/software/coreutils/basename>
       or available locally via: info '(coreutils) basename invocation'

GNU coreutils 9.4                April 2024                      BASENAME(1)
```

# Bash 的 basenc 命令

```
basenc [OPTION]... [FILE]...
```

**功能**

对数据进行编码或解码。

**类型**

可执行文件（`/usr/bin/basenc`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--base64` - BASE64 编解码，同 base64 命令
  - `--base64url` - URL 安全的 BASE64 编解码
  - `--base32` - BASE32 编解码，同 base32 命令
  - `--base32hex` - 扩展十六进制字母 BASE32 编码
  - `--base16` - 十六进制编码
  - `--base2msbf` - 高位优先二进制编码
  - `--base2lsbf` - 低位优先二进制编码
  - `-d`，`--decode` - 解码；不带此选项则为编码
  - `-i`，`--ignore-garbage` - 解码时忽略无效字符
  - `-w`，`--wrap=COLS` - 编码输出时一行的字符数；默认为 76，设为 0 则不换行
  - `--z85` - 使用类 ASCII85 编码；编码时输入必须是 4 的倍数，解码时输入必须是 5 的倍数
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

```
$ man basenc > 1.txt
$ cat 1.txt
BASENC(1)                        User Commands                        BASENC(1)

NAME
       basenc - Encode/decode data and print to standard output

SYNOPSIS
```

```
       basenc [OPTION]... [FILE]
...
$ basenc --base16 1.txt > 2.txt     # 编码
$ cat 2.txt
424153454E432831292020202020202020202020202020202020202020205573657220
436F6D6D616E64732020202020202020202020202020202020202020204424153454E4328
31290A0A4E414D450A2020202020202020626173656E63202D20456E636F64652F6465636F6465
206461746120616E64207072696E7420746F207374616E64617264206F75747075740A0A5359
...
$ basenc --base16 -d 2.txt          # 解码
BASENC(1)                       User Commands                       BASENC(1)

NAME
       basenc - Encode/decode data and print to standard output

SYNOPSIS
       basenc [OPTION]... [FILE]
...
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| base32 | 进行 BASE32 编码或解码 |
| base64 | 进行 BASE64 编码或解码 |

## 推荐阅读

- Transform data into printable data - GNU Coreutils
- basenc(1) — Linux manual page

## 手册

```
BASENC(1)                       User Commands                       BASENC(1)

NAME
       basenc - Encode/decode data and print to standard output

SYNOPSIS
       basenc [OPTION]... [FILE]

DESCRIPTION
       basenc encode or decode FILE, or standard input, to standard output.
```

```
       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       --base64
              same as 'base64' program (RFC4648 section 4)

       --base64url
              file- and url-safe base64 (RFC4648 section 5)

       --base32
              same as 'base32' program (RFC4648 section 6)

       --base32hex
              extended hex alphabet base32 (RFC4648 section 7)

       --base16
              hex encoding (RFC4648 section 8)

       --base2msbf
              bit string with most significant bit (msb) first

       --base2lsbf
              bit string with least significant bit (lsb) first

       -d, --decode
              decode data

       -i, --ignore-garbage
              when decoding, ignore non-alphabet characters

       -w, --wrap=COLS
              wrap encoded lines after COLS character (default 76).  Use 0  to
              disable line wrapping

       --z85  ascii85-like encoding (ZeroMQ spec:32/Z85); when encoding, input
              length must be a multiple of 4; when decoding, input length must
              be a multiple of 5

       --help display this help and exit

       --version
              output version information and exit

       When  decoding, the input may contain newlines in addition to the bytes
       of the formal alphabet.  Use --ignore-garbage  to  attempt  to  recover
       from any other non-alphabet bytes in the encoded stream.

ENCODING EXAMPLES
              $ printf '\376\117\202' | basenc --base64
              /k+C
```

```
$ printf '\376\117\202' | basenc --base64url
_k-C

$ printf '\376\117\202' | basenc --base32
7ZHYE===

$ printf '\376\117\202' | basenc --base32hex
VP7O4===

$ printf '\376\117\202' | basenc --base16
FE4F82

$ printf '\376\117\202' | basenc --base2lsbf
011111111111001001000001

$ printf '\376\117\202' | basenc --base2msbf
111111100100111110000010

$ printf '\376\117\202\000' | basenc --z85
@.FaC
```

AUTHOR

Written by Simon Josefsson and Assaf Gordon.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/basenc>
or available locally via: info '(coreutils) basenc invocation'

# Bash 的 blkdiscard 命令

```
blkdiscard [OPTION]... DEVICE
```

**功能**

丢弃磁盘设备的区块。

**类型**

可执行文件（`/usr/sbin/blkdiscard`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-f`，`--force` - 禁用所有检查，强制丢弃
  - `-o`，`--offset offset` - 设备中要开始丢弃数据的字节偏移量；必须与区块大小对齐，默认为 0
  - `-l`，`--length length` - 要丢弃的字节数；必须与区块大小对齐，默认直到设备末尾
  - `-p`，`--step length` - 一次迭代中要丢弃的字节数；默认为通过一次 ioctl 调用丢弃所有字节
  - `-q`，`--quiet` - 不显示警告
  - `-s`，`--secure` - 执行安全擦除操作
  - `-z`，`--zeroout` - 用零填充区块而非丢弃
  - `-v`，`--verbose` - 显示详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 要操作的设备

## 示例

```
$ blkdiscard /dev/sdb                    # 丢弃整个磁盘
$ blkdiscard /dev/sdb1                   # 丢弃整个分区
$ sudo blkdiscard -o 1G -l 2G /dev/sdb   # 丢弃从 1G 开始的 2G 数据
```

## 推荐阅读

- blkdiscard(8) — Linux manual page

# 手册

```
BLKDISCARD(8)                System Administration                BLKDISCARD(8)

NAME
       blkdiscard - discard sectors on a device

SYNOPSIS
       blkdiscard [options] [-o offset] [-l length] device

DESCRIPTION
       blkdiscard is used to discard device sectors. This is useful for
       solid-state drivers (SSDs) and thinly-provisioned storage. Unlike
       fstrim(8), this command is used directly on the block device.

       By default, blkdiscard will discard all blocks on the device. Options
       may be used to modify this behavior based on range or size, as
       explained below.

       The device argument is the pathname of the block device.

       WARNING: All data in the discarded region on the device will be lost!

OPTIONS
       The offset and length arguments may be followed by the multiplicative
       suffixes KiB (=1024), MiB (=1024*1024), and so on for GiB, TiB, PiB,
       EiB, ZiB and YiB (the "iB" is optional, e.g., "K" has the same meaning
       as "KiB") or the suffixes KB (=1000), MB (=1000*1000), and so on for
       GB, TB, PB, EB, ZB and YB.

       -f, --force
           Disable all checking. Since v2.36 the block device is open in
           exclusive mode (O_EXCL) by default to avoid collision with mounted
           filesystem or another kernel subsystem. The --force option disables
           the exclusive access mode.

       -o, --offset offset
           Byte offset into the device from which to start discarding. The
           provided value must be aligned to the device sector size. The
           default value is zero.

       -l, --length length
           The number of bytes to discard (counting from the starting point).
           The provided value must be aligned to the device sector size. If
           the specified value extends past the end of the device, blkdiscard
           will stop at the device size boundary. The default value extends to
           the end of the device.

       -p, --step length
           The number of bytes to discard within one iteration. The default is
           to discard all by one ioctl call.
```

-q, --quiet
Suppress warning messages.

-s, --secure
Perform a secure discard. A secure discard is the same as a regular discard except that all copies of the discarded blocks that were possibly created by garbage collection must also be erased. This requires support from the device.

-z, --zeroout
Zero-fill rather than discard.

-v, --verbose
Display the aligned values of offset and length. If the --step option is specified, it prints the discard progress every second.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

EXIT STATUS

blkdiscard has the following exit status values:

0
success

1
failure; incorrect invocation, permissions or any other generic error

2
failure; since v2.39, the device does not support discard operation

AUTHORS

Lukas Czerner <lczerner@redhat.com>, Karel Zak <kzak@redhat.com>

SEE ALSO

fstrim(8)

REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The blkdiscard command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 blkid 命令

```
blkid [OPTION]... DEVICE...
```

**功能**

获取块设备的属性。

**类型**

可执行文件（`/usr/sbin/blkid`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-c`，`--cache-file cachefile` - 从 `cachefile` 读取缓存数据（而不是默认缓存文件）
  - `-d`，`--no-encoding` - 不要对 非打印字符 进行编码
  - `-D`，`--no-part-details` - 不要打印分区表中的信息
  - `-g`，`--garbage-collect` - 对缓存执行垃圾回收，以删除不再存在的设备
  - `-H`，`--hint setting` - 设置探测提示
  - `-i`，`--info` - 显示 I/O 限制
  - `-k`，`--list-filesystems` - 列出所有已知的文件系统和 RAID
  - `-l`，`--list-one` - 仅查找与 `--match-token` 选项指定的搜索参数匹配的设备
  - `-L`，`--label label` - 查找使用此文件系统标签的设备
  - `-n`，`--match-types list` - 将探测函数限制为指定的（以逗号分隔的）超级块类型（名称）列表
  - `-o`，`--output format` - 使用指定的输出格式
    - `full` - 打印所有标签；**默认**
    - `value` - 打印标签的值
    - `list` - 以用户友好的格式打印设备信息
    - `device` - 仅打印设备名称
    - `udev` - 打印键值对，以便轻松导入到 udev 环境中
    - `export` - 打印键值对以便轻松导入到环境变量中
  - `-O`，`--offset offset` - 在给定偏移量处进行探测
  - `-p`，`--probe` - 切换到低级超级块探测模式
  - `-s`，`--match-tag tag` - 对于每个（指定的）设备，仅显示与标签匹配的标签
  - `-S`，`--size size` - 覆盖设备/文件的大小

    - `-t`, `--match-token NAME=value` - 搜索名称为 NAME 且值为 value 的块设备，并显示找到的所有设备
    - `-u`, `--usages list` - 将探测函数限制为指定的（以逗号分隔的）“用法”类型列表
    - `-U`, `--uuid uuid` - 查找文件系统 UUID 为 `uuid` 的设备
    - `-h`, `--help` - 显示帮助
    - `-V`, `--version` - 显示版本
- `DEVICE` - 要操作的设备列表；省略时查看所有设备

## 示例

```
$ blkid
/dev/nvme0n1p1: UUID="02A0-97AC" BLOCK_SIZE="512" TYPE="vfat" PARTUUID="dfedd060-593f-
48e6-937d-a573d4a95dc5"
/dev/nvme0n1p2: UUID="8b2c9bcc-94f2-4cf2-a9d0-581e58f9d676" BLOCK_SIZE="4096" TYPE="ext4"
PARTUUID="4c3ba215-b476-4c6c-a37b-fb13e02d9d46"
```

## 推荐阅读

- blkid(8) — Linux manual page

## 手册

```
BLKID(8)                    System Administration                    BLKID(8)

NAME
       blkid - locate/print block device attributes

SYNOPSIS
       blkid --label label | --uuid uuid

       blkid [--no-encoding --garbage-collect --list-one --cache-file file]
       [--output format] [--match-tag tag] [--match-token NAME=value]
       [device...]

       blkid --probe [--offset offset] [--output format] [--size size]
       [--match-tag tag] [--match-types list] [--usages list]
       [--no-part-details] device...

       blkid --info [--output format] [--match-tag tag] device...

DESCRIPTION
       The blkid program is the command-line interface to working with the
       libblkid(3) library. It can determine the type of content (e.g.,
       filesystem or swap) that a block device holds, and also the attributes
       (tokens, NAME=value pairs) from the content metadata (e.g., LABEL or
```

UUID fields).

It is recommended to use lsblk(8) command to get information about block devices, or lsblk --fs to get an overview of filesystems, or findmnt(8) to search in already mounted filesystems.

lsblk(8) provides more information, better control on output formatting, easy to use in scripts and it does not require root permissions to get actual information. blkid reads information directly from devices and for non-root users it returns cached unverified information. blkid is mostly designed for system services and to test libblkid(3) functionality.

When device is specified, tokens from only this device are displayed. It is possible to specify multiple device arguments on the command line. If none is given, all partitions or unpartitioned devices which appear in /proc/partitions are shown, if they are recognized.

blkid has two main forms of operation: either searching for a device with a specific NAME=value pair, or displaying NAME=value pairs for one or more specified devices.

For security reasons blkid silently ignores all devices where the probing result is ambivalent (multiple colliding filesystems are detected). The low-level probing mode (-p) provides more information and extra exit status in this case. It’s recommended to use wipefs(8) to get a detailed overview and to erase obsolete stuff (magic strings) from the device.

OPTIONS

The size and offset arguments may be followed by the multiplicative suffixes like KiB (=1024), MiB (=1024*1024), and so on for GiB, TiB, PiB, EiB, ZiB and YiB (the "iB" is optional, e.g., "K" has the same meaning as "KiB"), or the suffixes KB (=1000), MB (=1000*1000), and so on for GB, TB, PB, EB, ZB and YB.

-c, --cache-file cachefile
Read from cachefile instead of reading from the default cache file (see the CONFIGURATION FILE section for more details). If you want to start with a clean cache (i.e., don’t report devices previously scanned but not necessarily available at this time), specify /dev/null.

-d, --no-encoding
Don’t encode non-printing characters. The non-printing characters are encoded by ^ and M- notation by default. Note that the --output udev output format uses a different encoding which cannot be disabled.

-D, --no-part-details
Don’t print information (PART_ENTRY_* tags) from partition table in low-level probing mode.

-g, --garbage-collect
    Perform a garbage collection pass on the blkid cache to remove
    devices which no longer exist.

-H, --hint setting
    Set probing hint. The hints are an optional way to force probing
    functions to check, for example, another location. The currently
    supported is "session_offset=number" to set session offset on
    multi-session UDF.

-i, --info
    Display information about I/O Limits (aka I/O topology). The
    'export' output format is automatically enabled. This option can be
    used together with the --probe option.

-k, --list-filesystems
    List all known filesystems and RAIDs and exit.

-l, --list-one
    Look up only one device that matches the search parameter specified
    with the --match-token option. If there are multiple devices that
    match the specified search parameter, then the device with the
    highest priority is returned, and/or the first device found at a
    given priority (but see below note about udev). Device types in
    order of decreasing priority are: Device Mapper, EVMS, LVM, MD, and
    finally regular block devices. If this option is not specified,
    blkid will print all of the devices that match the search
    parameter.

    This option forces blkid to use udev when used for LABEL or UUID
    tokens in --match-token. The goal is to provide output consistent
    with other utils (like mount(8), etc.) on systems where the same
    tag is used for multiple devices.

-L, --label label
    Look up the device that uses this filesystem label; this is equal
    to --list-one --output device --match-token LABEL=label. This
    lookup method is able to reliably use /dev/disk/by-label udev
    symlinks (dependent on a setting in /etc/blkid.conf). Avoid using
    the symlinks directly; it is not reliable to use the symlinks
    without verification. The --label option works on systems with and
    without udev.

    Unfortunately, the original blkid(8) from e2fsprogs uses the -L
    option as a synonym for -o list. For better portability, use -l -o
    device -t LABEL=label and -o list in your scripts rather than the
    -L option.

-n, --match-types list
    Restrict the probing functions to the specified (comma-separated)
    list of superblock types (names). The list items may be prefixed
    with "no" to specify the types which should be ignored. For

```
        example:

        blkid --probe --match-types vfat,ext3,ext4 /dev/sda1

        probes for vfat, ext3 and ext4 filesystems, and

        blkid --probe --match-types nominix /dev/sda1

        probes for all supported formats except minix filesystems. This
        option is only useful together with --probe.

    -o, --output format
        Use the specified output format. Note that the order of variables
        and devices is not fixed. See also option -s. The format parameter
        may be:

        full
            print all tags (the default)

        value
            print the value of the tags

        list
            print the devices in a user-friendly format; this output format
            is unsupported for low-level probing (--probe or --info).

            This output format is DEPRECATED in favour of the lsblk(8)
            command.

        device
            print the device name only; this output format is always
            enabled for the --label and --uuid options

        udev
            print key="value" pairs for easy import into the udev
            environment; the keys are prefixed by ID_FS_ or ID_PART_
            prefixes. The value may be modified to be safe for udev
            environment; allowed is plain ASCII, hex-escaping and valid
            UTF-8, everything else (including whitespaces) is replaced with
            '_'. The keys with _ENC postfix use hex-escaping for unsafe
            chars.

            The udev output returns the ID_FS_AMBIVALENT tag if more
            superblocks are detected, and ID_PART_ENTRY_* tags are always
            returned for all partitions including empty partitions.

            This output format is DEPRECATED.

        export
            print key=value pairs for easy import into the environment;
            this output format is automatically enabled when I/O Limits
            (--info option) are requested.
```

The non-printing characters are encoded by ^ and M- notation and all potentially unsafe characters are escaped.

-O, --offset offset
Probe at the given offset (only useful with --probe). This option can be used together with the --info option.

-p, --probe
Switch to low-level superblock probing mode (bypassing the cache).

Note that low-level probing also returns information about partition table type (PTTYPE tag) and partitions (PART_ENTRY_* tags). The tag names produced by low-level probing are based on names used internally by libblkid and it may be different than when executed without --probe (for example PART_ENTRY_UUID= vs PARTUUID=). See also --no-part-details.

-s, --match-tag tag
For each (specified) device, show only the tags that match tag. It is possible to specify multiple --match-tag options. If no tag is specified, then all tokens are shown for all (specified) devices. In order to just refresh the cache without showing any tokens, use --match-tag none with no other options.

-S, --size size
Override the size of device/file (only useful with --probe).

-t, --match-token NAME=value
Search for block devices with tokens named NAME that have the value value, and display any devices which are found. Common values for NAME include TYPE, LABEL, and UUID. If there are no devices specified on the command line, all block devices will be searched; otherwise only the specified devices are searched.

-u, --usages list
Restrict the probing functions to the specified (comma-separated) list of "usage" types. Supported usage types are: filesystem, raid, crypto and other. The list items may be prefixed with "no" to specify the usage types which should be ignored. For example:

blkid --probe --usages filesystem,other /dev/sda1

probes for all filesystem and other (e.g., swap) formats, and

blkid --probe --usages noraid /dev/sda1

probes for all supported formats except RAIDs. This option is only useful together with --probe.

-U, --uuid uuid
Look up the device that uses this filesystem uuid. For more details see the --label option.

```
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       If the specified device or device addressed by specified token (option
       --match-token) was found and it's possible to gather any information
       about the device, an exit status 0 is returned. Note the option
       --match-tag filters output tags, but it does not affect exit status.

       If the specified token was not found, or no (specified) devices could
       be identified, or it is impossible to gather any information about the
       device identifiers or device content an exit status of 2 is returned.

       For usage or other errors, an exit status of 4 is returned.

       If an ambivalent probing result was detected by low-level probing mode
       (-p), an exit status of 8 is returned.

CONFIGURATION FILE
       The standard location of the /etc/blkid.conf config file can be
       overridden by the environment variable BLKID_CONF. The following
       options control the libblkid library:

       SEND_UEVENT=<yes|not>
           Sends uevent when /dev/disk/by-{label,uuid,partuuid,partlabel}/
           symlink does not match with LABEL, UUID, PARTUUID or PARTLABEL on
           the device. Default is "yes".

       CACHE_FILE=<path>
           Overrides the standard location of the cache file. This setting can
           be overridden by the environment variable BLKID_FILE. Default is
           /run/blkid/blkid.tab, or /etc/blkid.tab on systems without a /run
           directory.

       EVALUATE=<methods>
           Defines LABEL and UUID evaluation method(s). Currently, the
           libblkid library supports the "udev" and "scan" methods. More than
           one method may be specified in a comma-separated list. Default is
           "udev,scan". The "udev" method uses udev /dev/disk/by-* symlinks
           and the "scan" method scans all block devices from the
           /proc/partitions file.

ENVIRONMENT
       Setting LIBBLKID_DEBUG=all enables debug output.

AUTHORS
       blkid was written by Andreas Dilger for libblkid and improved by
       Theodore Ts'o and Karel Zak.

SEE ALSO
```

```
       libblkid(3), findfs(8), lsblk(8), wipefs(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The blkid command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                          BLKID(8)
```

# Bash 的 blkzone 命令

```
blkzone CMD [OPTION]... DEVICE
```

**功能**

管理和查看分区式（zoned）块设备的 zone 信息。

> 这种设备常见于：
>
> - SMR 硬盘（Shingled Magnetic Recording）
> - ZNS SSD（Zoned Namespace SSD）
>
> ```
> 普通硬盘
>    _______       _______       _______
> __/       _____/       _____/       __
>
> 叠瓦盘
>    _______       _______       _______
> __/       _____/       _____/       __
>      _______       _______       _______
>   __/       _____/       _____/       __
> ```
>
> 在这类设备上，随机写会覆相邻的磁道，磁盘被划分为多个 zone（区域），写入必须按顺序进行，因此需要专门工具管理。

**类型**

可执行文件（`/usr/sbin/blkzone`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-o`, `--offset sector` - 起始扇区
  - `-l`, `--length sectors` - 长度（以扇区为单位）
  - `-c`, `--count count` - 命令可操作的最大区域数
  - `-f`, `--force` - 强制执行更改系统所用块设备上的区域状态的命令
  - `-v`, `--verbose` - 显示报告中返回的区域数量或重置扇区范围
  - `--version` - 显示版本
  - `-h`, `--help` - 显示帮助
  - `-V`, `--version` - 显示版本

- `DEVICE` - 要操作的设备

## 区域命令

| 命令 | 说明 |
| --- | --- |
| `report` | 报告 zone 信息 |
| `capacity` | 报告容量信息 |
| `reset` | 将 zone 的写指针恢复到起始位置 |
| `open` | 打开 zone，使其可以写入 |
| `close` | 关闭 zone |
| `finish` | 标记 zone 写入完成，写指针移动到末尾 |

## 字段说明

| 字段 | 说明 |
| --- | --- |
| `start` | zone 的起始扇区 |
| `len` | zone 的大小 |
| `cap` | zone 的容量 |
| `wptr` | 写指针 |
| `reset` | 建议重置写指针 |
| `non-seq` | 非顺序写入资源处于活动状态 |
| `cond` | zone 的状态 |
| `type` | zone 的类型 |

## zone 的状态

| 状态 | 说明 |
| --- | --- |
| `cl` | 关闭 |
| `nw` | 未写入指针 |
| `em` | 空 |
| `fu` | 满 |

| 状态 | 说明 |
| --- | --- |
| `oe` | 明确打开 |
| `oi` | 隐式打开 |
| `ol` | 离线 |
| `ro` | 只读 |
| `x?` | 保留状态 (不应报告) |

## 示例

### 查看 zone 状态

```
$ sudo blkzone report /dev/sdb
zone 0
start: 0
size : 256MB
write pointer: 128MB
```

- 写指针在 128MB 处，只能在写指针之后写入

### 写入 zone

```
$ sudo blkzone open -o 0 -l 256M /dev/sdb       # 打开一块区域（可能包含多个 zone）进行顺序写入
$ sudo dd if=data.bin of=/dev/sdb bs=4K         # 写入数据
$ sudo blkzone finish -o 0 -l 256M /dev/sdb     # 标记完成
```

## 推荐阅读

- blkzone(8) — Linux manual page

## 手册

```
BLKZONE(8)                  System Administration                  BLKZONE(8)

NAME
       blkzone - run zone command on a device

SYNOPSIS
       blkzone command [options] device
```

DESCRIPTION

blkzone is used to run zone command on device that support the Zoned Block Commands (ZBC) or Zoned-device ATA Commands (ZAC). The zones to operate on can be specified using the offset, count and length options.

The device argument is the pathname of the block device.

COMMANDS

report

The command blkzone report is used to report device zone information.

By default, the command will report all zones from the start of the block device. Options may be used to modify this behavior, changing the starting zone or the size of the report, as explained below.

Report output:

| | |
|---|---|
| start | Zone start sector |
| len | Zone length in number of sectors |
| cap | Zone capacity in number of sectors |
| wptr | Zone write pointer position |
| reset | Reset write pointer recommended |
| non-seq | Non-sequential write resources active |
| cond | Zone condition |
| type | Zone type |

Zone conditions:

| | |
|---|---|
| cl | Closed |

```
| nw | Not write pointer        |
|----|--------------------------|
|    |                          |
| em | Empty                    |
|----|--------------------------|
|    |                          |
| fu | Full                     |
|----|--------------------------|
|    |                          |
| oe | Explicitly opened        |
|----|--------------------------|
|    |                          |
| oi | Implicitly opened        |
|----|--------------------------|
|    |                          |
| ol | Offline                  |
|----|--------------------------|
|    |                          |
| ro | Read only                |
|----|--------------------------|
|    |                          |
| x? | Reserved conditions      |
|    | (should not be reported) |
└────┴──────────────────────────┘
```

capacity
: The command blkzone capacity is used to report device capacity information.

  By default, the command will report the sum, in number of sectors, of all zone capacities on the device. Options may be used to modify this behavior, changing the starting zone or the size of the report, as explained below.

reset
: The command blkzone reset is used to reset one or more zones. Unlike sg_reset_wp(8), this command operates from the block layer and can reset a range of zones.

open
: The command blkzone open is used to explicitly open one or more zones. Unlike sg_zone(8), open action, this command operates from the block layer and can open a range of zones.

close
: The command blkzone close is used to close one or more zones. Unlike sg_zone(8), close action, this command operates from the block layer and can close a range of zones.

finish
: The command blkzone finish is used to finish (transition to full condition) one or more zones. Unlike sg_zone(8), finish action, this command operates from the block layer and can finish a range of zones.

By default, the reset, open, close and finish commands will operate from the zone at device sector 0 and operate on all zones. Options may be used to modify this behavior as explained below.

## OPTIONS

The offset and length option arguments may be followed by the multiplicative suffixes KiB (=1024), MiB (=1024*1024), and so on for GiB, TiB, PiB, EiB, ZiB and YiB (the "iB" is optional, e.g., "K" has the same meaning as "KiB") or the suffixes KB (=1000), MB (=1000*1000), and so on for GB, TB, PB, EB, ZB and YB. Additionally, the 0x prefix can be used to specify offset and length in hex.

-o, --offset sector
: The starting zone specified as a sector offset. The provided offset in sector units (512 bytes) should match the start of a zone. The default value is zero.

-l, --length sectors
: The maximum number of sectors the command should operate on. The default value is the number of sectors remaining after offset. This option cannot be used together with the option --count.

-c, --count count
: The maximum number of zones the command should operate on. The default value is the number of zones starting from offset. This option cannot be used together with the option --length.

-f, --force
: Enforce commands to change zone status on block devices used by the system.

-v, --verbose
: Display the number of zones returned in the report or the range of sectors reset.

-h, --help
: Display help text and exit.

-V, --version
: Print version and exit.

## AUTHORS

Shaun Tancheff <shaun@tancheff.com>, Karel Zak <kzak@redhat.com>

## SEE ALSO

sg_rep_zones(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

```
       The blkzone command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                         BLKZONE(8)
```

# Bash 的 blockdev 命令

```
blockdev [OPTION]... COMMAND... DEVICE...
```

**功能**

从命令行对块设备调用 `ioctl`。

**类型**

可执行文件（`/usr/sbin/blockdev`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-q` - 不打印任何信息
  - `-v` - 打印详细信息
  - `--report` - 打印指定设备的报告
  - `-h`, `--help` - 显示帮助
  - `-V`, `--version` - 显示版本
- `COMMAND` - 要执行的命令
  - `--flushbufs` - 冲洗缓冲区
  - `--getalignoff` - 获取设备对齐偏移量（alignment offset），用于确定读写是否对齐
  - `--getbsz` - 获取设备的块大小；单位是字节
  - `--getdiscardzeroes` - 检查设备在执行 discard 操作时是否自动清零
  - `--getdiskseq` - 获取设备的序列号（disk sequence number）
  - `--getfra` - 获取前置读请求数量（front read ahead）
  - `--getiomin` - 获取最小 I/O 单位（最小 I/O size）
  - `--getioopt` - 获取设备的最优 I/O 大小
  - `--getmaxsect` - 获取设备最大扇区数（每次 I/O 可处理的最大扇区数）
  - `--getpbsz` - 获取物理块大小（physical block size）
  - `--getra` - 获取读提前（read ahead）扇区数
  - `--getro` - 检查设备是否只读（read-only）
  - `--getsize64` - 获取设备大小（以字节为单位，64 位）
  - `--getsize` - 获取设备大小（以 512 字节扇区为单位）
  - `--getss` - 获取设备的扇区大小（sector size）
  - `--getsz` - 获取设备扇区数量（与 --getsize 类似）
  - `--rereadpt` - 重新读取分区表（re-read partition table）

    - `--setbsz bytes` - 设置设备块大小（block size）
    - `--setfra sectors` - 设置前置读（front read ahead）扇区数；以 512 字节的扇区为单位
    - `--setra sectors` - 设置预读（read ahead）扇区数；以 512 字节的扇区为单位
    - `--setro` - 将设备设置为只读
    - `--setrw` - 将设备设置为可读写
- `DEVICE` - 要操作的设备

## 示例

### 查看设备大小（扇区数和字节数）

```
# 查看 /dev/sda 的扇区数
sudo blockdev --getsz /dev/sda

# 查看 /dev/sda 的大小（字节）
sudo blockdev --getsize64 /dev/sda
```

### 查看设备的扇区大小和块大小

```
# 扇区大小
sudo blockdev --getss /dev/sda

# 块大小
sudo blockdev --getbsz /dev/sda
```

### 检查设备是否只读

```
sudo blockdev --getro /dev/sda
# 输出 1 表示只读，0 表示可读写
```

### 设置设备为只读或可读写

```
# 设置只读
sudo blockdev --setro /dev/sda

# 设置可读写
sudo blockdev --setrw /dev/sda
```

### 查看和设置读提前（read-ahead）

```
# 查看 read-ahead 扇区数
sudo blockdev --getra /dev/sda
```

```
# 设置 read-ahead 为 256 扇区
sudo blockdev --setra 256 /dev/sda
```

**冲洗缓冲区（flush buffers）**

```
sudo blockdev --flushbufs /dev/sda
```

- 这个命令会把所有缓存的写操作刷新到磁盘上，类似 `sync`，确保数据安全。

**重新读取分区表**

```
sudo blockdev --rereadpt /dev/sda
```

- 当你修改了分区表（如用 `fdisk` 或 `parted`）后，可以用这个命令让内核重新识别分区。

**综合示例：设置 read-ahead、查看大小并刷新缓冲区**

```
sudo blockdev --setra 128 /dev/sda
sudo blockdev --getsize64 /dev/sda
sudo blockdev --flushbufs /dev/sda
```

## 推荐阅读

- blockdev(8) — Linux manual page

## 手册

```
BLOCKDEV(8)                 System Administration                BLOCKDEV(8)

NAME
       blockdev - call block device ioctls from the command line

SYNOPSIS
       blockdev [-q] [-v] command [command...] device [device...]

       blockdev --report [device...]

       blockdev -h|-V

DESCRIPTION
       The utility blockdev allows one to call block device ioctls from the
       command line.

OPTIONS
       -q
           Be quiet.
```

```
       -v
           Be verbose.

       --report
           Print a report for the specified device. It is possible to give
           multiple devices. If none is given, all devices which appear in
           /proc/partitions are shown. Note that the partition StartSec is in
           512-byte sectors.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

COMMANDS
       It is possible to give multiple devices and multiple commands.

       --flushbufs
           Flush buffers.

       --getalignoff
           Get alignment offset.

       --getbsz
           Print the blocksize in bytes. This size does not describe device
           topology. It's the size used internally by the kernel and it may be
           modified (for example) by filesystem driver on mount.

       --getdiscardzeroes
           Get discard zeroes support status.

       --getdiskseq
           Get disk sequence number.

       --getfra
           Get filesystem readahead in 512-byte sectors.

       --getiomin
           Get minimum I/O size.

       --getioopt
           Get optimal I/O size.

       --getmaxsect
           Get max sectors per request.

       --getpbsz
           Get physical block (sector) size.

       --getra
           Print readahead (in 512-byte sectors).
```

```
       --getro
           Get read-only. Print 1 if the device is read-only, 0 otherwise.

       --getsize64
           Print device size in bytes.

       --getsize
           Print device size (32-bit!) in sectors. Deprecated in favor of the
           --getsz option.

       --getss
           Print logical sector size in bytes - usually 512.

       --getsz
           Get size in 512-byte sectors.

       --rereadpt
           Reread partition table

       --setbsz bytes
           Set blocksize. Note that the block size is specific to the current
           file descriptor opening the block device, so the change of block
           size only persists for as long as blockdev has the device open, and
           is lost once blockdev exits.

       --setfra sectors
           Set filesystem readahead (same as --setra on 2.6 kernels).

       --setra sectors
           Set readahead (in 512-byte sectors).

       --setro
           Set read-only. The currently active access to the device may not be
           affected by the change. For example, a filesystem already mounted
           in read-write mode will not be affected. The change applies after
           remount.

       --setrw
           Set read-write.

AUTHORS
       blockdev was written by Andries E. Brouwer and rewritten by Karel Zak.

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The blockdev command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.
```

# Bash 的 cat 命令

```
cat [OPTION]... [FILE]...
```

**功能**

拼接文件内容并打印到标准输出。

**类型**

可执行文件（`/usr/bin/cat`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-A`，`--show-all` - 显示所有特殊符号；同 `-vET`
  - `-b`，`--number-nonblank` - 非空行显示行号
  - `-e` - 行末显示 `$`，其它不可打印字符显示为 `^` 或 `M-` 前缀形式；同 `-vE`
  - `-E`，`--show-ends` - 行末显示 `$`
  - `-n`，`--number` - 显示行号
  - `-s`，`--squeeze-blank` - 抑制重复的空行
  - `-t` - 制表符显示为 `^I`，其它不可打印字符显示为 `^` 或 `M-` 前缀形式；同 `-vT`
  - `-T`，`--show-tabs` - 制表符显示为 `^I`
  - ~~`-u` - 不进行缓冲；已废弃，被忽略~~
  - `-v`，`--show-nonprinting` - 不可打印字符显示为 `^` 或 `M-` 前缀形式，参考 ASCII 编码表
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 要拼接的文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

**查看文件**

```
$ cat file          # 查看 file
$ cat file1 file2   # 查看 file1 和 file2
```

**写文件**

```
$ cat > file        # 将标准输入写入 file
```

- 将标准输入覆盖写入 `file` 文件
- 按 `Ctrl + D` 结束

```
$ cat >> file       # 将标准输入写入 file
```

- 将标准输入追加写入 `file` 文件
- 按 `Ctrl + D` 结束

```
$ cat src1 src2 > dest    # 将 src1 和 src2 拼接，写入 dest
```

- 拼接 `src1` 和 `src2`，覆盖写入 `dest`

## 相关命令

| 命令 | 说明 |
| --- | --- |
| tac | 逆序拼接文件内容并打印到标准输出 |
| rev | 反转行中的字符 |

## 推荐阅读

- Concatenate and write files - GNU Coreutils
- cat(1) — Linux manual page

## 手册

```
CAT(1)                          User Commands                          CAT(1)

NAME
       cat - concatenate files and print on the standard output

SYNOPSIS
       cat [OPTION]... [FILE]...

DESCRIPTION
       Concatenate FILE(s) to standard output.

       With no FILE, or when FILE is -, read standard input.

       -A, --show-all
```

```
              equivalent to -vET

       -b, --number-nonblank
              number nonempty output lines, overrides -n

       -e     equivalent to -vE

       -E, --show-ends
              display $ at end of each line

       -n, --number
              number all output lines

       -s, --squeeze-blank
              suppress repeated empty output lines

       -t     equivalent to -vT

       -T, --show-tabs
              display TAB characters as ^I

       -u     (ignored)

       -v, --show-nonprinting
              use ^ and M- notation, except for LFD and TAB

       --help display this help and exit

       --version
              output version information and exit

EXAMPLES
       cat f - g
              Output f's contents, then standard input, then g's contents.

       cat    Copy standard input to standard output.

AUTHOR
       Written by Torbjorn Granlund and Richard M. Stallman.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       tac(1)
```

```
       Full documentation <https://www.gnu.org/software/coreutils/cat>
       or available locally via: info '(coreutils) cat invocation'

GNU coreutils 9.4                April 2024                          CAT(1)
```

# Bash 的 chcon 命令

```
chcon [OPTION]... CONTEXT FILE...
```

**功能**

修改文件的 SELinux 安全上下文。

**类型**

可执行文件（`/usr/bin/chcon`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--dereference` - 影响符号链接引用的源文件，而非符号链接本身（默认）
  - `-h`，`--no-dereference` - 影响符号链接本身，而不是其引用的源文件
  - `-u`，`--user=USER` - 在目标安全上下文中设置用户为 `USER`
  - `-r`，`--role=ROLE` - 在目标安全上下文中设置角色为 `ROLE`
  - `-t`，`--type=TYPE` - 在目标安全上下文中设置类型为 `TYPE`
  - `-l`，`--range=RANGE` - 在目标安全上下文中设置范围为 `RANGE`
  - `--no-preserve-root` - 不保护 `/` 目录
  - `--preserve-root` - 保护 `/` 目录
  - `--reference=RFILE` - 引用 `RFILE` 的安全上下文
  - `-R`，`--recursive` - 对目录进行递归操作
  - `-v`，`--verbose` - 打印详细信息
  - `-H` - 使用 `-R` 选项进行递归操作时，指向目录的符号链接只有作为命令行参数时才遍历
  - `-L` - 使用 `-R` 选项进行递归操作时，遍历所有指向目录的符号链接
  - `-P` - 使用 `-R` 选项进行递归操作时，不遍历任何符号链接（默认）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CONTEXT` - 安全上下文；通常格式为 `USER:ROLE:TYPE:LEVEL`
- `FILE` - 文件路径

## 示例

```
$ sudo chcon USER:ROLE:TYPE:RANGE 1.txt     # 设置安全上下文
$ ls -Z 1.txt                               # 查看安全上下文
```

```
USER:ROLE:TYPE:RANGE 1.txt
$ sudo chcon --reference=1.txt 2.txt        # 将 2.txt 的安全上下文设为和 1.txt 一样
$ ls -Z 2.txt
USER:ROLE:TYPE:RANGE 2.txt
```

## 推荐阅读

- Change SELinux context of file - GNU Coreutils
- chcon(1) — Linux manual page

## 手册

```
CHCON(1)                          User Commands                          CHCON(1)

NAME
       chcon - change file security context

SYNOPSIS
       chcon [OPTION]... CONTEXT FILE...
       chcon [OPTION]... [-u USER] [-r ROLE] [-l RANGE] [-t TYPE] FILE...
       chcon [OPTION]... --reference=RFILE FILE...

DESCRIPTION
       Change  the  SELinux  security  context  of each FILE to CONTEXT.  With
       --reference, change the security context of each FILE to that of RFILE.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       --dereference
              affect the referent of each symbolic link (this is the default),
              rather than the symbolic link itself

       -h, --no-dereference
              affect symbolic links instead of any referenced file

       -u, --user=USER
              set user USER in the target security context

       -r, --role=ROLE
              set role ROLE in the target security context

       -t, --type=TYPE
              set type TYPE in the target security context

       -l, --range=RANGE
              set range RANGE in the target security context
```

```
       --no-preserve-root
              do not treat '/' specially (the default)

       --preserve-root
              fail to operate recursively on '/'

       --reference=RFILE
              use  RFILE's  security  context rather than specifying a CONTEXT
              value

       -R, --recursive
              operate on files and directories recursively

       -v, --verbose
              output a diagnostic for every file processed

       The following options modify how a hierarchy is traversed when  the  -R
       option  is also specified.  If more than one is specified, only the fi-
       nal one takes effect.

       -H     if a command line argument is a symbolic link  to  a  directory,
              traverse it

       -L     traverse every symbolic link to a directory encountered

       -P     do not traverse any symbolic links (default)

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Russell Coker and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/chcon>
       or available locally via: info '(coreutils) chcon invocation'

GNU coreutils 9.4                April 2024                          CHCON(1)
```

# Bash 的 chgrp 命令

```
chgrp [OPTION]... GROUP FILE...
```

**功能**

修改文件的所属组。

**类型**

可执行文件（`/usr/bin/chgrp`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--changes` - 仅对发生变化的文件打印详细信息
  - `-f`，`--silent`，`--quiet` - 忽略大部分错误信息
  - `-v`，`--verbose` - 打印详细信息
  - `--dereference` - 影响符号链接引用的源文件，而非符号链接本身（默认）
  - `-h`，`--no-dereference` - 影响符号链接本身，而不是其引用的源文件
  - `--no-preserve-root` - 不保护 `/` 目录
  - `--preserve-root` - 保护 `/` 目录
  - `--reference=RFILE` - 引用 `RFILE` 文件的配置
  - `-R`，`--recursive` - 对目录进行递归操作
  - `-H` - 使用 `-R` 选项进行递归操作时，指向目录的符号链接只有作为命令行参数时才遍历
  - `-L` - 使用 `-R` 选项进行递归操作时，遍历所有指向目录的符号链接
  - `-P` - 使用 `-R` 选项进行递归操作时，不遍历任何符号链接（默认）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `GROUP` - 要设为的所属组
- `FILE` - 文件列表

## 示例

**使用组名**

```
$ chgrp primers file.txt                    # 将 file 的所属组设为 primers
```

### 使用组 ID

```
$ chgrp :1001 file.txt                    # 将 file 的所属组 ID 设为 1001
```

### 其它

```
$ chgrp --reference=2.txt 1.txt           # 将 1.txt 的所属组设为和 2.txt 一样
$ chgrp primers -R dir                    # 递归设置目录 dir 中的文件
```

### 修改根目录（危险！）

```
$ chgrp primers --no-preserve-root -R /   # 危险！千万不要执行此命令！
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| chown | 修改文件的所有者和所属组 |
| chmod | 修改文件的访问权限 |

## 推荐阅读

- Change group ownership - GNU Coreutils
- chgrp(1) — Linux manual page

## 手册

```
CHGRP(1)                          User Commands                          CHGRP(1)

NAME
       chgrp - change group ownership

SYNOPSIS
       chgrp [OPTION]... GROUP FILE...
       chgrp [OPTION]... --reference=RFILE FILE...

DESCRIPTION
       Change  the  group of each FILE to GROUP.  With --reference, change the
       group of each FILE to that of RFILE.

       -c, --changes
              like verbose but report only when a change is made
```

```
       -f, --silent, --quiet
              suppress most error messages

       -v, --verbose
              output a diagnostic for every file processed

       --dereference
              affect the referent of each symbolic link (this is the default),
              rather than the symbolic link itself

       -h, --no-dereference
              affect symbolic links instead of  any  referenced  file  (useful
              only on systems that can change the ownership of a symlink)

       --no-preserve-root
              do not treat '/' specially (the default)

       --preserve-root
              fail to operate recursively on '/'

       --reference=RFILE
              use  RFILE's group rather than specifying a GROUP.  RFILE is al-
              ways dereferenced if a symbolic link.

       -R, --recursive
              operate on files and directories recursively

       The following options modify how a hierarchy is traversed when  the  -R
       option  is also specified.  If more than one is specified, only the fi-
       nal one takes effect.

       -H     if a command line argument is a symbolic link  to  a  directory,
              traverse it

       -L     traverse every symbolic link to a directory encountered

       -P     do not traverse any symbolic links (default)

       --help display this help and exit

       --version
              output version information and exit

EXAMPLES
       chgrp staff /u
              Change the group of /u to "staff".

       chgrp -hR staff /u
              Change the group of /u and subfiles to "staff".

AUTHOR
       Written by David MacKenzie and Jim Meyering.
```

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

chown(1), chown(2)

Full documentation <https://www.gnu.org/software/coreutils/chgrp>
or available locally via: info '(coreutils) chgrp invocation'

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 chmod 命令

```
chmod [OPTION]... MODE[,MODE]... FILE...
```

**功能**

修改文件的访问权限。

**类型**

可执行文件（`/usr/bin/chmod`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--changes` - 仅对发生变化的文件打印详细信息
  - `-f`，`--silent`，`--quiet` - 忽略大部分错误信息
  - `-v`，`--verbose` - 打印详细信息
  - `--no-preserve-root` - 不保护 `/` 目录
  - `--preserve-root` - 保护 `/` 目录
  - `--reference=RFILE` - 引用 `RFILE` 文件的配置
  - `-R`，`--recursive` - 对目录进行递归操作
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `MODE` - 访问权限
- `FILE` - 文件列表

## 权限的表示方法

| 符号 | 数字 | 含义 |
| --- | --- | --- |
| u | - | 所有者 (user) |
| g | - | 组 (group) |
| o | - | 其他人 (others) |
| a | - | 所有 (all) |
| + | - | 添加权限 |

| 符号 | 数字 | 含义 |
| --- | --- | --- |
| - | - | 移除权限 |
| = | - | 设置为指定权限 |
| r | 4 | 读（read） |
| w | 2 | 写（write） |
| x | 1 | 执行（execute） |
| s | 4 | setuid |
| s | 2 | setgid |
| t | 1 | 粘滞位（sticky bit） |

- `setuid` - 用户执行该文件时，可以将用户 ID 设为该文件的所有者 ID
- `setgid` - 用户执行该文件时，可以将组 ID 设为该文件的所属组 ID
- 粘滞位 - 目录中的文件只能被其所有者删除

## 示例

### 符号操作

```
$ chmod u=rwx file  # 对所有者设为可读、可写、可执行
$ chmod g+rx file   # 对所属组添加读取、执行权限
$ chmod o= file     # 对其他人设为没有任何权限
```

### setuid 和 setgid

```
$ chmod u+s file    # 添加 setuid
$ chmod g+s file    # 添加 setgid
$ chmod u-s file    # 取消 setuid
$ chmod g-s file    # 取消 setgid
```

### 粘滞位

```
$ chmod +t dir      # 添加粘滞位
$ chmod -t dir      # 取消粘滞位
```

### 数字操作

```
$ chmod 750 file    # 所有者=7 所属组=5 其他=0
```

```
$ chmod 7750 dir    # 特殊=7 所有者=7 所属组=5 其他=0
```

- 特殊权限为 `7 = 4 + 2 + 1` 即 `setuid`、`setgid`、`sticky` 均设置
- 所有者权限为 `7 = 4 + 2 + 1` 即可读、可写、可执行
- 所属组权限为 `5 = 4 + 1` 即可读、可执行，不可写
- 其他人权限为 `0` 即不可读、不可写、不可执行

## 相关命令

| 命令 | 说明 |
| --- | --- |
| chown | 修改文件的所有者和所属组 |
| chgrp | 修改文件的所有者 |

## 推荐阅读

- Change access permissions - GNU Coreutils
- chmod(1) — Linux manual page

## 手册

```
CHMOD(1)                          User Commands                          CHMOD(1)

NAME
       chmod - change file mode bits

SYNOPSIS
       chmod [OPTION]... MODE[,MODE]... FILE...
       chmod [OPTION]... OCTAL-MODE FILE...
       chmod [OPTION]... --reference=RFILE FILE...

DESCRIPTION
       This manual page documents the GNU version of chmod.  chmod changes the
       file  mode  bits of each given file according to mode, which can be ei-
       ther a symbolic representation of changes to make, or an  octal  number
       representing the bit pattern for the new mode bits.

       The  format  of a symbolic mode is [ugoa...][[-+=][perms...]...], where
       perms is either zero or more letters from the set rwxXst, or  a  single
       letter  from  the set ugo.  Multiple symbolic modes can be given, sepa-
       rated by commas.

       A combination of the letters ugoa controls which users' access  to  the
       file  will  be  changed:  the  user who owns it (u), other users in the
```

file's group (g), other users not in the file's group (o), or all users (a). If none of these are given, the effect is as if (a) were given, but bits that are set in the umask are not affected.

The operator + causes the selected file mode bits to be added to the existing file mode bits of each file; - causes them to be removed; and = causes them to be added and causes unmentioned bits to be removed except that a directory's unmentioned set user and group ID bits are not affected.

The letters rwxXst select file mode bits for the affected users: read (r), write (w), execute (or search for directories) (x), execute/search only if the file is a directory or already has execute permission for some user (X), set user or group ID on execution (s), restricted deletion flag or sticky bit (t). Instead of one or more of these letters, you can specify exactly one of the letters ugo: the permissions granted to the user who owns the file (u), the permissions granted to other users who are members of the file's group (g), and the permissions granted to users that are in neither of the two preceding categories (o).

A numeric mode is from one to four octal digits (0-7), derived by adding up the bits with values 4, 2, and 1. Omitted digits are assumed to be leading zeros. The first digit selects the set user ID (4) and set group ID (2) and restricted deletion or sticky (1) attributes. The second digit selects permissions for the user who owns the file: read (4), write (2), and execute (1); the third selects permissions for other users in the file's group, with the same values; and the fourth for other users not in the file's group, with the same values.

chmod never changes the permissions of symbolic links; the chmod system call cannot change their permissions. This is not a problem since the permissions of symbolic links are never used. However, for each symbolic link listed on the command line, chmod changes the permissions of the pointed-to file. In contrast, chmod ignores symbolic links encountered during recursive directory traversals.

SETUID AND SETGID BITS

chmod clears the set-group-ID bit of a regular file if the file's group ID does not match the user's effective group ID or one of the user's supplementary group IDs, unless the user has appropriate privileges. Additional restrictions may cause the set-user-ID and set-group-ID bits of MODE or RFILE to be ignored. This behavior depends on the policy and functionality of the underlying chmod system call. When in doubt, check the underlying system behavior.

For directories chmod preserves set-user-ID and set-group-ID bits unless you explicitly specify otherwise. You can set or clear the bits with symbolic modes like u+s and g-s. To clear these bits for directories with a numeric mode requires an additional leading zero like 00755, leading minus like -6000, or leading equals like =755.

RESTRICTED DELETION FLAG OR STICKY BIT

The restricted deletion flag or sticky bit is a single bit, whose interpretation depends on the file type. For directories, it prevents unprivileged users from removing or renaming a file in the directory unless they own the file or the directory; this is called the restricted deletion flag for the directory, and is commonly found on world-writable directories like /tmp. For regular files on some older systems, the bit saves the program's text image on the swap device so it will load more quickly when run; this is called the sticky bit.

## OPTIONS

Change the mode of each FILE to MODE. With --reference, change the mode of each FILE to that of RFILE.

-c, --changes
: like verbose but report only when a change is made

-f, --silent, --quiet
: suppress most error messages

-v, --verbose
: output a diagnostic for every file processed

--no-preserve-root
: do not treat '/' specially (the default)

--preserve-root
: fail to operate recursively on '/'

--reference=RFILE
: use RFILE's mode instead of specifying MODE values. RFILE is always dereferenced if a symbolic link.

-R, --recursive
: change files and directories recursively

--help display this help and exit

--version
: output version information and exit

Each MODE is of the form '[ugoa]*([-+=]([rwxXst]*|[ugo]))+|[-+=][0-7]+'.

## AUTHOR

Written by David MacKenzie and Jim Meyering.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.

```
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       chmod(2)

       Full documentation <https://www.gnu.org/software/coreutils/chmod>
       or available locally via: info '(coreutils) chmod invocation'

GNU coreutils 9.4                 April 2024                          CHMOD(1)
```

# Bash 的 choom 命令

```
choom [OPTION] [--] [COMMAND] [ARGS ...]
```

**功能**

查看或调整进程的 OOM（Out-Of-Memory）评分；或以指定的 OOM 调整评分执行命令。

Linux 内核在内存耗尽时会启动 OOM Killer，选择一个 OOM 评分最高的进程杀掉。

**类型**

可执行文件（`/usr/bin/choom`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-p`，`--pid` - 要查看或调整 OOM 评分的的进程ID
  - `-n`，`--adjust` - 要设置的 OOM 调整评分，取值范围是 0 到 1000，越大越容易被杀
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `COMMAND` - 要执行的命令
- `ARGS` - 命令的参数列表

## 示例

**查看和调整进程的 OOM 评分**

```
$ choom -p $$                                          # 查看当前 shell 进程的 oom 评分
pid 944375's current OOM score: 666
pid 944375's current OOM score adjust value: 0
$ choom -p $$ -n 500                                   # 将调整评分设为 500
pid 944375's OOM score adjust value changed from 0 to 500
$ choom -p $$                                          # 再次查看当前 shell 进程的 oom 评分
pid 944375's current OOM score: 1000
pid 944375's current OOM score adjust value: 500
```

- 原始评分为 `666`
- 调整评分被设为 `500`

- 调整后评分为 `666 + 500`，超过了上限 `1000` 因此缩限到 `1000`

**以指定的 OOM 调整评分执行命令**

```
$ choom -n 1000 -- tail -f /var/log/syslog                    # 以 1000 的调整评分执行
tail -f /var/log/syslog
```

- 以 `1000` 的调整评分执行 `tail -f /var/log/syslog`，内存不足时会优先杀死此进程

## 推荐阅读

- choom(1) — Linux manual page

## 手册

```
CHOOM(1)                        User Commands                        CHOOM(1)

NAME
       choom - display and adjust OOM-killer score.

       choom -p PID

       choom -p PID -n number

       choom -n number [--] command [argument ...]

DESCRIPTION
       The choom command displays and adjusts Out-Of-Memory killer score
       setting.

OPTIONS
       -p, --pid pid
           Specifies process ID.

       -n, --adjust value
           Specify the adjust score value.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

NOTES
       Linux kernel uses the badness heuristic to select which process gets
       killed in out of memory conditions.

       The badness heuristic assigns a value to each candidate task ranging
```

from 0 (never kill) to 1000 (always kill) to determine which process is targeted. The units are roughly a proportion along that range of allowed memory the process may allocate from based on an estimation of its current memory and swap use. For example, if a task is using all allowed memory, its badness score will be 1000. If it is using half of its allowed memory, its score will be 500.

There is an additional factor included in the badness score: the current memory and swap usage is discounted by 3% for root processes.

The amount of "allowed" memory depends on the context in which the oom killer was called. If it is due to the memory assigned to the allocating task’s cpuset being exhausted, the allowed memory represents the set of mems assigned to that cpuset. If it is due to a mempolicy’s node(s) being exhausted, the allowed memory represents the set of mempolicy nodes. If it is due to a memory limit (or swap limit) being reached, the allowed memory is that configured limit. Finally, if it is due to the entire system being out of memory, the allowed memory represents all allocatable resources.

The adjust score value is added to the badness score before it is used to determine which task to kill. Acceptable values range from -1000 to +1000. This allows userspace to polarize the preference for oom killing either by always preferring a certain task or completely disabling it. The lowest possible value, -1000, is equivalent to disabling oom killing entirely for that task since it will always report a badness score of 0.

Setting an adjust score value of +500, for example, is roughly equivalent to allowing the remainder of tasks sharing the same system, cpuset, mempolicy, or memory controller resources to use at least 50% more memory. A value of -500, on the other hand, would be roughly equivalent to discounting 50% of the task’s allowed memory from being considered as scoring against the task.

## AUTHORS

Karel Zak <kzak@redhat.com>

## SEE ALSO

proc(5)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The choom command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 chown 命令

```
chown [OPTION]... [OWNER][:[GROUP]] FILE...
```

**功能**

修改文件的所有者和所属组。

**类型**

可执行文件（`/usr/bin/chown`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--changes` - 仅对发生变化的文件打印详细信息
  - `-f`，`--silent`，`--quiet` - 忽略大部分错误信息
  - `-v`，`--verbose` - 打印详细信息
  - `--dereference` - 影响符号链接引用的源文件，而非符号链接本身（默认）
  - `-h`，`--no-dereference` - 影响符号链接本身，而不是其引用的源文件
  - `--from=CURRENT_OWNER:CURRENT_GROUP` - 仅当文件的所有者和所属主是 `CURRENT_OWNER` 和 `CURRENT_GROUP` 才进行修改
  - `--no-preserve-root` - 不保护 `/` 目录
  - `--preserve-root` - 保护 `/` 目录
  - `--reference=RFILE` - 引用 `RFILE` 文件的配置
  - `-R`，`--recursive` - 对目录进行递归操作
  - `-H` - 使用 `-R` 选项进行递归操作时，指向目录的符号链接只有作为命令行参数时才遍历
  - `-L` - 使用 `-R` 选项进行递归操作时，遍历所有指向目录的符号链接
  - `-P` - 使用 `-R` 选项进行递归操作时，不遍历任何符号链接（默认）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `OWNER` - 要设为的所有者
- `GROUP` - 要设为的所属组
- `FILE` - 文件列表

## 示例

**用户名或组名**

```
$ chown planc file.txt                  # 将 file 的所有者设为 planc
$ chown :primers file.txt               # 将 file 的所属组设为 primers
$ chown planc:primers file.txt          # 将 file 的所有者设为 planc 所属组设为 primers
```

### 用 ID 或组 ID

```
$ chown 1001 file.txt                   # 将 file 的所有者 ID 设为 1001
$ chown :1001 file.txt                  # 将 file 的所属组 ID 设为 1001
$ chown 1001:1001 file.txt              # 将 file 的所有者 ID 设为 1001 所属组 ID 设为 1001
```

### 其它

```
$ chown --reference=2.txt 1.txt         # 将 1.txt 的所有者和所属组设为和 2.txt 一样
$ chown planc:primers -R dir            # 递归设置目录 dir 中的文件
```

### 修改根目录（危险！）

```
$ chown planc --no-preserve-root -R /   # 危险！千万不要执行此命令！
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| chgrp | 修改文件的所属组 |
| chmod | 修改文件的访问权限 |

## 推荐阅读

- Change file owner and group - GNU Coreutils
- chown(1) — Linux manual page

## 手册

```
CHOWN(1)                        User Commands                        CHOWN(1)

NAME
       chown - change file owner and group

SYNOPSIS
       chown [OPTION]... [OWNER][:[GROUP]] FILE...
       chown [OPTION]... --reference=RFILE FILE...
```

DESCRIPTION

This manual page documents the GNU version of chown. chown changes the user and/or group ownership of each given file. If only an owner (a user name or numeric user ID) is given, that user is made the owner of each given file, and the files' group is not changed. If the owner is followed by a colon and a group name (or numeric group ID), with no spaces between them, the group ownership of the files is changed as well. If a colon but no group name follows the user name, that user is made the owner of the files and the group of the files is changed to that user's login group. If the colon and group are given, but the owner is omitted, only the group of the files is changed; in this case, chown performs the same function as chgrp. If only a colon is given, or if the entire operand is empty, neither the owner nor the group is changed.

OPTIONS

Change the owner and/or group of each FILE to OWNER and/or GROUP. With --reference, change the owner and group of each FILE to those of RFILE.

-c, --changes
like verbose but report only when a change is made

-f, --silent, --quiet
suppress most error messages

-v, --verbose
output a diagnostic for every file processed

--dereference
affect the referent of each symbolic link (this is the default), rather than the symbolic link itself

-h, --no-dereference
affect symbolic links instead of any referenced file (useful only on systems that can change the ownership of a symlink)

--from=CURRENT_OWNER:CURRENT_GROUP
change the owner and/or group of each file only if its current owner and/or group match those specified here. Either may be omitted, in which case a match is not required for the omitted attribute

--no-preserve-root
do not treat '/' specially (the default)

--preserve-root
fail to operate recursively on '/'

--reference=RFILE
use RFILE's owner and group rather than specifying OWNER:GROUP values. RFILE is always dereferenced.

```
       -R, --recursive
              operate on files and directories recursively

       The  following  options modify how a hierarchy is traversed when the -R
       option is also specified.  If more than one is specified, only the  fi-
       nal one takes effect.

       -H     if  a  command  line argument is a symbolic link to a directory,
              traverse it

       -L     traverse every symbolic link to a directory encountered

       -P     do not traverse any symbolic links (default)

       --help display this help and exit

       --version
              output version information and exit

       Owner is unchanged if missing.  Group  is  unchanged  if  missing,  but
       changed  to login group if implied by a ':' following a symbolic OWNER.
       OWNER and GROUP may be numeric as well as symbolic.

EXAMPLES
       chown root /u
              Change the owner of /u to "root".

       chown root:staff /u
              Likewise, but also change its group to "staff".

       chown -hR root /u
              Change the owner of /u and subfiles to "root".

AUTHOR
       Written by David MacKenzie and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       chown(2)

       Full documentation <https://www.gnu.org/software/coreutils/chown>
       or available locally via: info '(coreutils) chown invocation'
```

# Bash 的 chroot 命令

```
chroot [OPTION] NEWROOT [COMMAND [ARG]...]
```

**功能**

执行命令时指定根路径。

**类型**

可执行文件（`/usr/sbin/chroot`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--groups=G_LIST` - 以列表（`g1,g2,...,gn`）的格式补充组
  - `--userspec=USER:GROUP` - 指定用户和组
  - `--skip-chdir` - 不更改工作目录
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NEWROOT` 要作为根路径的路径
- `COMMAND` 要执行的命令；如果不指定此参数，则切换后运行当前 shell 程序
- `ARG` 命令的参数列表

**返回值**

- 125 - `chroot` 自身失败
- 126 - `COMMAND` 无法运行
- 127 - `COMMAND` 不存在
- 其它情况返回 `COMMAND` 的返回值

## 示例

```
$ ls /tmp/                                              # 查看 /tmp 目录
tmp.XG4RL3jbgF  tmp.Urtrw97emL  tmp.M8ln2SXW4p
$ mkdir -p ./fake_root/tmp                              # 创建目录
$ touch ./fake_root/tmp/1.txt                           # 创建文件
$ wget https://busybox.net/downloads/binaries/1.35.0-x86_64-linux-musl/busybox -O
./fake_root/busybox
$ sudo chroot ./fake_root/ ./busybox ls /tmp/           # 指定 ./fake_root/ 为根目录，查看 /tmp
```

```
目录
1.txt
```

> 因为将根目录设为 `./fake_root/` 之后，没有 `/bin` 目录下的各种命令，因此下载了 `busybox` 来执行 `ls` 命令。

## 推荐阅读

- Run a command with a different root directory - GNU Coreutils
- chroot(1) — Linux manual page

## 手册

```
CHROOT(8)                       User Commands                       CHROOT(8)

NAME
       chroot - run command or interactive shell with special root directory

SYNOPSIS
       chroot [OPTION] NEWROOT [COMMAND [ARG]...]
       chroot OPTION

DESCRIPTION
       Run COMMAND with root directory set to NEWROOT.

       --groups=G_LIST
              specify supplementary groups as g1,g2,..,gN

       --userspec=USER:GROUP
              specify user and group (ID or name) to use

       --skip-chdir
              do not change working directory to '/'

       --help display this help and exit

       --version
              output version information and exit

       If no command is given, run '"$SHELL" -i' (default: '/bin/sh -i').

   Exit status:
       125    if the chroot command itself fails

       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found
```

```
       -      the exit status of COMMAND otherwise

AUTHOR
       Written by Roland McGrath.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       chroot(2)

       Full documentation <https://www.gnu.org/software/coreutils/chroot>
       or available locally via: info '(coreutils) chroot invocation'

GNU coreutils 9.4                April 2024                        CHROOT(8)
```

# Bash 的 chrt 命令

```
chrt [OPTION] -p [PRIORITY] PID
chrt [OPTION] PRIORITY [COMMAND] [ARGS ...]
```

**功能**

查看或设置进程的调度策略和优先级；或以指定的进程优先级执行命令。

**类型**

可执行文件（`/usr/bin/chrt`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--all-tasks` - 设置检索进程下的所有任务（线程）
  - `-m`，`--max` - 查看优先级的取值范围
  - `-p`，`--pid` - 查看或设置的进程ID
  - `-v`，`--verbose` - 显示状态信息
  - `--help` - 显示帮助
  - `--version` - 显示版本
  - 策略选项：
    - `-o`，`--other` - 将调度策略设为 `SCHED_OTHER`（分时调度）；这是 Linux 的默认调度策略
    - `-f`，`--fifo` - 将调度策略设为 `SCHED_FIFO`（先进先出调度）；无时间片，可能饿死所有任务
    - `-r`，`--rr` - 将调度策略设为 `SCHED_RR`（轮流调度）；有时间片，可能饿死低优先级任务；不指定策略时默认使用此选项
    - `-b`，`--batch` - 将调度策略设为 `SCHED_BATCH`（批处理调度）
    - `-i`，`--idle` - 将调度策略设为 `SCHED_IDLE`（闲时调度）；用于低优先级任务
    - `-d`，`--deadline` - 将调度策略设为 `SCHED_DEADLINE`（截止时间调度）；在截止时间前完成的硬实时调度策略
  - 调度选项：
    - `-T`，`--sched-runtime nanoseconds` - 指定 `SCHED_DEADLINE` 策略的运行时参数
    - `-P`，`--sched-period nanoseconds` - 指定 `SCHED_DEADLINE` 策略的周期参数
    - `-D`，`--sched-deadline nanoseconds` - 指定 `SCHED_DEADLINE` 策略的截止时间参数
    - `-R`，`--reset-on-fork` - 使用 `SCHED_RESET_ON_FORK` 或 `SCHED_FLAG_RESET_ON_FORK` 标志

- `PRIORITY` - 优先级；数值越大，优先级越高
- `COMMAND` - 要执行的命令
- `ARGS` - 命令的参数列表

# 示例

### 查看优先级的取值范围

```
$ chrt -m
SCHED_OTHER min/max priority    : 0/0
SCHED_FIFO min/max priority     : 1/99
SCHED_RR min/max priority       : 1/99
SCHED_BATCH min/max priority    : 0/0
SCHED_IDLE min/max priority     : 0/0
SCHED_DEADLINE min/max priority : 0/0
```

- 可以看到只有 `SCHED_FIFO` 和 `SCHED_RR` 支持优先级

### 查看和设置进程调度策略和优先级

```
$ chrt -p $$                                    # 查看当前 bash 进程的调度策略和优先级
pid 944375's current scheduling policy: SCHED_OTHER
pid 944375's current scheduling priority: 0
$ sudo chrt -r -p 20 $$                         # 将策略设为 SCHED_RR，优先级设为 20
$ chrt -p $$                                    # 再次查看当前 bash 进程的调度策略和优先级
pid 944375's current scheduling policy: SCHED_RR
pid 944375's current scheduling priority: 20
```

### 以指定的进程调度策略和优先级执行命令

```
$ chrt -i 0 tail -f /var/log/syslog             # 以 SCHED_IDLE 策略执行 tail -f /var/log/syslog
```

- 虽然 `SCHED_IDLE` 不支持优先级，但不能省略优先级参数，必须设置为 `0`

# 推荐阅读

- chrt(1) — Linux manual page

# 手册

NAME

chrt - manipulate the real-time attributes of a process

SYNOPSIS

chrt [options] priority command argument ...

chrt [options] -p [priority] PID

DESCRIPTION

chrt sets or retrieves the real-time scheduling attributes of an existing PID, or runs command with the given attributes.

POLICIES

-o, --other
Set scheduling policy to SCHED_OTHER (time-sharing scheduling). This is the default Linux scheduling policy.

-f, --fifo
Set scheduling policy to SCHED_FIFO (first in-first out).

-r, --rr
Set scheduling policy to SCHED_RR (round-robin scheduling). When no policy is defined, the SCHED_RR is used as the default.

-b, --batch
Set scheduling policy to SCHED_BATCH (scheduling batch processes). Linux-specific, supported since 2.6.16. The priority argument has to be set to zero.

-i, --idle
Set scheduling policy to SCHED_IDLE (scheduling very low priority jobs). Linux-specific, supported since 2.6.23. The priority argument has to be set to zero.

-d, --deadline
Set scheduling policy to SCHED_DEADLINE (sporadic task model deadline scheduling). Linux-specific, supported since 3.14. The priority argument has to be set to zero. See also --sched-runtime, --sched-deadline and --sched-period. The relation between the options required by the kernel is runtime ⇐ deadline ⇐ period. chrt copies period to deadline if --sched-deadline is not specified and deadline to runtime if --sched-runtime is not specified. It means that at least --sched-period has to be specified. See sched(7) for more details.

SCHEDULING OPTIONS

-T, --sched-runtime nanoseconds
Specifies runtime parameter for SCHED_DEADLINE policy (Linux-specific).

-P, --sched-period nanoseconds
Specifies period parameter for SCHED_DEADLINE policy (Linux-specific). Note that the kernel's lower limit is 100 milliseconds.

-D, --sched-deadline nanoseconds
Specifies deadline parameter for SCHED_DEADLINE policy (Linux-specific).

-R, --reset-on-fork
Use SCHED_RESET_ON_FORK or SCHED_FLAG_RESET_ON_FORK flag. Linux-specific, supported since 2.6.31.

Each thread has a reset-on-fork scheduling flag. When this flag is set, children created by fork(2) do not inherit privileged scheduling policies. After the reset-on-fork flag has been enabled, it can be reset only if the thread has the CAP_SYS_NICE capability. This flag is disabled in child processes created by fork(2).

More precisely, if the reset-on-fork flag is set, the following rules apply for subsequently created children:

- If the calling thread has a scheduling policy of SCHED_FIFO or SCHED_RR, the policy is reset to SCHED_OTHER in child processes.

- If the calling process has a negative nice value, the nice value is reset to zero in child processes.

OPTIONS
-a, --all-tasks
Set or retrieve the scheduling attributes of all the tasks (threads) for a given PID.

-m, --max
Show minimum and maximum valid priorities, then exit.

-p, --pid
Operate on an existing PID and do not launch a new task.

-v, --verbose
Show status information.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

EXAMPLES
The default behavior is to run a new command:

chrt priority command [arguments]

You can also retrieve the real-time attributes of an existing task:

```
chrt -p PID
```

Or set them:

```
chrt -r -p priority PID
```

This, for example, sets real-time scheduling to priority 30 for the process PID with the SCHED_RR (round-robin) class:

```
chrt -r -p 30 PID
```

Reset priorities to default for a process:

```
chrt -o -p 0 PID
```

See sched(7) for a detailed discussion of the different scheduler classes and how they interact.

## PERMISSIONS

A user must possess CAP_SYS_NICE to change the scheduling attributes of a process. Any user can retrieve the scheduling information.

## NOTES

Only SCHED_FIFO, SCHED_OTHER and SCHED_RR are part of POSIX 1003.1b Process Scheduling. The other scheduling attributes may be ignored on some systems.

Linux' default scheduling policy is SCHED_OTHER.

## AUTHORS

Robert Love <rml@tech9.net>, Karel Zak <kzak@redhat.com>

## SEE ALSO

nice(1), renice(1), taskset(1), sched(7)

See sched_setscheduler(2) for a description of the Linux scheduling scheme.

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The chrt command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 cksum 命令

```
cksum [OPTION]... [FILE]...
```

**功能**

计算和校验文件的校验和。

**类型**

可执行文件（`/usr/bin/cksum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--algorithm=TYPE` - 校验和算法；支持：`sysv`，`crc`，`md5`，`sha1`，`sha224`，`sha256`，`sha384`，`sha512`，`blake2b`，`sm3`，默认为 `crc`
  - `--base64` - 输出 BASE64 编码而不是十六进制编码
  - `-c`，`--check` - 从文件中读取校验和进行校验
  - `-l`，`--length=BITS` - 结果保留的长度（二进制位），只能用于 `blake2b` 不能超过 512 且必须是 8 的倍数
  - `--raw` - 输出原始数据而不是十六进制编码；通常用于重定向
  - `--tag` - 生成 BSD 风格的输出（默认）
  - `--untagged` - 生成反向、无标签（不附带算法类型）风格的输出
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--debug` - 校验时显示所使用的实现
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算**

```
$ cksum 1.txt                          # 默认为 CRC
2503682644 3208 1.txt
$ cksum -a sha224 1.txt                # 指定 SHA224
SHA224 (1.txt) = e97bd3540da0cf24cf74a776a523ed7c89e89d435c9f59070f9dbb64
$ cksum -a sha224 --base64 1.txt       # 输出 BASE64
SHA224 (1.txt) = 6XvTVA2gzyTPdKd2pSPtfInonUNcn1kHD527ZA==
$ cksum -a sha224 --raw 1.txt          # 输出原始数据
 δ§v¥#킬蜃\Y»d
$ cksum -a sha224 --untagged 1.txt     # 输出反向无标签风格，无法用于校验
e97bd3540da0cf24cf74a776a523ed7c89e89d435c9f59070f9dbb64  1.txt
```

**校验**

```
$ cksum -a sha512 1.txt 2.txt 3.txt > sum.txt
$ cksum -c sum.txt                     # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

# 相关命令

| 命令 | 说明 |
|---|---|
| b2sum | 计算和校验 BLAKE2 值 |
| md5sum | 计算和校验 MD5 值 |
| sha1sum | 计算和校验 SHA1 值 |
| sha224sum | 计算和校验 SHA224 值 |
| sha256sum | 计算和校验 SHA256 值 |
| sha384sum | 计算和校验 SHA384 值 |
| sha512sum | 计算和校验 SHA512 值 |

# 推荐阅读

- Print and verify file checksums - GNU Coreutils
- cksum(1) — Linux manual page

# 手册

CKSUM(1) User Commands CKSUM(1)

NAME
cksum - compute and verify file checksums

SYNOPSIS
cksum [OPTION]... [FILE]...

DESCRIPTION
Print or verify checksums. By default use the 32 bit CRC algorithm.

With no FILE, or when FILE is -, read standard input.

Mandatory arguments to long options are mandatory for short options too.

-a, --algorithm=TYPE
select the digest type to use. See DIGEST below.

--base64
emit base64-encoded digests, not hexadecimal

-c, --check
read checksums from the FILEs and check them

-l, --length=BITS
digest length in bits; must not exceed the max for the blake2 algorithm and must be a multiple of 8

--raw emit a raw binary digest, not hexadecimal

--tag create a BSD-style checksum (the default)

--untagged
create a reversed style checksum, without digest type

-z, --zero
end each output line with NUL, not newline, and disable file name escaping

The following five options are useful only when verifying checksums:
--ignore-missing
don't fail or report status for missing files

--quiet
don't print OK for each successfully verified file

--status
don't output anything, status code shows success

--strict
exit non-zero for improperly formatted checksum lines

```
       -w, --warn
              warn about improperly formatted checksum lines

       --debug
              indicate which implementation used

       --help display this help and exit

       --version
              output version information and exit

   DIGEST determines the digest algorithm and default output format:
       sysv   (equivalent to sum -s)

       bsd    (equivalent to sum -r)

       crc    (equivalent to cksum)

       md5    (equivalent to md5sum)

       sha1   (equivalent to sha1sum)

       sha224 (equivalent to sha224sum)

       sha256 (equivalent to sha256sum)

       sha384 (equivalent to sha384sum)

       sha512 (equivalent to sha512sum)

       blake2b
              (equivalent to b2sum)

       sm3    (only available through cksum)

       When checking, the input should be a former output of this program,  or
       equivalent standalone program.

AUTHOR
       Written by Padraig Brady and Q. Frank Xia.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
```

```
       Full documentation <https://www.gnu.org/software/coreutils/cksum>
       or available locally via: info '(coreutils) cksum invocation'

GNU coreutils 9.4                 April 2024                          CKSUM(1)
```

# Bash 的 comm 命令

```
comm [OPTION]... FILE1 FILE2
```

**功能**

逐行比较已排序的文件 `FILE1` 和 `FILE2`。

产生三列输出，以此为：`FILE1` 的特有行，`FILE2` 的特有行，`FILE1` 和 `FILE2` 的共有行。

**类型**

可执行文件（`/usr/bin/comm`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-1` - 不显示第 1 列（`FILE1` 的特有行）
  - `-2` - 不显示第 2 列（`FILE2` 的特有行）
  - `-3` - 不显示第 3 列（`FILE1` 和 `FILE2` 的共有行）
  - `--check-order` - 即使所有行都匹配，也要检查排序是否正确
  - `--nocheck-order` - 不检查是否正确排序
  - `--output-delimiter=STR` - 使用 `STR` 分隔输出的列
  - `--total` - 输出摘要
  - `-z`, `--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE1` - 比较的文件；`-` 表示读取标准输入，`FILE1` 和 `FILE2` 不能同时为 `-`
- `FILE2` - 比较的文件；`-` 表示读取标准输入，`FILE1` 和 `FILE2` 不能同时为 `-`

## 示例

```
$ cat 1.txt
AAAAA
BBBBB
CCCCC
DDDDD
EEEEE
$ cat 2.txt
AAAAA
```

```
BBBBB
DDDDD
EEEEE
FFFFF
$ comm 1.txt 2.txt      # 比较
                AAAAA
                BBBBB
CCCCC
                DDDDD
                EEEEE
        FFFFF
```

## 推荐阅读

- Compare two sorted files line by line - GNU Coreutils
- comm(1) — Linux manual page

## 手册

```
COMM(1)                          User Commands                          COMM(1)

NAME
       comm - compare two sorted files line by line

SYNOPSIS
       comm [OPTION]... FILE1 FILE2

DESCRIPTION
       Compare sorted files FILE1 and FILE2 line by line.

       When FILE1 or FILE2 (not both) is -, read standard input.

       With  no  options,  produce  three-column  output.  Column one contains
       lines unique to FILE1, column two contains lines unique to  FILE2,  and
       column three contains lines common to both files.

       -1     suppress column 1 (lines unique to FILE1)

       -2     suppress column 2 (lines unique to FILE2)

       -3     suppress column 3 (lines that appear in both files)

       --check-order
              check  that  the  input  is  correctly sorted, even if all input
              lines are pairable

       --nocheck-order
```

```
              do not check that the input is correctly sorted

       --output-delimiter=STR
              separate columns with STR

       --total
              output a summary

       -z, --zero-terminated
              line delimiter is NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

       Note, comparisons honor the rules specified by 'LC_COLLATE'.

EXAMPLES
       comm -12 file1 file2
              Print only lines present in both file1 and file2.

       comm -3 file1 file2
              Print lines in file1 not in file2, and vice versa.

AUTHOR
       Written by Richard M. Stallman and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       join(1), uniq(1)

       Full documentation <https://www.gnu.org/software/coreutils/comm>
       or available locally via: info '(coreutils) comm invocation'

GNU coreutils 9.4                April 2024                          COMM(1)
```

# Bash 的 cp 命令

```
cp [OPTION]... SOURCE... DEST
```

**功能**

复制文件或目录。

**类型**

可执行文件（`/usr/bin/cp`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--archive` - 复制整个目录树并保持文件的原有属性；同 `-dR --preserve=all`
  - `--attributes-only` - 只复制文件属性，不复制文件数据
  - `--backup[=CONTROL]` - 选择怎样备份被覆盖的目标文件
  - `-b` - 备份被覆盖的目标文件
  - `--copy-contents` - 复制特殊文件（设备、FIFO、管道等）
  - `-d` - 复制符号链接本身，而不是复制源文件；同 `--no-dereference --preserve=links`
  - `--debug` - 输出调试信息（文件的复制方式）
  - `-f`，`--force` - 强制操作，目标存在且无法打开时会尝试将其删除
  - `-i`，`--interactive` - 覆盖前提示
  - `-H` - `SOURCE` 是符号链接时复制源文件而非符号链接，但递归复制目录中的符号链接时不生效
  - `-l`，`--link` - 创建文件的硬链接而不是复制
  - `-L`，`--dereference` - `SOURCE` 是符号链接时复制源文件而非符号链接，递归复制目录中的符号链接时也生效
  - `-n`，`--no-clobber` - 不覆盖现有文件，也不失败
  - `-P`，`--no-dereference` - `SOURCE` 是符号链接时复制符号链接而非源文件
  - `-p` - 保持文件原有的属性；同 `--preserve=mode,ownership,timestamps`
  - `--preserve[=ATTR_LIST]` - 保持文件原有的属性
  - `--no-preserve=ATTR_LIST` - 不保持文件原有的属性
  - `--parents` - 复制时保持路径
  - `-R`，`-r`，`--recursive` - 递归复制目录
  - `--reflink[=WHEN]` - 控制在支持 CoW（Copy-on-Write，写时复制）的文件系统上进行轻量复制

    - `--remove-destination` - 复制前删除已经存在的目标
    - `--sparse=WHEN` - 控制稀疏文件的创建
    - `--strip-trailing-slashes` - 删除每个 `SOURCE` 尾部的 `/`
    - `-s`, `--symbolic-link` - 创建文件的符号链接而不是复制
    - `-S`, `--suffix=SUFFIX` - 设置备份文件的后缀
    - `-t`, `--target-directory=DIRECTORY` - 指定目标目录
    - `-T`, `--no-target-directory` - 将 `DEST` 视为文件（如果 `DEST` 是目录，则进行合并而非复制到它内部）
    - `--update[=UPDATE]` - 控制哪些现有文件需要更新；取值：`all`, `none`, `older`
    - `-u` - 仅在源文件比目标新的时候复制；同 `--update[=older]`
    - `-v`, `--verbose` - 输出详细信息（进行的全部操作）
    - `-x`, `--one-file-system` - 跳过跨文件系统的复制
    - `-Z` - 将目标文件的 SELinux 安全上下文设为默认类型
    - `--context[=CTX]`
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `SOURCE` - 源文件
- `DEST` - 目标路径

## 示例

### 复制文件

```
$ ls                            # 空目录下操作
$ touch src.txt                 # 创建文件
$ cp src.txt dest.txt           # 复制文件
$ ls
src.txt dest.txt
$ mkdir dir                     # 创建目录
$ cp src.txt ./dir/dest.txt     # 复制到目录中
$ ls dir
dest.txt
```

### 复制目录

```
$ mkdir src dest1 dest2         # 创建目录
$ touch src/1.txt               # 创建文件
$ tree src
src
└── 1.txt
$ cp -r src dest1               # 将 src 复制到 dest1 内部
$ tree dest1
dest1
```

```
└── src
    └── 1.txt
$ cp -rT src dest2              # 复制 src 合并到 dest2
$ tree dest2
dest2
└── 1.txt
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| mv | 移动文件 |
| ln | 链接文件 |

## 推荐阅读

- Copy files and directories - GNU Coreutils
- cp(1) — Linux manual page

## 手册

```
CP(1)                            User Commands                           CP(1)

NAME
       cp - copy files and directories

SYNOPSIS
       cp [OPTION]... [-T] SOURCE DEST
       cp [OPTION]... SOURCE... DIRECTORY
       cp [OPTION]... -t DIRECTORY SOURCE...

DESCRIPTION
       Copy SOURCE to DEST, or multiple SOURCE(s) to DIRECTORY.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -a, --archive
              same as -dR --preserve=all

       --attributes-only
              don't copy the file data, just the attributes

       --backup[=CONTROL]
```

```
       make a backup of each existing destination file

-b     like --backup but does not accept an argument

--copy-contents
       copy contents of special files when recursive

-d     same as --no-dereference --preserve=links

--debug
       explain how a file is copied.  Implies -v

-f, --force
       if an existing destination file cannot be opened, remove it  and
       try  again  (this  option  is ignored when the -n option is also
       used)

-i, --interactive
       prompt before overwrite (overrides a previous -n option)

-H     follow command-line symbolic links in SOURCE

-l, --link
       hard link files instead of copying

-L, --dereference
       always follow symbolic links in SOURCE

-n, --no-clobber
       do not overwrite an existing file and do not fail  (overrides  a
       -u  or  previous  -i  option).  See also --update; equivalent to
       --update=none.

-P, --no-dereference
       never follow symbolic links in SOURCE

-p     same as --preserve=mode,ownership,timestamps

--preserve[=ATTR_LIST]
       preserve the specified attributes

--no-preserve=ATTR_LIST
       don't preserve the specified attributes

--parents
       use full source file name under DIRECTORY

-R, -r, --recursive
       copy directories recursively

--reflink[=WHEN]
       control clone/CoW copies. See below
```

```
--remove-destination
       remove each existing destination file before attempting to  open
       it (contrast with --force)

--sparse=WHEN
       control creation of sparse files. See below

--strip-trailing-slashes
       remove any trailing slashes from each SOURCE argument

-s, --symbolic-link
       make symbolic links instead of copying

-S, --suffix=SUFFIX
       override the usual backup suffix

-t, --target-directory=DIRECTORY
       copy all SOURCE arguments into DIRECTORY

-T, --no-target-directory
       treat DEST as a normal file

--update[=UPDATE]
       control     which    existing    files    are    updated;    UP-
       DATE={all,none,older(default)}.  See below

-u     equivalent to --update[=older]

-v, --verbose
       explain what is being done

-x, --one-file-system
       stay on this file system

-Z     set SELinux security context of destination file to default type

--context[=CTX]
       like -Z, or if CTX is specified then set the  SELinux  or  SMACK
       security context to CTX

--help display this help and exit

--version
       output version information and exit

ATTR_LIST  is  a  comma-separated  list  of  attributes. Attributes are
'mode' for permissions (including any ACL and xattr permissions), 'own-
ership' for user and group, 'timestamps' for file  timestamps,  'links'
for  hard  links,  'context' for security context, 'xattr' for extended
attributes, and 'all' for all attributes.

By default, sparse SOURCE files are detected by a crude  heuristic  and
the corresponding DEST file is made sparse as well.  That is the behav-
```

ior selected by --sparse=auto. Specify --sparse=always to create a sparse DEST file whenever the SOURCE file contains a long enough sequence of zero bytes. Use --sparse=never to inhibit creation of sparse files.

UPDATE controls which existing files in the destination are replaced. 'all' is the default operation when an --update option is not specified, and results in all existing files in the destination being replaced. 'none' is similar to the --no-clobber option, in that no files in the destination are replaced, but also skipped files do not induce a failure. 'older' is the default operation when --update is specified, and results in files being replaced if they're older than the corresponding source file.

When --reflink[=always] is specified, perform a lightweight copy, where the data blocks are copied only when modified. If this is not possible the copy fails, or if --reflink=auto is specified, fall back to a standard copy. Use --reflink=never to ensure a standard copy is performed.

The backup suffix is '~', unless set with --suffix or SIMPLE_BACKUP_SUFFIX. The version control method may be selected via the --backup option or through the VERSION_CONTROL environment variable. Here are the values:

none, off
: never make backups (even if --backup is given)

numbered, t
: make numbered backups

existing, nil
: numbered if numbered backups exist, simple otherwise

simple, never
: always make simple backups

As a special case, cp makes a backup of SOURCE when the force and backup options are given and SOURCE and DEST are the same name for an existing, regular file.

## AUTHOR

Written by Torbjorn Granlund, David MacKenzie, and Jim Meyering.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

```
SEE ALSO
       install(1)

       Full documentation <https://www.gnu.org/software/coreutils/cp>
       or available locally via: info '(coreutils) cp invocation'

GNU coreutils 9.4                 April 2024                              CP(1)
```

# Bash 的 csplit 命令

```
csplit [OPTION]... FILE PATTERN...
```

**功能**

将文件拆分为多个部分。

拆分后依次输出到文件 `xx00`， `xx01`， …，并在标准输出上依次打印输出文件的字节数。

**类型**

可执行文件（`/usr/bin/csplit`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--suffix-format=FORMAT` - 输出文件名的后缀格式
  - `-f`，`--prefix=PREFIX` - 输出文件名的前缀
  - `-k`，`--keep-files` - 产生错误时不要删除输出文件
  - `--suppress-matched` - 删除匹配的行
  - `-n`，`--digits=DIGITS` - 输出文件名的后缀数字位数
  - `-s`，`--quiet`，`--silent` - 不打印输出
  - `-z`，`--elide-empty-files` - 不输出空文件
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 要拆分的文件；`-` 表示读取标准输入
- `PATTERN` - 模式列表，每个模式由以下结构组成：
  - `整数` - 复制到 `INTEGER` 行（不含该行）
  - `/正则表达式/[行偏移量]` - 复制到匹配的行（不含该行）；行偏移量可以省略
  - `%正则表达式%[行偏移量]` - 跳过到匹配的行（不含该行）；行偏移量可以省略
  - `{整数}` - 重复前面的模式 `INTEGER` 次
  - `{*}` - 重复前面的模式经可能多次

## 示例

**行号**

```
$ man csplit | cat -n | csplit - 30 60
1012
1953
$ cat xx00                                          # xx00 从 0 到 29 行
...
    25               do not remove output files on errors
    26
    27         --suppress-matched
    28               suppress the lines matching PATTERN
    29
$ cat xx01                                          # xx01 从 30 到 59 行
    30         -n, --digits=DIGITS
    31               use specified number of digits instead of 2
    32
    33         -s, --quiet, --silent
...
    55               repeat the previous pattern specified number of times
    56
    57         {*}    repeat the previous pattern as many times as possible
    58
    59         A line OFFSET is an integer optionally preceded by '+' or '-'
$ cat xx02                                          # xx02 从 60 行 到结尾
    60
    61  AUTHOR
    62         Written by Stuart Kemp and David MacKenzie.
    63
    64  REPORTING BUGS
...
```

**正则表达式**

```
$ man csplit | csplit - '/^[A-Z]+$/' '{*}'          # 复制到匹配 /^[A-Z]+$/ 的行（不含该
行），重复尽可能多次
80
77
52
1451
228
531
$ cat xx0                                           # 按每个小节拆分
CSPLIT(1)                       User Commands                       CSPLIT(1)

$ cat xx01
NAME
       csplit - split a file into sections determined by context lines
$ cat xx02
SYNOPSIS
       csplit [OPTION]... FILE PATTERN...

$ cat xx03
DESCRIPTION
```

```
       Output  pieces of FILE separated by PATTERN(s) to files 'xx00', 'xx01',
       ..., and output byte counts of each piece to standard output.
...
$ cat xx04
AUTHOR
       Written by Stuart Kemp and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>
```

## 推荐阅读

- Split a file into context-determined pieces - GNU Coreutils
- csplit(1) — Linux manual page

## 手册

```
CSPLIT(1)                        User Commands                        CSPLIT(1)

NAME
       csplit - split a file into sections determined by context lines

SYNOPSIS
       csplit [OPTION]... FILE PATTERN...

DESCRIPTION
       Output  pieces of FILE separated by PATTERN(s) to files 'xx00', 'xx01',
       ..., and output byte counts of each piece to standard output.

       Read standard input if FILE is -

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -b, --suffix-format=FORMAT
              use sprintf FORMAT instead of %02d

       -f, --prefix=PREFIX
              use PREFIX instead of 'xx'

       -k, --keep-files
              do not remove output files on errors

       --suppress-matched
              suppress the lines matching PATTERN
```

```
       -n, --digits=DIGITS
              use specified number of digits instead of 2

       -s, --quiet, --silent
              do not print counts of output file sizes

       -z, --elide-empty-files
              suppress empty output files

       --help display this help and exit

       --version
              output version information and exit

   Each PATTERN may be:
       INTEGER
              copy up to but not including specified line number

       /REGEXP/[OFFSET]
              copy up to but not including a matching line

       %REGEXP%[OFFSET]
              skip to, but not including a matching line

       {INTEGER}
              repeat the previous pattern specified number of times

       {*}    repeat the previous pattern as many times as possible

       A line OFFSET is an integer optionally preceded by '+' or '-'

AUTHOR
       Written by Stuart Kemp and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/csplit>
       or available locally via: info '(coreutils) csplit invocation'

GNU coreutils 9.4               April 2024                        CSPLIT(1)
```

# Bash 的 ctrlaltdel 命令

```
ctrlaltdel [OPTION]... FUNC
```

**功能**

设置 `Ctrl + Alt + Del` 组合键的功能。

**类型**

可执行文件（`/usr/sbin/ctrlaltdel`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FUNC` - 按下 `Ctrl + Alt + Del` 组合键的功能：
  - `hard` - 立即重启计算机，无需调用 `sync(2)` 函数，也无需任何其他准备工作；**默认**
  - `soft` - 让内核向 init 进程（PID 为 1）发送 `SIGINT` 信号

## 推荐阅读

- ctrlaltdel(8) — Linux manual page

## 手册

```
CTRLALTDEL(8)               System Administration              CTRLALTDEL(8)

NAME
       ctrlaltdel - set the function of the Ctrl-Alt-Del combination

SYNOPSIS
       ctrlaltdel hard|soft

DESCRIPTION
       Based on examination of the linux/kernel/reboot.c code, it is clear
       that there are two supported functions that the <Ctrl-Alt-Del> sequence
       can perform.

       hard
           Immediately reboot the computer without calling sync(2) and without
```

```
            any other preparation. This is the default.

        soft
            Make the kernel send the SIGINT (interrupt) signal to the init
            process (this is always the process with PID 1). If this option is
            used, the init(8) program must support this feature. Since there
            are now several init(8) programs in the Linux community, please
            consult the documentation for the version that you are currently
            using.

        When the command is run without any argument, it will display the
        current setting.

        The function of ctrlaltdel is usually set in the /etc/rc.local file.

OPTIONS
        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

FILES
        /etc/rc.local

AUTHORS
        Peter Orbaek <poe@daimi.aau.dk>

SEE ALSO
        init(8), systemd(1)

REPORTING BUGS
        For bug reports, use the issue tracker at
        https://github.com/util-linux/util-linux/issues.

AVAILABILITY
        The ctrlaltdel command is part of the util-linux package which can be
        downloaded from Linux Kernel Archive
        <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                    CTRLALTDEL(8)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 cut 命令

```
cut OPTION... [FILE]...
```

**功能**

从文件的行中截取指定的部分。

**类型**

可执行文件（`/usr/bin/cut`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`, `--bytes=LIST` - 要截取的字节
  - `-c`, `--characters=LIST` - 要截取的字符
  - `-d`, `--delimiter=DELIM` - 使用 `DELIM` 作为字段分隔符，而不是 `TAB`
  - `-f`, `--fields=LIST` - 要截取的字段
  - ~~`-n` - 防止截断多字节字符（无效，被忽略）~~
  - `--complement` - 对选取的字节、字符、字段进行补全
  - `-s`, `--only-delimited` - 不要打印不含分隔符的行
  - `--output-delimiter=STRING` - 打印输出时用 `STRING` 代替分隔符
  - `-z`, `--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示当前帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；`-` 表示读取标准输入

## 示例

```
$ cat 1.txt
id,name,age,city
1,Alice,23,London
2,Bob,31,Paris
3,Carol,27,Tokyo
$ cut 1.txt -c 1-10 1.txt       # 截取 1 至 10 字符
id,name,ag
1,Alice,23
2,Bob,31,P
3,Carol,27
```

```
id,name,ag
1,Alice,23
2,Bob,31,P
3,Carol,27
$ cut -d ',' -f 2,3 1.txt       # 逗号作为分隔符，截取 2,3 字段
name,age
Alice,23
Bob,31
Carol,27
```

## 推荐阅读

- Print selected parts of lines - GNU Coreutils
- cut(1) — Linux manual page

## 手册

```
CUT(1)                          User Commands                          CUT(1)

NAME
       cut - remove sections from each line of files

SYNOPSIS
       cut OPTION... [FILE]...

DESCRIPTION
       Print selected parts of lines from each FILE to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -b, --bytes=LIST
              select only these bytes

       -c, --characters=LIST
              select only these characters

       -d, --delimiter=DELIM
              use DELIM instead of TAB for field delimiter

       -f, --fields=LIST
              select only these fields;  also print any line that contains  no
              delimiter character, unless the -s option is specified

       -n     (ignored)
```

```
       --complement
              complement the set of selected bytes, characters or fields

       -s, --only-delimited
              do not print lines not containing delimiters

       --output-delimiter=STRING
              use STRING as the output delimiter the default is to use the in-
              put delimiter

       -z, --zero-terminated
              line delimiter is NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

       Use  one,  and  only  one of -b, -c or -f.  Each LIST is made up of one
       range, or many ranges separated by commas.  Selected input  is  written
       in  the  same order that it is read, and is written exactly once.  Each
       range is one of:

       N      N'th byte, character or field, counted from 1

       N-     from N'th byte, character or field, to end of line

       N-M    from N'th to M'th (included) byte, character or field

       -M     from first to M'th (included) byte, character or field

AUTHOR
       Written by David M. Ihnat, David MacKenzie, and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/cut>
       or available locally via: info '(coreutils) cut invocation'

GNU coreutils 9.4                April 2024                           CUT(1)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 date 命令

```
date [OPTION]... [+FORMAT]
```

**功能**

查看或设置日期和时间。

**类型**

可执行文件（`/usr/bin/date`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-d`，`--date=STRING` - 查看指定的日期，而不是现在
  - `--debug` - 打印解析过程
  - `-f`，`--file=DATEFILE` - 使用文件 `DATEFILE` 指定日期，每行一次，类似 `-d` 选项
  - `-I[FMT]`，`--iso-8601[=FMT]` - 使用 ISO 8601 格式
  - `--resolution` - 查看时间戳的可用分辨率
  - `-R`，`--rfc-email` - 使用 RFC 5322 格式
  - `--rfc-3339=FMT` - 使用 RFC 3339 格式
  - `-r`，`--reference=FILE` - 查看文件 `FILE` 最后被修改的时间
  - `-s`，`--set=STRING` - 将时间设为字符串 `STRING` 表示的时间
  - `-u`，`--utc`，`--universal` - 查看或设置 UTC 时间；默认为本地时间
  - `--help` - 显示当前帮助
  - `--version` - 显示版本

## 格式

| 格式 | 含义 |
| --- | --- |
| `%%` | 字面值 `%` |
| `%a` | 本地化的星期简称（例如：Sun） |
| `%A` | 本地化的星期全称（例如：Sunday） |
| `%b` | 本地化的月份简称（例如：Jan） |
| `%B` | 本地化的月份全称（例如：January） |

| 格式 | 含义 |
| --- | --- |
| `%c` | 本地化的日期与时间（例如：Thu Mar 3 23:05:25 2005） |
| `%C` | 世纪（类似 `%Y`，但去掉最后两位数字，例如：20） |
| `%d` | 月中的第几天（例如：01） |
| `%D` | 日期，等同于 `%m/%d/%y` |
| `%e` | 月中的第几天（空格填充），等同于 `%_d` |
| `%F` | 完整日期，类似于 `%Y-%m-%d` |
| `%g` | ISO 周的年份的最后两位数字（参见 `%G`） |
| `%G` | ISO 周的年份（参见 `%V`，通常与 `%V` 一起使用） |
| `%h` | 等同于 `%b` |
| `%H` | 小时（00..23） |
| `%I` | 小时（01..12） |
| `%j` | 一年中的第几天（001..366） |
| `%k` | 小时（空格填充，0..23），等同于 `%_H` |
| `%l` | 小时（空格填充，1..12），等同于 `%_I` |
| `%m` | 月份（01..12） |
| `%M` | 分钟（00..59） |
| `%n` | 换行符 |
| `%N` | 纳秒（000000000..999999999） |
| `%p` | 本地化的 AM 或 PM（若未知则为空） |
| `%P` | 与 `%p` 相同，但为小写形式 |
| `%q` | 一年中的季度（1..4） |
| `%r` | 本地化的 12 小时时间（例如：11:11:04 PM） |
| `%R` | 24 小时制的小时与分钟，等同于 `%H:%M` |
| `%s` | 自纪元（1970-01-01 00:00 UTC）以来的秒数（时间戳） |
| `%S` | 秒（00..60） |
| `%t` | 制表符（tab） |
| `%T` | 时间，等同于 `%H:%M:%S` |

| 格式 | 含义 |
| --- | --- |
| `%u` | 星期几（1..7，1 表示星期一） |
| `%U` | 一年中的周数（以星期日为每周第一天，00..53） |
| `%V` | ISO 周数（以星期一为每周第一天，01..53） |
| `%w` | 星期几（0..6，0 表示星期日） |
| `%W` | 一年中的周数（以星期一为每周第一天，00..53） |
| `%x` | 本地化的日期表示（例如：12/31/99） |
| `%X` | 本地化的时间表示（例如：23:13:48） |
| `%y` | 年份的最后两位数字（00..99） |
| `%Y` | 完整年份 |
| `%z` | 数字形式的时区偏移 +hhmm（例如：-0400） |
| `%:z` | 数字形式的时区偏移 +hh:mm（例如：-04:00） |
| `%::z` | 数字形式的时区偏移 +hh:mm:ss（例如：-04:00:00） |
| `%:::z` | 数字形式的时区偏移，带必要的冒号精度（例如：-04，+05:30） |

默认情况下，数值字段会使用 0 进行填充，可以在 `%` 后添加以下标识修改此类行为：

| 标识 | 含义 |
| --- | --- |
| `-` | 不填充 |
| `_` | 使用空格填充 |
| `0` | 使用 0 填充 |
| `+` | 使用 0 填充，并在大于 4 位的年份前添加 `+` |
| `^` | 使用大写 |
| `#` | 使用相反的大小写 |

## 示例

### 查看时间

```
$ date
Thu Nov 13 11:33:44 AM UTC 2025
$ date +"%Y-%m-%d %H:%M:%S"              # 指定格式
2025-11-13 11:33:44
```

```
$ date -d @1763024722                    # 查看指定的时间戳
Thu Nov 13 09:05:22 AM UTC 2025
$ date -d "next Monday"                  # 查看下周一的时间
Mon Nov 17 12:00:00 AM UTC 2025
$ date -d "2025-12-25 +7 days"           # 查看 2025-12-25 的 7 天之后的日期
Thu Jan  1 12:00:00 AM UTC 2026
```

**设置时间**

```
$ sudo date -s "2012-12-21 08:30:00"
Fri Dec 21 08:30:00 AM UTC 2012
```

> 现代操作系统通常使用 NTP 时间同步服务，使用 `date` 命令设置时间后会立刻被 NTP 覆盖。

## 推荐阅读

- Print or set system date and time - GNU Coreutils
- date(1) — Linux manual page

## 手册

```
DATE(1)                          User Commands                          DATE(1)

NAME
       date - print or set the system date and time

SYNOPSIS
       date [OPTION]... [+FORMAT]
       date [-u|--utc|--universal] [MMDDhhmm[[CC]YY][.ss]]

DESCRIPTION
       Display  date  and  time  in  the given FORMAT.  With -s, or with [MMD-
       Dhhmm[[CC]YY][.ss]], set the date and time.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -d, --date=STRING
              display time described by STRING, not 'now'

       --debug
              annotate  the  parsed date, and warn about questionable usage to
              stderr

       -f, --file=DATEFILE
              like --date; once for each line of DATEFILE
```

```
-I[FMT], --iso-8601[=FMT]
       output date/time in ISO 8601 format.  FMT='date' for  date  only
       (the  default),  'hours', 'minutes', 'seconds', or 'ns' for date
       and    time    to    the    indicated    precision.     Example:
       2006-08-14T02:34:56-06:00

--resolution
       output   the   available   resolution   of  timestamps  Example:
       0.000000001

-R, --rfc-email
       output date and time in RFC 5322 format.  Example: Mon,  14  Aug
       2006 02:34:56 -0600

--rfc-3339=FMT
       output  date/time in RFC 3339 format.  FMT='date', 'seconds', or
       'ns' for date and time to  the  indicated  precision.   Example:
       2006-08-14 02:34:56-06:00

-r, --reference=FILE
       display the last modification time of FILE

-s, --set=STRING
       set time described by STRING

-u, --utc, --universal
       print or set Coordinated Universal Time (UTC)

--help display this help and exit

--version
       output version information and exit

All  options  that  specify the date to display are mutually exclusive.
I.e.: --date, --file, --reference, --resolution.

FORMAT controls the output.  Interpreted sequences are:

%%     a literal %

%a     locale's abbreviated weekday name (e.g., Sun)

%A     locale's full weekday name (e.g., Sunday)

%b     locale's abbreviated month name (e.g., Jan)

%B     locale's full month name (e.g., January)

%c     locale's date and time (e.g., Thu Mar  3 23:05:25 2005)

%C     century; like %Y, except omit last two digits (e.g., 20)
```

```
  %d     day of month (e.g., 01)
  %D     date; same as %m/%d/%y
  %e     day of month, space padded; same as %_d
  %F     full date; like %+4Y-%m-%d
  %g     last two digits of year of ISO week number (see %G)
  %G     year of ISO week number (see %V); normally useful only with %V
  %h     same as %b
  %H     hour (00..23)
  %I     hour (01..12)
  %j     day of year (001..366)
  %k     hour, space padded ( 0..23); same as %_H
  %l     hour, space padded ( 1..12); same as %_I
  %m     month (01..12)
  %M     minute (00..59)
  %n     a newline
  %N     nanoseconds (000000000..999999999)
  %p     locale's equivalent of either AM or PM; blank if not known
  %P     like %p, but lower case
  %q     quarter of year (1..4)
  %r     locale's 12-hour clock time (e.g., 11:11:04 PM)
  %R     24-hour hour and minute; same as %H:%M
  %s     seconds since the Epoch (1970-01-01 00:00 UTC)
  %S     second (00..60)
  %t     a tab
  %T     time; same as %H:%M:%S
  %u     day of week (1..7); 1 is Monday
  %U     week number of year, with Sunday as first day of week (00..53)
```

```
       %V     ISO week number, with Monday as first day of week (01..53)

       %w     day of week (0..6); 0 is Sunday

       %W     week number of year, with Monday as first day of week (00..53)

       %x     locale's date representation (e.g., 12/31/99)

       %X     locale's time representation (e.g., 23:13:48)

       %y     last two digits of year (00..99)

       %Y     year

       %z     +hhmm numeric time zone (e.g., -0400)

       %:z    +hh:mm numeric time zone (e.g., -04:00)

       %::z   +hh:mm:ss numeric time zone (e.g., -04:00:00)

       %:::z  numeric time zone with :  to  necessary  precision  (e.g.,  -04,
              +05:30)

       %Z     alphabetic time zone abbreviation (e.g., EDT)

       By  default,  date  pads numeric fields with zeroes.  The following op-
       tional flags may follow '%':

       -      (hyphen) do not pad the field

       _      (underscore) pad with spaces

       0      (zero) pad with zeros

       +      pad with zeros, and put '+' before future years with >4 digits

       ^      use upper case if possible

       #      use opposite case if possible

       After any flags comes an optional field width,  as  a  decimal  number;
       then an optional modifier, which is either E to use the locale's alter-
       nate  representations  if available, or O to use the locale's alternate
       numeric symbols if available.

EXAMPLES
       Convert seconds since the Epoch (1970-01-01 UTC) to a date

              $ date --date='@2147483647'

       Show the time on the west coast of the US (use tzselect(1) to find TZ)
```

```
              $ TZ='America/Los_Angeles' date

       Show the local time for 9AM next Friday on the west coast of the US

              $ date --date='TZ="America/Los_Angeles" 09:00 next Fri'

DATE STRING
       The --date=STRING is a mostly free format human  readable  date  string
       such  as  "Sun, 29 Feb 2004 16:21:42 -0800" or "2004-02-29 16:21:42" or
       even "next Thursday".  A date string may contain items indicating  cal-
       endar  date,  time of day, time zone, day of week, relative time, rela-
       tive date, and numbers.  An empty string indicates the beginning of the
       day.  The date string format is more complex than is easily  documented
       here but is fully described in the info documentation.

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/date>
       or available locally via: info '(coreutils) date invocation'

GNU coreutils 9.4                April 2024                          DATE(1)
```

# Bash 的 dd 命令

```
dd [OPERAND]...
```

**功能**

复制文件并进行格式转换。

**类型**

可执行文件（`/usr/bin/dd`），属于 coreutils。

**参数**

- `OPERAND` 操作数

## 操作数

| 操作数 | 说明 |
| --- | --- |
| `bs=BYTES` | 每次读写最多 BYTES 字节（默认 512）。会覆盖 `ibs` 与 `obs` 的设置。 |
| `cbs=BYTES` | 每次转换 BYTES 字节。 |
| `conv=CONVS` | 根据逗号分隔的符号列表对文件进行转换。 |
| `count=N` | 只复制 N 个输入块。 |
| `ibs=BYTES` | 每次读取最多 BYTES 字节（默认 512）。 |
| `if=FILE` | 从 FILE 读取，而不是从标准输入读取。 |
| `iflag=FLAGS` | 按逗号分隔的标志列表控制读取行为。 |
| `obs=BYTES` | 每次写入 BYTES 字节（默认 512）。 |
| `of=FILE` | 写入到 FILE，而不是标准输出。 |
| `oflag=FLAGS` | 按逗号分隔的标志列表控制写入行为。 |
| `seek=N`（或 `oseek=N`） | 跳过 N 个 obs 大小的输出块。 |
| `skip=N`（或 `iseek=N`） | 跳过 N 个 ibs 大小的输入块。 |

| 操作数 | 说明 |
| --- | --- |
| `status=LEVEL` | 输出到 stderr 的信息级别：`none` 只显示错误；`noxfer` 隐藏最终传输统计；`progress` 显示周期性传输统计。 |

> `N` 和 `BYTES` 可用使用以下单位：c=1, w=2, b=512, kB=1000, K=1024, MB=1000*1000, M=1024*1024, xM=M 以及 G, T, P, E, Z, Y, R, Q 等二进制单位前缀.

## 转换参数

| 转换参数（`CONV`） | 说明 |
| --- | --- |
| `ascii` | 从 **EBCDIC 转换为 ASCII**。 |
| `ebcdic` | 从 **ASCII 转换为 EBCDIC**。 |
| `ibm` | 从 **ASCII 转换为另一种 IBM EBCDIC 变体**。 |
| `block` | 将以换行符结尾的记录填充空格，使其达到 `cbs` 指定大小。 |
| `unblock` | 将 `cbs` 大小记录中的尾随空格替换为换行符。 |
| `lcase` | 将大写字母转换为小写字母。 |
| `ucase` | 将小写字母转换为大写字母。 |
| `sparse` | 遇到全 NUL 的输出块时尝试使用 seek，而不是实际写入。 |
| `swab` | 每两个字节交换顺序（字节对调）。 |
| `sync` | 将每个输入块填充 NUL，使其达到 ibs 大小；与 `block` 或 `unblock` 一起使用时填充空格而非 NUL。 |
| `excl` | 如果输出文件已存在则失败。 |
| `nocreat` | 不创建输出文件。 |
| `notrunc` | 不截断输出文件。 |
| `noerror` | 读取出错后继续。 |
| `fdatasync` | 结束前将输出文件的数据物理写入磁盘。 |
| `fsync` | 行为类似，但同时写入文件元数据。 |

## Flag 参数

| Flag | 说明 |
| --- | --- |
| `append` | 追加模式（仅对输出有意义，建议搭配 `conv=notrunc`）。 |
| `direct` | 使用直接 I/O（跳过缓存）。 |
| `directory` | 失败，除非目标是目录。 |
| `dsync` | 数据使用同步 I/O。 |
| `sync` | 同上，但也同步元数据。 |
| `fullblock` | 累积完整的输入块（仅适用于 iflag）。 |
| `nonblock` | 使用非阻塞 I/O。 |
| `noatime` | 不更新访问时间。 |
| `nocache` | 请求丢弃缓存（同时可参考 `oflag=sync`）。 |
| `noctty` | 不从文件获取控制终端。 |
| `nofollow` | 不跟随符号链接。 |

## 示例

创建 100MB 的文件

```
$ dd if=/dev/zero of=file.bin bs=1M count=100
```

创建 10GB 的稀疏文件（不实际占用硬盘空间）

```
$ dd if=/dev/zero of=file.bin bs=1M count=0 seek=10G
```

将硬盘保存到镜像文件

```
$ dd if=/dev/sda of=/path/disk.img bs=4M
```

测试硬盘读取速度

```
$ dd if=/dev/sda of=/dev/null bs=4M status=progress
```

清空硬盘（危险）

```
$ dd if=/dev/zero of=/dev/sdb bs=1M
```

制作可启动 U 盘（危险）

```
$ dd if=ubuntu.iso of=/dev/sdb bs=4M status=progress oflag=sync
```

## 推荐阅读

- Convert and copy a file - GNU Coreutils
- dd(1) — Linux manual page

## 手册

```
DD(1)                            User Commands                            DD(1)

NAME
       dd - convert and copy a file

SYNOPSIS
       dd [OPERAND]...
       dd OPTION

DESCRIPTION
       Copy a file, converting and formatting according to the operands.

       bs=BYTES
              read and write up to BYTES bytes at a time (default: 512); over-
              rides ibs and obs

       cbs=BYTES
              convert BYTES bytes at a time

       conv=CONVS
              convert the file as per the comma separated symbol list

       count=N
              copy only N input blocks

       ibs=BYTES
              read up to BYTES bytes at a time (default: 512)

       if=FILE
              read from FILE instead of stdin

       iflag=FLAGS
              read as per the comma separated symbol list

       obs=BYTES
              write BYTES bytes at a time (default: 512)

       of=FILE
              write to FILE instead of stdout
```

```
oflag=FLAGS
       write as per the comma separated symbol list

seek=N (or oseek=N) skip N obs-sized output blocks

skip=N (or iseek=N) skip N ibs-sized input blocks

status=LEVEL
       The  LEVEL  of information to print to stderr; 'none' suppresses
       everything but error messages,  'noxfer'  suppresses  the  final
       transfer  statistics, 'progress' shows periodic transfer statis-
       tics

N and BYTES may be followed by the following  multiplicative  suffixes:
c=1,  w=2,  b=512,  kB=1000,  K=1024,  MB=1000*1000, M=1024*1024, xM=M,
GB=1000*1000*1000, G=1024*1024*1024, and so on for T, P, E, Z, Y, R, Q.
Binary prefixes can be used, too: KiB=K, MiB=M, and so on.  If  N  ends
in 'B', it counts bytes not blocks.

Each CONV symbol may be:

ascii  from EBCDIC to ASCII

ebcdic from ASCII to EBCDIC

ibm    from ASCII to alternate EBCDIC

block  pad newline-terminated records with spaces to cbs-size

unblock
       replace trailing spaces in cbs-size records with newline

lcase  change upper case to lower case

ucase  change lower case to upper case

sparse try to seek rather than write all-NUL output blocks

swab   swap every pair of input bytes

sync   pad  every  input  block  with  NULs to ibs-size; when used with
       block or unblock, pad with spaces rather than NULs

excl   fail if the output file already exists

nocreat
       do not create the output file

notrunc
       do not truncate the output file

noerror
```

```
               continue after read errors

        fdatasync
               physically write output file data before finishing

        fsync  likewise, but also write metadata

        Each FLAG symbol may be:

        append append mode (makes sense  only  for  output;  conv=notrunc  sug-
               gested)

        direct use direct I/O for data

        directory
               fail unless a directory

        dsync  use synchronized I/O for data

        sync   likewise, but also for metadata

        fullblock
               accumulate full blocks of input (iflag only)

        nonblock
               use non-blocking I/O

        noatime
               do not update access time

        nocache
               Request to drop cache.  See also oflag=sync

        noctty do not assign controlling terminal from file

        nofollow
               do not follow symlinks

        Sending a USR1 signal to a running 'dd' process makes it print I/O sta-
        tistics to standard error and then resume copying.

        Options are:

        --help display this help and exit

        --version
               output version information and exit

AUTHOR
        Written by Paul Rubin, David MacKenzie, and Stuart Kemp.

REPORTING BUGS
        GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
```

```
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/dd>
       or available locally via: info '(coreutils) dd invocation'

GNU coreutils 9.4                April 2024                            DD(1)
```

# Bash 的 delpart 命令

```
delpart device partition
```

**功能**

通知内核移除分区（不会实际删除分区）。

**类型**

可执行文件（`/usr/bin/delpart`），属于 util-linux。

**参数**

- `device` - 硬盘设备
- `partition` - 分区号

## 示例

```
$ sudo delpart /dev/sda 3
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| addpart | 添加硬盘分区 |

## 推荐阅读

- delpart(8) — Linux manual page

## 手册

```
DELPART(8)                  System Administration                  DELPART(8)

NAME
       delpart - tell the kernel to forget about a partition

SYNOPSIS
```

```
       delpart device partition

DESCRIPTION
       delpart asks the Linux kernel to forget about the specified partition
       (a number) on the specified device. The command is a simple wrapper
       around the "del partition" ioctl.

       This command doesn’t manipulate partitions on a block device.

OPTIONS
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

SEE ALSO
       addpart(8), fdisk(8), parted(8), partprobe(8), partx(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The delpart command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                        DELPART(8)
```

# Bash 的 df 命令

```
df [OPTION]... [FILE]...
```

**功能**

查看文件系统的空间占用。

**类型**

可执行文件（`/usr/bin/df`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 包括伪文件系统、重复的文件系统以及无法访问的文件系统
  - `-B`，`--block-size=SIZE` - 指定单位，例如 `-B 4K` 表示以 4096 字节为单位
  - `-h`，`--human-readable` - 自动添加单位以便人类阅读，1K 表示 1024 字节
  - `-H`，`--si` - 自动添加单位以便人类阅读，1K 表示 1000 字节
  - `-i`，`--inodes` - 列出 inode 信息而不是块占用
  - `-k` - 等价于 `-B 1K`
  - `-l`，`--local` - 范围限制在本地文件系统
  - `--no-sync` - 不在获取占用信息前调用 sync（默认）
  - `--output[=FIELD_LIST]` - 使用 `FIELD_LIST` 指定的输出格式；有效值：
    - `source` - 文件系统源路径
    - `fstype` - 文件系统类型
    - `itotal` - inode 总数
    - `iused` - inode 占用数
    - `iavail` - inode 可用数
    - `ipcent` - inode 占用百分比
    - `size` - 总空间
    - `used` - 占用空间
    - `avail` - 可用空间
    - `pcent` - 空间占用百分比
    - `file` - 文件路径（对应参数中的 `FILE`）
    - `target` - 挂载路径
  - `-P`，`--portability` - 使用 POSIX 输出格式
  - `--sync` - 在获取占用信息前自动调用 sync

    - `--total` - 忽略对可用空间影响不大的条目并得出总计数据
    - `-t`, `--type=TYPE` - 范围限制在 `TYPE` 指定的文件系统类型
    - `-T`, `--print-type` - 打印文件系统类型
    - `-x`, `--exclude-type=TYPE` - 忽略 `TYPE` 指定的文件系统类型
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `FILE` - 文件列表；默认查看所有文件系统

## 示例

### 基本使用

```
$ df -h
Filesystem      Size  Used Avail Use% Mounted on
tmpfs           1.4G  1.7M  1.4G   1% /run
efivarfs        128K   13K  111K  10% /sys/firmware/efi/efivars
/dev/nvme0n1p2  915G  172G  697G  20% /
tmpfs           6.8G     0  6.8G   0% /dev/shm
tmpfs           5.0M     0  5.0M   0% /run/lock
/dev/nvme0n1p1  1.1G  6.2M  1.1G   1% /boot/efi
tmpfs           1.4G   16K  1.4G   1% /run/user/1000
```

### 指定格式

```
$ df -h --output=source,fstype,used,avail,size,pcent
Filesystem     Type      Used Avail  Size Use%
tmpfs          tmpfs     1.7M  1.4G  1.4G   1%
efivarfs       efivarfs   13K  111K  128K  10%
/dev/nvme0n1p2 ext4      172G  697G  915G  20%
tmpfs          tmpfs        0  6.8G  6.8G   0%
tmpfs          tmpfs        0  5.0M  5.0M   0%
/dev/nvme0n1p1 vfat      6.2M  1.1G  1.1G   1%
tmpfs          tmpfs      16K  1.4G  1.4G   1%
```

## 推荐阅读

- Report file system space usage - GNU Coreutils
- df(1) — Linux manual page

## 手册

```
DF(1)                          User Commands                          DF(1)
```

NAME

df - report file system space usage

SYNOPSIS

df [OPTION]... [FILE]...

DESCRIPTION

This manual page documents the GNU version of df. df displays the amount of space available on the file system containing each file name argument. If no file name is given, the space available on all currently mounted file systems is shown. Space is shown in 1K blocks by default, unless the environment variable POSIXLY_CORRECT is set, in which case 512-byte blocks are used.

If an argument is the absolute file name of a device node containing a mounted file system, df shows the space available on that file system rather than on the file system containing the device node. This version of df cannot show the space available on unmounted file systems, because on most kinds of systems doing so requires non-portable intimate knowledge of file system structures.

OPTIONS

Show information about the file system on which each FILE resides, or all file systems by default.

Mandatory arguments to long options are mandatory for short options too.

-a, --all
include pseudo, duplicate, inaccessible file systems

-B, --block-size=SIZE
scale sizes by SIZE before printing them; e.g., '-BM' prints sizes in units of 1,048,576 bytes; see SIZE format below

-h, --human-readable
print sizes in powers of 1024 (e.g., 1023M)

-H, --si
print sizes in powers of 1000 (e.g., 1.1G)

-i, --inodes
list inode information instead of block usage

-k like --block-size=1K

-l, --local
limit listing to local file systems

--no-sync
do not invoke sync before getting usage info (default)

--output[=FIELD_LIST]

```
              use the output format defined by FIELD_LIST, or print all fields
              if FIELD_LIST is omitted.

       -P, --portability
              use the POSIX output format

       --sync invoke sync before getting usage info

       --total
              elide all entries insignificant to available space, and  produce
              a grand total

       -t, --type=TYPE
              limit listing to file systems of type TYPE

       -T, --print-type
              print file system type

       -x, --exclude-type=TYPE
              limit listing to file systems not of type TYPE

       -v     (ignored)

       --help display this help and exit

       --version
              output version information and exit

       Display  values  are  in  units  of  the  first  available  SIZE  from
       --block-size, and the DF_BLOCK_SIZE, BLOCK_SIZE and BLOCKSIZE  environ-
       ment  variables.   Otherwise,  units  default  to 1024 bytes (or 512 if
       POSIXLY_CORRECT is set).

       The SIZE argument is an integer and  optional  unit  (example:  10K  is
       10*1024).   Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or KB,MB,...
       (powers of 1000).  Binary prefixes can be used, too: KiB=K, MiB=M,  and
       so on.

       FIELD_LIST  is a comma-separated list of columns to be included.  Valid
       field names are: 'source', 'fstype', 'itotal', 'iused', 'iavail',  'ip-
       cent',  'size', 'used', 'avail', 'pcent', 'file' and 'target' (see info
       page).

AUTHOR
       Written by Torbjorn Granlund, David MacKenzie, and Paul Eggert.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
```

```
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/df>
       or available locally via: info '(coreutils) df invocation'

GNU coreutils 9.4                April 2024                             DF(1)
```

# Bash 的 dir 命令

```
dir [OPTION]... [FILE]...
```

**功能**

列出目录内容。

和 ls 命令基本等价，区别在于默认使用 `-C` 和 `-b` 选项。

**类型**

可执行文件（`/usr/bin/dir`），属于 coreutils。

**参数**

- `OPTION` 选项：
- `-a`，`--all` - 显示隐藏文件（文件名以 `.` 开头）
  - `-A`，`--almost-all` - 显示隐藏文件，但不显示 `.` 和 `..`
  - `--author` - 显示文件的作者；和 `-l` 选项一起使用
  - `-b`，`--escape` - 非 图形字符 显示为 C 风格的 转义字符
  - `--block-size=SIZE` - 文件大小的单位；和 `-l` 选项一起使用
  - `-B`，`--ignore-backups` 不显示备份文件（文件名以 `~` 结束）
  - `-c` - 按 ctime（文件状态信息最后更新的时间）排序；和 `-lt` 选项一起使用
  - `-C` - 竖向排列
  - `--color[=WHEN]` - 何时使用彩色，有效值：`always`，`auto`，`never`
  - `-d`，`--directory` - 查看目录本身，而不是目录的内容
  - `-D`，`--dired` - 生成为 Emacs 的 dired 模式设计的输出
  - `-f` - 显示所有文件且不进行排序（按照目录顺序显示）
  - `-F`，`--classify[=WHEN]` - 何时添加 文件类型指示器，有效值：`always`，`auto`，`never`
  - `--file-type` - 显示文件类型指示器；同 `--classify` 但是不添加 `*`
  - `--format=WORD` - 格式，有效值：`across`，`commas`，`horizontal`，`long`，`single-column`，`verbose`，`vertical`
  - `--full-time` - 显示完整的时间；同 `-l --time-style=full-iso`
  - `-g` - 显示文件详情；同 `-l` 但不显示所有者
  - `--group-directories-first` - 分组时将目录排在文件前面
  - `-G`，`--no-group` - 不显示分组名
  - `-h`，`--human-readable` - 让详细信息更适合人类阅读；和 `-l` 或 `-s` 选择一起使用

- `--si` - 让详细信息更适合人类阅读；同 `-h` 但是文件大小以 1000 为单位，而不是 1024
- `-H`，`--dereference-command-line` - 参数是符号链接时，查看源文件
- `--dereference-command-line-symlink-to-dir` - 参数是目录的符号链接时，查看源目录
- `--hide=PATTERN` - 隐藏匹配的文件
- `--hyperlink[=WHEN]` - 何时将文件名显示为超链接（现代终端上点击可以打开文件管理器），有效值：`always`，`auto`，`never`
- `--indicator-style=WORD` - 添加 文件类型指示器 ，有效值：`none`，`slash`，`file-type`，`classify`
- `-i`，`--inode` - 显示文件的 inode 编号
- `-I`，`--ignore=PATTERN` - 不显示匹配的文件
- `-k`，`--kibibytes` - 文件大小以 1024 字节为单位
- `-l` - 显示详细信息
- `-L`，`--dereference` - 符号链接显示源文件的信息，而非符号链接本身的信息
- `-m` - 用逗号分隔条目
- `-n`，`--numeric-uid-gid` - 显示详细信息，但是所有者和所属组显示为 ID 而非名称
- `-N`，`--literal` - 含特殊字符的文件名不添加单引号
- `-o` - 显示详细信息，但不显示所属组信息
- `-p`，`--indicator-style=slash` - 添加 `/` 指示器
- `-q`，`--hide-control-chars` - 非 图形字符 显示为 `?`
- `--show-control-chars` - 直接显示非 图形字符
- `-Q`，`--quote-name` - 用双引号包裹文件名
- `--quoting-style=WORD` - 引号风格，有效值：`literal`，`locale`，`shell`，`shell-always`，`shell-escape`，`shell-escape-always`，`c`，`escape`
- `-r`，`--reverse` - 逆序
- `-R`，`--recursive` - 递归查看子目录
- `-s`，`--size` - 显示文件分配的块数
- `-S` - 按文件大小排序，最大的优先
- `--sort=WORD` - 排序规则，有效值：`none`，`size`，`time`，`version`，`extension`，`width`
- `--time=WORD` - 显示或排序的时间，有效值：`atime`，`ctime`，`mtime`，`btime`
- `--time-style=TIME_STYLE` - 日期和时间的格式
- `-t` - 按时间排序
- `-T`，`--tabsize=COLS` - 制表符的宽度
- `-u` - 显示 atime 或按 atime 排序
- `-U` - 不排序
- `-v` - 自然排序
- `-w`，`--width=COLS` - 一行的宽度，0 表示无限
- `-x` - 横向排列
- `-X` - 按扩展名排序

    - `-Z`, `--context` - 显示安全上下文
    - `--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
    - `-1` - 每行只显示一个文件
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `FILE` - 文件列表

## 推荐阅读

- Briefly list directory contents - GNU Coreutils
- dir(1) — Linux manual page

## 手册

```
DIR(1)                          User Commands                          DIR(1)

NAME
       dir - list directory contents

SYNOPSIS
       dir [OPTION]... [FILE]...

DESCRIPTION
       List  information  about  the FILEs (the current directory by default).
       Sort entries alphabetically if none of -cftuvSUX nor --sort  is  speci-
       fied.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -a, --all
              do not ignore entries starting with .

       -A, --almost-all
              do not list implied . and ..

       --author
              with -l, print the author of each file

       -b, --escape
              print C-style escapes for nongraphic characters

       --block-size=SIZE
              with  -l,  scale  sizes  by  SIZE  when  printing  them;  e.g.,
              '--block-size=M'; see SIZE format below

       -B, --ignore-backups
```

```
       do not list implied entries ending with ~

-c     with  -lt: sort by, and show, ctime (time of last change of file
       status information); with -l: show ctime and sort by name;  oth-
       erwise: sort by ctime, newest first

-C     list entries by columns

--color[=WHEN]
       color the output WHEN; more info below

-d, --directory
       list directories themselves, not their contents

-D, --dired
       generate output designed for Emacs' dired mode

-f     list all entries in directory order

-F, --classify[=WHEN]
       append indicator (one of */=>@|) to entries WHEN

--file-type
       likewise, except do not append '*'

--format=WORD
       across  -x, commas -m, horizontal -x, long -l, single-column -1,
       verbose -l, vertical -C

--full-time
       like -l --time-style=full-iso

-g     like -l, but do not list owner

--group-directories-first
       group directories before files; can be augmented with  a  --sort
       option, but any use of --sort=none (-U) disables grouping

-G, --no-group
       in a long listing, don't print group names

-h, --human-readable
       with -l and -s, print sizes like 1K 234M 2G etc.

--si   likewise, but use powers of 1000 not 1024

-H, --dereference-command-line
       follow symbolic links listed on the command line

--dereference-command-line-symlink-to-dir
       follow  each  command line symbolic link that points to a direc-
       tory
```

```
--hide=PATTERN
       do not list implied entries matching shell  PATTERN  (overridden
       by -a or -A)

--hyperlink[=WHEN]
       hyperlink file names WHEN

--indicator-style=WORD
       append indicator with style WORD to entry names: none (default),
       slash (-p), file-type (--file-type), classify (-F)

-i, --inode
       print the index number of each file

-I, --ignore=PATTERN
       do not list implied entries matching shell PATTERN

-k, --kibibytes
       default  to  1024-byte  blocks  for file system usage; used only
       with -s and per directory totals

-l     use a long listing format

-L, --dereference
       when showing file information for a symbolic link, show informa-
       tion for the file the link references rather than for  the  link
       itself

-m     fill width with a comma separated list of entries

-n, --numeric-uid-gid
       like -l, but list numeric user and group IDs

-N, --literal
       print entry names without quoting

-o     like -l, but do not list group information

-p, --indicator-style=slash
       append / indicator to directories

-q, --hide-control-chars
       print ? instead of nongraphic characters

--show-control-chars
       show nongraphic characters as-is (the default, unless program is
       'ls' and output is a terminal)

-Q, --quote-name
       enclose entry names in double quotes

--quoting-style=WORD
       use  quoting style WORD for entry names: literal, locale, shell,
```

```
       shell-always, shell-escape, shell-escape-always, c, escape
       (overrides QUOTING_STYLE environment variable)

-r, --reverse
       reverse order while sorting

-R, --recursive
       list subdirectories recursively

-s, --size
       print the allocated size of each file, in blocks

-S     sort by file size, largest first

--sort=WORD
       sort by WORD instead of name: none (-U), size (-S), time (-t),
       version (-v), extension (-X), width

--time=WORD
       select which timestamp used to display or sort; access time
       (-u): atime, access, use; metadata change time (-c): ctime, sta-
       tus; modified time (default): mtime, modification; birth time:
       birth, creation;

       with -l, WORD determines which time to show; with --sort=time,
       sort by WORD (newest first)

--time-style=TIME_STYLE
       time/date format with -l; see TIME_STYLE below

-t     sort by time, newest first; see --time

-T, --tabsize=COLS
       assume tab stops at each COLS instead of 8

-u     with -lt: sort by, and show, access time; with -l: show access
       time and sort by name; otherwise: sort by access time, newest
       first

-U     do not sort; list entries in directory order

-v     natural sort of (version) numbers within text

-w, --width=COLS
       set output width to COLS.  0 means no limit

-x     list entries by lines instead of by columns

-X     sort alphabetically by entry extension

-Z, --context
       print any security context of each file
```

```
       --zero end each output line with NUL, not newline

       -1     list one file per line

       --help display this help and exit

       --version
              output version information and exit

       The  SIZE  argument  is  an  integer and optional unit (example: 10K is
       10*1024).  Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or  KB,MB,...
       (powers  of 1000).  Binary prefixes can be used, too: KiB=K, MiB=M, and
       so on.

       The TIME_STYLE argument can be  full-iso,  long-iso,  iso,  locale,  or
       +FORMAT.   FORMAT  is  interpreted  like in date(1).  If FORMAT is FOR-
       MAT1<newline>FORMAT2, then FORMAT1 applies to non-recent files and FOR-
       MAT2 to recent files.  TIME_STYLE prefixed with 'posix-'  takes  effect
       only  outside  the POSIX locale.  Also the TIME_STYLE environment vari-
       able sets the default style to use.

       The WHEN argument defaults to  'always'  and  can  also  be  'auto'  or
       'never'.

       Using  color  to distinguish file types is disabled both by default and
       with --color=never.  With --color=auto, ls emits color codes only  when
       standard  output is connected to a terminal.  The LS_COLORS environment
       variable can change the settings.  Use the dircolors(1) command to  set
       it.

   Exit status:
       0      if OK,

       1      if minor problems (e.g., cannot access subdirectory),

       2      if serious trouble (e.g., cannot access command-line argument).

AUTHOR
       Written by Richard M. Stallman and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/dir>
       or available locally via: info '(coreutils) dir invocation'
```

# Bash 的 dircolors 命令

```
dircolors [OPTION]... [FILE]
```

**功能**

查看或设置 ls 和 dir 命令的颜色。

**类型**

可执行文件（`/usr/bin/dircolors`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--sh`，`--bourne-shell` - 以 Bash 代码的格式输出环境变量 `LS_COLORS`
  - `-c`，`--csh`，`--c-shell` - 以 C shell 代码的格式输出环境变量 `LS_COLORS`
  - `-p`，`--print-database` - 打印默认值
  - `--print-ls-colors` - 以彩色打印颜色配置
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 使用此文件指定颜色

## 示例

**查看环境变量**

```
$ dircolors -b      # 以 Bash 代码的格式输出
LS_COLORS='rs=0:di=01;34:ln=01;36:mh=00:pi=40;33:so=01;35:do=01;35:bd=40;33;01:cd=40;33;01:or=40;31;01:mi=00:su=37;41:sg=30;43:ca=00:tw=30
;42:ow=34;42:st=37;44:ex=01;32:*.tar=01;31:*.tgz=01;31:*.arc=01;31:*.arj=01;31:*.taz=01;31:*.lha=01;31:*.lz4=01;31:*.lzh=01;31:*.lzma=01;3
1:*.tlz=01;31:*.txz=01;31:*.tzo=01;31:*.t7z=01;31:*.zip=01;31:*.z=01;31:*.dz=01;31:*.gz=01;31:*.lrz=01;31:*.lz=01;31:*.lzo=01;31:*.xz=01;3
1:*.zst=01;31:*.tzst=01;31:*.bz2=01;31:*.bz=01;31:*.tbz=01;31:*.tbz2=01;31:*.tz=01;31:*.deb=01;31:*.rpm=01;31:*.jar=01;31:*.war=01;31:*.ea
r=01;31:*.sar=01;31:*.rar=01;31:*.alz=01;31:*.ace=01;31:*.zoo=01;31:*.cpio=01;31:*.7z=01;31:*.rz=01;31:*.cab=01;31:*.wim=01;31:*.swm=01;31
:*.dwm=01;31:*.esd=01;31:*.avif=01;35:*.jpg=01;35:*.jpeg=01;35:*.mjpg=01;35:*.mjpeg=01;35:*.gif=01;35:*.bmp=01;35:*.pbm=01;35:*.pgm=01;35:
*.ppm=01;35:*.tga=01;35:*.xbm=01;35:*.xpm=01;35:*.tif=01;35:*.tiff=01;35:*.png=01;35:*.svg=01;35:*.svgz=01;35:*.mng=01;35:*.pcx=01;35:*.mo
v=01;35:*.mpg=01;35:*.mpeg=01;35:*.m2v=01;35:*.mkv=01;35:*.webm=01;35:*.webp=01;35:*.ogm=01;35:*.mp4=01;35:*.m4v=01;35:*.mp4v=01;35:*.vob=
01;35:*.qt=01;35:*.nuv=01;35:*.wmv=01;35:*.asf=01;35:*.rm=01;35:*.rmvb=01;35:*.flc=01;35:*.avi=01;35:*.fli=01;35:*.flv=01;35:*.gl=01;35:*.
dl=01;35:*.xcf=01;35:*.xwd=01;35:*.yuv=01;35:*.cgm=01;35:*.emf=01;35:*.ogv=01;35:*.ogx=01;35:*.aac=00;36:*.au=00;36:*.flac=00;36:*.m4a=00;
36:*.mid=00;36:*.midi=00;36:*.mka=00;36:*.mp3=00;36:*.mpc=00;36:*.ogg=00;36:*.ra=00;36:*.wav=00;36:*.oga=00;36:*.opus=00;36:*.spx=00;36:*.
xspf=00;36:*~=00;90:*#=00;90:*.bak=00;90:*.crdownload=00;90:*.dpkg-dist=00;90:*.dpkg-new=00;90:*.dpkg-old=00;90:*.dpkg-
tmp=00;90:*.old=00;90:*.orig=00;90:*.part=00;90:*.rej=00;90:*.rpmnew=00;90:*.rpmorig=00;90:*.rpmsave=00;90:*.swp=00;90:*.tmp=00;90:*.ucf-
dist=00;90:*.ucf-new=00;90:*.ucf-old=00;90:';
export LS_COLORS
$ dircolors -c      # 以 C shell 代码的格式输出
setenv LS_COLORS
'rs=0:di=01;34:ln=01;36:mh=00:pi=40;33:so=01;35:do=01;35:bd=40;33;01:cd=40;33;01:or=40;31;01:mi=00:su=37;41:sg=30;43:ca=00:tw=30;42:ow=34;
42:st=37;44:ex=01;32:*.tar=01;31:*.tgz=01;31:*.arc=01;31:*.arj=01;31:*.taz=01;31:*.lha=01;31:*.lz4=01;31:*.lzh=01;31:*.lzma=01;31:*.tlz=01
;31:*.txz=01;31:*.tzo=01;31:*.t7z=01;31:*.zip=01;31:*.z=01;31:*.dz=01;31:*.gz=01;31:*.lrz=01;31:*.lz=01;31:*.lzo=01;31:*.xz=01;31:*.zst=01
;31:*.tzst=01;31:*.bz2=01;31:*.bz=01;31:*.tbz=01;31:*.tbz2=01;31:*.tz=01;31:*.deb=01;31:*.rpm=01;31:*.jar=01;31:*.war=01;31:*.ear=01;31:*.
sar=01;31:*.rar=01;31:*.alz=01;31:*.ace=01;31:*.zoo=01;31:*.cpio=01;31:*.7z=01;31:*.rz=01;31:*.cab=01;31:*.wim=01;31:*.swm=01;31:*.dwm=01;
31:*.esd=01;31:*.avif=01;35:*.jpg=01;35:*.jpeg=01;35:*.mjpg=01;35:*.mjpeg=01;35:*.gif=01;35:*.bmp=01;35:*.pbm=01;35:*.pgm=01;35:*.ppm=01;3
5:*.tga=01;35:*.xbm=01;35:*.xpm=01;35:*.tif=01;35:*.tiff=01;35:*.png=01;35:*.svg=01;35:*.svgz=01;35:*.mng=01;35:*.pcx=01;35:*.mov=01;35:*.
mpg=01;35:*.mpeg=01;35:*.m2v=01;35:*.mkv=01;35:*.webm=01;35:*.webp=01;35:*.ogm=01;35:*.mp4=01;35:*.m4v=01;35:*.mp4v=01;35:*.vob=01;35:*.qt
=01;35:*.nuv=01;35:*.wmv=01;35:*.asf=01;35:*.rm=01;35:*.rmvb=01;35:*.flc=01;35:*.avi=01;35:*.fli=01;35:*.flv=01;35:*.gl=01;35:*.dl=01;35:*
.xcf=01;35:*.xwd=01;35:*.yuv=01;35:*.cgm=01;35:*.emf=01;35:*.ogv=01;35:*.ogx=01;35:*.aac=00;36:*.au=00;36:*.flac=00;36:*.m4a=00;36:*.mid=0
0;36:*.midi=00;36:*.mka=00;36:*.mp3=00;36:*.mpc=00;36:*.ogg=00;36:*.ra=00;36:*.wav=00;36:*.oga=00;36:*.opus=00;36:*.spx=00;36:*.xspf=00;36
:*~=00;90:*#=00;90:*.bak=00;90:*.crdownload=00;90:*.dpkg-dist=00;90:*.dpkg-new=00;90:*.dpkg-old=00;90:*.dpkg-
tmp=00;90:*.old=00;90:*.orig=00;90:*.part=00;90:*.rej=00;90:*.rpmnew=00;90:*.rpmorig=00;90:*.rpmsave=00;90:*.swp=00;90:*.tmp=00;90:*.ucf-
dist=00;90:*.ucf-new=00;90:*.ucf-old=00;90:'
```

**查看颜色**

```
$ dircolors --print-ls-colors
```

输出：

```
3333333333333333333333333333333rrrrrrrrrrrrrrrrrrrrrrrrrrrrrr
rs      0
di      01;34
ln      01;36
mh      00
pi      40;33
so      01;35
do      01;35
bd      40;33;01
cd      40;33;01
or      40;31;01
```

推荐阅

- Briefly list directory contents - GNU Coreutils
- dir(1) — Linux manual page

## 手册

```
DIRCOLORS(1)                    User Commands                    DIRCOLORS(1)

NAME
       dircolors - color setup for ls

SYNOPSIS
       dircolors [OPTION]... [FILE]

DESCRIPTION
       Output commands to set the LS_COLORS environment variable.

   Determine format of output:
       -b, --sh, --bourne-shell
              output Bourne shell code to set LS_COLORS

       -c, --csh, --c-shell
              output C shell code to set LS_COLORS

       -p, --print-database
              output defaults

       --print-ls-colors
              output fully escaped colors for display

       --help display this help and exit

       --version
              output version information and exit

       If  FILE  is  specified,  read  it to determine which colors to use for
       which file types and extensions.  Otherwise, a precompiled database  is
       used.   For  details  on  the  format  of  these  files, run 'dircolors
       --print-database'.

AUTHOR
       Written by H. Peter Anvin.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/dircolors>
       or available locally via: info '(coreutils) dircolors invocation'

GNU coreutils 9.4                April 2024                     DIRCOLORS(1)
```

# Bash 的 dirname 命令

```
dirname OPTION... NAME...
```

### 功能

提取文件路径中的目录（从文件路径中删除最后一个组成部分）。

### 类型

可执行文件（`/usr/bin/dirname`），属于 coreutils。

### 参数

- `OPTION` 选项：
  - `-z`, `--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NAME` - 文件路径列表；不要求文件存在

## 示例

### 绝对路径

```
$ dirname /path/to/file.txt
/path/to
$ dirname /path/to
/path
$ dirname /path
/
$ dirname /         # 根目录会得到自身
/
```

### 相对路径

```
$ dirname ./path/to/file.txt
./path/to
$ dirname path/to/file.txt
path/to
$ dirname ./path/../path/to/file.txt
./path/../path/to
```

- 只会删除路径的最后一个组成部分，不会简化路径

## 相关命令

| 命令 | 说明 |
| --- | --- |
| basename | 从文件路径中去除目录和后缀 |
| realpath | 解析路径 |

## 推荐阅读

- Strip last file name component - GNU Coreutils
- dirname(1) — Linux manual page

## 手册

```
DIRNAME(1)                      User Commands                      DIRNAME(1)

NAME
       dirname - strip last component from file name

SYNOPSIS
       dirname [OPTION] NAME...

DESCRIPTION
       Output each NAME with its last non-slash component and trailing slashes
       removed;  if  NAME contains no /'s, output '.' (meaning the current di-
       rectory).

       -z, --zero
              end each output line with NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

EXAMPLES
       dirname /usr/bin/
              -> "/usr"

       dirname dir1/str dir2/str
              -> "dir1" followed by "dir2"

       dirname stdio.h
              -> "."
```

AUTHOR
       Written by David MacKenzie and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       basename(1), readlink(1)

       Full documentation <https://www.gnu.org/software/coreutils/dirname>
       or available locally via: info '(coreutils) dirname invocation'

# Bash 的 dmesg 命令

```
dmesg [OPTIONS]
```

**功能**

查看或控制内核消息缓冲。

**类型**

可执行文件（`/usr/bin/dmesg`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-C`，`--clear` - 清空消息缓冲
  - `-c`，`--read-clear` - 显示消息后清空缓存
  - `-D`，`--console-off` - 停止向控制台（通常是串口）打印消息
  - `-d`，`--show-delta` - 显示时间戳及消息之间的时间间隔
  - `-E`，`--console-on` - 向控制台（通常是串口）打印消息
  - `-e`，`--reltime` - 以易于理解的格式打印本地时间
  - `-F`，`--file FILE` - 从 `FILE` 文件中读取内核消息
  - `-f`，`--facility LIST` - 仅查看 `LIST` 列出来源的消息，以逗号分隔
  - `-H`，`--human` - 启用人类可读输出
  - `-J`，`--json` - 以 JSON 格式输出
  - `-k`，`--kernel` - 仅显示内核消息
  - `-L`，`--color[=WHEN]` - 根据 `WHEN` 设置输出颜色；取值为：`auto`，`never`，`always`
  - `-l`，`--level LIST` - 仅查看 `LIST` 列出级别的消息，以逗号分隔，加号（`+`）同时包含更高级别；从低到高依次伪：
    - `debug` - 调试
    - `info` - 常规信息
    - `notice` - 重要信息
    - `warn` - 警告
    - `err` - 错误
    - `crit` - 严重错误
    - `alert` - 警报
    - `emerg` - 紧急
  - `-n`，`--console-level LEVEL` - 向控制台（通常是串口）打印的消息级别

- `--noescape` - 停止转义不可打印字符
- `-P`, `--nopager` - 不分页
- `-p`, `--force-prefix` - 强制显示前缀
- `-r`, `--raw` - 打印原始消息缓冲区
- `-S`, `--syslog` - 通过 [syslog(2)] 接口读取内核消息
- `-s`, `--buffer-size SIZE` - 使用大小伪 `SIZE` 的缓冲查询内核消息缓冲；默认为 16392
- `-T`, `--ctime` - 显示人类可读的时间
- `--since TIME` - 显示 `TIME` （含）以后的消息
- `--until time` - 显示 `TIME` （不含）之前的消息
- `-t`, `--notime` - 不打印内核时间戳
- `--time-format FORMAT` - 使用 `FORMAT` 指定的格式打印时间
- `-u`, `--userspace` - 打印用户空间的消息
- `-w`, `--follow` - 等待新消息
- `-W`, `--follow-new` - 等待且仅显示新消息
- `-x`, `--decode` - 将来源和级别的数值解析为任何可读的前缀
- `--help` - 显示帮助
- `--version` - 显示版本

## 示例

```
$ dmesg
[    0.000000] Linux version 6.8.0-87-generic (buildd@lcy02-amd64-034) (x86_64-linux-gnu-gcc-13 (Ubuntu 13.3.0-6ubuntu2~24.04) 13.3.0, GNU ld (GNU Binutils for Ubuntu) 2.42) #88-Ubuntu SMP PREEMPT_DYNAMIC Sat Oct 11 09:28:41 UTC 2025 (Ubuntu 6.8.0-87.88-generic 6.8.12)
[    0.000000] Command line: BOOT_IMAGE=/boot/vmlinuz-6.8.0-87-generic root=UUID=8b2c9bcc-94f2-4cf2-a9d0-581e58f9d676 ro
[    0.000000] KERNEL supported cpus:
[    0.000000]   Intel GenuineIntel
[    0.000000]   AMD AuthenticAMD
[    0.000000]   Hygon HygonGenuine
[    0.000000]   Centaur CentaurHauls
[    0.000000]   zhaoxin   Shanghai
[    0.000000] BIOS-provided physical RAM map:
[    0.000000] BIOS-e820: [mem 0x0000000000000000-0x000000000009ffff] usable
...
```

## 推荐阅读

- dmesg(1) — Linux manual page

# 手册

```
DMESG(1)                        User Commands                        DMESG(1)

NAME
       dmesg - print or control the kernel ring buffer

SYNOPSIS
       dmesg [options]

       dmesg --clear

       dmesg --read-clear [options]

       dmesg --console-level level

       dmesg --console-on

       dmesg --console-off

DESCRIPTION
       dmesg is used to examine or control the kernel ring buffer.

       The default action is to display all messages from the kernel ring
       buffer.

OPTIONS
       The --clear, --read-clear, --console-on, --console-off, and
       --console-level options are mutually exclusive.

       -C, --clear
           Clear the ring buffer.

       -c, --read-clear
           Clear the ring buffer after first printing its contents.

       -D, --console-off
           Disable the printing of messages to the console.

       -d, --show-delta
           Display the timestamp and the time delta spent between messages. If
           used together with --notime then only the time delta without the
           timestamp is printed.

       -E, --console-on
           Enable printing messages to the console.

       -e, --reltime
           Display the local time and the delta in human-readable format. Be
           aware that conversion to the local time could be inaccurate (see -T
           for more details).
```

```
-F, --file file
    Read the syslog messages from the given file. Note that -F does not
    support messages in kmsg format. The old syslog format is supported
    only.

-f, --facility list
    Restrict output to the given (comma-separated) list of facilities.
    For example:

    dmesg --facility=daemon

    will print messages from system daemons only. For all supported
    facilities see the --help output.

-H, --human
    Enable human-readable output. See also --color, --reltime and
    --nopager.

-J, --json
    Use JSON output format. The time output format is in "sec.usec"
    format only, log priority level is not decoded by default (use
    --decode to split into facility and priority), the other options to
    control the output format or time format are silently ignored.

-k, --kernel
    Print kernel messages.

-L, --color[=when]
    Colorize the output. The optional argument when can be auto, never
    or always. If the when argument is omitted, it defaults to auto.
    The colors can be disabled; for the current built-in default see
    the --help output. See also the COLORS section below.

-l, --level list
    Restrict output to the given (comma-separated) list of levels. For
    example:

    dmesg --level=err,warn

    will print error and warning messages only. For all supported
    levels see the --help output.

    Appending a plus + to a level name also includes all higher levels.
    For example:

    dmesg --level=err+

    will print levels err, crit, alert and emerg.

    Prepending it will include all lower levels.

-n, --console-level level
```

Set the level at which printing of messages is done to the console. The level is a level number or abbreviation of the level name. For all supported levels see the --help output.

For example, -n 1 or -n emerg prevents all messages, except emergency (panic) messages, from appearing on the console. All levels of messages are still written to /proc/kmsg, so syslogd(8) can still be used to control exactly where kernel messages appear. When the -n option is used, dmesg will not print or clear the kernel ring buffer.

--noescape
The unprintable and potentially unsafe characters (e.g., broken multi-byte sequences, terminal controlling chars, etc.) are escaped in format \x<hex> for security reason by default. This option disables this feature at all. It’s usable for example for debugging purpose together with --raw. Be careful and don’t use it by default.

-P, --nopager
Do not pipe output into a pager. A pager is enabled by default for --human output.

-p, --force-prefix
Add facility, level or timestamp information to each line of a multi-line message.

-r, --raw
Print the raw message buffer, i.e., do not strip the log-level prefixes, but all unprintable characters are still escaped (see also --noescape).

Note that the real raw format depends on the method how dmesg reads kernel messages. The /dev/kmsg device uses a different format than syslog(2). For backward compatibility, dmesg returns data always in the syslog(2) format. It is possible to read the real raw data from /dev/kmsg by, for example, the command 'dd if=/dev/kmsg iflag=nonblock'.

-S, --syslog
Force dmesg to use the syslog(2) kernel interface to read kernel messages. The default is to use /dev/kmsg rather than syslog(2) since kernel 3.5.0.

-s, --buffer-size size
Use a buffer of size to query the kernel ring buffer. This is 16392 by default. (The default kernel syslog buffer size was 4096 at first, 8192 since 1.3.54, 16384 since 2.1.113.) If you have set the kernel buffer to be larger than the default, then this option can be used to view the entire buffer.

-T, --ctime
Print human-readable timestamps.

```
    Be aware that the timestamp could be inaccurate! The time source
    used for the logs is not updated after system SUSPEND/RESUME.
    Timestamps are adjusted according to current delta between boottime
    and monotonic clocks, this works only for messages printed after
    last resume.

--since time
    Display record since the specified time. Supported is the subsecond
    granularity. The time is possible to specify in absolute way as
    well as by relative notation (e.g. '1 hour ago'). Be aware that the
    timestamp could be inaccurate and see --ctime for more details.

--until time
    Display record until the specified time. Supported is the subsecond
    granularity. The time is possible to specify in absolute way as
    well as by relative notation (e.g. '1 hour ago'). Be aware that the
    timestamp could be inaccurate and see --ctime for more details.

-t, --notime
    Do not print kernel’s timestamps.

--time-format format
    Print timestamps using the given format, which can be ctime,
    reltime, delta or iso. The first three formats are aliases of the
    time-format-specific options. The iso format is a dmesg
    implementation of the ISO-8601 timestamp format. The purpose of
    this format is to make the comparing of timestamps between two
    systems, and any other parsing, easy. The definition of the iso
    timestamp is: YYYY-MM-DD<T>HH:MM:SS,<microseconds>←+><timezone
    offset from UTC>.

    The iso format has the same issue as ctime: the time may be
    inaccurate when a system is suspended and resumed.

-u, --userspace
    Print userspace messages.

-w, --follow
    Wait for new messages. This feature is supported only on systems
    with a readable /dev/kmsg (since kernel 3.5.0).

-W, --follow-new
    Wait and print only new messages.

-x, --decode
    Decode facility and level (priority) numbers to human-readable
    prefixes.

-h, --help
    Display help text and exit.

-V, --version
```

```
            Print version and exit.

COLORS
       The output colorization is implemented by terminal-colors.d(5)
       functionality. Implicit coloring can be disabled by an empty file

          /etc/terminal-colors.d/dmesg.disable

       for the dmesg command or for all tools by

          /etc/terminal-colors.d/disable

       The user-specific $XDG_CONFIG_HOME/terminal-colors.d or
       $HOME/.config/terminal-colors.d overrides the global setting.

       Note that the output colorization may be enabled by default, and in
       this case terminal-colors.d directories do not have to exist yet.

       The logical color names supported by dmesg are:

       subsys
           The message sub-system prefix (e.g., "ACPI:").

       time
           The message timestamp.

       timebreak
           The message timestamp in short ctime format in --reltime or --human
           output.

       alert
           The text of the message with the alert log priority.

       crit
           The text of the message with the critical log priority.

       err
           The text of the message with the error log priority.

       warn
           The text of the message with the warning log priority.

       segfault
           The text of the message that inform about segmentation fault.

EXIT STATUS
       dmesg can fail reporting permission denied error. This is usually
       caused by dmesg_restrict kernel setting, please see syslog(2) for more
       details.

AUTHORS
       Karel Zak <kzak@redhat.com>
```

```
       dmesg was originally written by Theodore Ts’o <tytso@athena.mit.edu>.

SEE ALSO
       terminal-colors.d(5), syslogd(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The dmesg command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-23                          DMESG(1)
```

# Bash 的 du 命令

```
du [OPTION]... [FILE]...
```

**功能**

估算文件空间占用。

**类型**

可执行文件（`/usr/bin/du`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-0`，`--null` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `-a`，`--all` - 所有文件写入计数，而不只是目录
  - `--apparent-size` - 打印实际大小而非设备使用量
  - `-B`，`--block-size=SIZE` - 指定单位，例如 `-B 4K` 表示以 4096 字节为单位
  - `-b`，`--bytes` - 等价于 `--apparent-size --block-size=1`
  - `-c`，`--total` - 得出总计数据
  - `-D`，`--dereference-args` - 参数是符号链接时，查看源文件
  - `-d`，`--max-depth=N` - 查看目录下最大深度不超过 `N` 的文件；参数目录自身深度为 0
  - `--files0-from=F` - 从文件 `F` 中读取没有了列表，空字符（`\0`）分隔
  - `-H` - 等价于 `--dereference-args`
  - `-h`，`--human-readable` - 自动添加单位以便人类阅读， 1K 表示 1024 字节
  - `--inodes` - 列出 inode 信息而不是块占用
  - `-k` - 等价于 `-B 1K`
  - `-L`，`--dereference` - 符号链接查看源文件，而非符号链接本身
  - `-l`，`--count-links` - 硬链接多次计算大小
  - `-m` - 等价于 `-B 1M`
  - `-P`，`--no-dereference` - 符号链接查看本身，而非源文件
  - `-S`，`--separate-dirs` - 目录的大小不包括子目录的大小
  - `--si` - 自动添加单位以便人类阅读，1K 表示 1000 字节
  - `-s`，`--summarize` - 仅显示每个参数的总计数据
  - `-t`，`--threshold=SIZE` - 忽略小于 `SIZE` 的条目
  - `--time` - 显示文件的最后修改时间

    - `--time=WORD` - 显示 `WORD` 格式的时间：`atime`，`access`，`use`，`ctime` 或 `status`
    - `--time-style=STYLE` - 显示 `STYLE` 风格的时间：`full-iso`，`long-iso`，`iso`，`+FORMAT`（参考 date）
    - `-X`，`--exclude-from=FILE` - 排除与 `FILE` 文件中的模式匹配的文件
    - `--exclude=PATTERN` - 排除匹配模式 `PATTERN` 的文件
    - `-x`，`--one-file-system` - 跳过不同文件系统中的目录
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `FILE` - 文件列表；默认查看当前工作目录

## 示例

```
$ sudo du -d 1 -h .     # 当前目录，最大深度为 1
25M     ./dir1
4.0K    ./dir2
26M     .
```

## 推荐阅读

- Estimate file space usage - GNU Coreutils
- du(1) — Linux manual page

## 手册

```
DU(1)                           User Commands                           DU(1)

NAME
       du - estimate file space usage

SYNOPSIS
       du [OPTION]... [FILE]...
       du [OPTION]... --files0-from=F

DESCRIPTION
       Summarize  device  usage  of the set of FILEs, recursively for directo-
       ries.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -0, --null
              end each output line with NUL, not newline
```

```
-a, --all
       write counts for all files, not just directories

--apparent-size
       print  apparent sizes rather than device usage; although the ap-
       parent size is usually smaller, it may be larger due to holes in
       ('sparse') files, internal fragmentation, indirect  blocks,  and
       the like

-B, --block-size=SIZE
       scale  sizes  by  SIZE  before printing them; e.g., '-BM' prints
       sizes in units of 1,048,576 bytes; see SIZE format below

-b, --bytes
       equivalent to '--apparent-size --block-size=1'

-c, --total
       produce a grand total

-D, --dereference-args
       dereference only symlinks that are listed on the command line

-d, --max-depth=N
       print the total for a directory (or file, with --all) only if it
       is  N  or  fewer  levels  below  the  command   line   argument;
       --max-depth=0 is the same as --summarize

--files0-from=F
       summarize  device  usage of the NUL-terminated file names speci-
       fied in file F; if F is -, then read names from standard input

-H     equivalent to --dereference-args (-D)

-h, --human-readable
       print sizes in human readable format (e.g., 1K 234M 2G)

--inodes
       list inode usage information instead of block usage

-k     like --block-size=1K

-L, --dereference
       dereference all symbolic links

-l, --count-links
       count sizes many times if hard linked

-m     like --block-size=1M

-P, --no-dereference
       don't follow any symbolic links (this is the default)

-S, --separate-dirs
```

```
              for directories do not include size of subdirectories

       --si   like -h, but use powers of 1000 not 1024

       -s, --summarize
              display only a total for each argument

       -t, --threshold=SIZE
              exclude entries  smaller  than  SIZE  if  positive,  or  entries
              greater than SIZE if negative

       --time show time of the last modification of any file in the directory,
              or any of its subdirectories

       --time=WORD
              show  time  as WORD instead of modification time: atime, access,
              use, ctime or status

       --time-style=STYLE
              show times using STYLE, which can be: full-iso,  long-iso,  iso,
              or +FORMAT; FORMAT is interpreted like in 'date'

       -X, --exclude-from=FILE
              exclude files that match any pattern in FILE

       --exclude=PATTERN
              exclude files that match PATTERN

       -x, --one-file-system
              skip directories on different file systems

       --help display this help and exit

       --version
              output version information and exit

       Display  values  are  in  units  of  the  first  available  SIZE  from
       --block-size, and the DU_BLOCK_SIZE, BLOCK_SIZE and BLOCKSIZE  environ-
       ment  variables.   Otherwise,  units  default  to 1024 bytes (or 512 if
       POSIXLY_CORRECT is set).

       The SIZE argument is an integer and  optional  unit  (example:  10K  is
       10*1024).   Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or KB,MB,...
       (powers of 1000).  Binary prefixes can be used, too: KiB=K, MiB=M,  and
       so on.

PATTERNS
       PATTERN  is  a shell pattern (not a regular expression).  The pattern ?
       matches any one character, whereas * matches any  string  (composed  of
       zero,  one  or  multiple  characters).  For example, *.o will match any
       files whose names end in .o.  Therefore, the command

              du --exclude='*.o'
```

will skip all files and subdirectories ending in .o (including the file .o itself).

## AUTHOR

Written by Torbjorn Granlund, David MacKenzie, Paul Eggert, and Jim Meyering.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/du>
or available locally via: info '(coreutils) du invocation'

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 echo 命令

```
echo [OPTION]... [STRING]...
```

**功能**

显示一行文本。

**类型**

Bash 内置命令。

> `echo` 也是 coreutils 的一个可执行文件；为了保持 POSIX 标准的兼容性，确保在其他 shell（比如一些极简的或古老的 shell）或者某些特定环境下，命令依然可用。

**说明**

- 选项
  - `-n` - 不在末尾自动添加换行
  - `-e` - 启用转义字符
  - `-E` - 禁用转义字符
- 转义字符（开启 `-e` 选项时有效）
  - `\\` - 表示反斜杠 `\` 本身
  - `\a` - 响铃
  - `\b` - 退格
  - `\c` - 不再产生输出
  - `\e` - ANSI 转义代码
  - `\f` - 换页符
  - `\n` - 换行符
  - `\r` - 回车符
  - `\t` - 水平制表符
  - `\v` - 垂直制表符
  - `\0NNN` - 八进制字符编码
  - `\xXXX` - 十六进制字符编码

## ANSI 转义代码

使用 ANSI 转义代码的格式为 `\e[代码;代码;代码m`：可以同时包含任意数量的代码，以分号（`;`）分隔，并以 `m` 结束。

| 代码 | 说明 | 备注 |
| --- | --- | --- |
| 0 | 重置 | 清除所有效果，恢复默认 |
| 1 | 强调（加粗） | - |
| 2 | 弱化 | - |
| 3 | 斜体 | - |
| 4 | 下划线 | - |
| 5 | 缓慢闪烁 | - |
| 6 | 快速闪烁 | 未被广泛支持 |
| 7 | 反显（前景色与背景色交换） | - |
| 8 | 隐藏 | 未被广泛支持 |
| 9 | 删除线 | 未被广泛支持 |
| 10 | 主要字体 | - |
| 11-19 | 替代字体 | - |
| 20 | 德文尖角体 | 未被广泛支持 |
| 21 | 关闭粗体或双下划线 | 未被广泛支持 |
| 22 | 正常颜色强度 | - |
| 23 | 非斜体、非尖角体 | - |
| 24 | 关闭下划线 | 消除 4 和 21 |
| 25 | 关闭闪烁 | 消除 5 和 6 |
| 27 | 关闭反显 | 消除 7 |
| 28 | 关闭隐藏 | 消除 8 |
| 29 | 关闭删除线 | 消除 9 |
| 30-37 | 设置前景色 | 见 颜色代码 |
| 38 | 设置前景色 | 后跟 `5;n`（256 色模式）或 `2;r;g;b`（24 位 RGB 模式） |
| 39 | 恢复默认前景色 | 消除 30-38 |
| 40-47 | 设置背景色 | 见 颜色代码 |

| 代码 | 说明 | 备注 |
| --- | --- | --- |
| 48 | 设置背景色 | 后跟 `5;n`(256 色模式) 或 `2;r;g;b`（24 位 RGB 模式） |
| 49 | 恢复默认背景色 | 消除 40-48 |
| 51 | 边框 | 未被广泛支持 |
| 52 | 包围 | 未被广泛支持 |
| 53 | 上划线 | 未被广泛支持 |
| 55 | 关闭上划线 | 消除 53 |
| 60 | 表意文字下划线或右边线 | 未被广泛支持 |
| 61 | 表意文字双下划线或双右边线 | 未被广泛支持 |
| 62 | 表意文字上划线或左边线 | 未被广泛支持 |
| 63 | 表意文字双上划线或双左边线 | 未被广泛支持 |
| 64 | 表意文字强调 | 未被广泛支持 |
| 65 | 关闭表意文字属性 | 消除 60-64 |
| 90-97 | 设置明亮的前景色 | aixterm 非标准，见 颜色代码 |
| 100–107 | 设置明亮的背景色 | aixterm 非标准，见 颜色代码 |

## 颜色代码

| 颜色 | 前景色代码 | 背景色代码 | 备注 |
| --- | --- | --- | --- |
| 黑 | 30 | 40 | |
| 红 | 31 | 41 | |
| 绿 | 32 | 42 | |
| 黄 | 33 | 43 | |
| 蓝 | 34 | 44 | |
| 紫 | 35 | 45 | |
| 青 | 36 | 46 | |

| 颜色 | 前景色代码 | 背景色代码 | 备注 |
| --- | --- | --- | --- |
| 白 | 37 | 47 | |
| 亮黑（灰） | 90 | 100 | |
| 亮红 | 91 | 101 | |
| 亮绿 | 92 | 102 | |
| 亮黄 | 93 | 103 | |
| 亮蓝 | 94 | 104 | |
| 亮紫 | 95 | 105 | |
| 亮青 | 96 | 106 | |
| 亮白 | 97 | 107 | |

## 256 色模式

## 示例

```
# 显示文本
echo https://xplanc.org/
```

```
# 加粗 前景黄色 背景红色
echo -e "\e[1;33;41mhttps://xplanc.org/\e[0m"

# 使用 256 色模式，前景色是 201 背景色是 51
echo -e "\e[38;5;201;48;5;51mhttps://xplanc.org/\e[0m"

# 使用 24 位 RGB 模式，前景色是 255,255,0 背景色是 0,0,255
echo -e "\e[38;2;255;255;0;48;2;0;0;255mhttps://xplanc.org/\e[0m"
```

## 推荐阅读

- Print a line of text - GNU Coreutils
- echo(1) — Linux manual page
- ANSI escape code - Wikipedia

## 手册

```
ECHO(1)                          User Commands                          ECHO(1)

NAME
       echo - display a line of text

SYNOPSIS
       echo [SHORT-OPTION]... [STRING]...
       echo LONG-OPTION

DESCRIPTION
       Echo the STRING(s) to standard output.

       -n     do not output the trailing newline

       -e     enable interpretation of backslash escapes

       -E     disable interpretation of backslash escapes (default)

       --help display this help and exit

       --version
              output version information and exit

       If -e is in effect, the following sequences are recognized:

       \\     backslash

       \a     alert (BEL)
```

```
       \b     backspace

       \c     produce no further output

       \e     escape

       \f     form feed

       \n     new line

       \r     carriage return

       \t     horizontal tab

       \v     vertical tab

       \0NNN  byte with octal value NNN (1 to 3 digits)

       \xHH   byte with hexadecimal value HH (1 to 2 digits)

       NOTE: your shell may have its own version of echo, which usually super-
       sedes  the  version described here.  Please refer to your shell's docu-
       mentation for details about the options it supports.

       NOTE: printf(1) is a preferred alternative, which does not have  issues
       outputting option-like strings.

AUTHOR
       Written by Brian Fox and Chet Ramey.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       printf(1)

       Full documentation <https://www.gnu.org/software/coreutils/echo>
       or available locally via: info '(coreutils) echo invocation'

GNU coreutils 9.4                April 2024                          ECHO(1)
```

# Bash 的 env 命令

```
env [OPTION]... [-] [NAME=VALUE]... [COMMAND [ARG]...]
```

### 功能

创建新的环境运行程序。

### 类型

可执行文件（`/usr/bin/env`），属于 coreutils。

### 参数

- `OPTION` 选项：
  - `-i`，`--ignore-environment` - 忽略原本的环境变量，从空的环境开始
  - `-0`，`--null` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `-u`，`--unset=NAME` - 删除环境变量 `NAME`
  - `-C`，`--chdir=DIR` - 切换工作目录
  - `-S`，`--split-string=S` - 将 `S` 按空格拆分为独立的参数
  - `--block-signal[=SIG]` - 阻止向执行的命令发送 `SIG` 信号
  - `--default-signal[=SIG]` - 将 `SIG` 信号的处理方式重置为默认
  - `--ignore-signal[=SIG]` - 将 `SIG` 信号的处理方式设为忽略
  - `--list-signal-handling` - 将处理方式不为默认的所有信号在标准错误中列出
  - `-v`，`--debug` - 打印每个处理步骤的详细信息
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `-` - 在所有选项之后，忽略原本的环境变量；同 `-i` 选项
- `NAME=VALUE` - 要设置的环境变量
- `COMMAND` - 要执行的命令，如果为空则打印所有环境变量
- `ARG` - 命令 `COMMAND` 的参数，在 shebang 行中需要使用 `-S` 选项

### 返回值

- 125 - `env` 命令自身失败
- 126 - `COMMAND` 无法被调用
- 127 - `COMMAND` 无法被找到
- 其它 - `COMMAND` 的返回值

# 示例

### 设置环境变量

设置环境变量：

```
$ env NAME=Primers bash -c 'echo $NAME use $LANG'
Primers use en_US.UTF-8
$ NAME=Primers bash -c 'echo $NAME use $LANG'       # 可以不使用 env
Primers use en_US.UTF-8
```

- 注意，这里不能直接写作 `env NAME=Primers echo Hello $NAME`，因为这样会在执行命令之前就将 `$NAME` 展开为空
- 单纯设置环境变量可以不使用 `env`，直接在命令前设置即可

清空并设置环境变量：

```
$ env -i NAME=Primers bash -c 'echo $NAME use $LANG'
Primers use
```

- 原有的环境变量被清空，因此 `$LANG` 为空

### 查看环境

```
$ env | grep PATH
PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin
$ env | grep -i proxy
ALL_PROXY=http://localhost:7890
```

### #!/usr/bin/env -S

```
#!/usr/bin/env -S bash -c
echo Hello World
```

如果不使用 `-S` 选项，`bash -c` 会被视为一个整体命令而失败：

```
/usr/bin/env: ‘bash -c’: No such file or directory
/usr/bin/env: use -[v]S to pass options in shebang lines
```

# 推荐阅读

- Run a command in a modified environment - GNU Coreutils
- env(1) — Linux manual page

# 手册

```
ENV(1)                          User Commands                          ENV(1)

NAME
       env - run a program in a modified environment

SYNOPSIS
       env [OPTION]... [-] [NAME=VALUE]... [COMMAND [ARG]...]

DESCRIPTION
       Set each NAME to VALUE in the environment and run COMMAND.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -i, --ignore-environment
              start with an empty environment

       -0, --null
              end each output line with NUL, not newline

       -u, --unset=NAME
              remove variable from the environment

       -C, --chdir=DIR
              change working directory to DIR

       -S, --split-string=S
              process and split S into separate arguments; used to pass multi-
              ple arguments on shebang lines

       --block-signal[=SIG]
              block delivery of SIG signal(s) to COMMAND

       --default-signal[=SIG]
              reset handling of SIG signal(s) to the default

       --ignore-signal[=SIG]
              set handling of SIG signal(s) to do nothing

       --list-signal-handling
              list non default signal handling to stderr

       -v, --debug
              print verbose information for each processing step

       --help display this help and exit

       --version
              output version information and exit
```

```
       A mere - implies -i.  If no COMMAND, print the resulting environment.

       SIG may be a signal name like 'PIPE', or a  signal  number  like  '13'.
       Without  SIG,  all known signals are included.  Multiple signals can be
       comma-separated.  An empty SIG argument is a no-op.

   Exit status:
       125    if the env command itself fails

       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found

       -      the exit status of COMMAND otherwise

OPTIONS
   -S/--split-string usage in scripts
       The -S option allows specifying multiple parameters in a script.   Run-
       ning a script named 1.pl containing the following first line:

              #!/usr/bin/env -S perl -w -T
              ...

       Will execute perl -w -T 1.pl .

       Without the '-S' parameter the script will likely fail with:

              /usr/bin/env: 'perl -w -T': No such file or directory

       See the full documentation for more details.

   --default-signal[=SIG] usage
       This  option  allows  setting  a  signal handler to its default action,
       which is not possible using the traditional shell  trap  command.   The
       following  example  ensures  that  seq will be terminated by SIGPIPE no
       matter how this signal is being handled in  the  process  invoking  the
       command.

              sh -c 'env --default-signal=PIPE seq inf | head -n1'

NOTES
       POSIX's exec(3p) pages says:
              "many  existing applications wrongly assume that they start with
              certain signals set to the default action  and/or  unblocked....
              Therefore,  it is best not to block or ignore signals across ex-
              ecs without explicit reason to do  so,  and  especially  not  to
              block  signals  across execs of arbitrary (not closely cooperat-
              ing) programs."

AUTHOR
       Written by Richard Mlynarik, David MacKenzie, and Assaf Gordon.
```

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

sigaction(2), sigprocmask(2), signal(7)

Full documentation <https://www.gnu.org/software/coreutils/env>
or available locally via: info '(coreutils) env invocation'

# Bash 的 expand 命令

```
expand [OPTION]... [FILE]...
```

**功能**

将制表符（`\t`）转换为空格。

**类型**

可执行文件（`/usr/bin/expand`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-i`，`--initial` - 非空格后的制表符不要转换
  - `-t`，`--tabs=N` - 制表位的宽度设为 `N`；默认为 8
  - `-t`，`--tabs=LIST` - 使用 `LIST` 指定制表位的位置
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ echo -e "A\tB\tC" > 1.txt      # 创建一个文件
$ cat 1.txt                      # 查看文件
A       B       C
$ expand 1.txt                   # 默认制表位宽度为 8
A       B       C
$ expand -t 4 1.txt              # 指定制表位宽度为 4
A   B   C
$ expand -t 4,16 1.txt           # 指定制表位的位置为 4,16
A   B           C
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| unexpand | 将空格转换为制表符（`\t`） |

## 推荐阅读

- Convert tabs to spaces - GNU Coreutils
- expand(1) — Linux manual page

## 手册

```
EXPAND(1)                       User Commands                       EXPAND(1)

NAME
       expand - convert tabs to spaces

SYNOPSIS
       expand [OPTION]... [FILE]...

DESCRIPTION
       Convert tabs in each FILE to spaces, writing to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -i, --initial
              do not convert tabs after non blanks

       -t, --tabs=N
              have tabs N characters apart, not 8

       -t, --tabs=LIST
              use comma separated list of tab positions.  The  last  specified
              position  can  be prefixed with '/' to specify a tab size to use
              after the last explicitly specified tab stop.  Also a prefix  of
              '+'  can  be  used  to align remaining tab stops relative to the
              last specified tab stop instead of the first column

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
```

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

unexpand(1)

Full documentation <https://www.gnu.org/software/coreutils/expand>
or available locally via: info '(coreutils) expand invocation'

# Bash 的 expr 命令

```
expr EXPRESSION
```

**功能**

计算表达式。

**类型**

可执行文件（`/usr/bin/expr`），属于 coreutils。

**参数**

- `EXPRESSION` 表达式：
  - `ARG1 | ARG2` - 如果 `ARG1` 不是 0 或 null，则结果为 `ARG1`，否则为 `ARG2`
  - `ARG1 & ARG2` - 如果 `ARG1` 和 `ARG2` 都不是 0 或 null，则结果为 `ARG1`，否则为 0
  - `ARG1 < ARG2` - `ARG1` $<$ `ARG2`
  - `ARG1 <= ARG2` - `ARG1` $\leq$ `ARG2`
  - `ARG1 = ARG2` - `ARG1` $=$ `ARG2`
  - `ARG1 != ARG2` - `ARG1` $\neq$ `ARG2`
  - `ARG1 >= ARG2` - `ARG1` $\geq$ `ARG2`
  - `ARG1 > ARG2` - `ARG1` $>$ `ARG2`
  - `ARG1 + ARG2` - `ARG1` $+$ `ARG2`；仅支持整数，下同
  - `ARG1 - ARG2` - `ARG1` $-$ `ARG2`
  - `ARG1 * ARG2` - `ARG1` $\times$ `ARG2`
  - `ARG1 / ARG2` - `ARG1` $\div$ `ARG2`
  - `ARG1 % ARG2` - `ARG1` $\mathbf{mod}$ `ARG2`
  - `STRING : REGEXP` - 在字符串 `STRING` 中使用正则表达式 `REGEXP` 进行锚定匹配
  - `match STRING REGEXP` - 同 `STRING : REGEXP`
  - `substr STRING POS LENGTH` - 从字符串 `STRING` 中 `POS` 位置（从 1 开始）截取长度为 `LENGTH` 的子串
  - `index STRING CHARS` - 从字符串 `STRING` 中查找 `CHARS` 中任意字符首次出现的位置，找不到则得到 0
  - `length STRING` - 获取字符串 `STRING` 的长度
  - `( EXPRESSION )` - 括号嵌套表达式
- `OPTION` 选项：
  - `--help` - 显示帮助

- `--version` - 显示版本

## 示例

```
$ expr 1 '>' 3                      # 判断大于
0
$ expr 1 + 1                        # 加法
2
$ expr 2 - 3                        # 减法
-1
$ expr 10 / 3                       # 除法
3
$ expr 10 % 3                       # 取余
1
$ expr length primers               # 获取 primers 的长度
7
$ expr index primers aeiou          # 在 primers 中查找 aeiou 中的任意一个字符位置
3
$ expr substr primers 3 3           # 从 primers 中 3 位置截取长度为 3 的子串
ime
$ expr match 'https://xplanc.org/primers/' 'https\?://[^/]*/\(.*\)'
primers/
```

## 推荐阅读

- Evaluate expressions - GNU Coreutils
- expr(1) — Linux manual page

## 手册

```
EXPR(1)                          User Commands                          EXPR(1)

NAME
       expr - evaluate expressions

SYNOPSIS
       expr EXPRESSION
       expr OPTION

DESCRIPTION
       --help display this help and exit

       --version
              output version information and exit

       Print  the  value of EXPRESSION to standard output.  A blank line below
```

```
separates increasing precedence groups.  EXPRESSION may be:

ARG1 | ARG2
       ARG1 if it is neither null nor 0, otherwise ARG2

ARG1 & ARG2
       ARG1 if neither argument is null or 0, otherwise 0

ARG1 < ARG2
       ARG1 is less than ARG2

ARG1 <= ARG2
       ARG1 is less than or equal to ARG2

ARG1 = ARG2
       ARG1 is equal to ARG2

ARG1 != ARG2
       ARG1 is unequal to ARG2

ARG1 >= ARG2
       ARG1 is greater than or equal to ARG2

ARG1 > ARG2
       ARG1 is greater than ARG2

ARG1 + ARG2
       arithmetic sum of ARG1 and ARG2

ARG1 - ARG2
       arithmetic difference of ARG1 and ARG2

ARG1 * ARG2
       arithmetic product of ARG1 and ARG2

ARG1 / ARG2
       arithmetic quotient of ARG1 divided by ARG2

ARG1 % ARG2
       arithmetic remainder of ARG1 divided by ARG2

STRING : REGEXP
       anchored pattern match of REGEXP in STRING

match STRING REGEXP
       same as STRING : REGEXP

substr STRING POS LENGTH
       substring of STRING, POS counted from 1

index STRING CHARS
       index in STRING where any CHARS is found, or 0
```

```
       length STRING
              length of STRING

       + TOKEN
              interpret TOKEN as a string, even if it is a

              keyword like 'match' or an operator like '/'

       ( EXPRESSION )
              value of EXPRESSION

       Beware that many operators need to be escaped  or  quoted  for  shells.
       Comparisons are arithmetic if both ARGs are numbers, else lexicographi-
       cal.   Pattern  matches  return the string matched between \( and \) or
       null; if \( and \) are not used, they return the number  of  characters
       matched or 0.

       Exit  status  is 0 if EXPRESSION is neither null nor 0, 1 if EXPRESSION
       is null or 0, 2 if EXPRESSION is syntactically invalid, and 3 if an er-
       ror occurred.

AUTHOR
       Written by Mike Parker, James Youngman, and Paul Eggert.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/expr>
       or available locally via: info '(coreutils) expr invocation'

GNU coreutils 9.4                April 2024                          EXPR(1)
```

# Bash 的 factor 命令

```
factor [OPTION] [NUMBER]...
```

**功能**

质因数分解。

**类型**

可执行文件（`/usr/bin/factor`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-h`，`--exponents` - 以 `x^n` 的形式打印重复的因子
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NUMBER` 要分解的整数列表

## 示例

```
$ factor 1995 2000 2005 2010 2015 2020 2025
1995: 3 5 7 19
2000: 2 2 2 2 5 5 5
2005: 5 401
2010: 2 3 5 67
2015: 5 13 31
2020: 2 2 5 101
2025: 3 3 3 3 5 5
$ factor -h 1995 2000 2005 2010 2015 2020 2025
1995: 3 5 7 19
2000: 2^4 5^3
2005: 5 401
2010: 2 3 5 67
2015: 5 13 31
2020: 2^2 5 101
2025: 3^4 5^2
```

## 推荐阅读

- Print prime factors - GNU Coreutils

- factor(1) — Linux manual page

## 手册

```
FACTOR(1)                        User Commands                        FACTOR(1)

NAME
       factor - factor numbers

SYNOPSIS
       factor [OPTION] [NUMBER]...

DESCRIPTION
       Print  the prime factors of each specified integer NUMBER.  If none are
       specified on the command line, read them from standard input.

       -h, --exponents
              print repeated factors in form p^e unless e is 1

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Paul Rubin, Torbjorn Granlund, and Niels Moller.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/factor>
       or available locally via: info '(coreutils) factor invocation'

GNU coreutils 9.4                 April 2024                        FACTOR(1)
```

# Bash 的 fallocate 命令

```
fallocate [OPTIONS] FILE
```

**功能**

为文件预分配磁盘空间，或取消预分配的空间。

**类型**

可执行文件（`/usr/bin/fallocate`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-c`，`--collapse-range` - 折叠（删除）指定区间，并将其后的数据向前移动
  - `-d`，`--dig-holes` - 检测和释放空洞，但保留文件大小
  - `-i`，`--insert-range` - 在指定偏移处插入一段空洞，并将后续数据向后移动
  - `-l`，`--length LENGTH` - 在 `LENGTH` 长度内进行操作
  - `-n`，`--keep-size` - 保持文件逻辑大小不变
  - `-o`，`--offset OFFSET` - 从 `OFFSET` 位置之后开始操作
  - `-p`，`--punch-hole` - 将指定区域作为空洞释放，此后读取该返回将得到零
  - `-v`，`--verbose` - 显示详细的操作信息
  - `-x`，`--posix` - 使用 POSIX 兼容的空间预分配方式
  - `-z`，`--zero-range` - 将指定区间置零
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 要操作的问题

> 空洞指文件占据但不实际使用的空间，例如一个文件里面有很多连续的零，则可以将该空间释放以作他用。

## 示例

```
$ fallocate -d file.data                    # 检测并释放文件中的空洞空间
$ fallocate -p -o 100 -l 100  file.data     # 将文件 100 到 200 自己作为空洞释放
```

## 推荐阅读

- fallocate(1) — Linux manual page

# 手册

```
FALLOCATE(1)                    User Commands                    FALLOCATE(1)

NAME
       fallocate - preallocate or deallocate space to a file

SYNOPSIS
       fallocate [-c|-p|-z] [-o offset] -l length [-n] filename

       fallocate -d [-o offset] [-l length] filename

       fallocate -x [-o offset] -l length filename

DESCRIPTION
       fallocate is used to manipulate the allocated disk space for a file,
       either to deallocate or preallocate it. For filesystems which support
       the fallocate(2) system call, preallocation is done quickly by
       allocating blocks and marking them as uninitialized, requiring no IO to
       the data blocks. This is much faster than creating a file by filling it
       with zeroes.

       The exit status returned by fallocate is 0 on success and 1 on failure.

OPTIONS
       The length and offset arguments may be followed by the multiplicative
       suffixes KiB (=1024), MiB (=1024*1024), and so on for GiB, TiB, PiB,
       EiB, ZiB, and YiB (the "iB" is optional, e.g., "K" has the same meaning
       as "KiB") or the suffixes KB (=1000), MB (=1000*1000), and so on for
       GB, TB, PB, EB, ZB, and YB.

       The options --collapse-range, --dig-holes, --punch-hole, and
       --zero-range are mutually exclusive.

       -c, --collapse-range
           Removes a byte range from a file, without leaving a hole. The byte
           range to be collapsed starts at offset and continues for length
           bytes. At the completion of the operation, the contents of the file
           starting at the location offset+length will be appended at the
           location offset, and the file will be length bytes smaller. The
           option --keep-size may not be specified for the collapse-range
           operation.

           Available since Linux 3.15 for ext4 (only for extent-based files)
           and XFS.

           A filesystem may place limitations on the granularity of the
           operation, in order to ensure efficient implementation. Typically,
```

offset and length must be a multiple of the filesystem logical block size, which varies according to the filesystem type and configuration. If a filesystem has such a requirement, the operation will fail with the error EINVAL if this requirement is violated.

-d, --dig-holes
Detect and dig holes. This makes the file sparse in-place, without using extra disk space. The minimum size of the hole depends on filesystem I/O block size (usually 4096 bytes). Also, when using this option, --keep-size is implied. If no range is specified by --offset and --length, then the entire file is analyzed for holes.

You can think of this option as doing a "cp --sparse" and then renaming the destination file to the original, without the need for extra disk space.

See --punch-hole for a list of supported filesystems.

-i, --insert-range
Insert a hole of length bytes from offset, shifting existing data.

-l, --length length
Specifies the length of the range, in bytes.

-n, --keep-size
Do not modify the apparent length of the file. This may effectively allocate blocks past EOF, which can be removed with a truncate.

-o, --offset offset
Specifies the beginning offset of the range, in bytes.

-p, --punch-hole
Deallocates space (i.e., creates a hole) in the byte range starting at offset and continuing for length bytes. Within the specified range, partial filesystem blocks are zeroed, and whole filesystem blocks are removed from the file. After a successful call, subsequent reads from this range will return zeroes. This option may not be specified at the same time as the --zero-range option. Also, when using this option, --keep-size is implied.

Supported for XFS (since Linux 2.6.38), ext4 (since Linux 3.0), Btrfs (since Linux 3.7), tmpfs (since Linux 3.5) and gfs2 (since Linux 4.16).

-v, --verbose
Enable verbose mode.

-x, --posix
Enable POSIX operation mode. In that mode allocation operation always completes, but it may take longer time when fast allocation is not supported by the underlying filesystem.

```
       -z, --zero-range
           Zeroes space in the byte range starting at offset and continuing
           for length bytes. Within the specified range, blocks are
           preallocated for the regions that span the holes in the file. After
           a successful call, subsequent reads from this range will return
           zeroes.

           Zeroing is done within the filesystem preferably by converting the
           range into unwritten extents. This approach means that the
           specified range will not be physically zeroed out on the device
           (except for partial blocks at the either end of the range), and I/O
           is (otherwise) required only to update metadata.

           Option --keep-size can be specified to prevent file length
           modification.

           Available since Linux 3.14 for ext4 (only for extent-based files)
           and XFS.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

AUTHORS
       Eric Sandeen <sandeen@redhat.com>, Karel Zak <kzak@redhat.com>

SEE ALSO
       truncate(1), fallocate(2), posix_fallocate(3)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The fallocate command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                      FALLOCATE(1)
```

# Bash 的 false 命令

```
false OPTION
```

**功能**

什么也不做，直接返回失败（1），通常用于 `if` 语句。

**类型**

Bash 内置命令。

> `false` 也是 coreutils 的一个可执行文件；为了保持 POSIX 标准的兼容性，确保在其他 shell（比如一些极简的或古老的 shell）或者某些特定环境下，命令依然可用。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
if [ false ]; then
    echo 吃饭
else
    echo 睡觉
fi
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| true | 直接返回失败 |
| test | 评估表达式的布尔值 |

## 推荐阅读

- false(1) — Linux manual page

# 手册

```
FALSE(1)                        User Commands                        FALSE(1)

NAME
       false - do nothing, unsuccessfully

SYNOPSIS
       false [ignored command line arguments]
       false OPTION

DESCRIPTION
       Exit with a status code indicating failure.

       --help display this help and exit

       --version
              output version information and exit

       NOTE:  your  shell may have its own version of false, which usually su-
       persedes the version described here.  Please refer to your shell's doc-
       umentation for details about the options it supports.

AUTHOR
       Written by Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/false>
       or available locally via: info '(coreutils) false invocation'

GNU coreutils 9.4                 April 2024                        FALSE(1)
```

# Bash 的 findfs 命令

```
findfs [OPTION]... NAME=value
```

**功能**

通过标签或 UUID 查找文件系统。

**类型**

可执行文件（`/usr/sbin/findfs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `NAME=value` - 查找的字段和值：
  - `LABEL=<label>` - 通过标签查找文件系统
  - `UUID=<uuid>` - 通过 UUID 查找文件系统
  - `PARTLABEL=<label>` - 通过分区标签查找
  - `PARTUUID=<uuid>` - 通过分区 UUID 查找

## 示例

```
$ blkid                     # 查看所有块设备的 UUID
/dev/nvme0n1p1: UUID="02A0-97AC" BLOCK_SIZE="512" TYPE="vfat" PARTUUID="dfedd060-593f-
48e6-937d-a573d4a95dc5"
/dev/nvme0n1p2: UUID="8b2c9bcc-94f2-4cf2-a9d0-581e58f9d676" BLOCK_SIZE="4096" TYPE="ext4"
PARTUUID="4c3ba215-b476-4c6c-a37b-fb13e02d9d46"
$ findfs UUID=02A0-97AC     # 查找设备
/dev/nvme0n1p1
$ findfs UUID=8b2c9bcc-94f2-4cf2-a9d0-581e58f9d676
/dev/nvme0n1p2
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| blkid | 获取块设备的属性 |

| 命令 | 说明 |
| --- | --- |
| lsblk | 列出块设备 |

## 推荐阅读

- findfs(8) — Linux manual page

## 手册

```
FINDFS(8)                   System Administration                   FINDFS(8)

NAME
       findfs - find a filesystem by label or UUID

SYNOPSIS
       findfs NAME=value

DESCRIPTION
       findfs will search the block devices in the system looking for a
       filesystem or partition with specified tag. The currently supported
       tags are:

       LABEL=<label>
           Specifies filesystem label.

       UUID=<uuid>
           Specifies filesystem UUID.

       PARTUUID=<uuid>
           Specifies partition UUID. This partition identifier is supported
           for example for GUID Partition Table (GPT) partition tables.

       PARTLABEL=<label>
           Specifies partition label (name). The partition labels are
           supported for example for GUID Partition Table (GPT) or MAC
           partition tables.

       If the filesystem or partition is found, the device name will be
       printed on stdout.

       The complete overview about filesystems and partitions you can get for
       example by

          lsblk --fs

          partx --show <disk>

          blkid
```

```
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       0
           success

       1
           label or uuid cannot be found

       2
           usage error, wrong number of arguments or unknown option

ENVIRONMENT
       LIBBLKID_DEBUG=all
           enables libblkid debug output.

AUTHORS
       findfs was originally written by Theodore Ts’o <tytso@mit.edu> and
       re-written for the util-linux package by Karel Zak <kzak@redhat.com>.

SEE ALSO
       blkid(8), lsblk(8), partx(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The findfs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          FINDFS(8)
```

# Bash 的 findmnt 命令

```
findmnt [OPTIONS]
```

**功能**

查询文件系统及挂载点。

**类型**

可执行文件（`/usr/bin/findmnt`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-A, --all` 禁用所有内置过滤器，显示所有文件系统。
  - `-a, --ascii` 使用 ASCII 字符显示树状格式。
  - `-b, --bytes` 以字节为单位显示文件系统大小，而不是人类可读格式。 默认单位是字节，前缀以 2^10 为基数（KiB、MiB 等）。为了可读性，符号缩写会只显示首字母，例如 “1 KiB” 显示为 “1 K”，省略了 “iB”。
  - `-C, --nocanonicalize` 不进行路径规范化。此选项影响路径比较和标签（LABEL、UUID 等）的解析。
  - `-c, --canonicalize` 对所有输出路径进行规范化。
  - `--deleted` 显示内核标记为已删除的挂载目标（mountpoint）的文件系统。
  - `-D, --df` 模拟 `df(1)` 输出。等价于 `-o SOURCE,FSTYPE,SIZE,USED,AVAIL,USE%,TARGET`，但排除伪文件系统。使用 `--all` 可以显示所有文件系统。
  - `-d, --direction word` 搜索方向，可选值为 forward（向前）或 backward（向后）。
  - `-e, --evaluate` 将所有标签（LABEL、UUID、PARTUUID、PARTLABEL）转换为 SOURCE 列对应的设备名。可能存在标签重复的情况，此时会显示 SOURCES 列（多行显示所有匹配设备）。此选项主要对 `fstab` 有意义。
  - `-F, --tab-file path` 指定替代挂载表文件。如果与 `--fstab`、`--mtab` 或 `--kernel` 一起使用，则覆盖默认路径。指定多次会禁用树状输出（改为列表模式）。
  - `-f, --first-only` 只显示第一个匹配的文件系统。

- `-i, --invert` 反转匹配逻辑（显示不匹配的项）。
- `-J, --json` 使用 JSON 格式输出。
- `-k, --kernel` 在 `/proc/self/mountinfo` 中搜索。输出为树状格式，默认选项，仅显示内核维护的挂载选项（参见 `--mtab`）。
- `-l, --list` 使用列表格式输出。当输出受 `-t`、`-O`、`-S`、`-T` 限制且未使用 `--submounts`，或指定多个 `-F` 文件时，自动启用列表格式。
- `-M, --mountpoint path` 明确指定挂载点文件或目录。可与 `--target` 互换使用。
- `-m, --mtab` 在 `/etc/mtab` 中搜索。默认输出列表格式，可能包含用户空间挂载选项。
- `-N, --task tid` 使用替代命名空间 `/proc/<tid>/mountinfo` 而非默认 `/proc/self/mountinfo`。指定多次会禁用树状输出。参考 `unshare(1)`。
- `-n, --noheadings` 不显示表头行。
- `-O, --options list` 限制输出的文件系统选项。多个选项用逗号分隔，-t 与 -O 可以叠加。每个选项精确匹配，可在单个选项前加 `no` 表示排除。`+` 前缀可禁用 `no` 解释。
- `-o, --output list` 自定义输出列。列名可通过 `--help` 查看。默认 TARGET 列会显示树状格式（除非使用 `--list` 或 `--raw`）。可用 `+列名` 扩展默认列。
- `--output-all` 输出几乎所有可用列，排除需要 `--poll` 的列。
- `-P, --pairs` 以 `key="value"` 形式输出，值中的特殊字符以十六进制转义。可与 `--shell` 配合使用。
- `-p, --poll[=list]` 监控 `/proc/self/mountinfo` 的变化，支持动作：mount、umount、remount、move。动作可用逗号分隔，默认监控所有动作。可配合 `--timeout` 或 `--first-only` 限制阻塞时间。
- `--pseudo` 仅显示伪文件系统。
- `--shadow` 仅显示被其他文件系统覆盖的挂载点。
- `-R, --submounts` 递归显示选中文件系统的子挂载点。子挂载点总是树状输出。对 `--mtab` 或 `--fstab` 无效。
- `-r, --raw` 使用原始格式输出，所有潜在危险字符以十六进制转义。
- `--real` 仅显示真实文件系统。
- `-S, --source spec` 明确指定挂载源，可用 device、maj:min、LABEL、UUID、PARTLABEL、PARTUUID 等。
- `-s, --fstab` 在 `/etc/fstab` 中搜索，列表格式输出。

- `-T, --target path` 指定挂载目标路径。如果不是挂载点，会反向检查路径元素获取挂载点（仅支持 kernel 搜索，不支持 fstab）。
- `-t, --types list` 限制输出的文件系统类型，可用 `no` 前缀排除类型。
- `--tree` 启用树状输出（仅当存在父子关系时有效，fstab 等可能被忽略）。
- `--shadowed` 仅显示被覆盖的挂载点（同 `--shadow`）。
- `-U, --uniq` 忽略重复挂载目标的文件系统（跳过被覆盖挂载点）。
- `-u, --notruncate` 不截断列文本。默认只截断非 TARGET、SOURCE、UUID、LABEL、PARTUUID、PARTLABEL 列。
- `-v, --nofsroot` 对 bind mount 或 btrfs 子卷，SOURCE 列不显示 [/dir]。
- `-w, --timeout milliseconds` 指定 `--poll` 阻塞时间上限（毫秒）。
- `-x, --verify` 检查挂载表内容。默认验证 `/etc/fstab` 可解析性，可与 `--tab-file` 一起使用。可指定源或目标过滤，`--verbose` 可显示详细信息。
- `--verbose` 强制显示更多信息（当前仅对 `--verify` 有效）。
- `--vfs-all` 当与 VFS-OPTIONS 列一起使用时，显示所有 VFS（与文件系统无关）标志，用于审计，显示默认内核挂载选项。
- `-y, --shell` 将列名修改为仅包含 shell 变量允许的字符，可与 `--pairs` 配合使用。
- `--help` - 显示帮助
- `--version` - 显示版本

## 示例

```
$ findmnt
TARGET                                SOURCE FSTYPE      OPTIONS
/                                     /dev/nvme0n1p2
                                             ext4        rw,relatime
├─/sys                                sysfs  sysfs
rw,nosuid,nodev,noexec,relatime
│ ├─/sys/firmware/efi/efivars         efivarfs
│ │                                           efivarfs
rw,nosuid,nodev,noexec,relatime
│ ├─/sys/kernel/security              securityfs
│ │                                           securityfs
rw,nosuid,nodev,noexec,relatime
│ ├─/sys/fs/cgroup                    cgroup2
...
```

# 推荐阅读

- findmnt(8) — Linux manual page

# 手册

```
FINDMNT(8)                  System Administration                  FINDMNT(8)

NAME
       findmnt - find a filesystem

SYNOPSIS
       findmnt [options]

       findmnt [options] device|mountpoint

       findmnt [options] [--source] device [--target path|--mountpoint
       mountpoint]

DESCRIPTION
       findmnt will list all mounted filesystems or search for a filesystem.
       The findmnt command is able to search in /etc/fstab, /etc/mtab or
       /proc/self/mountinfo. If device or mountpoint is not given, all
       filesystems are shown.

       The device may be specified by device name, major:minor numbers,
       filesystem label or UUID, or partition label or UUID. Note that findmnt
       follows mount(8) behavior where a device name may be interpreted as a
       mountpoint (and vice versa) if the --target, --mountpoint or --source
       options are not specified.

       The command-line option --target accepts any file or directory and then
       findmnt displays the filesystem for the given path.

       The command prints all mounted filesystems in the tree-like format by
       default. The default output, is subject to change. So whenever
       possible, you should avoid using default output in your scripts. Always
       explicitly define expected columns by using --output columns-list in
       environments where a stable output is required.

       The relationship between block devices and filesystems is not always
       one-to-one. The filesystem may use more block devices. This is why
       findmnt provides  SOURCE and SOURCES (pl.) columns. The column SOURCES
       displays all devices where it is possible to find the same filesystem
       UUID (or another tag specified in fstab when executed with --fstab and
       --evaluate).

OPTIONS
       -A, --all
           Disable all built-in filters and print all filesystems.
```

-a, --ascii
    Use ascii characters for tree formatting.

-b, --bytes
    Print the sizes in bytes rather than in a human-readable format.

    By default, the unit, sizes are expressed in, is byte, and unit prefixes are in power of 2^10 (1024). Abbreviations of symbols are exhibited truncated in order to reach a better readability, by exhibiting alone the first letter of them; examples: "1 KiB" and "1 MiB" are respectively exhibited as "1 K" and "1 M", then omitting on purpose the mention "iB", which is part of these abbreviations.

-C, --nocanonicalize
    Do not canonicalize paths at all. This option affects the comparing of paths and the evaluation of tags (LABEL, UUID, etc.).

-c, --canonicalize
    Canonicalize all printed paths.

--deleted
    Print filesystems where target (mountpoint) is marked as deleted by kernel.

-D, --df
    Imitate the output of df(1). This option is equivalent to -o SOURCE,FSTYPE,SIZE,USED,AVAIL,USE%,TARGET but excludes all pseudo filesystems. Use --all to print all filesystems.

-d, --direction word
    The search direction, either forward or backward.

-e, --evaluate
    Convert all tags (LABEL, UUID, PARTUUID, or PARTLABEL) to the corresponding device names for the SOURCE column. It’s an unusual situation, but the same tag may be duplicated (used for more devices). For this purpose, there is SOURCES (pl.) column. This column displays by multi-line cell all devices where the tag is detected by libblkid. This option makes sense for fstab only.

-F, --tab-file path
    Search in an alternative file. If used with --fstab, --mtab or --kernel, then it overrides the default paths. If specified more than once, then tree-like output is disabled (see the --list option).

-f, --first-only
    Print the first matching filesystem only.

-i, --invert
    Invert the sense of matching.

```
-J, --json
    Use JSON output format.

-k, --kernel
    Search in /proc/self/mountinfo. The output is in the tree-like
    format. This is the default. The output contains only mount options
    maintained by kernel (see also --mtab).

-l, --list
    Use the list output format. This output format is automatically
    enabled if the output is restricted by the -t, -O, -S or -T option
    and the option --submounts is not used or if more that one source
    file (the option -F) is specified.

-M, --mountpoint path
    Explicitly define the mountpoint file or directory. See also
    --target.

-m, --mtab
    Search in /etc/mtab. The output is in the list format by default
    (see --tree). The output may include user space mount options.

-N, --task tid
    Use alternative namespace /proc/<tid>/mountinfo rather than the
    default /proc/self/mountinfo. If the option is specified more than
    once, then tree-like output is disabled (see the --list option).
    See also the unshare(1) command.

-n, --noheadings
    Do not print a header line.

-O, --options list
    Limit the set of printed filesystems. More than one option may be
    specified in a comma-separated list. The -t and -O options are
    cumulative in effect. It is different from -t in that each option
    is matched exactly; a leading no at the beginning does not have
    global meaning. The "no" can used for individual items in the list.
    The "no" prefix interpretation can be disabled by "+" prefix.

-o, --output list
    Define output columns. See the --help output to get a list of the
    currently supported columns. The TARGET column contains tree
    formatting if the --list or --raw options are not specified.

    The default list of columns may be extended if list is specified in
    the format +list (e.g., findmnt -o +PROPAGATION).

--output-all
    Output almost all available columns. The columns that require
    --poll are not included.

-P, --pairs
    Produce output in the form of key="value" pairs. All potentially
```

unsafe value characters are hex-escaped (\x<code>). See also option --shell.

-p, --poll[=list]
Monitor changes in the /proc/self/mountinfo file. Supported actions are: mount, umount, remount and move. More than one action may be specified in a comma-separated list. All actions are monitored by default.

The time for which --poll will block can be restricted with the --timeout or --first-only options.

The standard columns always use the new version of the information from the mountinfo file, except the umount action which is based on the original information cached by findmnt. The poll mode allows using extra columns:

ACTION
mount, umount, move or remount action name; this column is enabled by default

OLD-TARGET
available for umount and move actions

OLD-OPTIONS
available for umount and remount actions

--pseudo
Print only pseudo filesystems.

--shadow
Print only filesystems over-mounted by another filesystem.

-R, --submounts
Print recursively all submounts for the selected filesystems. The restrictions defined by options -t, -O, -S, -T and --direction are not applied to submounts. All submounts are always printed in tree-like order. The option enables the tree-like output format by default. This option has no effect for --mtab or --fstab.

-r, --raw
Use raw output format. All potentially unsafe characters are hex-escaped (\x<code>).

--real
Print only real filesystems.

-S, --source spec
Explicitly define the mount source. Supported specifications are device, maj:min, LABEL=label, UUID=uuid, PARTLABEL=label and PARTUUID=uuid.

-s, --fstab

Search in /etc/fstab. The output is in the list format (see --list).

-T, --target path
Define the mount target. If path is not a mountpoint file or directory, then findmnt checks the path elements in reverse order to get the mountpoint (this feature is supported only when searching in kernel files and unsupported for --fstab). It’s recommended to use the option --mountpoint when checks of path elements are unwanted and path is a strictly specified mountpoint.

-t, --types list
Limit the set of printed filesystems. More than one type may be specified in a comma-separated list. The list of filesystem types can be prefixed with no to specify the filesystem types on which no action should be taken. For more details see mount(8).

--tree
Enable tree-like output if possible. The options is silently ignored for tables where is missing child-parent relation (e.g., fstab).

--shadowed
Print only filesystems over-mounted by another filesystem.

-U, --uniq
Ignore filesystems with duplicate mount targets, thus effectively skipping over-mounted mount points.

-u, --notruncate
Do not truncate text in columns. The default is to not truncate the TARGET, SOURCE, UUID, LABEL, PARTUUID, PARTLABEL columns. This option disables text truncation also in all other columns.

-v, --nofsroot
Do not print a [/dir] in the SOURCE column for bind mounts or btrfs subvolumes.

-w, --timeout milliseconds
Specify an upper limit on the time for which --poll will block, in milliseconds.

-x, --verify
Check mount table content. The default is to verify /etc/fstab parsability and usability. It’s possible to use this option also with --tab-file. It’s possible to specify source (device) or target (mountpoint) to filter mount table. The option --verbose forces findmnt to print more details.

--verbose
Force findmnt to print more information (--verify only for now).

--vfs-all

```
            When used with VFS-OPTIONS column, print all VFS (fs-independent)
            flags. This option is designed for auditing purposes to list also
            default VFS kernel mount options which are normally not listed.

        -y, --shell
            The column name will be modified to contain only characters allowed
            for shell variable identifiers. This is usable, for example, with
            --pairs. Note that this feature has been automatically enabled for
            --pairs in version 2.37, but due to compatibility issues, now it's
            necessary to request this behavior by --shell.

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

EXIT STATUS
        The exit value is 0 if there is something to display, or 1 on any error
        (for example if no filesystem is found based on the user's filter
        specification, or the device path or mountpoint does not exist).

ENVIRONMENT
        LIBMOUNT_FSTAB=<path>
            overrides the default location of the fstab file

        LIBMOUNT_MTAB=<path>
            overrides the default location of the mtab file

        LIBMOUNT_DEBUG=all
            enables libmount debug output

        LIBSMARTCOLS_DEBUG=all
            enables libsmartcols debug output

        LIBSMARTCOLS_DEBUG_PADDING=on
            use visible padding characters.

EXAMPLES
        findmnt --fstab -t nfs
            Prints all NFS filesystems defined in /etc/fstab.

        findmnt --fstab /mnt/foo
            Prints all /etc/fstab filesystems where the mountpoint directory is
            /mnt/foo. It also prints bind mounts where /mnt/foo is a source.

        findmnt --fstab --target /mnt/foo
            Prints all /etc/fstab filesystems where the mountpoint directory is
            /mnt/foo.

        findmnt --fstab --evaluate
            Prints all /etc/fstab filesystems and converts LABEL= and UUID=
            tags to the real device names.
```

```
       findmnt -n --raw --evaluate --output=target LABEL=/boot
           Prints only the mountpoint where the filesystem with label "/boot"
           is mounted.

       findmnt --poll --mountpoint /mnt/foo
           Monitors mount, unmount, remount and move on /mnt/foo.

       findmnt --poll=umount --first-only --mountpoint /mnt/foo
           Waits for /mnt/foo unmount.

       findmnt --poll=remount -t ext3 -O ro
           Monitors remounts to read-only mode on all ext3 filesystems.

AUTHORS
       Karel Zak <kzak@redhat.com>

SEE ALSO
       fstab(5), mount(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The findmnt command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                        FINDMNT(8)
```

# Bash 的 flock 命令

```
flock [OPTIONS]
```

**功能**

管理文件的锁。

**类型**

可执行文件（`/usr/bin/flock`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-c, --command command` 通过 `-c` 将一个单独命令（不带参数）传递给 shell 执行。
  - `-E, --conflict-exit-code number` 当使用 `-n`（非阻塞）且锁冲突，或使用 `-w`（等待超时）且超时未获取锁时，指定返回的退出状态码。
  - 默认值为 1
  - 可用范围：0~255
  - `-F, --no-fork` 执行命令前 **不 fork**，flock 进程直接被命令替换，并继续持有锁。
  - 与 `--close` 不兼容，否则没有进程持锁。
  - `-e, -x, --exclusive` 获取独占锁（也称写锁），这是默认模式。
  - `-n, --nb, --nonblock` 如果锁无法立即获取，则直接失败而不等待。
  - 返回值可通过 `-E` 指定
  - `-o, --close` 在执行命令前关闭持锁的文件描述符。
  - 适用于命令生成子进程时，让子进程不持锁
  - `-s, --shared` 获取共享锁（也称读锁）。
  - `-u, --unlock` 释放锁。通常不需要手动调用，因为文件关闭时锁会自动释放。
  - 特殊情况需要，例如命令内部可能 fork 背景进程，不希望其持锁
  - `-w, --wait, --timeout seconds` 在指定秒数内未获取锁则失败。

- 支持小数秒
- 返回值可通过 `-E` 指定
- 秒数为 0 时等同于 `--nonblock`
- `--verbose` 输出获取锁所花的时间，或者说明无法获取锁的原因。
- `--help` - 显示帮助
- `--version` - 显示版本

## 示例

### 阻塞式加锁

```
$ flock /tmp/lockfile -c "echo '任务开始'; sleep 5; echo '任务结束'"    # 加锁然后执行命令
任务开始
任务结束
```

- 如果文件已经被其它进程加锁，则会阻塞到锁释放
- `flock` 结束时自动释放锁

### 非阻塞式加锁

```
$ flock -n /tmp/lockfile -c "echo '任务开始'; sleep 5; echo '任务结束'"    # 加锁然后执行命令
任务开始
任务结束
```

- 如果文件已经被其它进程加锁，则会直接返回失败（不等的也不执行命令）
- `flock` 结束时自动释放锁

### 在脚本内加锁一段时间

```
#!/bin/bash
exec 200 > /tmp/lockfile        # 打开文件并指定文件描述符为 200
flock 200                       # 加锁

echo "任务开始"
sleep 10
echo "任务结束"
```

- 脚本退出时，文件描述符 `200` 关闭，自动释放锁
- 也可以通过 `flock -u 200` 提前释放锁，或通过 `exec 200>&-` 关闭文件来释放锁

## 推荐阅读

- flock(1) — Linux manual page

## 手册

```
FLOCK(1)                        User Commands                        FLOCK(1)

NAME
       flock - manage locks from shell scripts

SYNOPSIS
       flock [options] file|directory command [arguments]

       flock [options] file|directory -c command

       flock [options] number

DESCRIPTION
       This utility manages flock(2) locks from within shell scripts or from
       the command line.

       The first and second of the above forms wrap the lock around the
       execution of a command, in a manner similar to su(1) or newgrp(1). They
       lock a specified file or directory, which is created (assuming
       appropriate permissions) if it does not already exist. By default, if
       the lock cannot be immediately acquired, flock waits until the lock is
       available.

       The third form uses an open file by its file descriptor number. See the
       examples below for how that can be used.

OPTIONS
       -c, --command command
           Pass a single command, without arguments, to the shell with -c.

       -E, --conflict-exit-code number
           The exit status used when the -n option is in use, and the
           conflicting lock exists, or the -w option is in use, and the
           timeout is reached. The default value is 1. The number has to be in
           the range of 0 to 255.

       -F, --no-fork
           Do not fork before executing command. Upon execution the flock
           process is replaced by command which continues to hold the lock.
           This option is incompatible with --close as there would otherwise
           be nothing left to hold the lock.

       -e, -x, --exclusive
           Obtain an exclusive lock, sometimes called a write lock. This is
```

the default.

-n, --nb, --nonblock
Fail rather than wait if the lock cannot be immediately acquired. See the -E option for the exit status used.

-o, --close
Close the file descriptor on which the lock is held before executing command. This is useful if command spawns a child process which should not be holding the lock.

-s, --shared
Obtain a shared lock, sometimes called a read lock.

-u, --unlock
Drop a lock. This is usually not required, since a lock is automatically dropped when the file is closed. However, it may be required in special cases, for example if the enclosed command group may have forked a background process which should not be holding the lock.

-w, --wait, --timeout seconds
Fail if the lock cannot be acquired within seconds. Decimal fractional values are allowed. See the -E option for the exit status used. The zero number of seconds is interpreted as --nonblock.

--verbose
Report how long it took to acquire the lock, or why the lock could not be obtained.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

EXIT STATUS

The command uses <sysexits.h> exit status values for everything, except when using either of the options -n or -w which report a failure to acquire the lock with an exit status given by the -E option, or 1 by default. The exit status given by -E has to be in the range of 0 to 255.

When using the command variant, and executing the child worked, then the exit status is that of the child command.

NOTES

flock does not detect deadlock. See flock(2) for details.

Some file systems (e. g. NFS and CIFS) have a limited implementation of flock(2) and flock may always fail. For details see flock(2), nfs(5) and mount.cifs(8). Depending on mount options, flock can always fail

there.

EXAMPLES

Note that "shell> " in examples is a command line prompt.

shell1> flock /tmp -c cat; shell2> flock -w .007 /tmp -c echo; /bin/echo $?

Set exclusive lock to directory /tmp and the second command will fail.

shell1> flock -s /tmp -c cat; shell2> flock -s -w .007 /tmp -c echo; /bin/echo $?

Set shared lock to directory /tmp and the second command will not fail. Notice that attempting to get exclusive lock with second command would fail.

shell> flock -x local-lock-file echo 'a b c'

Grab the exclusive lock "local-lock-file" before running echo with 'a b c'.

(; flock -n 9 || exit 1; # ... commands executed under lock ...; ) 9>/var/lock/mylockfile

The form is convenient inside shell scripts. The mode used to open the file doesn’t matter to flock; using > or >> allows the lockfile to be created if it does not already exist, however, write permission is required. Using < requires that the file already exists but only read permission is required.

[ "${FLOCKER}" != "$0" ] && exec env FLOCKER="$0" flock -en "$0" "$0" "$@" || :

This is useful boilerplate code for shell scripts. Put it at the top of the shell script you want to lock and it’ll automatically lock itself on the first run. If the environment variable $FLOCKER is not set to the shell script that is being run, then execute flock and grab an exclusive non-blocking lock (using the script itself as the lock file) before re-execing itself with the right arguments. It also sets the FLOCKER environment variable to the right value so it doesn’t run again.

shell> exec 4<>/var/lock/mylockfile; shell> flock -n 4

This form is convenient for locking a file without spawning a subprocess. The shell opens the lock file for reading and writing as file descriptor 4, then flock is used to lock the descriptor.

AUTHORS

H. Peter Anvin <hpa@zytor.com>

COPYRIGHT

Copyright © 2003-2006 H. Peter Anvin. This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SEE ALSO

```
       flock(2)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The flock command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          FLOCK(1)
```

# Bash 的 fmt 命令

```
fmt [-WIDTH] [OPTION]... [FILE]...
```

**功能**

简单文本格式化。

**类型**

可执行文件（`/usr/bin/fmt`），属于 coreutils。

**参数**

- `WIDTH` - 每行宽度；同 `-w` 选项，默认为 75
- `OPTION` 选项：
  - `-c`，`--crown-margin` - 保留前两行的缩进
  - `-p`，`--prefix=STRING` - 仅格式化以 `STRING` 开头的行，并将前缀重新附加到格式化的行
  - `-s`，`--split-only` - 拆分长行，但不补充
  - `-t`，`--tagged-paragraph` - 第一行的缩进与第二行不同
  - `-u`，`--uniform-spacing` - 单词之间空一格，句子之后空两格
  - `-w`，`--width=WIDTH` - 每行宽度；默认为 75
  - `-g`，`--goal=WIDTH` - 目标宽度（单词提前换行）；默认为行宽的 93%
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ echo 'Primers 编程伙伴 https://xplanc.org/primers/' > 1.txt    # 创建文件
$ cat 1.txt                                                      # 查看文件
Primers 编程伙伴 https://xplanc.org/primers/
$ fmt -w 40 1.txt                                                # 行宽度 40
Primers 编程伙伴
https://xplanc.org/primers/
$ fmt -w 40 -g 10 1.txt                                          # 行宽度 40 目标宽度 10
Primers
编程伙伴
https://xplanc.org/primers/
```

# 推荐阅读

- Reformat paragraph text - GNU Coreutils
- fmt(1) — Linux manual page

# 手册

```
FMT(1)                          User Commands                          FMT(1)

NAME
       fmt - simple optimal text formatter

SYNOPSIS
       fmt [-WIDTH] [OPTION]... [FILE]...

DESCRIPTION
       Reformat  each  paragraph  in  the FILE(s), writing to standard output.
       The option -WIDTH is an abbreviated form of --width=DIGITS.

       With no FILE, or when FILE is -, read standard input.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -c, --crown-margin
              preserve indentation of first two lines

       -p, --prefix=STRING
              reformat  only lines beginning with STRING, reattaching the pre-
              fix to reformatted lines

       -s, --split-only
              split long lines, but do not refill

       -t, --tagged-paragraph
              indentation of first line different from second

       -u, --uniform-spacing
              one space between words, two after sentences

       -w, --width=WIDTH
              maximum line width (default of 75 columns)

       -g, --goal=WIDTH
              goal width (default of 93% of width)

       --help display this help and exit

       --version
```

```
              output version information and exit

AUTHOR
       Written by Ross Paterson.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/fmt>
       or available locally via: info '(coreutils) fmt invocation'

GNU coreutils 9.4                April 2024                            FMT(1)
```

# Bash 的 fold 命令

```
fold [OPTION]... [FILE]...
```

**功能**

自动换行以适应宽度。

**类型**

可执行文件（`/usr/bin/fold`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--bytes` - 计算字节数而不是列数
  - `-s`，`--spaces` - 在空格处换行
  - `-w`，`--width=WIDTH` - 行的宽度；默认为 80
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ echo 'Primers 编程伙伴 https://xplanc.org/primers/' | fold -w 20
Primers 编程伙伴
 https://xplanc.org/
primers/
```

## 推荐阅读

- Wrap input lines to fit in specified width - GNU Coreutils
- fold(1) — Linux manual page

## 手册

```
FOLD(1)                          User Commands                          FOLD(1)

NAME
```

```
       fold - wrap each input line to fit in specified width

SYNOPSIS
       fold [OPTION]... [FILE]...

DESCRIPTION
       Wrap input lines in each FILE, writing to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -b, --bytes
              count bytes rather than columns

       -s, --spaces
              break at spaces

       -w, --width=WIDTH
              use WIDTH columns instead of 80

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       fmt(1)

       Full documentation <https://www.gnu.org/software/coreutils/fold>
       or available locally via: info '(coreutils) fold invocation'

GNU coreutils 9.4                April 2024                          FOLD(1)
```

# Bash 的 fsck.cramfs 命令

```
fsck.cramfs [OPTION]... [FILESYSTEM]
```

**功能**

检查并修复压缩 ROM 文件系统。

**类型**

可执行文件（`/usr/sbin/fsck.cramfs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-v`，`--verbose` - 显示详细信息
  - `-b`，`--blocksize blocksize` - 块大小
  - `--extract[=directory]` - 测试解压缩整个文件系统
  - `-a` - ~~自动修复~~；此选项将被静默忽略
  - `-y` - ~~总是 yes~~；此选项将被静默忽略
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FILESYSTEM` - 要检查的文件系统

**返回值**

- 0 - 成功
- 4 - 文件系统错误未得到纠正
- 8 - 操作错误
- 16 - 打印使用信息

## 示例

```
$ fsck.cramfs -v dir.img
cramfs endianness is little
dir.img: OK
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| fsck | 检查并修复 Linux 文件系统 |
| mkfs.cramfs | 创建 cramfs 文件系统 |

## 推荐阅读

- fsck.cramfs(8) — Linux manual page

## 手册

```
FSCK.CRAMFS(8)              System Administration             FSCK.CRAMFS(8)

NAME
       fsck.cramfs - fsck compressed ROM file system

SYNOPSIS
       fsck.cramfs [options] file

DESCRIPTION
       fsck.cramfs is used to check the cramfs file system.

OPTIONS
       -v, --verbose
           Enable verbose messaging.

       -b, --blocksize blocksize
           Use this blocksize, defaults to page size. Must be equal to what
           was set at creation time. Only used for --extract.

       --extract[=directory]
           Test to uncompress the whole file system. Optionally extract
           contents of the file to directory.

       -a
           This option is silently ignored.

       -y
           This option is silently ignored.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       0
```

```
           success

       4
           file system was left uncorrected

       8
           operation error, such as unable to allocate memory

       16
           usage information was printed

SEE ALSO
       mount(8), mkfs.cramfs(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The fsck.cramfs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                   FSCK.CRAMFS(8)
```

# Bash 的 fsck.minix 命令

```
fsck.minix [OPTION]... FILESYSTEM
```

**功能**

检查并修复 minix 文件系统。

**类型**

可执行文件（`/usr/sbin/fsck.minix`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-l`，`--list` - 列出所有文件名
  - `-r`，`--repair`
  - `-a`，`--auto`
  - `-v`，`--verbose`
  - `-s`，`--super`
  - `-m`，`--uncleared`
  - `-f`，`--force`
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FILESYSTEM` - 要检查的文件系统

## 示例

```
$ fsck.minix -v minix.img       # 检查文件系统

     5 inodes used (0%)
   178 zones used (1%)

     3 regular files
     2 directories
     0 character device files
     0 block device files
     0 links
     0 symbolic links
------
     5 files
```

```
$ fsck.minix -l minix.img       # 查看文件系统
/1.txt
/2.txt
/dir:
/dir/file
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| fsck | 检查并修复 Linux 文件系统 |
| mkfs.minix | 创建 minix 文件系统 |

## 推荐阅读

- fsck.minix(8) — Linux manual page

## 手册

```
FSCK.MINIX(8)                System Administration                FSCK.MINIX(8)

NAME
       fsck.minix - check consistency of Minix filesystem

SYNOPSIS
       fsck.minix [options] device

DESCRIPTION
       fsck.minix performs a consistency check for the Linux MINIX filesystem.

       The program assumes the filesystem is quiescent. fsck.minix should not
       be used on a mounted device unless you can be sure nobody is writing to
       it. Remember that the kernel can write to device when it searches for
       files.

       The device name will usually have the following form:
           ┌──────────────────┬──────────────┐
           │                  │              │
           │ /dev/hda[1-63]   │ IDE disk 1   │
           ├──────────────────┼──────────────┤
           │                  │              │
           │ /dev/hdb[1-63]   │ IDE disk 2   │
           ├──────────────────┼──────────────┤
           │                  │              │
           │ /dev/sda[1-15]   │ SCSI disk 1  │
           ├──────────────────┼──────────────┤
```

```
          │                │             │
          │ /dev/sdb[1-15] │ SCSI disk 2 │
          └────────────────┴─────────────┘

       If the filesystem was changed, i.e., repaired, then fsck.minix
       will print "FILE SYSTEM HAS BEEN CHANGED" and will sync(2)
       three times before exiting. There is no need to reboot after
       check.

WARNING
       fsck.minix should not be used on a mounted filesystem. Using
       fsck.minix on a mounted filesystem is very dangerous, due to
       the possibility that deleted files are still in use, and can
       seriously damage a perfectly good filesystem! If you absolutely
       have to run fsck.minix on a mounted filesystem, such as the
       root filesystem, make sure nothing is writing to the disk, and
       that no files are "zombies" waiting for deletion.

OPTIONS
       -l, --list
           List all filenames.

       -r, --repair
           Perform interactive repairs.

       -a, --auto
           Perform automatic repairs. This option implies --repair and
           serves to answer all of the questions asked with the
           default. Note that this can be extremely dangerous in the
           case of extensive filesystem damage.

       -v, --verbose
           Be verbose.

       -s, --super
           Output super-block information.

       -m, --uncleared
           Activate MINIX-like "mode not cleared" warnings.

       -f, --force
           Force a filesystem check even if the filesystem was marked
           as valid. Marking is done by the kernel when the filesystem
           is unmounted.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

DIAGNOSTICS
       There are numerous diagnostic messages. The ones mentioned here
```

are the most commonly seen in normal usage.

If the device does not exist, fsck.minix will print "unable to read super block". If the device exists, but is not a MINIX filesystem, fsck.minix will print "bad magic number in super-block".

## EXIT STATUS

The exit status returned by fsck.minix is the sum of the following:

0
: No errors

3
: Filesystem errors corrected, system should be rebooted if filesystem was mounted

4
: Filesystem errors left uncorrected

7
: Combination of exit statuses 3 and 4

8
: Operational error

16
: Usage or syntax error

## AUTHORS

Linus Torvalds <torvalds@cs.helsinki.fi>. Exit status values by Rik Faith <faith@cs.unc.edu> Added support for filesystem valid flag: Dr. Wettstein <greg%wind.uucp@plains.nodak.edu>. Check to prevent fsck of mounted filesystem added by Daniel Quinlan <quinlan@yggdrasil.com>. Minix v2 fs support by Andreas Schwab <schwab@issan.informatik.uni-dortmund.de>, updated by Nicolai Langfeldt <janl@math.uio.no>. Portability patch by Russell King <rmk@ecs.soton.ac.uk>.

## SEE ALSO

fsck(8), fsck.ext2(8), mkfs(8), mkfs.ext2(8), mkfs.minix(8), reboot(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The fsck.minix command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 fsck 命令

```
fsck [OPTION]... [FILESYSTEM]...
```

**功能**

检查并修复 Linux 文件系统。

> `fsck` 会扫描文件系统结构，检查 inode、block、目录结构、链接关系等是否一致，必要时尝试修复。
>
> 主要用于以下场景：
>
> - 系统异常断电
> - 磁盘损坏
> - 文件系统错误
> - 启动时自检

**类型**

可执行文件（`/usr/sbin/fsck`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-l` - 为整个磁盘设备创建独占的 `flock(2)` 锁定文件（`/run/fsck/<diskname>.lock`）
  - `-r [fd]` - 每次 fsck 完成后，报告某些统计信息，如果指定了文件描述符 `fd`，则将进度信息写入该文件
  - `-s` - 序列化 fsck 操作
  - `-t fslist` - 指定要检查的文件系统类型
  - `-A` - 遍历 `/etc/fstab` 文件，并尝试一次性检查所有文件系统
  - `-C [fd]` - 显示文件系统检查器的进度条，如果指定了文件描述符 `fd`，则将进度信息写入该文件
  - `-M` - 不要检查已挂载的文件系统
  - `-N` - 显示要执行的操作而不实际执行
  - `-P` - 当 `-A` 标志被设置时，根文件系统将与其他文件系统并行检查
  - `-R` - 当 `-A` 标志被设置时，跳过根文件系统
  - `-T` - 启动时不显示标题
  - `-V` - 生成详细输出
  - `-?`, `--help` - 显示帮助

    - `--version` - 显示版本
- `FILESYSTEM` - 要检查的文件系统

**返回值**

- 0 - 没有错误
- 1 - 文件系统错误已修复
- 2 - 系统需要重启
- 4 - 文件系统错误未得到纠正
- 8 - 操作错误
- 16 - 用法或语法错误
- 32 - 检查被用户请求取消
- 128 - 共享库错误

## 示例

```
$ sudo fsck /dev/sda1
fsck from util-linux 2.39.3
fsck.fat 4.2 (2021-01-31)
There is no label in boot sector, but there is volume label 'ESD-USB' stored in root
directory
1) Copy volume label from root directory to boot sector
2) Remove volume label from root directory
[12?q]? 2

*** Filesystem was changed ***
The changes have not yet been written, you can still choose to leave the
filesystem unmodified:
1) Write changes
2) Leave filesystem unchanged
[12?q]? 1
/dev/sda1: 8052 files, 321598/1919936 clusters
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| mkfs | 创建 Linux 文件系统 |
| fsck.btrfs | 检查并修复 btrfs 文件系统 |
| fsck.cramfs | 检查并修复 cramfs 文件系统 |
| fsck.ext2 | 检查并修复 ext2 文件系统 |

| 命令 | 说明 |
| --- | --- |
| fsck.ext3 | 检查并修复 ext3 文件系统 |
| fsck.ext4 | 检查并修复 ext4 文件系统 |
| fsck.fat | 检查并修复 fat 文件系统 |
| fsck.minix | 检查并修复 minix 文件系统 |
| fsck.msdos | 检查并修复 msdos 文件系统 |
| fsck.vfat | 检查并修复 vfat 文件系统 |
| fsck.xfs | 检查并修复 xfs 文件系统 |

## 推荐阅读

- fsck(8) — Linux manual page

## 手册

```
FSCK(8)                       System Administration                       FSCK(8)

NAME
       fsck - check and repair a Linux filesystem

SYNOPSIS
       fsck [-lsAVRTMNP] [-r [fd]] [-C [fd]] [-t fstype] [filesystem...] [--]
       [fs-specific-options]

DESCRIPTION
       fsck is used to check and optionally repair one or more Linux
       filesystems. filesystem can be a device name (e.g., /dev/hdc1,
       /dev/sdb2), a mount point (e.g., /, /usr, /home), or a filesystem label
       or UUID specifier (e.g., UUID=8868abf6-88c5-4a83-98b8-bfc24057f7bd or
       LABEL=root). Normally, the fsck program will try to handle filesystems
       on different physical disk drives in parallel to reduce the total
       amount of time needed to check all of them.

       If no filesystems are specified on the command line, and the -A option
       is not specified, fsck will default to checking filesystems in
       /etc/fstab serially. This is equivalent to the -As options.

       The exit status returned by fsck is the sum of the following
       conditions:

       0
           No errors
```

```
       1
           Filesystem errors corrected

       2
           System should be rebooted

       4
           Filesystem errors left uncorrected

       8
           Operational error

       16
           Usage or syntax error

       32
           Checking canceled by user request

       128
           Shared-library error

       The exit status returned when multiple filesystems are checked is the
       bit-wise OR of the exit statuses for each filesystem that is checked.

       In actuality, fsck is simply a front-end for the various filesystem
       checkers (fsck.fstype) available under Linux. The filesystem-specific
       checker is searched for in the PATH environment variable. If the PATH
       is undefined then fallback to /sbin.

       Please see the filesystem-specific checker manual pages for further
       details.

OPTIONS
       -l
           Create an exclusive flock(2) lock file (/run/fsck/<diskname>.lock)
           for whole-disk device. This option can be used with one device only
           (this means that -A and -l are mutually exclusive). This option is
           recommended when more fsck instances are executed in the same time.
           The option is ignored when used for multiple devices or for
           non-rotating disks. fsck does not lock underlying devices when
           executed to check stacked devices (e.g. MD or DM) - this feature is
           not implemented yet.

       -r [fd]
           Report certain statistics for each fsck when it completes. These
           statistics include the exit status, the maximum run set size (in
           kilobytes), the elapsed all-clock time and the user and system CPU
           time used by the fsck run. For example:

           /dev/sda1: status 0, rss 92828, real 4.002804, user 2.677592, sys
           0.86186
```

GUI front-ends may specify a file descriptor fd, in which case the progress bar information will be sent to that file descriptor in a machine parsable format. For example:

/dev/sda1 0 92828 4.002804 2.677592 0.86186

-s

Serialize fsck operations. This is a good idea if you are checking multiple filesystems and the checkers are in an interactive mode. (Note: e2fsck(8) runs in an interactive mode by default. To make e2fsck(8) run in a non-interactive mode, you must either specify the -p or -a option, if you wish for errors to be corrected automatically, or the -n option if you do not.)

-t fslist

Specifies the type(s) of filesystem to be checked. When the -A flag is specified, only filesystems that match fslist are checked. The fslist parameter is a comma-separated list of filesystems and options specifiers. All of the filesystems in this comma-separated list may be prefixed by a negation operator 'no' or '!', which requests that only those filesystems not listed in fslist will be checked. If none of the filesystems in fslist is prefixed by a negation operator, then only those listed filesystems will be checked.

Options specifiers may be included in the comma-separated fslist. They must have the format opts=fs-option. If an options specifier is present, then only filesystems which contain fs-option in their mount options field of /etc/fstab will be checked. If the options specifier is prefixed by a negation operator, then only those filesystems that do not have fs-option in their mount options field of /etc/fstab will be checked.

For example, if opts=ro appears in fslist, then only filesystems listed in /etc/fstab with the ro option will be checked.

For compatibility with Mandrake distributions whose boot scripts depend upon an unauthorized UI change to the fsck program, if a filesystem type of loop is found in fslist, it is treated as if opts=loop were specified as an argument to the -t option.

Normally, the filesystem type is deduced by searching for filesys in the /etc/fstab file and using the corresponding entry. If the type cannot be deduced, and there is only a single filesystem given as an argument to the -t option, fsck will use the specified filesystem type. If this type is not available, then the default filesystem type (currently ext2) is used.

-A

Walk through the /etc/fstab file and try to check all filesystems in one run. This option is typically used from the /etc/rc system initialization file, instead of multiple commands for checking a single filesystem.

The root filesystem will be checked first unless the -P option is specified (see below). After that, filesystems will be checked in the order specified by the fs_passno (the sixth) field in the /etc/fstab file. Filesystems with a fs_passno value of 0 are skipped and are not checked at all. Filesystems with a fs_passno value of greater than zero will be checked in order, with filesystems with the lowest fs_passno number being checked first. If there are multiple filesystems with the same pass number, fsck will attempt to check them in parallel, although it will avoid running multiple filesystem checks on the same physical disk.

fsck does not check stacked devices (RAIDs, dm-crypt, ...) in parallel with any other device. See below for FSCK_FORCE_ALL_PARALLEL setting. The /sys filesystem is used to determine dependencies between devices.

Hence, a very common configuration in /etc/fstab files is to set the root filesystem to have a fs_passno value of 1 and to set all other filesystems to have a fs_passno value of 2. This will allow fsck to automatically run filesystem checkers in parallel if it is advantageous to do so. System administrators might choose not to use this configuration if they need to avoid multiple filesystem checks running in parallel for some reason - for example, if the machine in question is short on memory so that excessive paging is a concern.

fsck normally does not check whether the device actually exists before calling a filesystem specific checker. Therefore non-existing devices may cause the system to enter filesystem repair mode during boot if the filesystem specific checker returns a fatal error. The /etc/fstab mount option nofail may be used to have fsck skip non-existing devices. fsck also skips non-existing devices that have the special filesystem type auto.

-C [fd]
Display completion/progress bars for those filesystem checkers (currently only for ext[234]) which support them. fsck will manage the filesystem checkers so that only one of them will display a progress bar at a time. GUI front-ends may specify a file descriptor fd, in which case the progress bar information will be sent to that file descriptor.

-M
Do not check mounted filesystems and return an exit status of 0 for mounted filesystems.

-N
Don’t execute, just show what would be done.

-P
When the -A flag is set, check the root filesystem in parallel with the other filesystems. This is not the safest thing in the world to

do, since if the root filesystem is in doubt things like the e2fsck(8) executable might be corrupted! This option is mainly provided for those sysadmins who don’t want to repartition the root filesystem to be small and compact (which is really the right solution).

-R
When checking all filesystems with the -A flag, skip the root filesystem. (This is useful in case the root filesystem has already been mounted read-write.)

-T
Don’t show the title on startup.

-V
Produce verbose output, including all filesystem-specific commands that are executed.

-?, --help
Display help text and exit.

--version
Display version information and exit.

FILESYSTEM SPECIFIC OPTIONS

Options which are not understood by fsck are passed to the filesystem-specific checker!

These options must not take arguments, as there is no way for fsck to be able to properly guess which options take arguments and which don’t.

Options and arguments which follow the -- are treated as filesystem-specific options to be passed to the filesystem-specific checker.

Please note that fsck is not designed to pass arbitrarily complicated options to filesystem-specific checkers. If you’re doing something complicated, please just execute the filesystem-specific checker directly. If you pass fsck some horribly complicated options and arguments, and it doesn’t do what you expect, don’t bother reporting it as a bug. You’re almost certainly doing something that you shouldn’t be doing with fsck. Options to different filesystem-specific fsck’s are not standardized.

ENVIRONMENT

The fsck program’s behavior is affected by the following environment variables:

FSCK_FORCE_ALL_PARALLEL
If this environment variable is set, fsck will attempt to check all of the specified filesystems in parallel, regardless of whether the filesystems appear to be on the same device. (This is useful for RAID systems or high-end storage systems such as those sold by

companies such as IBM or EMC.) Note that the fs_passno value is still used.

FSCK_MAX_INST
This environment variable will limit the maximum number of filesystem checkers that can be running at one time. This allows configurations which have a large number of disks to avoid fsck starting too many filesystem checkers at once, which might overload CPU and memory resources available on the system. If this value is zero, then an unlimited number of processes can be spawned. This is currently the default, but future versions of fsck may attempt to automatically determine how many filesystem checks can be run based on gathering accounting data from the operating system.

PATH
The PATH environment variable is used to find filesystem checkers.

FSTAB_FILE
This environment variable allows the system administrator to override the standard location of the /etc/fstab file. It is also useful for developers who are testing fsck.

LIBBLKID_DEBUG=all
enables libblkid debug output.

LIBMOUNT_DEBUG=all
enables libmount debug output.

## FILES

/etc/fstab

## AUTHORS

Theodore Ts’o <tytso@mit.edu>>, Karel Zak <kzak@redhat.com>

## SEE ALSO

fstab(5), mkfs(8), fsck.ext2(8) or fsck.ext3(8) or e2fsck(8), fsck.cramfs(8), fsck.jfs(8), fsck.nfs(8), fsck.minix(8), fsck.msdos(8), fsck.vfat(8), fsck.xfs(8), reiserfsck(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The fsck command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 fsfreeze 命令

```
fsfreeze [OPTION]... MOUNTPOINT
```

**功能**

暂停对文件系统（Ext3/4、ReiserFS、JFS、XFS）的访问。

**类型**

可执行文件（`/usr/sbin/fsfreeze`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-f`，`--freeze` - 冻结指定的文件系统
  - `-u`，`--unfreeze` - 解冻指定的文件系统
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `MOUNTPOINT` - 要操作的挂载点

## 示例

```
$ sudo fsfreeze --freeze ./mnt          # 冻结挂载点
$ sudo fsfreeze --unfreeze ./mnt        # 解冻挂载点
```

```
$ sudo fsfreeze --freeze ./mnt          # 冻结挂载点
$ sudo echo hello > ./mnt/file.txt      # 尝试写入；会阻塞，直到挂载点解冻
```

## 推荐阅读

- fsfreeze(8) — Linux manual page

## 手册

```
FSFREEZE(8)                System Administration                FSFREEZE(8)
```

NAME

fsfreeze - suspend access to a filesystem (Ext3/4, ReiserFS, JFS, XFS)

SYNOPSIS

fsfreeze --freeze|--unfreeze mountpoint

DESCRIPTION

fsfreeze suspends or resumes access to a filesystem.

fsfreeze halts any new access to the filesystem and creates a stable image on disk. fsfreeze is intended to be used with hardware RAID devices that support the creation of snapshots.

fsfreeze is unnecessary for device-mapper devices. The device-mapper (and LVM) automatically freezes a filesystem on the device when a snapshot creation is requested. For more details see the dmsetup(8) man page.

The mountpoint argument is the pathname of the directory where the filesystem is mounted. The filesystem must be mounted to be frozen (see mount(8)).

Note that access-time updates are also suspended if the filesystem is mounted with the traditional atime behavior (mount option strictatime, for more details see mount(8)).

OPTIONS

-f, --freeze

This option requests the specified filesystem to be frozen from new modifications. When this is selected, all ongoing transactions in the filesystem are allowed to complete, new write(2) system calls are halted, other calls which modify the filesystem are halted, and all dirty data, metadata, and log information are written to disk. Any process attempting to write to the frozen filesystem will block waiting for the filesystem to be unfrozen.

Note that even after freezing, the on-disk filesystem can contain information on files that are still in the process of unlinking. These files will not be unlinked until the filesystem is unfrozen or a clean mount of the snapshot is complete.

-u, --unfreeze

This option is used to un-freeze the filesystem and allow operations to continue. Any filesystem modifications that were blocked by the freeze are unblocked and allowed to complete.

-h, --help

Display help text and exit.

-V, --version

Print version and exit.

FILESYSTEM SUPPORT

This command will work only if filesystem supports has support for freezing. List of these filesystems include (2016-12-18) btrfs, ext2/3/4, f2fs, jfs, nilfs2, reiserfs, and xfs. Previous list may be incomplete, as more filesystems get support. If in doubt easiest way to know if a filesystem has support is create a small loopback mount and test freezing it.

## NOTES

This man page is based on xfs_freeze(8).

## AUTHORS

Written by Hajime Taira.

## SEE ALSO

mount(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The fsfreeze command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 fstrim 命令

```
fstrim [OPTION]... [MOUNTPOINT]
```

**功能**

丢弃已挂载文件系统上未使用的数据块。

**类型**

可执行文件（`/usr/sbin/fstrim`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-A`，`--fstab` - 修剪 `/etc/fstab` 中的所有支持此操作的已挂载文件系统
  - `-a`，`--all` - 修剪所有支持此操作的已挂载文件系统
  - `-n`，`--dry-run` - 执行整个过程，但不进行实际修剪
  - `-o`，`--offset offset` - 从文件系统的 `offset` 字节处开始搜索；默认为 0
  - `-l`，`--length length` - 从 `offset` 开始搜索 `length` 字节
  - `-I`，`--listed-in list` - 修剪 `list` 列表中文件里指示的文件系统；多个文件用冒号（`:`）分隔
  - `-m`，`--minimum minimum-size` - 要丢弃的最小连续空闲范围；单位是字节
  - `-t`，`--types list` - 允许或禁止修剪此类型的文件系统；多个类型用逗号（`,`）分隔，`no` 前缀表示禁止修剪的类型
  - `-v`，`--verbose` - 打印详细信息
  - `--quiet-unsupported` - 遇到不支持的文件系统时，不打印提示
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `MOUNTPOINT` - 要操作的挂载点

## 示例

```
$ sudo umount ./mnt
```

## 推荐阅读

- fstrim(8) — Linux manual page

# 手册

```
FSTRIM(8)                   System Administration                   FSTRIM(8)

NAME
       fstrim - discard unused blocks on a mounted filesystem

SYNOPSIS
       fstrim [-Aav] [-o offset] [-l length] [-m minimum-size] [mountpoint]

DESCRIPTION
       fstrim is used on a mounted filesystem to discard (or "trim") blocks
       which are not in use by the filesystem. This is useful for solid-state
       drives (SSDs) and thinly-provisioned storage.

       By default, fstrim will discard all unused blocks in the filesystem.
       Options may be used to modify this behavior based on range or size, as
       explained below.

       The mountpoint argument is the pathname of the directory where the
       filesystem is mounted and is required when -A, -a, --fstab, or --all
       are unspecified.

       Running fstrim frequently, or even using mount -o discard, might
       negatively affect the lifetime of poor-quality SSD devices. For most
       desktop and server systems a sufficient trimming frequency is once a
       week. Note that not all devices support a queued trim, so each trim
       command incurs a performance penalty on whatever else might be trying
       to use the disk at the time.

OPTIONS
       The offset, length, and minimum-size arguments may be followed by the
       multiplicative suffixes KiB (=1024), MiB (=1024*1024), and so on for
       GiB, TiB, PiB, EiB, ZiB and YiB (the "iB" is optional, e.g., "K" has
       the same meaning as "KiB") or the suffixes KB (=1000), MB (=1000*1000),
       and so on for GB, TB, PB, EB, ZB and YB.

       -A, --fstab
           Trim all mounted filesystems mentioned in /etc/fstab on devices
           that support the discard operation. The root filesystem is
           determined from kernel command line if missing in the file. The
           other supplied options, like --offset, --length and --minimum, are
           applied to all these devices. Errors from filesystems that do not
           support the discard operation, read-only devices, autofs and
           read-only filesystems are silently ignored. Filesystems with
           "X-fstrim.notrim" mount option are skipped.

       -a, --all
           Trim all mounted filesystems on devices that support the discard
           operation. The other supplied options, like --offset, --length and
           --minimum, are applied to all these devices. Errors from
```

filesystems that do not support the discard operation, read-only devices and read-only filesystems are silently ignored.

-n, --dry-run
This option does everything apart from actually call FITRIM ioctl.

-o, --offset offset
Byte offset in the filesystem from which to begin searching for free blocks to discard. The default value is zero, starting at the beginning of the filesystem.

-l, --length length
The number of bytes (after the starting point) to search for free blocks to discard. If the specified value extends past the end of the filesystem, fstrim will stop at the filesystem size boundary. The default value extends to the end of the filesystem.

-I, --listed-in list
Specifies a colon-separated list of files in fstab or kernel mountinfo format. All missing or empty files are silently ignored. The evaluation of the list stops after first non-empty file. For example:

--listed-in /etc/fstab:/proc/self/mountinfo.

Filesystems with "X-fstrim.notrim" mount option in fstab are skipped.

-m, --minimum minimum-size
Minimum contiguous free range to discard, in bytes. (This value is internally rounded up to a multiple of the filesystem block size.) Free ranges smaller than this will be ignored and fstrim will adjust the minimum if it’s smaller than the device’s minimum, and report that (fstrim_range.minlen) back to userspace. By increasing this value, the fstrim operation will complete more quickly for filesystems with badly fragmented freespace, although not all blocks will be discarded. The default value is zero, discarding every free block.

-t, --types list
Specifies allowed or forbidden filesystem types when used with --all or --fstab. The list is a comma-separated list of the filesystem names. The list follows how mount -t evaluates type patterns. Only specified filesystem types are allowed. All specified types are forbidden if the list is prefixed by "no" or each filesystem prefixed by "no" is forbidden. If the option is not used, then all filesystems (except "autofs") are allowed.

-v, --verbose
Verbose execution. With this option fstrim will output the number of bytes passed from the filesystem down the block stack to the device for potential discard. This number is a maximum discard amount from the storage device’s perspective, because FITRIM ioctl

```
            called repeated will keep sending the same sectors for discard
            repeatedly.

            fstrim will report the same potential discard bytes each time, but
            only sectors which had been written to between the discards would
            actually be discarded by the storage device. Further, the kernel
            block layer reserves the right to adjust the discard ranges to fit
            raid stripe geometry, non-trim capable devices in a LVM setup, etc.
            These reductions would not be reflected in fstrim_range.len (the
            --length option).

       --quiet-unsupported
            Suppress error messages if trim operation (ioctl) is unsupported.
            This option is meant to be used in systemd service file or in
            cron(8) scripts to hide warnings that are result of known problems,
            such as NTFS driver reporting Bad file descriptor when device is
            mounted read-only, or lack of file system support for ioctl FITRIM
            call. This option also cleans exit status when unsupported
            filesystem specified on fstrim command line.

       -h, --help
            Display help text and exit.

       -V, --version
            Print version and exit.

EXIT STATUS
       0
            success

       1
            failure

       32
            all failed

       64
            some filesystem discards have succeeded, some failed

       The command fstrim --all returns 0 (all succeeded), 32 (all failed) or
       64 (some failed, some succeeded).

AUTHORS
       Lukas Czerner <lczerner@redhat.com>, Karel Zak <kzak@redhat.com>

SEE ALSO
       blkdiscard(8), mount(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
```

```
       The fstrim command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-11-21                          FSTRIM(8)
```

# Bash 的 getopt 命令

```
getopt [OPTIONS] -- PARAMETERS
```

**功能**

解析命令行参数。

**类型**

可执行文件（`/usr/bin/getopt`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-a`，`--alternative` - 允许长选项使用单个横线
  - `-l`，`--longoptions OPTIONS` - 要识别的长选项列表，用逗号分隔，用冒号表示选项需要附带参数
  - `-n`，`--name NAME` - 打印错误消息时以 `NAME` 作为程序的名称
  - `-o`，`--options OPTIONS` - 要识别的短选项列表，用冒号表示选项需要附带参数；无此选项则以首个参数作为此选项的值
  - `-q`，`--quiet` - 不打印错误消息
  - `-Q`，`--quiet-output` - 不打印常规输出
  - `-s`，`--shell SHELL` - 按照 `SEHLL` 的需要生成引号；默认为 BASH
  - `-T`，`--test` - 检测 `getopt` 命令是否为增强版本
  - `-u`，`--unquoted` - 不对输出添加引号
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `PARAMETERS` - 要解析的参数列表

## 示例

**短选项**

```
$ getopt -o f:hv -- -f file other_args
 -f 'file' -- 'other_args'
$ getopt -o f:hv -- -f file other_args -h
 -f 'file' -h -- 'other_args'
```

- `f:hv` 表示接受 `-f`，`-h`，`-v` 三个选项，其中 `-f` 需要附带参数

**长选项**

```
$ getopt -o '' -l input:,output: -- --input file1 --output file2 other_args
 --input 'file1' --output 'file2' -- 'other_args'
```

- `-o ''` 表示没有短选项，如果不这样设置，则会将 `--input` 当作短选项列表

**在脚本文件中使用**

```
#!/bin/bash

# 解析参数，报告错误时以 $0 为名
PARSED=$(getopt -n "$0" -o hvf: -l help,verbose,file: -- "$@") || exit 1

# 将解析后的参数设为 $1 $2 ...
eval set -- "$PARSED"

while true; do
  case "$1" in
    -h|--help)
      echo "show help"
      shift                # 从 $@ 移除一个参数
      ;;
    -v|--verbose)
      echo "verbose mode"
      shift                # 从 $@ 移除一个参数
      ;;
    -f|--file)
      echo "file = $2"
      shift 2              # 从 $@ 移除两个参数
      ;;
    --)
      shift                # 从 $@ 移除一个参数
      break
      ;;
  esac
done

echo "剩余参数: $@"
```

# 推荐阅读

- getopt(1) — Linux manual page

# 手册

NAME

getopt - parse command options (enhanced)

SYNOPSIS

```
getopt optstring parameters

getopt [options] [--] optstring parameters

getopt [options] -o|--options optstring [options] [--] parameters
```

DESCRIPTION

getopt is used to break up (parse) options in command lines for easy parsing by shell procedures, and to check for valid options. It uses the GNU getopt(3) routines to do this.

The parameters getopt is called with can be divided into two parts: options which modify the way getopt will do the parsing (the options and the optstring in the SYNOPSIS), and the parameters which are to be parsed (parameters in the SYNOPSIS). The second part will start at the first non-option parameter that is not an option argument, or after the first occurrence of '--'. If no '-o' or '--options' option is found in the first part, the first parameter of the second part is used as the short options string.

If the environment variable GETOPT_COMPATIBLE is set, or if the first parameter is not an option (does not start with a '-', the first format in the SYNOPSIS), getopt will generate output that is compatible with that of other versions of getopt(1). It will still do parameter shuffling and recognize optional arguments (see the COMPATIBILITY section for more information).

Traditional implementations of getopt(1) are unable to cope with whitespace and other (shell-specific) special characters in arguments and non-option parameters. To solve this problem, this implementation can generate quoted output which must once again be interpreted by the shell (usually by using the eval command). This has the effect of preserving those characters, but you must call getopt in a way that is no longer compatible with other versions (the second or third format in the SYNOPSIS). To determine whether this enhanced version of getopt(1) is installed, a special test option (-T) can be used.

OPTIONS

-a, --alternative
Allow long options to start with a single '-'.

-l, --longoptions longopts
The long (multi-character) options to be recognized. More than one option name may be specified at once, by separating the names with commas. This option may be given more than once, the longopts are cumulative. Each long option name in longopts may be followed by

one colon to indicate it has a required argument, and by two colons to indicate it has an optional argument.

-n, --name progname
The name that will be used by the getopt(3) routines when it reports errors. Note that errors of getopt(1) are still reported as coming from getopt.

-o, --options shortopts
The short (one-character) options to be recognized. If this option is not found, the first parameter of getopt that does not start with a '-' (and is not an option argument) is used as the short options string. Each short option character in shortopts may be followed by one colon to indicate it has a required argument, and by two colons to indicate it has an optional argument. The first character of shortopts may be '+' or '-' to influence the way options are parsed and output is generated (see the SCANNING MODES section for details).

-q, --quiet
Disable error reporting by getopt(3).

-Q, --quiet-output
Do not generate normal output. Errors are still reported by getopt(3), unless you also use -q.

-s, --shell shell
Set quoting conventions to those of shell. If the -s option is not given, the BASH conventions are used. Valid arguments are currently 'sh', 'bash', 'csh', and 'tcsh'.

-T, --test
Test if your getopt(1) is this enhanced version or an old version. This generates no output, and sets the error status to 4. Other implementations of getopt(1), and this version if the environment variable GETOPT_COMPATIBLE is set, will return '--' and error status 0.

-u, --unquoted
Do not quote the output. Note that whitespace and special (shell-dependent) characters can cause havoc in this mode (like they do with other getopt(1) implementations).

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

PARSING
This section specifies the format of the second part of the parameters of getopt (the parameters in the SYNOPSIS). The next section (OUTPUT) describes the output that is generated. These parameters were typically

the parameters a shell function was called with. Care must be taken that each parameter the shell function was called with corresponds to exactly one parameter in the parameter list of getopt (see the EXAMPLES). All parsing is done by the GNU getopt(3) routines.

The parameters are parsed from left to right. Each parameter is classified as a short option, a long option, an argument to an option, or a non-option parameter.

A simple short option is a '-' followed by a short option character. If the option has a required argument, it may be written directly after the option character or as the next parameter (i.e., separated by whitespace on the command line). If the option has an optional argument, it must be written directly after the option character if present.

It is possible to specify several short options after one '-', as long as all (except possibly the last) do not have required or optional arguments.

A long option normally begins with '--' followed by the long option name. If the option has a required argument, it may be written directly after the long option name, separated by '=', or as the next argument (i.e., separated by whitespace on the command line). If the option has an optional argument, it must be written directly after the long option name, separated by '=', if present (if you add the '=' but nothing behind it, it is interpreted as if no argument was present; this is a slight bug, see the BUGS). Long options may be abbreviated, as long as the abbreviation is not ambiguous.

Each parameter not starting with a '-', and not a required argument of a previous option, is a non-option parameter. Each parameter after a '--' parameter is always interpreted as a non-option parameter. If the environment variable POSIXLY_CORRECT is set, or if the short option string started with a '+', all remaining parameters are interpreted as non-option parameters as soon as the first non-option parameter is found.

OUTPUT

Output is generated for each element described in the previous section. Output is done in the same order as the elements are specified in the input, except for non-option parameters. Output can be done in compatible (unquoted) mode, or in such way that whitespace and other special characters within arguments and non-option parameters are preserved (see QUOTING). When the output is processed in the shell script, it will seem to be composed of distinct elements that can be processed one by one (by using the shift command in most shell languages). This is imperfect in unquoted mode, as elements can be split at unexpected places if they contain whitespace or special characters.

If there are problems parsing the parameters, for example because a required argument is not found or an option is not recognized, an error

will be reported on stderr, there will be no output for the offending element, and a non-zero error status is returned.

For a short option, a single '-' and the option character are generated as one parameter. If the option has an argument, the next parameter will be the argument. If the option takes an optional argument, but none was found, the next parameter will be generated but be empty in quoting mode, but no second parameter will be generated in unquoted (compatible) mode. Note that many other getopt(1) implementations do not support optional arguments.

If several short options were specified after a single '-', each will be present in the output as a separate parameter.

For a long option, '--' and the full option name are generated as one parameter. This is done regardless whether the option was abbreviated or specified with a single '-' in the input. Arguments are handled as with short options.

Normally, no non-option parameters output is generated until all options and their arguments have been generated. Then '--' is generated as a single parameter, and after it the non-option parameters in the order they were found, each as a separate parameter. Only if the first character of the short options string was a '-', non-option parameter output is generated at the place they are found in the input (this is not supported if the first format of the SYNOPSIS is used; in that case all preceding occurrences of '-' and '+' are ignored).

QUOTING

In compatibility mode, whitespace or 'special' characters in arguments or non-option parameters are not handled correctly. As the output is fed to the shell script, the script does not know how it is supposed to break the output into separate parameters. To circumvent this problem, this implementation offers quoting. The idea is that output is generated with quotes around each parameter. When this output is once again fed to the shell (usually by a shell eval command), it is split correctly into separate parameters.

Quoting is not enabled if the environment variable GETOPT_COMPATIBLE is set, if the first form of the SYNOPSIS is used, or if the option '-u' is found.

Different shells use different quoting conventions. You can use the '-s' option to select the shell you are using. The following shells are currently supported: 'sh', 'bash', 'csh' and 'tcsh'. Actually, only two 'flavors' are distinguished: sh-like quoting conventions and csh-like quoting conventions. Chances are that if you use another shell script language, one of these flavors can still be used.

SCANNING MODES

The first character of the short options string may be a '-' or a '+' to indicate a special scanning mode. If the first calling form in the SYNOPSIS is used they are ignored; the environment variable

```
       POSIXLY_CORRECT is still examined, though.

       If the first character is '+', or if the environment variable
       POSIXLY_CORRECT is set, parsing stops as soon as the first non-option
       parameter (i.e., a parameter that does not start with a '-') is found
       that is not an option argument. The remaining parameters are all
       interpreted as non-option parameters.

       If the first character is a '-', non-option parameters are outputted at
       the place where they are found; in normal operation, they are all
       collected at the end of output after a '--' parameter has been
       generated. Note that this '--' parameter is still generated, but it
       will always be the last parameter in this mode.

COMPATIBILITY
       This version of getopt(1) is written to be as compatible as possible to
       other versions. Usually you can just replace them with this version
       without any modifications, and with some advantages.

       If the first character of the first parameter of getopt is not a '-',
       getopt goes into compatibility mode. It will interpret its first
       parameter as the string of short options, and all other arguments will
       be parsed. It will still do parameter shuffling (i.e., all non-option
       parameters are output at the end), unless the environment variable
       POSIXLY_CORRECT is set, in which case, getopt will prepend a '+' before
       short options automatically.

       The environment variable GETOPT_COMPATIBLE forces getopt into
       compatibility mode. Setting both this environment variable and
       POSIXLY_CORRECT offers 100% compatibility for 'difficult' programs.
       Usually, though, neither is needed.

       In compatibility mode, leading '-' and '+' characters in the short
       options string are ignored.

RETURN CODES
       getopt returns error code 0 for successful parsing, 1 if getopt(3)
       returns errors, 2 if it does not understand its own parameters, 3 if an
       internal error occurs like out-of-memory, and 4 if it is called with
       -T.

EXAMPLES
       Example scripts for (ba)sh and (t)csh are provided with the getopt(1)
       distribution, and are installed in /usr/share/doc/util-linux/examples/
       directory.

ENVIRONMENT
       POSIXLY_CORRECT
           This environment variable is examined by the getopt(3) routines. If
           it is set, parsing stops as soon as a parameter is found that is
           not an option or an option argument. All remaining parameters are
           also interpreted as non-option parameters, regardless whether they
           start with a '-'.
```

```
       GETOPT_COMPATIBLE
           Forces getopt to use the first calling format as specified in the
           SYNOPSIS.

BUGS
       getopt(3) can parse long options with optional arguments that are given
       an empty optional argument (but cannot do this for short options). This
       getopt(1) treats optional arguments that are empty as if they were not
       present.

       The syntax if you do not want any short option variables at all is not
       very intuitive (you have to set them explicitly to the empty string).

AUTHOR
       Frodo Looijaard <frodo@frodo.looijaard.name>

SEE ALSO
       bash(1), tcsh(1), getopt(3)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The getopt command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                         GETOPT(1)
```

# Bash 的 getty 命令

```
getty [OPTION]... PORT [BAUD_RATE]... [TERM]
```

**功能**

在终端设备上启动登录提示并等待用户登录。

**类型**

可执行文件（`/usr/sbin/getty`），属于 util-linux。

> 通常是 /usr/sbin/agetty 的软链接。

**参数**

- `OPTIONS` 选项：
  - `-8`，`--8bits` - 允许 8 位字符输入（禁用奇偶校验的第 8 位剥离）
  - `-a`，`--autologin username` - 自动以指定用户登录（无需输入用户名）
  - `-c`，`--noreset` - 不重置终端控制模式
  - `-E`，`--remote` - 表示连接来自远程终端（通常用于远程线路或终端服务器）
  - `-f`，`--issue-file path` - 指定显示登录提示信息的文件（默认通常为 /etc/issue）
  - `--show-issue` - 仅显示 issue 文件内容后退出
  - `-h`，`--flow-control` - 启用硬件流控制（RTS/CTS）
  - `-H`，`--host fakehost` - 指定伪造的远程主机名并传递给 login
  - `-i`，`--noissue` - 不显示 issue 文件内容
  - `-I`，`--init-string initstring` - 向终端或调制解调器发送初始化字符串
  - `-J`，`--noclear` - 不清屏直接显示登录提示
  - `-l`，`--login-program login_program` - 指定要执行的登录程序（默认是 /bin/login）
  - `-L`，`--local-line[=mode]` - 设置为本地线路（忽略 modem 控制信号）
  - `-m`，`--extract-baud` - 从终端连接字符串中提取波特率
  - `--list-speeds` - 显示支持的串口波特率列表
  - `-n`，`--skip-login` - 不提示用户名，直接执行登录程序
  - `-N`，`--nonewline` - 不在登录提示后输出换行符
  - `-o`，`--login-options login_options` - 向 login 程序传递额外参数
  - `-p`，`--login-pause` - 在执行 login 前等待用户按键
  - `-r`，`--chroot directory` - 在执行 login 前切换到指定的 chroot 目录
  - `-R`，`--hangup` - 在退出时对终端执行挂断操作

    - `-s`, `--keep-baud` - 保持当前串口波特率设置
    - `-t`, `--timeout timeout` - 指定等待用户名输入的超时时间（秒）
    - `-U`, `--detect-case` - 自动检测用户名输入的大小写并转换
    - `-w`, `--wait-cr` - 等待回车字符后再显示登录提示
    - `--nohints` - 不显示登录提示中的帮助信息
    - `--nohostname` - 不在登录提示中显示主机名
    - `--long-hostname` - 在登录提示中显示完整主机名（FQDN）
    - `--erase-chars string` - 设置退格（erase）字符
    - `--kill-chars string` - 设置删除整行（kill）字符
    - `--chdir directory` - 在执行 login 前切换到指定目录
    - `--delay number` - 在启动后延迟指定秒数再显示登录提示
    - `--nice number` - 设置进程的 nice 值（调整优先级）
    - `--reload` - 重新加载配置（通常用于 systemd 环境）
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `PORT` - 端口；例如 `tty1`, `ttyS0`, `ttyUSB0` 等
- `BAUD_RATE` - 波特率；
- `TERM` - 终端类型；例如 `vt100`, `linux` 等

## 示例

```
$ getty -t 60 tty1 38400     # 在 tty1 上以 38400 波特率启动，60秒内未操作则自动断开
```

## 推荐阅读

- agetty(8) — Linux manual page

## 手册

```
AGETTY(8)                   System Administration                   AGETTY(8)

NAME
       agetty - alternative Linux getty

SYNOPSIS
       agetty [options] port [baud_rate...] [term]

DESCRIPTION
       agetty opens a tty port, prompts for a login name and invokes the
       /bin/login command. It is normally invoked by init(8).
```

agetty has several non-standard features that are useful for hardwired and for dial-in lines:

- Adapts the tty settings to parity bits and to erase, kill, end-of-line and uppercase characters when it reads a login name. The program can handle 7-bit characters with even, odd, none or space parity, and 8-bit characters with no parity. The following special characters are recognized: Control-U (kill); DEL and backspace (erase); carriage return and line feed (end of line). See also the --erase-chars and --kill-chars options.

- Optionally deduces the baud rate from the CONNECT messages produced by Hayes(tm)-compatible modems.

- Optionally does not hang up when it is given an already opened line (useful for call-back applications).

- Optionally does not display the contents of the /etc/issue file.

- Optionally displays an alternative issue files or directories instead of /etc/issue or /etc/issue.d.

- Optionally does not ask for a login name.

- Optionally invokes a non-standard login program instead of /bin/login.

- Optionally turns on hardware flow control.

- Optionally forces the line to be local with no need for carrier detect.

This program does not use the /etc/gettydefs (System V) or /etc/gettytab (SunOS 4) files.

ARGUMENTS

port
: A path name relative to the /dev directory. If a "-" is specified, agetty assumes that its standard input is already connected to a tty port and that a connection to a remote user has already been established.

  Under System V, a "-" port argument should be preceded by a "--".

baud_rate,...
: A comma-separated list of one or more baud rates. Each time agetty receives a BREAK character it advances through the list, which is treated as if it were circular.

  Baud rates should be specified in descending order, so that the null character (Ctrl-@) can also be used for baud-rate switching.

```
            This argument is optional and unnecessary for virtual terminals.

            The default for serial terminals is keep the current baud rate (see
            --keep-baud) and if unsuccessful then default to '9600'.

        term
            The value to be used for the TERM environment variable. This
            overrides whatever init(1) may have set, and is inherited by login
            and the shell.

            The default is 'vt100', or 'linux' for Linux on a virtual terminal,
            or 'hurd' for GNU Hurd on a virtual terminal.

OPTIONS
        -8, --8bits
            Assume that the tty is 8-bit clean, hence disable parity detection.

        -a, --autologin username
            Automatically log in the specified user without asking for a
            username or password. Using this option causes an -f username
            option and argument to be added to the /bin/login command line. See
            --login-options, which can be used to modify this option’s
            behavior.

            Note that --autologin may affect the way in which getty initializes
            the serial line, because on auto-login agetty does not read from
            the line and it has no opportunity optimize the line setting.

        -c, --noreset
            Do not reset terminal cflags (control modes). See termios(3) for
            more details.

        -E, --remote
            Typically the login(1) command is given a remote hostname when
            called by something such as telnetd(8). This option allows agetty
            to pass what it is using for a hostname to login(1) for use in
            utmp(5). See --host, login(1), and utmp(5).

            If the --host fakehost option is given, then an -h fakehost option
            and argument are added to the /bin/login command line.

            If the --nohostname option is given, then an -H option is added to
            the /bin/login command line.

            See --login-options.

        -f, --issue-file path
            Specifies a ":" delimited list of files and directories to be
            displayed instead of /etc/issue (or other). All specified files and
            directories are displayed, missing or empty files are silently
            ignored. If the specified path is a directory then display all
            files with .issue file extension in version-sort order from the
            directory. This allows custom messages to be displayed on different
```

terminals. The --noissue option will override this option.

--show-issue
Display the current issue file (or other) on the current terminal and exit. Use this option to review the current setting, it is not designed for any other purpose. Note that output may use some default or incomplete information as proper output depends on terminal and agetty command line.

-h, --flow-control
Enable hardware (RTS/CTS) flow control. It is left up to the application to disable software (XON/XOFF) flow protocol where appropriate.

-H, --host fakehost
Write the specified fakehost into the utmp file. Normally, no login host is given, since agetty is used for local hardwired connections and consoles. However, this option can be useful for identifying terminal concentrators and the like.

-i, --noissue
Do not display the contents of /etc/issue (or other) before writing the login prompt. Terminals or communications hardware may become confused when receiving lots of text at the wrong baud rate; dial-up scripts may fail if the login prompt is preceded by too much text.

-I, --init-string initstring
Set an initial string to be sent to the tty or modem before sending anything else. This may be used to initialize a modem.
Non-printable characters may be sent by writing their octal code preceded by a backslash (\). For example, to send a linefeed character (ASCII 10, octal 012), write \12.

-J, --noclear
Do not clear the screen before prompting for the login name. By default the screen is cleared.

-l, --login-program login_program
Invoke the specified login_program instead of /bin/login. This allows the use of a non-standard login program. Such a program could, for example, ask for a dial-up password or use a different password file. See --login-options.

-L, --local-line[=mode]
Control the CLOCAL line flag. The optional mode argument is auto, always or never. If the mode argument is omitted, then the default is always. If the --local-line option is not given at all, then the default is auto.

always
Forces the line to be a local line with no need for carrier detect. This can be useful when you have a locally attached

terminal where the serial line does not set the carrier-detect signal.

never
Explicitly clears the CLOCAL flag from the line setting and the carrier-detect signal is expected on the line.

auto
The agetty default. Does not modify the CLOCAL setting and follows the setting enabled by the kernel.

-m, --extract-baud
Try to extract the baud rate from the CONNECT status message produced by Hayes(tm)-compatible modems. These status messages are of the form: "<junk><speed><junk>". agetty assumes that the modem emits its status message at the same speed as specified with (the first) baud_rate value on the command line.

Since the --extract-baud feature may fail on heavily-loaded systems, you still should enable BREAK processing by enumerating all expected baud rates on the command line.

--list-speeds
Display supported baud rates. These are determined at compilation time.

-n, --skip-login
Do not prompt the user for a login name. This can be used in connection with the --login-program option to invoke a non-standard login process such as a BBS system. Note that with the --skip-login option, agetty gets no input from the user who logs in and therefore will not be able to figure out parity, character size, and newline processing of the connection. It defaults to space parity, 7 bit characters, and ASCII CR (13) end-of-line character. Beware that the program that agetty starts (usually /bin/login) is run as root.

-N, --nonewline
Do not print a newline before writing out /etc/issue.

-o, --login-options login_options
Options and arguments that are passed to login(1). Where \u is replaced by the login name. For example:

--login-options '-h darkstar -- \u'

See --autologin, --login-program and --remote.

Please read the SECURITY NOTICE below before using this option.

-p, --login-pause
Wait for any key before dropping to the login prompt. Can be combined with --autologin to save memory by lazily spawning shells.

```
-r, --chroot directory
    Change root to the specified directory.

-R, --hangup
    Call vhangup(2) to do a virtual hangup of the specified terminal.

-s, --keep-baud
    Try to keep the existing baud rate. The baud rates from the command
    line are used when agetty receives a BREAK character. If another
    baud rates specified then the original baud rate is also saved to
    the end of the wanted baud rates list. This can be used to return
    to the original baud rate after unexpected BREAKs.

-t, --timeout timeout
    Terminate if no user name could be read within timeout seconds. Use
    of this option with hardwired terminal lines is not recommended.

-U, --detect-case
    Turn on support for detecting an uppercase-only terminal. This
    setting will detect a login name containing only capitals as
    indicating an uppercase-only terminal and turn on some
    upper-to-lower case conversions. Note that this has no support for
    any Unicode characters.

-w, --wait-cr
    Wait for the user or the modem to send a carriage-return or a
    linefeed character before sending the /etc/issue file (or others)
    and the login prompt. This is useful with the --init-string option.

--nohints
    Do not print hints about Num, Caps and Scroll Locks.

--nohostname
    By default the hostname will be printed. With this option enabled,
    no hostname at all will be shown. This setting is also possible to
    able by LOGIN_PLAIN_PROMPT option in the /etc/login.defs
    configuration file (see below for more details).

--long-hostname
    By default the hostname is only printed until the first dot. With
    this option enabled, the fully qualified hostname by
    gethostname(3P) or (if not found) by getaddrinfo(3) is shown.

--erase-chars string
    This option specifies additional characters that should be
    interpreted as a backspace ("ignore the previous character") when
    the user types the login name. The default additional 'erase' has
    been '#', but since util-linux 2.23 no additional erase characters
    are enabled by default.

--kill-chars string
    This option specifies additional characters that should be
```

```
            interpreted as a kill ("ignore all previous characters") when the
            user types the login name. The default additional 'kill' has been
            '@', but since util-linux 2.23 no additional kill characters are
            enabled by default.

        --chdir directory
            Change directory before the login.

        --delay number
            Sleep seconds before open tty.

        --nice number
            Run login with this priority.

        --reload
            Ask all running agetty instances to reload and update their
            displayed prompts, if the user has not yet commenced logging in.
            After doing so the command will exit. This feature might be
            unsupported on systems without Linux inotify(7).

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

CONFIG FILE ITEMS
        agetty reads the /etc/login.defs configuration file (see
        login.defs(5)). Note that the configuration file could be distributed
        with another package (usually shadow-utils). The following
        configuration items are relevant for agetty:

        LOGIN_PLAIN_PROMPT (boolean)
            Tell agetty that printing the hostname should be suppressed in the
            login: prompt. This is an alternative to the --nohostname command
            line option. The default value is no.

EXAMPLE
        This section shows examples for the process field of an entry in the
        /etc/inittab file. You'll have to prepend appropriate values for the
        other fields. See inittab(5) for more details.

        For a hardwired line or a console tty:

           /sbin/agetty 9600 ttyS1

        For a directly connected terminal without proper carrier-detect wiring
        (try this if your terminal just sleeps instead of giving you a
        password: prompt):

           /sbin/agetty --local-line 9600 ttyS1 vt100

        For an old-style dial-in line with a 9600/2400/1200 baud modem:
```

```
   /sbin/agetty --extract-baud --timeout 60 ttyS1 9600,2400,1200
```

For a Hayes modem with a fixed 115200 bps interface to the machine (the example init string turns off modem echo and result codes, makes modem/computer DCD track modem/modem DCD, makes a DTR drop cause a disconnection, and turns on auto-answer after 1 ring):

```
   /sbin/agetty --wait-cr --init-string 'ATE0Q1&D2&C1S0=1\015'
   115200 ttyS1
```

## SECURITY NOTICE

If you use the --login-program and --login-options options, be aware that a malicious user may try to enter lognames with embedded options, which then get passed to the used login program. agetty does check for a leading "-" and makes sure the logname gets passed as one parameter (so embedded spaces will not create yet another parameter), but depending on how the login binary parses the command line that might not be sufficient. Check that the used login program cannot be abused this way.

Some programs use "--" to indicate that the rest of the command line should not be interpreted as options. Use this feature if available by passing "--" before the username gets passed by \u.

## ISSUE FILES

The default issue file is /etc/issue. If the file exists, then agetty also checks for /etc/issue.d directory. The directory is optional extension to the default issue file and content of the directory is printed after /etc/issue content. If the /etc/issue does not exist, then the directory is ignored. All files with .issue extension from the directory are printed in version-sort order. The directory can be used to maintain 3rd-party messages independently on the primary system /etc/issue file.

Since version 2.35 additional locations for issue file and directory are supported. If the default /etc/issue does not exist, then agetty checks for /run/issue and /run/issue.d, thereafter for /usr/lib/issue and /usr/lib/issue.d. The directory /etc is expected for host specific configuration, /run is expected for generated stuff and /usr/lib for static distribution maintained configuration.

The default path maybe overridden by --issue-file option. In this case specified path has to be file or directory and all the default issue file and directory locations are ignored.

The issue file feature can be completely disabled by --noissue option.

It is possible to review the current issue file by agetty --show-issue on the current terminal.

The issue files may contain certain escape codes to display the system name, date, time et cetera. All escape codes consist of a backslash (\)

immediately followed by one of the characters listed below.

4 or 4{interface}
    Insert the IPv4 address of the specified network interface (for example: \4{eth0}). If the interface argument is not specified, then select the first fully configured (UP, non-LOCALBACK, RUNNING) interface. If no configured interface is found, fall back to the IP address of the machine's hostname.

6 or 6{interface}
    The same as \4 but for IPv6.

b
    Insert the baudrate of the current line.

d
    Insert the current date.

e or e{name}
    Translate the human-readable name to an escape sequence and insert it (for example: \e{red}Alert text.\e{reset}). If the name argument is not specified, then insert \033. The currently supported names are: black, blink, blue, bold, brown, cyan, darkgray, gray, green, halfbright, lightblue, lightcyan, lightgray, lightgreen, lightmagenta, lightred, magenta, red, reset, reverse, yellow and white. All unknown names are silently ignored.

s
    Insert the system name (the name of the operating system). Same as 'uname -s'. See also the \S escape code.

S or S{VARIABLE}
    Insert the VARIABLE data from /etc/os-release. If this file does not exist then fall back to /usr/lib/os-release. If the VARIABLE argument is not specified, then use PRETTY_NAME from the file or the system name (see \s). This escape code can be used to keep /etc/issue distribution and release independent. Note that \S{ANSI_COLOR} is converted to the real terminal escape sequence.

l
    Insert the name of the current tty line.

m
    Insert the architecture identifier of the machine. Same as uname -m.

n
    Insert the nodename of the machine, also known as the hostname. Same as uname -n.

o
    Insert the NIS domainname of the machine. Same as hostname -d.

O
    Insert the DNS domainname of the machine.

r
    Insert the release number of the OS. Same as uname -r.

t
    Insert the current time.

u
    Insert the number of current users logged in.

U
    Insert the string "1 user" or "<n> users" where <n> is the number
    of current users logged in.

v
    Insert the version of the OS, that is, the build-date and such.

An example. On my system, the following /etc/issue file:

    This is \n.\o (\s \m \r) \t

displays as:

    This is thingol.orcan.dk (Linux i386 1.1.9) 18:29:30

FILES
/var/run/utmp
    the system status file.

/etc/issue
    printed before the login prompt.

/etc/os-release /usr/lib/os-release
    operating system identification data.

/dev/console
    problem reports (if syslog(3) is not used).

/etc/inittab
    init(8) configuration file for SysV-style init daemon.

BUGS
The baud-rate detection feature (the --extract-baud option) requires
that agetty be scheduled soon enough after completion of a dial-in call
(within 30 ms with modems that talk at 2400 baud). For robustness,
always use the --extract-baud option in combination with a multiple
baud rate command-line argument, so that BREAK processing is enabled.

The text in the /etc/issue file (or other) and the login prompt are
always output with 7-bit characters and space parity.

The baud-rate detection feature (the --extract-baud option) requires that the modem emits its status message after raising the DCD line.

## DIAGNOSTICS

Depending on how the program was configured, all diagnostics are written to the console device or reported via the syslog(3) facility. Error messages are produced if the port argument does not specify a terminal device; if there is no utmp entry for the current process (System V only); and so on.

## AUTHORS

Werner Fink <werner@suse.de>, Karel Zak <kzak@redhat.com>

The original agetty for serial terminals was written by W.Z. Venema <wietse@wzv.win.tue.nl> and ported to Linux by Peter Orbaek <poe@daimi.aau.dk>.

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The agetty command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 groups 命令

```
groups [OPTION]... [USERNAME]...
```

**功能**

查看用户的所属组。

**类型**

可执行文件（`/usr/bin/groups`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `USERNAME` - 用户名；默认为当前用户

## 示例

```
$ groups
primers adm cdrom sudo dip plugdev lxd
$ groups root
root : root
```

## 推荐阅读

- Print group names a user is in - GNU Coreutils
- groups(1) — Linux manual page

## 手册

```
GROUPS(1)                          User Commands                          GROUPS(1)

NAME
       groups - print the groups a user is in

SYNOPSIS
       groups [OPTION]... [USERNAME]...
```

DESCRIPTION

Print group memberships for each USERNAME or, if no USERNAME is specified, for the current process (which may differ if the groups database has changed).

--help display this help and exit

--version

output version information and exit

AUTHOR

Written by David MacKenzie and James Youngman.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>

Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.

This is free software: you are free to change and redistribute it.

There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

getent(1)

Full documentation <https://www.gnu.org/software/coreutils/groups>

or available locally via: info '(coreutils) groups invocation'

# Bash 的 hardlink 命令

```
hardlink [OPTIONS] [FILE...]
```

### 功能

通过硬链接合并多个相同的文件以节省磁盘空间。

### 类型

可执行文件（`/usr/bin/hardlink`），属于 util-linux。

### 参数

- `OPTIONS` 选项：
  - `-c`，`--content` - 仅根据文件内容判断是否相同，忽略文件属性；等价于 `-pot`
  - `-b`，`--io-size SIZE` - IO 缓冲区大小
  - `-d`，`--respect-dir` - 仅链接相同目录里的重复文件
  - `-f`，`--respect-name` - 仅链接同名的重复文件
  - `-i`，`--include REGEXP` - 仅扫描 `REGEXP` 匹配的文件
  - `-m`，`--maximize` - 合并时保留链接数最多的文件
  - `-M`，`--minimize` - 合并时保留链接数最少的文件
  - `-n`，`--dry-run` - 仅显示将要执行的操作，而不实际执行
  - `-o`，`--ignore-owner` - 允许链接所有者不同的重复文件
  - `-O`，`--keep-oldest` - 合并时保留最旧（mtime 最早）的文件
  - `-p`，`--ignore-mode` - 允许链接权限不同的重复文件
  - `-q`，`--quiet` - 不打印消息
  - `-r`，`--cache-size SIZE` - 校验和缓存大小
  - `-s`，`--minimum-size SIZE` - 不链接小于 `SIZE` 的文件
  - `-S`，`--maximum-size SIZE` - 不链接大于 `SIZE` 的文件
  - `-t`，`--ignore-time` - 允许链接时间不同的重复文件
  - `-v`，`--verbose` - 打印详细信息
  - `-x`，`--exclude REGEXP` - 排除 `REGEXP` 匹配的文件
  - `-X`，`--respect-xattrs` - 仅链接扩展属性也相同的重复文件
  - `-y`，`--method NAME` - 设置比较内容的方法，支持：`sha256`（默认），`sha1`，`crc32c` 和 `memcmp`
  - `--reflink[=WHEN]` - 创建写时复制克隆（引用链接），而非硬链接
  - `--skip-reflinks` - 忽略引用链接

- `--help` - 显示帮助
- `--version` - 显示版本

## 示例

```
$ echo Primers 编程伙伴 | tee 1.txt 2.txt       # 同时向两个文件写入相同内容
Primers 编程伙伴
$ ls -i                                         # 查看文件的 inode 号
9851498 1.txt  9851499 2.txt
$ hardlink .                                    # 合并相同文件
Mode:                   real
Method:                 sha256
Files:                  4
Linked:                 1 files
Compared:               0 xattrs
Compared:               1 files
Saved:                  21 B
Duration:               0.000573 seconds
$ ls -i                                         # 查看文件的 inode 号
9851498 1.txt  9851498 2.txt
```

## 推荐阅读

- hardlink(1) — Linux manual page

## 手册

```
HARDLINK(1)                     User Commands                     HARDLINK(1)

NAME
       hardlink - link multiple copies of a file

SYNOPSIS
       hardlink [options] [directory|file]...

DESCRIPTION
       hardlink is a tool that replaces copies of a file with either hardlinks
       or copy-on-write clones, thus saving space.

       hardlink first creates a binary tree of file sizes and then compares
       the content of files that have the same size. There are two basic
       content comparison methods. The memcmp method directly reads data
       blocks from files and compares them. The other method is based on
       checksums (like SHA256); in this case for each data block a checksum is
       calculated by the Linux kernel crypto API, and this checksum is stored
       in userspace and used for file comparisons.
```

For each file also an "intro" buffer (32 bytes) is cached. This buffer is used independently from the comparison method and requested cache-size and io-size. The "intro" buffer dramatically reduces operations with data content as files are very often different from the beginning.

OPTIONS

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

-c, --content
Consider only file content, not attributes, when determining whether two files are equal. Same as -pot.

-b, --io-size size
The size of the read(2) or sendfile(2) buffer used when comparing file contents. The size argument may be followed by the multiplicative suffixes KiB, MiB, etc. The "iB" is optional, e.g., "K" has the same meaning as "KiB". The default is 8KiB for memcmp method and 1MiB for the other methods. The only memcmp method uses process memory for the buffer, other methods use zero-copy way and I/O operation is done in the kernel. The size may be altered on the fly to fit a number of cached content checksums.

-d, --respect-dir
Only try to link files with the same directory name. The top-level directory (as specified on the hardlink command line) is ignored. For example, hardlink --respect-dir /foo /bar will link /foo/some/file with /bar/some/file, but not /bar/other/file. If combined with --respect-name, then entire paths (except the top-level directory) are compared.

-f, --respect-name
Only try to link files with the same (base)name. It’s strongly recommended to use long options rather than -f which is interpreted in a different way by other hardlink implementations.

-i, --include regex
A regular expression to include files. If the option --exclude has been given, this option re-includes files which would otherwise be excluded. If the option is used without --exclude, only files matched by the pattern are included.

-m, --maximize
Among equal files, keep the file with the highest link count.

-M, --minimize
Among equal files, keep the file with the lowest link count.

```
-n, --dry-run
    Do not act, just print what would happen.

-o, --ignore-owner
    Link and compare files even if their owner information (user and
    group) differs. Results may be unpredictable.

-O, --keep-oldest
    Among equal files, keep the oldest file (least recent modification
    time). By default, the newest file is kept. If --maximize or
    --minimize is specified, the link count has a higher precedence
    than the time of modification.

-p, --ignore-mode
    Link and compare files even if their mode is different. Results may
    be slightly unpredictable.

-q, --quiet
    Quiet mode, don’t print anything.

-r, --cache-size size
    The size of the cache for content checksums. All non-memcmp methods
    calculate checksum for each file content block (see --io-size),
    these checksums are cached for the next comparison. The size is
    important for large files or a large sets of files of the same
    size. The default is 10MiB.

-s, --minimum-size size
    The minimum size to consider. By default this is 1, so empty files
    will not be linked. The size argument may be followed by the
    multiplicative suffixes KiB (=1024), MiB (=1024*1024), and so on
    for GiB, TiB, PiB, EiB, ZiB and YiB (the "iB" is optional, e.g.,
    "K" has the same meaning as "KiB").

-S, --maximum-size size
    The maximum size to consider. By default this is 0 and 0 has the
    special meaning of unlimited. The size argument may be followed by
    the multiplicative suffixes KiB (=1024), MiB (=1024*1024), and so
    on for GiB, TiB, PiB, EiB, ZiB and YiB (the "iB" is optional, e.g.,
    "K" has the same meaning as "KiB").

-t, --ignore-time
    Link and compare files even if their time of modification is
    different. This is usually a good choice.

-v, --verbose
    Verbose output, explain to the user what is being done. If
    specified once, every hardlinked file is displayed. If specified
    twice, it also shows every comparison.

-x, --exclude regex
    A regular expression which excludes files from being compared and
    linked.
```

-X, --respect-xattrs
Only try to link files with the same extended attributes.

-y, --method name
Set the file content comparison method. The currently supported methods are sha256, sha1, crc32c and memcmp. The default is sha256, or memcmp if Linux Crypto API is not available. The methods based on checksums are implemented in zero-copy way, in this case file contents are not copied to the userspace and all calculation is done in kernel.

--reflink[=when]
Create copy-on-write clones (aka reflinks) rather than hardlinks. The reflinked files share only on-disk data, but the file mode and owner can be different. It’s recommended to use it with --ignore-owner and --ignore-mode options. This option implies --skip-reflinks to ignore already cloned files.

The optional argument when can be never, always, or auto. If the when argument is omitted, it defaults to auto, in this case, hardlink checks filesystem type and uses reflinks on BTRFS and XFS only, and fallback to hardlinks when creating reflink is impossible. The argument always disables filesystem type detection and fallback to hardlinks, in this case, only reflinks are allowed.

--skip-reflinks
Ignore already cloned files. This option may be used without --reflink when creating classic hardlinks.

ARGUMENTS
hardlink takes one or more directories which will be searched for files to be linked.

BUGS
The original hardlink implementation uses the option -f to force hardlinks creation between filesystem. This very rarely usable feature is no more supported by the current hardlink.

hardlink assumes that the trees it operates on do not change during operation. If a tree does change, the result is undefined and potentially dangerous. For example, if a regular file is replaced by a device, hardlink may start reading from the device. If a component of a path is replaced by a symbolic link or file permissions change, security may be compromised. Do not run hardlink on a changing tree or on a tree controlled by another user.

AUTHOR
There are multiple hardlink implementations. The very first implementation is from Jakub Jelinek for Fedora distribution, this implementation has been used in util-linux between versions v2.34 to v2.36. The current implementations is based on Debian version from Julian Andres Klode.

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The hardlink command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                      HARDLINK(1)
```

# Bash 的 head 命令

```
head [OPTION]... [FILE]...
```

**功能**

查看文件的开头。

**类型**

可执行文件（`/usr/bin/head`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--bytes=[-]NUM` - 查看文件开头的 `NUM` 字节，如果使用负号，则表示查看最后 `NUM` 字节之前的内容
  - `-n`，`--lines=[-]NUM` - 查看文件开头的 `NUM` 行，如果使用负号，则表示查看最后 `NUM` 行之前的内容
  - `-q`，`--quiet`，`--silent` - 不打印包含文件名在内的头部信息
  - `-v`，`--verbose` - 打印包含文件名在内的头部信息
  - `-z`，`--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ man head | head -n 5              # 显示前 5 行
HEAD(1)                         User Commands                         HEAD(1)

NAME
       head - output the first part of files

$ man head | head -n -10            # 显示除了后 10 行以外的内容
HEAD(1)                         User Commands                         HEAD(1)

NAME
       head - output the first part of files
...
COPYRIGHT
```

```
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| tail | 查看文件末尾 |
| more | 分页查看文件内容 |
| less | 增强版文件分页阅读命令 |

## 推荐阅读

- Output the first part of files - GNU Coreutils
- head(1) — Linux manual page

## 手册

```
HEAD(1)                          User Commands                          HEAD(1)

NAME
       head - output the first part of files

SYNOPSIS
       head [OPTION]... [FILE]...

DESCRIPTION
       Print the first 10 lines of each FILE to standard output. With more
       than one FILE, precede each with a header giving the file name.

       With no FILE, or when FILE is -, read standard input.

       Mandatory arguments to long options are mandatory for short options
       too.

       -c, --bytes=[-]NUM
              print the first NUM bytes of each file; with the leading '-',
              print all but the last NUM bytes of each file

       -n, --lines=[-]NUM
              print the first NUM lines instead of the first 10; with the
              leading '-', print all but the last NUM lines of each file
```

```
       -q, --quiet, --silent
              never print headers giving file names

       -v, --verbose
              always print headers giving file names

       -z, --zero-terminated
              line delimiter is NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

       NUM may have a multiplier suffix: b 512, kB 1000, K 1024, MB 1000*1000,
       M  1024*1024,  GB 1000*1000*1000, G 1024*1024*1024, and so on for T, P,
       E, Z, Y, R, Q.  Binary prefixes can be used, too: KiB=K, MiB=M, and  so
       on.

AUTHOR
       Written by David MacKenzie and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       tail(1)

       Full documentation <https://www.gnu.org/software/coreutils/head>
       or available locally via: info '(coreutils) head invocation'

GNU coreutils 9.4                April 2024                          HEAD(1)
```

# Bash 的 hostid 命令

```
hostid [OPTION]
```

**功能**

获取当前主机的数字 ID。

**类型**

可执行文件（`/usr/bin/hostid`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ hostid
007f0101
```

## 推荐阅读

- Print numeric host identifier - GNU Coreutils
- hostid(1) — Linux manual page

## 手册

```
HOSTID(1)                          User Commands                          HOSTID(1)

NAME
       hostid - print the numeric identifier for the current host

SYNOPSIS
       hostid [OPTION]

DESCRIPTION
       Print the numeric identifier (in hexadecimal) for the current host.
```

```
       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       gethostid(3)

       Full documentation <https://www.gnu.org/software/coreutils/hostid>
       or available locally via: info '(coreutils) hostid invocation'

GNU coreutils 9.4                April 2024                        HOSTID(1)
```

# Bash 的 i386 命令

```
i386 [OPTION]... [CMD [ARG]...]
```

**功能**

使用 32 位架构运行程序。

**类型**

可执行文件（`/usr/bin/i386`），属于 util-linux。

> 实际上是 `/usr/bin/setarch` 的符号链接，并默认将架构设为 `i386`。

**参数**

- `OPTION` 选项：
  - `--uname-2.6` - 使程序检测到的内核版本号为 `2.6`
  - `-v`，`--verbose` - 显示详细信息
  - `-3`，`--3gb` - 指定程序最多使用 3GB 的地址空间（3GB 用户空间，1GB 内核空间）
  - `--4gb` - ~~指定程序最多使用 4GB 的地址空间（4GB 用户空间，没有内核空间）~~；此选项已废弃不再生效
  - `-B`，`--32bit` - 将地址空间设为 32 位以模拟硬件；支持 ARM 和 Alpha 架构
  - `-F`，`--fdpic-funcptrs` - 将用户空间函数指向信号处理程序的指针视为指向地址描述符的指针；仅支持 ARM, Blackfin, Fujitsu FR-V 和 SuperH 架构
  - `-I`，`--short-inode` - ~~使用 16 位 inode 号~~；此选项已废弃不再生效
  - `-L`，`--addr-compat-layout` - 提供传统虚拟地址空间布局
  - `-R`，`--addr-no-randomize` - 禁用虚拟地址空间的随机化
  - `-S`，`--whole-seconds` - ~~时间值截断为整数秒~~；此选项已废弃不再生效
  - `-T`，`--sticky-timeouts` - 这使得 `select(2)`、`pselect(2)` 和 `ppoll(2)` 系统调用保留超时值
  - `-X`，`--read-implies-exec` - 如果设置此项，则 `mmap(2)` `PROT_READ` 还会添加 `PROT_EXEC`` 位
  - `-Z`，`--mmap-page-zero` - SVr4 错误模拟，会将 `mmap(2)` 第 0 页设置为只读。
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CMD` - 要运行的命令
- `ARG` - 传递给 `CMD` 的参数列表

## 示例

```
$ i386 arch
i686
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| setarch | 使用指定的架构或个性化标志运行程序 |

## 推荐阅读

- setarch(8) — Linux manual page

## 手册

```
SETARCH(8)                  System Administration                  SETARCH(8)

NAME
       setarch - change reported architecture in new program environment
       and/or set personality flags

SYNOPSIS
       setarch [arch] [options] [program [argument...]]

       setarch --list|-h|-V

       arch [options] [program [argument...]]

DESCRIPTION
       setarch modifies execution domains and process personality flags.

       The execution domains currently only affects the output of uname -m.
       For example, on an AMD64 system, running setarch i386 program will
       cause program to see i686 instead of x86_64 as the machine type. It can
       also be used to set various personality options. The default program is
       /bin/sh.

       Since version 2.33 the arch command line argument is optional and
       setarch may be used to change personality flags (ADDR_LIMIT_*,
       SHORT_INODE, etc) without modification of the execution domain.

OPTIONS
       --list
           List the architectures that setarch knows about. Whether setarch
```

```
    can actually set each of these architectures depends on the running
    kernel.

--show[=personality]
    Show the currently active personality and flags. If the personality
    argument is provided, it is shown instead of the current one.
    personality is a hexadecimal number with values was described in
    sys/personality.h.

--uname-2.6
    Causes the program to see a kernel version number beginning with
    2.6. Turns on UNAME26.

-v, --verbose
    Be verbose.

-3, --3gb
    Specifies program should use a maximum of 3GB of address space.
    Supported on x86. Turns on ADDR_LIMIT_3GB.

--4gb
    This option has no effect. It is retained for backward
    compatibility only, and may be removed in future releases.

-B, --32bit
    Limit the address space to 32 bits to emulate hardware. Supported
    on ARM and Alpha. Turns on ADDR_LIMIT_32BIT.

-F, --fdpic-funcptrs
    Treat user-space function pointers to signal handlers as pointers
    to address descriptors. This option has no effect on architectures
    that do not support FDPIC ELF binaries. In kernel v4.14 support is
    limited to ARM, Blackfin, Fujitsu FR-V, and SuperH CPU
    architectures.

-I, --short-inode
    Obsolete bug emulation flag. Turns on SHORT_INODE.

-L, --addr-compat-layout
    Provide legacy virtual address space layout. Use when the program
    binary does not have PT_GNU_STACK ELF header. Turns on
    ADDR_COMPAT_LAYOUT.

-R, --addr-no-randomize
    Disables randomization of the virtual address space. Turns on
    ADDR_NO_RANDOMIZE.

-S, --whole-seconds
    Obsolete bug emulation flag. Turns on WHOLE_SECONDS.

-T, --sticky-timeouts
    This makes select(2), pselect(2), and ppoll(2) system calls
    preserve the timeout value instead of modifying it to reflect the
```

amount of time not slept when interrupted by a signal handler. Use when program depends on this behavior. For more details see the timeout description in select(2) manual page. Turns on STICKY_TIMEOUTS.

-X, --read-implies-exec
If this is set then mmap(2) PROT_READ will also add the PROT_EXEC bit - as expected by legacy x86 binaries. Notice that the ELF loader will automatically set this bit when it encounters a legacy binary. Turns on READ_IMPLIES_EXEC.

-Z, --mmap-page-zero
SVr4 bug emulation that will set mmap(2) page zero as read-only. Use when program depends on this behavior, and the source code is not available to be fixed. Turns on MMAP_PAGE_ZERO.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

## EXAMPLE

```
setarch --addr-no-randomize mytestprog
setarch ppc32 rpmbuild --target=ppc --rebuild foo.src.rpm
setarch ppc32 -v -vL3 rpmbuild --target=ppc --rebuild bar.src.rpm
setarch ppc32 --32bit rpmbuild --target=ppc --rebuild foo.src.rpm
```

## AUTHORS

Elliot Lee <sopwith@redhat.com>, Jindrich Novy <jnovy@redhat.com>, Karel Zak <kzak@redhat.com>

## SEE ALSO

personality(2), select(2)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The setarch command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 id 命令

```
id [OPTION]... [USER]...
```

**功能**

获取用户的 ID。

**类型**

可执行文件（`/usr/bin/id`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-Z`，`--context` - 仅打印进程的安全上下文
  - `-g`，`--group` - 仅打印用户的有效组 ID
  - `-G`，`--groups` - 打印用户的所有组 ID
  - `-n`，`--name` - 打印名称而非 ID
  - `-r`，`--real` - 打印实际 ID 而非有效 ID
  - `-u`，`--user` - 仅打印有效用户 ID
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `USER` - 用户列表，可用是用户名或用户 ID

## 示例

```
$ id
uid=1000(primers) gid=1000(primers)
groups=1000(primers),4(adm),24(cdrom),27(sudo),30(dip),46(plugdev),101(lxd)
$ id -g         # 仅打印用户的有效组 ID
1000
$ id -G         # 打印用户的所有组 ID
1000 4 24 27 30 46 101
$ id -nG        # 打印用户的所有组名
primers adm cdrom sudo dip plugdev lxd
```

## 推荐阅读

- Print user identity - GNU Coreutils
- id(1) — Linux manual page

## 手册

```
ID(1)                           User Commands                           ID(1)

NAME
       id - print real and effective user and group IDs

SYNOPSIS
       id [OPTION]... [USER]...

DESCRIPTION
       Print user and group information for each specified USER, or (when USER
       omitted) for the current process.

       -a     ignore, for compatibility with other versions

       -Z, --context
              print only the security context of the process

       -g, --group
              print only the effective group ID

       -G, --groups
              print all group IDs

       -n, --name
              print a name instead of a number, for -ugG

       -r, --real
              print the real ID instead of the effective ID, with -ugG

       -u, --user
              print only the effective user ID

       -z, --zero
              delimit entries with NUL characters, not whitespace;

              not permitted in default format

       --help display this help and exit

       --version
              output version information and exit

       Without any OPTION, print some useful set of identified information.

AUTHOR
```

Written by Arnold Robbins and David MacKenzie.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/id>
or available locally via: info '(coreutils) id invocation'

# Bash 的 install 命令

```
install [OPTION]... SOURCE DEST
```

**功能**

安装文件。

**类型**

可执行文件（`/usr/bin/install`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--backup[=CONTROL]` - 制作备份
    - `none`，`off` - 始终不备份
    - `numbered`，`t` - 制作编号备份
    - `existing`，`nil` - 如果存在编号备份则制作编号备份，否则制作简单备份
    - `simple`，`never` - 始终制作简单备份
  - `-b` - 制作备份，由环境变量 `VERSION_CONTROL` 决定备份方式
  - `-c` - ~~复制而不是移动文件；现在始终是复制，忽略此选项~~
  - `-C`，`--compare` - 比较源文件和目标文件，如果内容和元信息完全一致则不修改目标文件
  - `-d`，`--directory` - 将所有参数视为目录（源文件复制到目录内部）
  - `-D` - 自动创建父目录
  - `--debug` - 打印详细步骤
  - `-g`，`--group=GROUP` - 设置目标文件的所属组
  - `-m`，`--mode=MODE` - 设置目标文件的访问权限
  - `-o`，`--owner=OWNER` - 设置目标文件的所有者
  - `-p`，`--preserve-timestamps` - 目标文件保留源文件的时间戳
  - `-s`，`--strip` - 剥离符号表
  - `--strip-program=PROGRAM` - 使用 `PROGRAM` 对二进制文件进行剥离
  - `-S`，`--suffix=SUFFIX` - 设置备份文件的后缀
  - `-t`，`--target-directory=DIRECTORY` - 指定目标目录
  - `-T`，`--no-target-directory` - 将 `DEST` 视为文件（如果 `DEST` 是目录，则进行合并而非复制到它内部）
  - `-v`，`--verbose` - 打印详细信息
  - `--preserve-context`

    - `-Z`
    - `--context[=CTX]`
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `USESOURCER` - 源文件
- `DEST` - 目标

## 示例

```
$ sudo install -u root -g root -m755 -s my_cmd /usr/bin
```

- 将可执行文件 `my_cmd` 剥离符号表后安装到 `/usr/bin`，所有者和所属组设为 `root`，权限设为 `755`

## 推荐阅读

- Copy files and set attributes - GNU Coreutils
- install(1) — Linux manual page

## 手册

```
INSTALL(1)                       User Commands                       INSTALL(1)

NAME
       install - copy files and set attributes

SYNOPSIS
       install [OPTION]... [-T] SOURCE DEST
       install [OPTION]... SOURCE... DIRECTORY
       install [OPTION]... -t DIRECTORY SOURCE...
       install [OPTION]... -d DIRECTORY...

DESCRIPTION
       This  install  program copies files (often just compiled) into destina-
       tion locations you choose.  If you  want  to  download  and  install  a
       ready-to-use package on a GNU/Linux system, you should instead be using
       a package manager like yum(1) or apt-get(1).

       In  the first three forms, copy SOURCE to DEST or multiple SOURCE(s) to
       the existing DIRECTORY, while setting permission modes and owner/group.
       In the 4th form, create all components of the given DIRECTORY(ies).

       Mandatory arguments to long options are  mandatory  for  short  options
       too.
```

```
--backup[=CONTROL]
       make a backup of each existing destination file

-b     like --backup but does not accept an argument

-c     (ignored)

-C, --compare
       compare  content  of  source  and  destination  files, and if no
       change to content, ownership, and permissions, do not modify the
       destination at all

-d, --directory
       treat all arguments as directory names; create all components of
       the specified directories

-D     create all leading components of DEST except the  last,  or  all
       components of --target-directory, then copy SOURCE to DEST

--debug
       explain how a file is copied.  Implies -v

-g, --group=GROUP
       set group ownership, instead of process' current group

-m, --mode=MODE
       set permission mode (as in chmod), instead of rwxr-xr-x

-o, --owner=OWNER
       set ownership (super-user only)

-p, --preserve-timestamps
       apply access/modification times of SOURCE files to corresponding
       destination files

-s, --strip
       strip symbol tables

--strip-program=PROGRAM
       program used to strip binaries

-S, --suffix=SUFFIX
       override the usual backup suffix

-t, --target-directory=DIRECTORY
       copy all SOURCE arguments into DIRECTORY

-T, --no-target-directory
       treat DEST as a normal file

-v, --verbose
       print the name of each created file or directory
```

```
       --preserve-context
              preserve SELinux security context

       -Z     set  SELinux  security context of destination file and each cre-
              ated directory to default type

       --context[=CTX]
              like -Z, or if CTX is specified then set the  SELinux  or  SMACK
              security context to CTX

       --help display this help and exit

       --version
              output version information and exit

       The   backup   suffix   is   '~',   unless   set   with   --suffix   or   SIM-
       PLE_BACKUP_SUFFIX.  The version control method may be selected via  the
       --backup  option  or  through the VERSION_CONTROL environment variable.
       Here are the values:

       none, off
              never make backups (even if --backup is given)

       numbered, t
              make numbered backups

       existing, nil
              numbered if numbered backups exist, simple otherwise

       simple, never
              always make simple backups

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       cp(1)

       Full documentation <https://www.gnu.org/software/coreutils/install>
       or available locally via: info '(coreutils) install invocation'
```

# Bash 的 ionice 命令

```
ionice [OPTIONS] [COMMAND [ARG]...]
```

**功能**

运行命令时附带 IO 调度优先级。

**类型**

可执行文件（`/usr/bin/ionice`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-c`，`--class CLASS` - 指定调度器
    - `0` - 无调度
    - `1` - 实时调度
    - `2` - 尽力调度
    - `3` - 空闲调度
  - `-n`，`--classdata LEVEL` 对应实时调度和尽力调度的优先级；取值 0（最高） 到 7（最低）
  - `-p`，`--pid PID...` - 指定目标进程的进程 ID
  - `-P`，`--pgid PGID...` - - 指定目标进程的进程组 ID
  - `-t`，`--ignore` - 忽略设置优先级的失败
  - `-u`，`--uid UID...` - 指定目标进程的用户 ID
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `COMMAND` - 要执行的命令
- `ARG` - 参数列表

## 示例

```
$ ionice                        # 查看直接运行命令的 IO 优先级（通常是 0）
none: prio 0
$ ionice -c 1 -n 0 my_cmd       # 以实时调度 0 优先级运行 my_cmd
```

## 推荐阅读

- ionice(1) — Linux manual page

## 手册

```
IONICE(1)                       User Commands                       IONICE(1)

NAME
       ionice - set or get process I/O scheduling class and priority

SYNOPSIS
       ionice [-c class] [-n level] [-t] -p PID

       ionice [-c class] [-n level] [-t] -P PGID

       ionice [-c class] [-n level] [-t] -u UID

       ionice [-c class] [-n level] [-t] command [argument] ...

DESCRIPTION
       This program sets or gets the I/O scheduling class and priority for a
       program. If no arguments or just -p is given, ionice will query the
       current I/O scheduling class and priority for that process.

       When command is given, ionice will run this command with the given
       arguments. If no class is specified, then command will be executed with
       the "best-effort" scheduling class. The default priority level is 4.

       As of this writing, a process can be in one of three scheduling
       classes:

       Idle
           A program running with idle I/O priority will only get disk time
           when no other program has asked for disk I/O for a defined grace
           period. The impact of an idle I/O process on normal system activity
           should be zero. This scheduling class does not take a priority
           argument. Presently, this scheduling class is permitted for an
           ordinary user (since kernel 2.6.25).

       Best-effort
           This is the effective scheduling class for any process that has not
           asked for a specific I/O priority. This class takes a priority
           argument from 0-7, with a lower number being higher priority.
           Programs running at the same best-effort priority are served in a
           round-robin fashion.

           Note that before kernel 2.6.26 a process that has not asked for an
           I/O priority formally uses "none" as scheduling class, but the I/O
           scheduler will treat such processes as if it were in the
           best-effort class. The priority within the best-effort class will
           be dynamically derived from the CPU nice level of the process:
```

```
            io_priority = (cpu_nice + 20) / 5.

            For kernels after 2.6.26 with the CFQ I/O scheduler, a process that
            has not asked for an I/O priority inherits its CPU scheduling
            class. The I/O priority is derived from the CPU nice level of the
            process (same as before kernel 2.6.26).

        Realtime
            The RT scheduling class is given first access to the disk,
            regardless of what else is going on in the system. Thus the RT
            class needs to be used with some care, as it can starve other
            processes. As with the best-effort class, 8 priority levels are
            defined denoting how big a time slice a given process will receive
            on each scheduling window. This scheduling class is not permitted
            for an ordinary (i.e., non-root) user.

OPTIONS
        -c, --class class
            Specify the name or number of the scheduling class to use; 0 for
            none, 1 for realtime, 2 for best-effort, 3 for idle.

        -n, --classdata level
            Specify the scheduling class data. This only has an effect if the
            class accepts an argument. For realtime and best-effort, 0-7 are
            valid data (priority levels), and 0 represents the highest priority
            level.

        -p, --pid PID...
            Specify the process IDs of running processes for which to get or
            set the scheduling parameters.

        -P, --pgid PGID...
            Specify the process group IDs of running processes for which to get
            or set the scheduling parameters.

        -t, --ignore
            Ignore failure to set the requested priority. If command was
            specified, run it even in case it was not possible to set the
            desired scheduling priority, which can happen due to insufficient
            privileges or an old kernel version.

        -u, --uid UID...
            Specify the user IDs of running processes for which to get or set
            the scheduling parameters.

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

NOTES
        Linux supports I/O scheduling priorities and classes since 2.6.13 with
```

the CFQ I/O scheduler.

EXAMPLES

- # ionice -c 3 -p 89

  Sets process with PID 89 as an idle I/O process.

- # ionice -c 2 -n 0 bash

  Runs 'bash' as a best-effort program with highest priority.

- # ionice -p 89 91

  Prints the class and priority of the processes with PID 89 and 91.

AUTHORS

Jens Axboe <jens@axboe.dk>, Karel Zak <kzak@redhat.com>

SEE ALSO

ioprio_set(2)

REPORTING BUGS

For bug reports, use the issue tracker at
https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The ionice command is part of the util-linux package which can be
downloaded from Linux Kernel Archive
<https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 ipcmk 命令

```
ipcmk [OPTIONS]
```

### 功能

创建进程间通信（IPC）资源。

### 类型

可执行文件（`/usr/bin/ipcmk`），属于 util-linux。

### 参数

- `OPTIONS` 选项：
  - `-M`，`--shmem SIZE` - 创建 `SIZE` 大小的共享内存；支持 KiB（1024）、KB（1000） 等单位后缀
  - `-Q`，`--queue` - 创建消息队列
  - `-S`，`--semaphore NUMBER` - 创建包含 `NUMBER` 个信号量的数组
  - `-p`，`--mode MODE` - 资源的访问权限；默认为 0644
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ ipcmk -M 256              # 创建 256 字节的共享内存
Shared memory id: 0
$ ipcmk -Q                  # 创建消息队列
Message queue id: 0
$ ipcmk -S 8                # 创建 8 个信号量的数组
Semaphore id: 0
$ ipcs                      # 查看所有 IPC 资源

------ Message Queues --------
key        msqid      owner      perms      used-bytes   messages
0xb8bca3e8 0          primers    644        0            0

------ Shared Memory Segments --------
key        shmid      owner      perms      bytes      nattch     status
0xcbff4ae1 0          primers    644        256        0

------ Semaphore Arrays --------
```

```
key        semid      owner      perms      nsems
0x768e487c 0          primers    644        8
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| ipcrm | 删除 IPC 资源 |
| ipcs | 查看 IPC 资源 |

## 推荐阅读

- ipcmk(1) — Linux manual page

## 手册

```
IPCMK(1)                        User Commands                        IPCMK(1)

NAME
       ipcmk - make various IPC resources

SYNOPSIS
       ipcmk [options]

DESCRIPTION
       ipcmk allows you to create System V inter-process communication (IPC)
       objects: shared memory segments, message queues, and semaphore arrays.

OPTIONS
       Resources can be specified with these options:

       -M, --shmem size
           Create a shared memory segment of size bytes. The size argument may
           be followed by the multiplicative suffixes KiB (=1024), MiB
           (=1024*1024), and so on for GiB, etc. (the "iB" is optional, e.g.,
           "K" has the same meaning as "KiB") or the suffixes KB (=1000), MB
           (=1000*1000), and so on for GB, etc.

       -Q, --queue
           Create a message queue.

       -S, --semaphore number
           Create a semaphore array with number of elements.

       Other options are:
```

```
       -p, --mode mode
           Access permissions for the resource. Default is 0644.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

AUTHORS
       Hayden A. James <hayden.james@gmail.com>

SEE ALSO
       ipcrm(1), ipcs(1), sysvipc(7)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The ipcmk command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          IPCMK(1)
```

# Bash 的 ipcrm 命令

```
ipcrm [OPTIONS]
```

**功能**

删除进程间通信（IPC）资源。

**类型**

可执行文件（`/usr/bin/ipcrm`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-a`，`--all [TYPE]` - 删除所有 `TYPE` 类型的 IPC 资源；`TYPE` 取值为 `shm`， `msg`，`sem`，省略时删除所有类型
  - `-M`，`--shmem-key KEY` - 根据 `KEY` 删除共享内存
  - `-m`，`--shmem-id ID` - 根据 `ID` 删除共享内存
  - `-Q`，`--queue-key KEY` - 根据 `KEY` 删除消息队列
  - `-q`，`--queue-id ID` - 根据 `ID` 删除消息队列
  - `-S`，`--semaphore-key KEY` - 根据 `KEY` 删除信号量数组
  - `-s`，`--semaphore-id ID` - 根据 `ID` 删除信号量数组
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ ipcmk -S 8            # 创建信号量数组
Semaphore id: 2
$ ipcrm -s 2            # 删除信号量数组
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| ipcmk | 创建 IPC 资源 |
| ipcs | 查看 IPC 资源 |

# 推荐阅读

- ipcrm(1) — Linux manual page

# 手册

```
IPCRM(1)                        User Commands                        IPCRM(1)

NAME
       ipcrm - remove certain IPC resources

SYNOPSIS
       ipcrm [options]

       ipcrm [shm|msg|sem] ID ...

DESCRIPTION
       ipcrm removes System V inter-process communication (IPC) objects and
       associated data structures from the system. In order to delete such
       objects, you must be superuser, or the creator or owner of the object.

       System V IPC objects are of three types: shared memory, message queues,
       and semaphores. Deletion of a message queue or semaphore object is
       immediate (regardless of whether any process still holds an IPC
       identifier for the object). A shared memory object is only removed
       after all currently attached processes have detached (shmdt(2)) the
       object from their virtual address space.

       Two syntax styles are supported. The old Linux historical syntax
       specifies a three-letter keyword indicating which class of object is to
       be deleted, followed by one or more IPC identifiers for objects of this
       type.

       The SUS-compliant syntax allows the specification of zero or more
       objects of all three types in a single command line, with objects
       specified either by key or by identifier (see below). Both keys and
       identifiers may be specified in decimal, hexadecimal (specified with an
       initial '0x' or '0X'), or octal (specified with an initial '0').

       The details of the removes are described in shmctl(2), msgctl(2), and
       semctl(2). The identifiers and keys can be found by using ipcs(1).

OPTIONS
       -a, --all [shm] [msg] [sem]
           Remove all resources. When an option argument is provided, the
           removal is performed only for the specified resource types.

           Warning! Do not use -a if you are unsure how the software using the
           resources might react to missing objects. Some programs create
           these resources at startup and may not have any code to deal with
```

an unexpected disappearance.

-M, --shmem-key shmkey
Remove the shared memory segment created with shmkey after the last detach is performed.

-m, --shmem-id shmid
Remove the shared memory segment identified by shmid after the last detach is performed.

-Q, --queue-key msgkey
Remove the message queue created with msgkey.

-q, --queue-id msgid
Remove the message queue identified by msgid.

-S, --semaphore-key semkey
Remove the semaphore created with semkey.

-s, --semaphore-id semid
Remove the semaphore identified by semid.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

## NOTES

In its first Linux implementation, ipcrm used the deprecated syntax shown in the second line of the SYNOPSIS. Functionality present in other *nix implementations of ipcrm has since been added, namely the ability to delete resources by key (not just identifier), and to respect the same command-line syntax. For backward compatibility the previous syntax is still supported.

## SEE ALSO

ipcmk(1), ipcs(1), msgctl(2), msgget(2), semctl(2), semget(2), shmctl(2), shmdt(2), shmget(2), ftok(3), sysvipc(7)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The ipcrm command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 ipcs 命令

```
ipcs [OPTIONS]
```

**功能**

查看进程间通信（IPC）资源。

**类型**

可执行文件（`/usr/bin/ipcs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-i`，`--id ID` - 要查看的资源 ID
  - `-m`，`--shmems` - 查看共享内存
  - `-q`，`--queues` - 查看消息队列
  - `-s`，`--semaphores` - 查看信号量数组
  - `-a`，`--all` - 查看所有资源（默认）
  - `-c`，`--creator` - 显示资源的创建者
  - `-l`，`--limits` - 显示资源的限制
  - `-p`，`--pid` - 显示资源的 PID
  - `-t`，`--time` - 查看资源的时间信息
  - `-u`，`--summary` - 查看资源的状态摘要
  - `-b`，`--bytes` - 以字节为单位打印大小
  - `--human` - 以人类可读的格式打印大小
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ ipcs                    # 查看所有 IPC 资源

------ Message Queues --------
key        msqid      owner      perms      used-bytes   messages
0xb8bca3e8 0          primers    644        0            0

------ Shared Memory Segments --------
key        shmid      owner      perms      bytes      nattch     status
```

```
0xcbff4ae1 0          primers    644        256        0

------ Semaphore Arrays --------
key        semid      owner      perms      nsems
0x768e487c 0          primers    644        8

$ ipcs -l                  # 查看资源的限制

------ Messages Limits --------
max queues system wide = 32000
max size of message (bytes) = 8192
default max size of queue (bytes) = 16384

------ Shared Memory Limits --------
max number of segments = 4096
max seg size (kbytes) = 18014398509465599
max total shared memory (kbytes) = 18446744073709551612
min seg size (bytes) = 1

------ Semaphore Limits --------
max number of arrays = 32000
max semaphores per array = 32000
max semaphores system wide = 1024000000
max ops per semop call = 500
semaphore max value = 32767

$ ipcs -p                  # 查看资源的 PID

------ Message Queues PIDs --------
msqid      owner      lspid      lrpid
0          primers    0          0

------ Shared Memory Creator/Last-op PIDs --------
shmid      owner      cpid       lpid
0          planc      883345     0
```

- `lspid` 是最后一次向消息队列发送消息的进程的 PID
- `lrpid` 是最后一次从消息队列接收消息的进程的 PID
- `cpid` 是创建共享内存的进程的 PID
- `lpid` 是最后一次操作共享内存的进程的 PID

## 相关命令

| 命令 | 说明 |
| --- | --- |
| ipcmk | 创建 IPC 资源 |

| 命令 | 说明 |
| --- | --- |
| ipcrm | 删除 IPC 资源 |

## 推荐阅读

- ipcs(1) — Linux manual page

## 手册

```
IPCS(1)                         User Commands                         IPCS(1)

NAME
       ipcs - show information on IPC facilities

SYNOPSIS
       ipcs [options]

DESCRIPTION
       ipcs shows information on System V inter-process communication
       facilities. By default it shows information about all three resources:
       shared memory segments, message queues, and semaphore arrays.

OPTIONS
       -i, --id id
           Show full details on just the one resource element identified by
           id. This option needs to be combined with one of the three resource
           options: -m, -q or -s.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

   Resource options
       -m, --shmems
           Write information about active shared memory segments.

       -q, --queues
           Write information about active message queues.

       -s, --semaphores
           Write information about active semaphore sets.

       -a, --all
           Write information about all three resources (default).

   Output formats
```

Of these options only one takes effect: the last one specified.

-c, --creator
Show creator and owner.

-l, --limits
Show resource limits.

-p, --pid
Show PIDs of creator and last operator.

-t, --time
Write time information. The time of the last control operation that changed the access permissions for all facilities, the time of the last msgsnd(2) and msgrcv(2) operations on message queues, the time of the last shmat(2) and shmdt(2) operations on shared memory, and the time of the last semop(2) operation on semaphores.

-u, --summary
Show status summary.

Representation
These affect only the -l (--limits) option.

-b, --bytes
Print the sizes in bytes rather than in a human-readable format.

By default, the unit, sizes are expressed in, is byte, and unit prefixes are in power of 2^10 (1024). Abbreviations of symbols are exhibited truncated in order to reach a better readability, by exhibiting alone the first letter of them; examples: "1 KiB" and "1 MiB" are respectively exhibited as "1 K" and "1 M", then omitting on purpose the mention "iB", which is part of these abbreviations.

--human
Print sizes in human-readable format.

CONFORMING TO
The Linux ipcs utility is not fully compatible to the POSIX ipcs utility. The Linux version does not support the POSIX -a, -b and -o options, but does support the -l and -u options not defined by POSIX. A portable application shall not use the -a, -b, -o, -l, and -u options.

NOTES
The current implementation of ipcs obtains information about available IPC resources by parsing the files in /proc/sysvipc. Before util-linux version v2.23, an alternate mechanism was used: the IPC_STAT command of msgctl(2), semctl(2), and shmctl(2). This mechanism is also used in later util-linux versions in the case where /proc is unavailable. A limitation of the IPC_STAT mechanism is that it can only be used to retrieve information about IPC resources for which the user has read permission.

```
AUTHORS
       Krishna Balasubramanian <balasub@cis.ohio-state.edu>

SEE ALSO
       ipcmk(1), ipcrm(1), msgrcv(2), msgsnd(2), semget(2), semop(2),
       shmat(2), shmdt(2), shmget(2), sysvipc(7)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The ipcs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                          IPCS(1)
```

# Bash 的 isosize 命令

```
isosize [OPTION]...
```

**功能**

输出 ISO9660 文件系统的长度。

**类型**

可执行文件（`/usr/sbin/isosize`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-x`，`--sectors` - 以易于理解的形式显示块数和块大小；以“扇区”为单位
  - `-d`，`--divisor number` - 以 `number` 字节为单位；
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `MOUNTPOINT` - 要操作的挂载点

## 示例

```
$ isosize ./test.iso
358400
$ isosize -d 100 ./test.iso
3584
$ isosize -x ./test.iso
sector count: 175, sector size: 2048
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| genisoimage | 创建 ISO 镜像 |

## 推荐阅读

- isosize(8) — Linux manual page

# 手册

```
ISOSIZE(8)                  System Administration                  ISOSIZE(8)

NAME
       isosize - output the length of an iso9660 filesystem

SYNOPSIS
       isosize [options] iso9660_image_file

DESCRIPTION
       This command outputs the length of an iso9660 filesystem that is
       contained in the specified file. This file may be a normal file or a
       block device (e.g. /dev/hdd or /dev/sr0). In the absence of any options
       (and errors), it will output the size of the iso9660 filesystem in
       bytes. This can now be a large number (>> 4 GB).

OPTIONS
       -x, --sectors
           Show the block count and block size in human-readable form. The
           output uses the term "sectors" for "blocks".

       -d, --divisor number
           Only has an effect when -x is not given. The value shown (if no
           errors) is the iso9660 file size in bytes divided by number. So if
           number is the block size then the shown value will be the block
           count.

           The size of the file (or block device) holding an iso9660
           filesystem can be marginally larger than the actual size of the
           iso9660 filesystem. One reason for this is that cd writers are
           allowed to add "run out" sectors at the end of an iso9660 image.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       0
           success

       1
           generic failure, such as invalid usage

       32
           all failed

       64
           some failed
```

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The isosize command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                        ISOSIZE(8)
```

# Bash 的 df 命令

```
join [OPTION]... FILE1 FILE2
```

**功能**

将两个文件中的行按照公共字段连接。

**类型**

可执行文件（`/usr/bin/join`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-1 FIELD` - `FILE1` 根据 `FIELD` 字段进行匹配
  - `-2 FIELD` - `FILE2` 根据 `FIELD` 字段进行匹配
  - `-a FILENUM` - `FILENUM` 取 1 或 2，表示 `FILE1` 或 `FILE2`，额外打印对应文件中不可配对的行
  - `-e STRING` - 使用 `STRING` 填补缺失的空行
  - `-i`, `--ignore-case` - 比较字段时忽略大小写
  - `-j FIELD` - `FILE1` 和 `FILE2` 都根据 `FIELD` 字段进行匹配
  - `-o FORMAT` - 输出格式
  - `-t CHAR` - 使用 `CHAR` 作为输入和输出的分隔符
  - `-v FILENUM` - `FILENUM` 取 1 或 2，表示 `FILE1` 或 `FILE2`，仅打印对应文件中不可配对的行
  - `--check-order` - 检查输入是否正确排序
  - `--nocheck-order` - 不检查输入是否正确排序
  - `--header` - 将每个文件的第一行视作字段标题，直接打印，不进行匹配
  - `-z`, `--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE1` - 要连接的文件1
- `FILE2` - 要连接的文件2

## 示例

```
$ cat 1.txt
1 Alice
2 Bob
3 Carol
$ cat 2.txt
1 90
2 85
4 70
$ join 1.txt 2.txt          # 默认按照第一列合并
1 Alice 90
2 Bob 85
$ join -a 1 1.txt 2.txt     # 保留文件 1 中未匹配的行
1 Alice 90
2 Bob 85
3 Carol
$ join -a 1 1.txt 2.txt     # 仅打印文件 1 中未匹配的行
3 Carol
```

## 推荐阅读

- Join lines on a common field - GNU Coreutils
- join(1) — Linux manual page

## 手册

```
JOIN(1)                          User Commands                          JOIN(1)

NAME
       join - join lines of two files on a common field

SYNOPSIS
       join [OPTION]... FILE1 FILE2

DESCRIPTION
       For  each  pair of input lines with identical join fields, write a line
       to standard output.  The default join field is the first, delimited  by
       blanks.

       When FILE1 or FILE2 (not both) is -, read standard input.

       -a FILENUM
              also  print unpairable lines from file FILENUM, where FILENUM is
              1 or 2, corresponding to FILE1 or FILE2

       -e STRING
              replace missing (empty) input fields with STRING; I.e.,  missing
              fields specified with '-12jo' options
```

```
-i, --ignore-case
       ignore differences in case when comparing fields

-j FIELD
       equivalent to '-1 FIELD -2 FIELD'

-o FORMAT
       obey FORMAT while constructing output line

-t CHAR
       use CHAR as input and output field separator

-v FILENUM
       like -a FILENUM, but suppress joined output lines

-1 FIELD
       join on this FIELD of file 1

-2 FIELD
       join on this FIELD of file 2

--check-order
       check  that  the  input  is  correctly sorted, even if all input
       lines are pairable

--nocheck-order
       do not check that the input is correctly sorted

--header
       treat the first line in each file as field headers,  print  them
       without trying to pair them

-z, --zero-terminated
       line delimiter is NUL, not newline

--help display this help and exit

--version
       output version information and exit

Unless  -t  CHAR  is  given, leading blanks separate fields and are ig-
nored, else fields are separated by CHAR.  Any FIELD is a field  number
counted  from 1.  FORMAT is one or more comma or blank separated speci-
fications, each being 'FILENUM.FIELD' or '0'.  Default  FORMAT  outputs
the  join  field, the remaining fields from FILE1, the remaining fields
from FILE2, all separated by CHAR.  If FORMAT is  the  keyword  'auto',
then the first line of each file determines the number of fields output
for each line.

Important:  FILE1  and  FILE2 must be sorted on the join fields.  E.g.,
use "sort -k 1b,1" if 'join' has no options, or use  "join  -t  ''"  if
'sort'  has no options.  Note, comparisons honor the rules specified by
'LC_COLLATE'.  If the input is not sorted  and  some  lines  cannot  be
```

```
       joined, a warning message will be given.

AUTHOR
       Written by Mike Haertel.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       comm(1), uniq(1)

       Full documentation <https://www.gnu.org/software/coreutils/join>
       or available locally via: info '(coreutils) join invocation'

GNU coreutils 9.4                April 2024                          JOIN(1)
```

# Bash 的 last 命令

```
last [OPTIONS] [USERNAME...] [TTY...]
```

**功能**

显示最近登录的用户列表。

**类型**

可执行文件（`/usr/bin/last`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-a`，`--hostlast` - 在最后一列显示主机名
  - `-d`，`--dns` - 将远程登录的 IP 地址转换为主机名
  - `-f`，`--file FILE` - 使用 `FILE` 文件代替默认的 `/var/log/wtmp`
  - `-F`，`--fulltimes` - 显示完整的日期和时间
  - `-i`，`--ip` - 显示远程登录的 IP 地址
  - `-n`，`--limit NUMBER` - 显示最近的 `NUMBER` 条记录
  - `-p`，`--present TIME` - 显示指定时间段内的记录；相当于同时使用 `--since` 和 `--until`
  - `-R`，`--nohostname` - 隐藏主机名
  - `-s`，`--since TIME` - 显示自指定时间 `TIME`（含） 以来的记录
  - `-t`，`--until TIME` - 显示自指定时间 `TIME`（不含） 之前的记录
  - `--time-format FORMAT` - 指定时间的显示格式：
    - `notime` - 不显示时间
    - `short` - 短格式（默认）
    - `full` - 长格式
    - `iso` - ISO-8601
  - `-w`，`--fullnames` - 显示完整的用户名
  - `-x`，`--system` - 显示系统关机条目和运行级别更改
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `USERNAME` 指定用户列表
- `TTY` 指定终端列表

# 示例

```
$ last -n 10                    # 查看最近的 10 条登录记录
...
$ last -n 10 -s '2026-01-01'    # 查看 2026-01-01 以来的 10 条登录记录
...
```

# 推荐阅读

- last(1) — Linux manual page

# 手册

```
LAST(1)                         User Commands                         LAST(1)

NAME
       last, lastb - show a listing of last logged in users

SYNOPSIS
       last [options] [username...] [tty...]

       lastb [options] [username...] [tty...]

DESCRIPTION
       last searches back through the /var/log/wtmp file (or the file
       designated by the -f option) and displays a list of all users logged in
       (and out) since that file was created. One or more usernames and/or
       ttys can be given, in which case last will show only the entries
       matching those arguments. Names of ttys can be abbreviated, thus last 0
       is the same as last tty0.

       When catching a SIGINT signal (generated by the interrupt key, usually
       control-C) or a SIGQUIT signal, last will show how far it has searched
       through the file; in the case of the SIGINT signal last will then
       terminate.

       The pseudo user reboot logs in each time the system is rebooted. Thus
       last reboot will show a log of all the reboots since the log file was
       created.

       lastb is the same as last, except that by default it shows a log of the
       /var/log/btmp file, which contains all the bad login attempts.

OPTIONS
       -a, --hostlast
           Display the hostname in the last column. Useful in combination with
           the --dns option.
```

```
-d, --dns
    For non-local logins, Linux stores not only the host name of the
    remote host, but its IP number as well. This option translates the
    IP number back into a hostname.

-f, --file file
    Tell last to use a specific file instead of /var/log/wtmp. The
    --file option can be given multiple times, and all of the specified
    files will be processed.

-F, --fulltimes
    Print full login and logout times and dates.

-i, --ip
    Like --dns , but displays the host’s IP number instead of the name.

-number; -n, --limit number
    Tell last how many lines to show.

-p, --present time
    Display the users who were present at the specified time. This is
    like using the options --since and --until together with the same
    time.

-R, --nohostname
    Suppresses the display of the hostname field.

-s, --since time
    Display the state of logins since the specified time. This is
    useful, e.g., to easily determine who was logged in at a particular
    time. The option is often combined with --until.

-t, --until time
    Display the state of logins until the specified time.

--time-format format
    Define the output timestamp format to be one of notime, short,
    full, or iso. The notime variant will not print any timestamps at
    all, short is the default, and full is the same as the --fulltimes
    option. The iso variant will display the timestamp in ISO-8601
    format. The ISO format contains timezone information, making it
    preferable when printouts are investigated outside of the system.

-w, --fullnames
    Display full user names and domain names in the output.

-x, --system
    Display the system shutdown entries and run level changes.

-h, --help
    Display help text and exit.
```

-V, --version
Print version and exit.

TIME FORMATS

The options that take the time argument understand the following formats:

| | |
|---|---|
| YYYYMMDDhhmmss | |
| YYYY-MM-DD hh:mm:ss | |
| YYYY-MM-DD hh:mm | (seconds will be set to 00) |
| YYYY-MM-DD | (time will be set to 00:00:00) |
| hh:mm:ss | (date will be set to today) |
| hh:mm | (date will be set to today, seconds to 00) |
| now | |
| yesterday | (time is set to 00:00:00) |
| today | (time is set to 00:00:00) |
| tomorrow | (time is set to 00:00:00) |
| +5min | |
| -5days | |

FILES

/var/log/wtmp, /var/log/btmp

NOTES

The files wtmp and btmp might not be found. The system only logs

information in these files if they are present. This is a local configuration issue. If you want the files to be used, they can be created with a simple touch(1) command (for example, touch /var/log/wtmp).

An empty entry is a valid type of wtmp entry. It means that an empty file or file with zeros is not interpreted as an error.

The utmp file format uses fixed sizes of strings, which means that very long strings are impossible to store in the file and impossible to display by last. The usual limits are 32 bytes for a user and line name and 256 bytes for a hostname.

## AUTHORS

Miquel van Smoorenburg <miquels@cistron.nl>

## SEE ALSO

login(1), wtmp(5), init(8), shutdown(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The last command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 lastb 命令

```
lastb [OPTION]... [TTY]...
```

**功能**

显示最近登录的用户列表。

**类型**

- `OPTION` 选项：
  - `-a`，`--hostlast` - 在最后一列显示主机名
  - `-d`，`--dns` - 将 IP 转换为主机名
  - `-f`，`--file file` - 使用指定的文件替代 `/var/log/wtmp`
  - `-F`，`--fulltimes` - 显示完整的登录、注销日期和时间
  - `-i`，`--ip` - 显示 IP 地址
  - `-number`，`-n`，`--limit number` - 显示 `number` 行
  - `-p`，`--present time` - 显示指定时间段内的日志
  - `-R`，`--nohostname` - 不显示主机名
  - `-s`，`--since time` - 显示 `time` 之后的日志
  - `-t`，`--until time` - 显示 `time` 之前的日志
  - `--time-format format` - 指定时间格式
  - `-w`，`--fullnames` - 显示完整的用户名和域名
  - `-x`，`--system` - 显示系统关机条目和运行级别更改
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `NAME` - 要定位的文件列表

**参数**

## 示例

```
$ sudo lastb
root     pts/4                        Mon Mar  9 06:35 - 06:35  (00:00)

btmp begins Mon Mar  9 06:35:51 2026
```

# 推荐阅读

- lastb(1) — Linux manual page

# 手册

```
LAST(1)                         User Commands                         LAST(1)

NAME
       last, lastb - show a listing of last logged in users

SYNOPSIS
       last [options] [username...] [tty...]

       lastb [options] [username...] [tty...]

DESCRIPTION
       last searches back through the /var/log/wtmp file (or the file
       designated by the -f option) and displays a list of all users logged in
       (and out) since that file was created. One or more usernames and/or
       ttys can be given, in which case last will show only the entries
       matching those arguments. Names of ttys can be abbreviated, thus last 0
       is the same as last tty0.

       When catching a SIGINT signal (generated by the interrupt key, usually
       control-C) or a SIGQUIT signal, last will show how far it has searched
       through the file; in the case of the SIGINT signal last will then
       terminate.

       The pseudo user reboot logs in each time the system is rebooted. Thus
       last reboot will show a log of all the reboots since the log file was
       created.

       lastb is the same as last, except that by default it shows a log of the
       /var/log/btmp file, which contains all the bad login attempts.

OPTIONS
       -a, --hostlast
           Display the hostname in the last column. Useful in combination with
           the --dns option.

       -d, --dns
           For non-local logins, Linux stores not only the host name of the
           remote host, but its IP number as well. This option translates the
           IP number back into a hostname.

       -f, --file file
           Tell last to use a specific file instead of /var/log/wtmp. The
           --file option can be given multiple times, and all of the specified
           files will be processed.
```

-F, --fulltimes
Print full login and logout times and dates.

-i, --ip
Like --dns , but displays the host’s IP number instead of the name.

-number; -n, --limit number
Tell last how many lines to show.

-p, --present time
Display the users who were present at the specified time. This is like using the options --since and --until together with the same time.

-R, --nohostname
Suppresses the display of the hostname field.

-s, --since time
Display the state of logins since the specified time. This is useful, e.g., to easily determine who was logged in at a particular time. The option is often combined with --until.

-t, --until time
Display the state of logins until the specified time.

--time-format format
Define the output timestamp format to be one of notime, short, full, or iso. The notime variant will not print any timestamps at all, short is the default, and full is the same as the --fulltimes option. The iso variant will display the timestamp in ISO-8601 format. The ISO format contains timezone information, making it preferable when printouts are investigated outside of the system.

-w, --fullnames
Display full user names and domain names in the output.

-x, --system
Display the system shutdown entries and run level changes.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

TIME FORMATS
The options that take the time argument understand the following formats:

| | |
|---|---|
| YYYYMMDDhhmmss | |

| | |
|---|---|
| YYYY-MM-DD hh:mm:ss | |
| YYYY-MM-DD hh:mm | (seconds will be set to 00) |
| YYYY-MM-DD | (time will be set to 00:00:00) |
| hh:mm:ss | (date will be set to today) |
| hh:mm | (date will be set to today, seconds to 00) |
| now | |
| yesterday | (time is set to 00:00:00) |
| today | (time is set to 00:00:00) |
| tomorrow | (time is set to 00:00:00) |
| +5min | |
| -5days | |

FILES

/var/log/wtmp, /var/log/btmp

NOTES

The files wtmp and btmp might not be found. The system only logs information in these files if they are present. This is a local configuration issue. If you want the files to be used, they can be created with a simple touch(1) command (for example, touch /var/log/wtmp).

An empty entry is a valid type of wtmp entry. It means that an empty file or file with zeros is not interpreted as an error.

The utmp file format uses fixed sizes of strings, which means that very long strings are impossible to store in the file and impossible to

```
       display by last. The usual limits are 32 bytes for a user and line name
       and 256 bytes for a hostname.

AUTHORS
       Miquel van Smoorenburg <miquels@cistron.nl>

SEE ALSO
       login(1), wtmp(5), init(8), shutdown(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The last command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                          LAST(1)
```

# Bash 的 ldattach 命令

```
ldattach [OPTION]... LDISC DEVICE
```

**功能**

给终端设备（通常是串口）附加一个“线路规程（line discipline）”。

**类型**

可执行文件（`/usr/sbin/ldattach`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-1`，`--onestopbit` - 将串口的停止位数设置为 1
  - `-2`，`--twostopbits` - 将串口的停止位数设置为 2
  - `-7`，`--sevenbits` - 将串口的字符大小设置为 7 位
  - `-8`，`--eightbits` - 将串口的字符大小设置为 8 位
  - `-d`，`--debug` - 保持在前台运行，以便可以中断或调试
  - `-e`，`--evenparity` - 将串口的校验设置为偶校验
  - `-i`，`--iflag value...` - 设置串口 `c_iflag` 字中的指定位；可以是数字或名称，以减号（`-`）开头则清除指定的位
  - `-n`，`--noparity` - 将串口的校验设置为无
  - `-o`，`--oddparity` - 将串口的校验设置为奇校验
  - `-s`，`--speed value` - 设置串口的速率（波特率）
  - `-c`，`--intro-command string` - 定义一个在调用 `ldattach` 之前通过串口发送的引导命令
  - `-p`，`--pause value` - 在调用 `ldattach` 之前休眠指定秒数；默认为 1
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `LDISC` - 线路规程
- `DEVICE` - 要操作的设备

## 线路规程

| 名称 (标识) | 说明 (Description) |
| --- | --- |
| TTY(0) | 默认线路规程，提供透明操作（raw 模式）以及习惯的终端线路编辑功能（cooked 模式）。 |
| SLIP(1) | 串行线路 IP (SLIP) 协议处理器，用于通过串行线路传输 TCP/IP 数据包。 |
| MOUSE(2) | 用于连接 RS232 的指点设备（串行鼠标）的设备驱动程序。 |
| PPP(3) | 点对点协议 (PPP) 处理器，用于通过串行线路传输网络数据包。 |
| STRIP(4); AX25(5); X25(6) | 线路驱动程序，用于通过异步串行线路传输 X.25 数据包。 |
| 6PACK(7); R3964(9) | Simatic R3964 模块的驱动程序。 |
| IRDA(11) | Linux IrDa（红外数据传输）驱动程序 - 参见 http://irda.sourceforge.net/ |
| HDLC(13) | 同步 HDLC 驱动程序。 |
| SYNC_PPP(14) | 同步 PPP 驱动程序。 |
| HCI(15) | 蓝牙 HCI UART 驱动程序。 |
| GIGASET_M101(16) | 西门子 Gigaset M101 串行 DECT 适配器的驱动程序。 |
| PPS(18) | 串行线路每秒脉冲 (PPS) 源的驱动程序。 |
| GSM0710(21) | GSM 07.10 多路复用协议调制解调器 (CMUX) 的驱动程序。 |

## 示例

### GPS（NMEA）

```
$ ldattach nmea /dev/ttyS0
```

### 蓝牙 HCI（UART）

```
$ ldattach hci /dev/ttyS1
```

### 拨号网络 PPP

```
$ ldattach ppp /dev/ttyS0
```

### 调制解调器 GSM

```
$ ldattach gsm /dev/ttyUSB0
```

## 推荐阅读

- ldattach(8) — Linux manual page

## 手册

```
LDATTACH(8)                 System Administration                LDATTACH(8)

NAME
       ldattach - attach a line discipline to a serial line

SYNOPSIS
       ldattach [-1278denoVh] [-i iflag] [-s speed] ldisc device

DESCRIPTION
       The ldattach daemon opens the specified device file (which should refer
       to a serial device) and attaches the line discipline ldisc to it for
       processing of the sent and/or received data. It then goes into the
       background keeping the device open so that the line discipline stays
       loaded.

       The line discipline ldisc may be specified either by name or by number.

       In order to detach the line discipline, kill(1) the ldattach process.

       With no arguments, ldattach prints usage information.

LINE DISCIPLINES
       Depending on the kernel release, the following line disciplines are
       supported:

       TTY(0)
           The default line discipline, providing transparent operation (raw
           mode) as well as the habitual terminal line editing capabilities
           (cooked mode).

       SLIP(1)
           Serial Line IP (SLIP) protocol processor for transmitting TCP/IP
           packets over serial lines.

       MOUSE(2)
           Device driver for RS232 connected pointing devices (serial mice).

       PPP(3)
           Point to Point Protocol (PPP) processor for transmitting network
           packets over serial lines.
```

STRIP(4); AX25(5); X25(6)
Line driver for transmitting X.25 packets over asynchronous serial lines.

6PACK(7); R3964(9)
Driver for Simatic R3964 module.

IRDA(11)
Linux IrDa (infrared data transmission) driver - see http://irda.sourceforge.net/

HDLC(13)
Synchronous HDLC driver.

SYNC_PPP(14)
Synchronous PPP driver.

HCI(15)
Bluetooth HCI UART driver.

GIGASET_M101(16)
Driver for Siemens Gigaset M101 serial DECT adapter.

PPS(18)
Driver for serial line Pulse Per Second (PPS) source.

GSM0710(21)
Driver for GSM 07.10 multiplexing protocol modem (CMUX).

OPTIONS

-1, --onestopbit
Set the number of stop bits of the serial line to one.

-2, --twostopbits
Set the number of stop bits of the serial line to two.

-7, --sevenbits
Set the character size of the serial line to 7 bits.

-8, --eightbits
Set the character size of the serial line to 8 bits.

-d, --debug
Keep ldattach in the foreground so that it can be interrupted or debugged, and to print verbose messages about its progress to standard error output.

-e, --evenparity
Set the parity of the serial line to even.

-i, --iflag value...
Set the specified bits in the c_iflag word of the serial line. The

```
           given value may be a number or a symbolic name. If value is
           prefixed by a minus sign, the specified bits are cleared instead.
           Several comma-separated values may be given in order to set and
           clear multiple bits.

       -n, --noparity
           Set the parity of the serial line to none.

       -o, --oddparity
           Set the parity of the serial line to odd.

       -s, --speed value
           Set the speed (the baud rate) of the serial line to the specified
           value.

       -c, --intro-command string
           Define an intro command that is sent through the serial line before
           the invocation of ldattach. E.g. in conjunction with line
           discipline GSM0710, the command 'AT+CMUX=0\r' is commonly suitable
           to switch the modem into the CMUX mode.

       -p, --pause value
           Sleep for value seconds before the invocation of ldattach. Default
           is one second.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

AUTHORS
       Tilman Schmidt <tilman@imap.cc>

SEE ALSO
       inputattach(1), ttys(4)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The ldattach command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                       LDATTACH(8)
```

# Bash 的 link 命令

```
link [OPTION] SRC DEST
```

**功能**

创建文件的硬链接。

**类型**

可执行文件（`/usr/bin/link`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `SRC` - 源文件
- `DEST` - 目标文件（硬链接）

## 示例

```
$ link /opt/xxxx/cmd /usr/bin/cmd
```

- 将文件 `/opt/xxxx/cmd` 链接到 `/usr/bin/cmd`

## 相关命令

| 命令 | 说明 |
| --- | --- |
| ln | 创建文件的链接（硬链接或符号链接） |

## 推荐阅读

- Make a hard link via the link syscall - GNU Coreutils
- link(1) — Linux manual page

## 手册

NAME

link - call the link function to create a link to a file

SYNOPSIS

link FILE1 FILE2
link OPTION

DESCRIPTION

Call the link function to create a link named FILE2 to an existing FILE1.

--help display this help and exit

--version
output version information and exit

AUTHOR

Written by Michael Stone.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

link(2)

Full documentation <https://www.gnu.org/software/coreutils/link>
or available locally via: info '(coreutils) link invocation'

# Bash 的 linux32 命令

```
linux32 [OPTION]... [CMD [ARG]...]
```

**功能**

**类型**

可执行文件（`/usr/bin/linux32`），属于 util-linux。

> 实际上是 `/usr/bin/setarch` 的符号链接，并默认将架构设为 `linux32`。

**参数**

- `OPTION` 选项：
  - `--uname-2.6` - 使程序检测到的内核版本号为 `2.6`
  - `-v`, `--verbose` - 显示详细信息
  - `-3`, `--3gb` - 指定程序最多使用 3GB 的地址空间（3GB 用户空间，1GB 内核空间）
  - `--4gb` - ~~指定程序最多使用 4GB 的地址空间（4GB 用户空间，没有内核空间）~~；此选项已废弃不再生效
  - `-B`, `--32bit` - 将地址空间设为 32 位以模拟硬件；支持 ARM 和 Alpha 架构
  - `-F`, `--fdpic-funcptrs` - 将用户空间函数指向信号处理程序的指针视为指向地址描述符的指针；仅支持 ARM, Blackfin, Fujitsu FR-V 和 SuperH 架构
  - `-I`, `--short-inode` - ~~使用 16 位 inode 号~~；此选项已废弃不再生效
  - `-L`, `--addr-compat-layout` - 提供传统虚拟地址空间布局
  - `-R`, `--addr-no-randomize` - 禁用虚拟地址空间的随机化
  - `-S`, `--whole-seconds` - ~~时间值截断为整数秒~~；此选项已废弃不再生效
  - `-T`, `--sticky-timeouts` - 这使得 `select(2)`、`pselect(2)` 和 `ppoll(2)` 系统调用保留超时值
  - `-X`, `--read-implies-exec` - 如果设置此项，则 `mmap(2)` `PROT_READ` 还会添加 `PROT_EXEC`` 位
  - `-Z`, `--mmap-page-zero` - SVr4 错误模拟，会将 `mmap(2)` 第 0 页设置为只读。
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CMD` - 要运行的命令
- `ARG` - 传递给 `CMD` 的参数列表

## 示例

```
$ linux32 arch
i686
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| setarch | 使用指定的架构或个性化标志运行程序 |

## 推荐阅读

- setarch(8) — Linux manual page

## 手册

```
SETARCH(8)                  System Administration                  SETARCH(8)

NAME
       setarch - change reported architecture in new program environment
       and/or set personality flags

SYNOPSIS
       setarch [arch] [options] [program [argument...]]

       setarch --list|-h|-V

       arch [options] [program [argument...]]

DESCRIPTION
       setarch modifies execution domains and process personality flags.

       The execution domains currently only affects the output of uname -m.
       For example, on an AMD64 system, running setarch i386 program will
       cause program to see i686 instead of x86_64 as the machine type. It can
       also be used to set various personality options. The default program is
       /bin/sh.

       Since version 2.33 the arch command line argument is optional and
       setarch may be used to change personality flags (ADDR_LIMIT_*,
       SHORT_INODE, etc) without modification of the execution domain.

OPTIONS
       --list
           List the architectures that setarch knows about. Whether setarch
           can actually set each of these architectures depends on the running
           kernel.
```

```
--show[=personality]
    Show the currently active personality and flags. If the personality
    argument is provided, it is shown instead of the current one.
    personality is a hexadecimal number with values was described in
    sys/personality.h.

--uname-2.6
    Causes the program to see a kernel version number beginning with
    2.6. Turns on UNAME26.

-v, --verbose
    Be verbose.

-3, --3gb
    Specifies program should use a maximum of 3GB of address space.
    Supported on x86. Turns on ADDR_LIMIT_3GB.

--4gb
    This option has no effect. It is retained for backward
    compatibility only, and may be removed in future releases.

-B, --32bit
    Limit the address space to 32 bits to emulate hardware. Supported
    on ARM and Alpha. Turns on ADDR_LIMIT_32BIT.

-F, --fdpic-funcptrs
    Treat user-space function pointers to signal handlers as pointers
    to address descriptors. This option has no effect on architectures
    that do not support FDPIC ELF binaries. In kernel v4.14 support is
    limited to ARM, Blackfin, Fujitsu FR-V, and SuperH CPU
    architectures.

-I, --short-inode
    Obsolete bug emulation flag. Turns on SHORT_INODE.

-L, --addr-compat-layout
    Provide legacy virtual address space layout. Use when the program
    binary does not have PT_GNU_STACK ELF header. Turns on
    ADDR_COMPAT_LAYOUT.

-R, --addr-no-randomize
    Disables randomization of the virtual address space. Turns on
    ADDR_NO_RANDOMIZE.

-S, --whole-seconds
    Obsolete bug emulation flag. Turns on WHOLE_SECONDS.

-T, --sticky-timeouts
    This makes select(2), pselect(2), and ppoll(2) system calls
    preserve the timeout value instead of modifying it to reflect the
    amount of time not slept when interrupted by a signal handler. Use
    when program depends on this behavior. For more details see the
```

timeout description in select(2) manual page. Turns on
STICKY_TIMEOUTS.

-X, --read-implies-exec
If this is set then mmap(2) PROT_READ will also add the PROT_EXEC
bit - as expected by legacy x86 binaries. Notice that the ELF
loader will automatically set this bit when it encounters a legacy
binary. Turns on READ_IMPLIES_EXEC.

-Z, --mmap-page-zero
SVr4 bug emulation that will set mmap(2) page zero as read-only.
Use when program depends on this behavior, and the source code is
not available to be fixed. Turns on MMAP_PAGE_ZERO.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

EXAMPLE

```
setarch --addr-no-randomize mytestprog
setarch ppc32 rpmbuild --target=ppc --rebuild foo.src.rpm
setarch ppc32 -v -vL3 rpmbuild --target=ppc --rebuild bar.src.rpm
setarch ppc32 --32bit rpmbuild --target=ppc --rebuild foo.src.rpm
```

AUTHORS

Elliot Lee <sopwith@redhat.com>, Jindrich Novy <jnovy@redhat.com>,
Karel Zak <kzak@redhat.com>

SEE ALSO

personality(2), select(2)

REPORTING BUGS

For bug reports, use the issue tracker at
https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The setarch command is part of the util-linux package which can be
downloaded from Linux Kernel Archive
<https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 linux64 命令

```
linux64 [OPTION]... [CMD [ARG]...]
```

**功能**

**类型**

可执行文件（`/usr/bin/linux64`），属于 util-linux。

> 实际上是 `/usr/bin/setarch` 的符号链接，并默认将架构设为 `linux64`。

**参数**

- `OPTION` 选项：
  - `--uname-2.6` - 使程序检测到的内核版本号为 `2.6`
  - `-v`，`--verbose` - 显示详细信息
  - `-3`，`--3gb` - 指定程序最多使用 3GB 的地址空间（3GB 用户空间，1GB 内核空间）
  - `--4gb` - ~~指定程序最多使用 4GB 的地址空间（4GB 用户空间，没有内核空间）~~；此选项已废弃不再生效
  - `-B`，`--32bit` - 将地址空间设为 32 位以模拟硬件；支持 ARM 和 Alpha 架构
  - `-F`，`--fdpic-funcptrs` - 将用户空间函数指向信号处理程序的指针视为指向地址描述符的指针；仅支持 ARM, Blackfin, Fujitsu FR-V 和 SuperH 架构
  - `-I`，`--short-inode` - ~~使用 16 位 inode 号~~；此选项已废弃不再生效
  - `-L`，`--addr-compat-layout` - 提供传统虚拟地址空间布局
  - `-R`，`--addr-no-randomize` - 禁用虚拟地址空间的随机化
  - `-S`，`--whole-seconds` - ~~时间值截断为整数秒~~；此选项已废弃不再生效
  - `-T`，`--sticky-timeouts` - 这使得 `select(2)`、`pselect(2)` 和 `ppoll(2)` 系统调用保留超时值
  - `-X`，`--read-implies-exec` - 如果设置此项，则 `mmap(2)` `PROT_READ` 还会添加 `PROT_EXEC`` 位
  - `-Z`，`--mmap-page-zero` - SVr4 错误模拟，会将 `mmap(2)` 第 0 页设置为只读。
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CMD` - 要运行的命令
- `ARG` - 传递给 `CMD` 的参数列表

## 示例

```
$ linux64 arch
x86_64
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| setarch | 使用指定的架构或个性化标志运行程序 |

## 推荐阅读

- setarch(8) — Linux manual page

## 手册

```
SETARCH(8)                  System Administration                  SETARCH(8)

NAME
       setarch - change reported architecture in new program environment
       and/or set personality flags

SYNOPSIS
       setarch [arch] [options] [program [argument...]]

       setarch --list|-h|-V

       arch [options] [program [argument...]]

DESCRIPTION
       setarch modifies execution domains and process personality flags.

       The execution domains currently only affects the output of uname -m.
       For example, on an AMD64 system, running setarch i386 program will
       cause program to see i686 instead of x86_64 as the machine type. It can
       also be used to set various personality options. The default program is
       /bin/sh.

       Since version 2.33 the arch command line argument is optional and
       setarch may be used to change personality flags (ADDR_LIMIT_*,
       SHORT_INODE, etc) without modification of the execution domain.

OPTIONS
       --list
           List the architectures that setarch knows about. Whether setarch
           can actually set each of these architectures depends on the running
           kernel.
```

```
--show[=personality]
    Show the currently active personality and flags. If the personality
    argument is provided, it is shown instead of the current one.
    personality is a hexadecimal number with values was described in
    sys/personality.h.

--uname-2.6
    Causes the program to see a kernel version number beginning with
    2.6. Turns on UNAME26.

-v, --verbose
    Be verbose.

-3, --3gb
    Specifies program should use a maximum of 3GB of address space.
    Supported on x86. Turns on ADDR_LIMIT_3GB.

--4gb
    This option has no effect. It is retained for backward
    compatibility only, and may be removed in future releases.

-B, --32bit
    Limit the address space to 32 bits to emulate hardware. Supported
    on ARM and Alpha. Turns on ADDR_LIMIT_32BIT.

-F, --fdpic-funcptrs
    Treat user-space function pointers to signal handlers as pointers
    to address descriptors. This option has no effect on architectures
    that do not support FDPIC ELF binaries. In kernel v4.14 support is
    limited to ARM, Blackfin, Fujitsu FR-V, and SuperH CPU
    architectures.

-I, --short-inode
    Obsolete bug emulation flag. Turns on SHORT_INODE.

-L, --addr-compat-layout
    Provide legacy virtual address space layout. Use when the program
    binary does not have PT_GNU_STACK ELF header. Turns on
    ADDR_COMPAT_LAYOUT.

-R, --addr-no-randomize
    Disables randomization of the virtual address space. Turns on
    ADDR_NO_RANDOMIZE.

-S, --whole-seconds
    Obsolete bug emulation flag. Turns on WHOLE_SECONDS.

-T, --sticky-timeouts
    This makes select(2), pselect(2), and ppoll(2) system calls
    preserve the timeout value instead of modifying it to reflect the
    amount of time not slept when interrupted by a signal handler. Use
    when program depends on this behavior. For more details see the
```

timeout description in select(2) manual page. Turns on
STICKY_TIMEOUTS.

-X, --read-implies-exec
If this is set then mmap(2) PROT_READ will also add the PROT_EXEC
bit - as expected by legacy x86 binaries. Notice that the ELF
loader will automatically set this bit when it encounters a legacy
binary. Turns on READ_IMPLIES_EXEC.

-Z, --mmap-page-zero
SVr4 bug emulation that will set mmap(2) page zero as read-only.
Use when program depends on this behavior, and the source code is
not available to be fixed. Turns on MMAP_PAGE_ZERO.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

EXAMPLE

```
setarch --addr-no-randomize mytestprog
setarch ppc32 rpmbuild --target=ppc --rebuild foo.src.rpm
setarch ppc32 -v -vL3 rpmbuild --target=ppc --rebuild bar.src.rpm
setarch ppc32 --32bit rpmbuild --target=ppc --rebuild foo.src.rpm
```

AUTHORS

Elliot Lee <sopwith@redhat.com>, Jindrich Novy <jnovy@redhat.com>,
Karel Zak <kzak@redhat.com>

SEE ALSO

personality(2), select(2)

REPORTING BUGS

For bug reports, use the issue tracker at
https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The setarch command is part of the util-linux package which can be
downloaded from Linux Kernel Archive
<https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 ln 命令

```
ln [OPTION] SRC DEST
```

**功能**

创建文件的硬链接或软链接（符号链接）。

**类型**

可执行文件（`/usr/bin/ln`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--backup[=CONTROL]` - 选择怎样备份被覆盖的目标文件
  - `-b` - 备份被覆盖的目标文件
  - `-d`，`-F`，`--directory` - 允许超级用户尝试创建目录的硬链接
  - `-f`，`--force` - 强制操作，目标存在且无法打开时会尝试将其删除
  - `-i`，`--interactive` - 覆盖前提示
  - `-L`，`--logical` - 源文件是符号链接时，链接该链接的源文件，而不是该链接自身
  - `-n`，`--no-dereference` - 目标是指向目录的符号链接时将其视作普通文件进行覆盖
  - `-P`，`--physical` - 源文件是符号链接时，直接创建指向该链接的硬链接
  - `-r`，`--relative` - 将源文件的路径转换为参照符号链接所在目录的相对路径（与 `-s` 一起使用）
  - `-s`，`--symbolic` - 创建符号链接，而不是硬链接
  - `-S`，`--suffix=SUFFIX` - 使用 `SUFFIX` 作为备份文件的后缀；默认后缀为 `~`
  - `-t`，`--target-directory=DIRECTORY` - 指定要创建链接的目录
  - `-T`，`--no-target-directory` - 将 `DEST` 视为文件（如果 `DEST` 是目录，则进行覆盖而非复制到它内部）
  - `-v`，`--verbose` - 打印详细信息
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `SRC` - 源文件
- `DEST` - 目标文件（硬链接）

## 示例

**绝对路径**

```
$ ln -s /opt/xxxx/cmd /usr/bin/        # 将 /opt/xxxx/cmd 链接到 /usr/bin/cmd
$ ls -l /usr/bin/cmd
lrwx-wxr-x 1 root root 5 Nov 17 12:12 /usr/bin/cmd -> /opt/xxxx/cmd
```

- `/usr/bin/cmd` 是链接到 `/opt/xxxx/cmd` 的软连接

**相对路径**

```
$ ln -sr ./cmd ./dir        # 将 ./cmd 链接到 ./dir/cmd
$ ls -l ./dir/cmd
lrwx-wxr-x 1 root root 5 Nov 17 12:12 ./dir/cmd -> ../cmd
```

- `./dir/cmd` 是链接到 `../cmd` 的软连接

错误示范

```
$ ln -s ./cmd ./dir        # 将 ./cmd 链接到 ./dir/cmd
$ ls -l ./dir/cmd
lrwx-wxr-x 1 root root 5 Nov 17 12:12 ./dir/cmd -> ./cmd
```

- `./dir/cmd` 是链接到 `./cmd` 的软连接

软连接在解析时相对路径参照自身所在目录，因此无法找到源文件，`./cmd` 指向自身导致无限递归错误。

## 相关命令

| 命令 | 说明 |
| --- | --- |
| link | 创建文件的硬链接 |
| readlink | 获取符号链接的源文件 |

## 推荐阅读

- Make links between files - GNU Coreutils
- ln(1) — Linux manual page

## 手册

```
LN(1)                          User Commands                          LN(1)
```

NAME

ln - make links between files

SYNOPSIS

ln [OPTION]... [-T] TARGET LINK_NAME
ln [OPTION]... TARGET
ln [OPTION]... TARGET... DIRECTORY
ln [OPTION]... -t DIRECTORY TARGET...

DESCRIPTION

In the 1st form, create a link to TARGET with the name LINK_NAME. In the 2nd form, create a link to TARGET in the current directory. In the 3rd and 4th forms, create links to each TARGET in DIRECTORY. Create hard links by default, symbolic links with --symbolic. By default, each destination (name of new link) should not already exist. When creating hard links, each TARGET must exist. Symbolic links can hold arbitrary text; if later resolved, a relative link is interpreted in relation to its parent directory.

Mandatory arguments to long options are mandatory for short options too.

--backup[=CONTROL]
make a backup of each existing destination file

-b like --backup but does not accept an argument

-d, -F, --directory
allow the superuser to attempt to hard link directories (note: will probably fail due to system restrictions, even for the superuser)

-f, --force
remove existing destination files

-i, --interactive
prompt whether to remove destinations

-L, --logical
dereference TARGETs that are symbolic links

-n, --no-dereference
treat LINK_NAME as a normal file if it is a symbolic link to a directory

-P, --physical
make hard links directly to symbolic links

-r, --relative
with -s, create links relative to link location

-s, --symbolic
make symbolic links instead of hard links

```
       -S, --suffix=SUFFIX
              override the usual backup suffix

       -t, --target-directory=DIRECTORY
              specify the DIRECTORY in which to create the links

       -T, --no-target-directory
              treat LINK_NAME as a normal file always

       -v, --verbose
              print name of each linked file

       --help display this help and exit

       --version
              output version information and exit

       The  backup  suffix  is  '~', unless set with --suffix or SIM-
       PLE_BACKUP_SUFFIX.  The version control method may be selected via  the
       --backup  option  or  through the VERSION_CONTROL environment variable.
       Here are the values:

       none, off
              never make backups (even if --backup is given)

       numbered, t
              make numbered backups

       existing, nil
              numbered if numbered backups exist, simple otherwise

       simple, never
              always make simple backups

       Using -s ignores -L and -P.  Otherwise, the last option specified  con-
       trols behavior when a TARGET is a symbolic link, defaulting to -P.

AUTHOR
       Written by Mike Parker and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       link(2), symlink(2)
```

```
       Full documentation <https://www.gnu.org/software/coreutils/ln>
       or available locally via: info '(coreutils) ln invocation'

GNU coreutils 9.4                 April 2024                              LN(1)
```

# Bash 的 logname 命令

```
logname [OPTION]
```

**功能**

查看当前登录的用户名。

**类型**

可执行文件（`/usr/bin/logname`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ logname
primers
```

## 推荐阅读

- Print current login name - GNU Coreutils
- logname(1) — Linux manual page

## 手册

```
LOGNAME(1)                      User Commands                     LOGNAME(1)

NAME
       logname - print user´s login name

SYNOPSIS
       logname [OPTION]

DESCRIPTION
       Print the user's login name.
```

```
       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by FIXME: unknown.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       getlogin(3)

       Full documentation <https://www.gnu.org/software/coreutils/logname>
       or available locally via: info '(coreutils) logname invocation'

GNU coreutils 9.4                April 2024                        LOGNAME(1)
```

# Bash 的 ls 命令

```
ls [OPTION]... [FILE]...
```

**功能**

列出文件信息。

**类型**

可执行文件（`/usr/bin/ls`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 显示隐藏文件（文件名以 `.` 开头）
  - `-A`，`--almost-all` - 显示隐藏文件，但不显示 `.` 和 `..`
  - `--author` - 显示文件的作者；和 `-l` 选项一起使用
  - `-b`，`--escape` - 非 图形字符 显示为 C 风格的 转义字符
  - `--block-size=SIZE` - 文件大小的单位；和 `-l` 选项一起使用
  - `-B`，`--ignore-backups` 不显示备份文件（文件名以 `~` 结束）
  - `-c` - 按 ctime（文件状态信息最后更新的时间）排序；和 `-lt` 选项一起使用
  - `-C` - 竖向排列
  - `--color[=WHEN]` - 何时使用彩色，有效值：`always`，`auto`，`never`
  - `-d`，`--directory` - 查看目录本身，而不是目录的内容
  - `-D`，`--dired` - 生成为 Emacs 的 dired 模式设计的输出
  - `-f` - 显示所有文件且不进行排序（按照目录顺序显示）
  - `-F`，`--classify[=WHEN]` - 何时添加 文件类型指示器，有效值：`always`，`auto`，`never`
  - `--file-type` - 显示文件类型指示器；同 `--classify` 但是不添加 `*`
  - `--format=WORD` - 格式，有效值：`across`，`commas`，`horizontal`，`long`，`single-column`，`verbose`，`vertical`
  - `--full-time` - 显示完整的时间；同 `-l --time-style=full-iso`
  - `-g` - 显示文件详情；同 `-l` 但不显示所有者
  - `--group-directories-first` - 分组时将目录排在文件前面
  - `-G`，`--no-group` - 不显示分组名
  - `-h`，`--human-readable` - 让详细信息更适合人类阅读；和 `-l` 或 `-s` 选择一起使用
  - `--si` - 让详细信息更适合人类阅读；同 `-h` 但是文件大小以 1000 为单位，而不是 1024
  - `-H`，`--dereference-command-line` - 参数是符号链接时，查看源文件

- `--dereference-command-line-symlink-to-dir` - 参数是目录的符号链接时，查看源目录
- `--hide=PATTERN` - 隐藏匹配的文件
- `--hyperlink[=WHEN]` - 何时将文件名显示为超链接（现代终端上点击可以打开文件管理器），有效值：`always`，`auto`，`never`
- `--indicator-style=WORD` - 添加 文件类型指示器 ，有效值：`none`，`slash`，`file-type`，`classify`
- `-i`，`--inode` - 显示文件的 inode 编号
- `-I`，`--ignore=PATTERN` - 不显示匹配的文件
- `-k`，`--kibibytes` - 文件大小以 1024 字节为单位
- `-l` - 显示详细信息
- `-L`，`--dereference` - 符号链接显示源文件的信息，而非符号链接本身的信息
- `-m` - 用逗号分隔条目
- `-n`，`--numeric-uid-gid` - 显示详细信息，但是所有者和所属组显示为 ID 而非名称
- `-N`，`--literal` - 含特殊字符的文件名不添加单引号
- `-o` - 显示详细信息，但不显示所属组信息
- `-p`，`--indicator-style=slash` - 添加 `/` 指示器
- `-q`，`--hide-control-chars` - 非 图形字符 显示为 `?`
- `--show-control-chars` - 直接显示非 图形字符
- `-Q`，`--quote-name` - 用双引号包裹文件名
- `--quoting-style=WORD` - 引号风格，有效值：`literal`，`locale`，`shell`，`shell-always`，`shell-escape`，`shell-escape-always`，`c`，`escape`
- `-r`，`--reverse` - 逆序
- `-R`，`--recursive` - 递归查看子目录
- `-s`，`--size` - 显示文件分配的块数
- `-S` - 按文件大小排序，最大的优先
- `--sort=WORD` - 排序规则，有效值：`none`，`size`，`time`，`version`，`extension`，`width`
- `--time=WORD` - 显示或排序的时间，有效值：`atime`，`ctime`，`mtime`，`btime`
- `--time-style=TIME_STYLE` - 日期和时间的格式
- `-t` - 按时间排序
- `-T`，`--tabsize=COLS` - 制表符的宽度
- `-u` - 显示 atime 或按 atime 排序
- `-U` - 不排序
- `-v` - 自然排序
- `-w`，`--width=COLS` - 一行的宽度，0 表示无限
- `-x` - 横向排列
- `-X` - 按扩展名排序
- `-Z`，`--context` - 显示安全上下文
- `--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）

    - `-1` - 每行只显示一个文件
    - `--help` - 显示当前帮助
    - `--version` - 显示版本
- `FILE` 文件，可以省略，默认为当前工作目录

**返回值**

- `0`：成功
- `1`：失败，小问题
- `2`：失败，严重问题

## 文件类型指示器

| 符号 | 含义 | 示例 |
| --- | --- | --- |
| `/` | 目录（directory） | `src/` |
| `@` | 符号链接（symbolic link） | `readme@` |
| `*` | 可执行文件（executable file） | `run.sh*` |
| `=` | UNIX 域套接字（socket） | `mysock=` |
| `\|` | 命名管道（FIFO） | `myfifo\|` |
| `>` | Solaris Door File | `door>` |

## 示例

**基本列出文件**

```
ls
```

- 列出当前目录下的文件（不包括隐藏文件）。
- 输出按字母顺序排序。

```
ls -a
```

- 包括隐藏文件（以 `.` 开头的文件）。

```
ls -1
```

- 每行只显示一个文件名，方便脚本处理。

**彩色与类型指示**

```
ls --color=auto
```

- 根据文件类型显示颜色（目录、可执行文件、链接等）。

```
ls -F
```

- 给文件添加类型符号：
  - `/` 目录
  - `*` 可执行文件
  - `@` 符号链接
  - `|` FIFO
  - `=` socket

```
ls -lF --color=auto
```

- 长列表 + 类型符号 + 彩色显示。

**多列与排序**

```
ls -C
```

- 多列竖向显示（默认行为）。

```
ls -x
```

- 多列横向显示（行填充）。

```
ls -lt
```

- 按 **修改时间** 排序，最近修改的在前。

```
ls -ltu
```

- 按 **访问时间** 排序（-u 选项），最近访问的在前。

```
ls -ltc
```

- 按 **状态改变时间** 排序（-c 选项）。

**文件名原样**

```
ls -N
```

- 文件名按原样显示，不加引号、不转义。

```
ls --show-control-chars
```

- 显示文件名中的控制字符（如换行、制表符等），不替换为 `?`。

**超链接显示（现代终端）**

```
ls --hyperlink=auto
```

- 支持 OSC 8 超链接，在终端中点击文件名可直接打开。

```
ls --hyperlink=always
```

- 强制生成超链接，即使终端不支持也输出控制序列。

**组合示例**

```
ls -lF --color=auto --hyperlink=auto
```

- 长列表 + 类型符号 + 彩色显示 + 可点击链接。
- 最适合现代终端使用。

```
ls -laN --show-control-chars
```

- 查看所有文件（包括隐藏）
- 原样显示文件名
- 显示控制字符

```
ls -ltuF --color=auto
```

- 按访问时间排序
- 显示类型符号
- 彩色显示

```
ls -1U
```

- 每行显示一个文件名
- 不排序（按照目录存储顺序），适合脚本处理。

```
ls -l --quoting-style=c
```

- 长列表
- 文件名用 **C 风格转义**（`\n`、`\t` 等），便于分析特殊字符。

## 推荐阅读

- List directory contents - GNU Coreutils
- ls(1) — Linux manual page

## 手册

```
LS(1)                            User Commands                           LS(1)

NAME
       ls - list directory contents

SYNOPSIS
       ls [OPTION]... [FILE]...

DESCRIPTION
       List  information  about  the FILEs (the current directory by default).
       Sort entries alphabetically if none of -cftuvSUX nor --sort  is  speci-
       fied.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -a, --all
              do not ignore entries starting with .

       -A, --almost-all
              do not list implied . and ..

       --author
              with -l, print the author of each file

       -b, --escape
              print C-style escapes for nongraphic characters

       --block-size=SIZE
              with  -l,  scale  sizes  by  SIZE  when  printing  them;  e.g.,
              '--block-size=M'; see SIZE format below
```

```
-B, --ignore-backups
       do not list implied entries ending with ~

-c     with  -lt: sort by, and show, ctime (time of last change of file
       status information); with -l: show ctime and sort by name;  oth-
       erwise: sort by ctime, newest first

-C     list entries by columns

--color[=WHEN]
       color the output WHEN; more info below

-d, --directory
       list directories themselves, not their contents

-D, --dired
       generate output designed for Emacs' dired mode

-f     list all entries in directory order

-F, --classify[=WHEN]
       append indicator (one of */=>@|) to entries WHEN

--file-type
       likewise, except do not append '*'

--format=WORD
       across  -x, commas -m, horizontal -x, long -l, single-column -1,
       verbose -l, vertical -C

--full-time
       like -l --time-style=full-iso

-g     like -l, but do not list owner

--group-directories-first
       group directories before files; can be augmented with  a  --sort
       option, but any use of --sort=none (-U) disables grouping

-G, --no-group
       in a long listing, don't print group names

-h, --human-readable
       with -l and -s, print sizes like 1K 234M 2G etc.

--si   likewise, but use powers of 1000 not 1024

-H, --dereference-command-line
       follow symbolic links listed on the command line

--dereference-command-line-symlink-to-dir
       follow  each  command line symbolic link that points to a direc-
       tory
```

```
--hide=PATTERN
       do not list implied entries matching shell  PATTERN  (overridden
       by -a or -A)

--hyperlink[=WHEN]
       hyperlink file names WHEN

--indicator-style=WORD
       append indicator with style WORD to entry names: none (default),
       slash (-p), file-type (--file-type), classify (-F)

-i, --inode
       print the index number of each file

-I, --ignore=PATTERN
       do not list implied entries matching shell PATTERN

-k, --kibibytes
       default  to  1024-byte  blocks  for file system usage; used only
       with -s and per directory totals

-l     use a long listing format

-L, --dereference
       when showing file information for a symbolic link, show informa-
       tion for the file the link references rather than for  the  link
       itself

-m     fill width with a comma separated list of entries

-n, --numeric-uid-gid
       like -l, but list numeric user and group IDs

-N, --literal
       print entry names without quoting

-o     like -l, but do not list group information

-p, --indicator-style=slash
       append / indicator to directories

-q, --hide-control-chars
       print ? instead of nongraphic characters

--show-control-chars
       show nongraphic characters as-is (the default, unless program is
       'ls' and output is a terminal)

-Q, --quote-name
       enclose entry names in double quotes

--quoting-style=WORD
```

```
       use  quoting style WORD for entry names: literal, locale, shell,
       shell-always,  shell-escape,  shell-escape-always,  c,  escape
       (overrides QUOTING_STYLE environment variable)

-r, --reverse
       reverse order while sorting

-R, --recursive
       list subdirectories recursively

-s, --size
       print the allocated size of each file, in blocks

-S     sort by file size, largest first

--sort=WORD
       sort  by  WORD instead of name: none (-U), size (-S), time (-t),
       version (-v), extension (-X), width

--time=WORD
       select which timestamp used to  display  or  sort;  access  time
       (-u): atime, access, use; metadata change time (-c): ctime, sta-
       tus;  modified  time (default): mtime, modification; birth time:
       birth, creation;

       with -l, WORD determines which time to show;  with  --sort=time,
       sort by WORD (newest first)

--time-style=TIME_STYLE
       time/date format with -l; see TIME_STYLE below

-t     sort by time, newest first; see --time

-T, --tabsize=COLS
       assume tab stops at each COLS instead of 8

-u     with  -lt:  sort by, and show, access time; with -l: show access
       time and sort by name; otherwise: sort by  access  time,  newest
       first

-U     do not sort; list entries in directory order

-v     natural sort of (version) numbers within text

-w, --width=COLS
       set output width to COLS.  0 means no limit

-x     list entries by lines instead of by columns

-X     sort alphabetically by entry extension

-Z, --context
       print any security context of each file
```

--zero end each output line with NUL, not newline

-1 list one file per line

--help display this help and exit

--version
output version information and exit

The SIZE argument is an integer and optional unit (example: 10K is 10*1024). Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or KB,MB,... (powers of 1000). Binary prefixes can be used, too: KiB=K, MiB=M, and so on.

The TIME_STYLE argument can be full-iso, long-iso, iso, locale, or +FORMAT. FORMAT is interpreted like in date(1). If FORMAT is FORMAT1<newline>FORMAT2, then FORMAT1 applies to non-recent files and FORMAT2 to recent files. TIME_STYLE prefixed with 'posix-' takes effect only outside the POSIX locale. Also the TIME_STYLE environment variable sets the default style to use.

The WHEN argument defaults to 'always' and can also be 'auto' or 'never'.

Using color to distinguish file types is disabled both by default and with --color=never. With --color=auto, ls emits color codes only when standard output is connected to a terminal. The LS_COLORS environment variable can change the settings. Use the dircolors(1) command to set it.

Exit status:
0 if OK,

1 if minor problems (e.g., cannot access subdirectory),

2 if serious trouble (e.g., cannot access command-line argument).

AUTHOR
Written by Richard M. Stallman and David MacKenzie.

REPORTING BUGS
GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
dircolors(1)

```
       Full documentation <https://www.gnu.org/software/coreutils/ls>
       or available locally via: info '(coreutils) ls invocation'

GNU coreutils 9.4                 April 2024                              LS(1)
```

# Bash 的 lsblk 命令

```
lsblk [OPTION]... [DEVICE]...
```

**功能**

列出块设备。

**类型**

可执行文件（`/usr/bin/lsblk`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-A`，`--noempty` - 不显示空的块设备
  - `-a`，`--all` - 显示所有块设备
  - `-b`，`--bytes` - 使用字节作为单位
  - `-D`，`--discard` - 显示设备的丢弃能力（TRIM，UNMAP）
  - `-d`，`--nodeps` - 不显示从设备
  - `-E`，`--dedup column` - 通过 `column` 列进行去重
  - `-e`，`--exclude list` - 排除 `list` 列出的设备，以逗号分隔
  - `-f`，`--fs` - 显示文件系统信息
  - `-I`，`--include list` - 包含 `list` 列出的设备，以逗号分隔
  - `-i`，`--ascii` - 使用 ASCII 字符显示树状结构；默认为以 制表符 显示树状结构
  - `-J`，`--json` - 以 JSON 格式输出
  - `-l`，`--list` - 以列表（而非树状）格式输出
  - `-M`，`--merge` - 将子树的父节点分组
  - `-m`，`--perms` - 显示设备的权限信息
  - `-N`，`--nvme` - 只显示 NVMe 设备
  - `-v`，`--virtio` - 只显示 VirtIO 设备
  - `-n`，`--noheadings` - 不显示标题行（列名）
  - `-o`，`--output list` - 指定要输出的列
  - `-O`，`--output-all` - 输出所有有效列
  - `-P`，`--pairs` - 以键值对的格式输出
  - `-p`，`--paths` - 显示设备路径
  - `-r`，`--raw` - 以原始格式输出

    - `-S`, `--scsi` - 仅显示 SCSI 设备
    - `-s`, `--inverse` - 逆序输出
    - `-T`, `--tree[=column]` - 显示树状结构，`column` 指定显示为树状的列
    - `-t`, `--topology` - 显示拓扑信息
    - `-w`, `--width number` - 指定输出宽度
    - `-x`, `--sort column` - 以 `column` 指定的列排序
    - `-y`, `--shell` - 使列名只包含 SHELL 变量标识符允许的字符
    - `-z`, `--zoned` - 显示设备的区域相关信息
    - `--sysroot directory` - 指定根目录，可用于查看来自其它系统的 udev 信息
    - `--help` - 显示当前帮助
    - `--version` - 显示版本
- `DEVICE` - 要查看的设备列表，默认为查看所有设备

**返回值**

- 0 - 成功
- 1 - 失败
- 32 - 未找到指定设备
- 64 - 未找到部分指定设备

## 示例

```
$ lsblk             # 查看块设备
NAME        MAJ:MIN RM   SIZE RO TYPE MOUNTPOINTS
nvme0n1     259:0    0 931.5G  0 disk
├─nvme0n1p1 259:1    0     1G  0 part /boot/efi
└─nvme0n1p2 259:2    0 930.5G  0 part /
$ lsblk -d          # 不显示从设备
NAME    MAJ:MIN RM   SIZE RO TYPE MOUNTPOINTS
nvme0n1 259:0    0 931.5G  0 disk
$ lsblk -a          # 显示所有设备
NAME        MAJ:MIN RM   SIZE RO TYPE MOUNTPOINTS
loop0         7:0    0     0B  0 loop
loop1         7:1    0     0B  0 loop
loop2         7:2    0     0B  0 loop
loop3         7:3    0     0B  0 loop
loop4         7:4    0     0B  0 loop
loop5         7:5    0     0B  0 loop
loop6         7:6    0     0B  0 loop
loop7         7:7    0     0B  0 loop
nvme0n1     259:0    0 931.5G  0 disk
├─nvme0n1p1 259:1    0     1G  0 part /boot/efi
└─nvme0n1p2 259:2    0 930.5G  0 part /
```

# 推荐阅读

- lsblk(8) — Linux manual page

# 手册

```
LSBLK(8)                    System Administration                    LSBLK(8)

NAME
       lsblk - list block devices

SYNOPSIS
       lsblk [options] [device...]

DESCRIPTION
       lsblk lists information about all available or the specified block
       devices. The lsblk command reads the sysfs filesystem and udev db to
       gather information. If the udev db is not available or lsblk is
       compiled without udev support, then it tries to read LABELs, UUIDs and
       filesystem types from the block device. In this case root permissions
       are necessary.

       By default, the command prints all block devices (except RAM disks) in
       a tree-like format. The same device can be repeated in the tree if it
       relates to other devices. The --merge option is recommended for more
       complicated setups to gather groups of devices and describe complex N:M
       relationships.

       The default output, as well as the default output from options like
       --fs and --topology, is subject to change. So whenever possible, you
       should avoid using default outputs in your scripts. Always explicitly
       define expected columns by using --output columns-list and --list in
       environments where a stable output is required.

       Use lsblk --help to get a list of all available columns.

       Note that lsblk might be executed in time when udev does not have all
       information about recently added or modified devices yet. In this case
       it is recommended to use udevadm settle before lsblk to synchronize
       with udev.

       The relationship between block devices and filesystems is not always
       one-to-one. The filesystem may use more block devices, or the same
       filesystem may be accessible by more paths. This is the reason why
       lsblk provides MOUNTPOINT and MOUNTPOINTS (pl.) columns. The column
       MOUNTPOINT displays only one mount point (usually the last mounted
       instance of the filesystem), and the column MOUNTPOINTS displays by
       multi-line cell all mount points associated with the device.

OPTIONS
```

```
-A, --noempty
    Don’t print empty devices.

-a, --all
    Disable all built-in filters and list all empty devices and RAM
    disk devices too.

-b, --bytes
    Print the sizes in bytes rather than in a human-readable format.

    By default, the unit, sizes are expressed in, is byte, and unit
    prefixes are in power of 2^10 (1024). Abbreviations of symbols are
    exhibited truncated in order to reach a better readability, by
    exhibiting alone the first letter of them; examples: "1 KiB" and "1
    MiB" are respectively exhibited as "1 K" and "1 M", then omitting
    on purpose the mention "iB", which is part of these abbreviations.

-D, --discard
    Print information about the discarding capabilities (TRIM, UNMAP)
    for each device.

-d, --nodeps
    Do not print holder devices or slaves. For example, lsblk --nodeps
    /dev/sda prints information about the sda device only.

-E, --dedup column
    Use column as a de-duplication key to de-duplicate output tree. If
    the key is not available for the device, or the device is a
    partition and parental whole-disk device provides the same key than
    the device is always printed.

    The usual use case is to de-duplicate output on system multi-path
    devices, for example by -E WWN.

-e, --exclude list
    Exclude the devices specified by the comma-separated list of major
    device numbers. Note that RAM disks (major=1) are excluded by
    default if --all is not specified. The filter is applied to the
    top-level devices only. This may be confusing for --list output
    format where hierarchy of the devices is not obvious.

-f, --fs
    Output info about filesystems. This option is equivalent to -o
    NAME,FSTYPE,FSVER,LABEL,UUID,FSAVAIL,FSUSE%,MOUNTPOINTS. The
    authoritative information about filesystems and raids is provided
    by the blkid(8) command.

-I, --include list
    Include devices specified by the comma-separated list of major
    device numbers. The filter is applied to the top-level devices
    only. This may be confusing for --list output format where
    hierarchy of the devices is not obvious.
```

```
-i, --ascii
    Use ASCII characters for tree formatting.

-J, --json
    Use JSON output format. It’s strongly recommended to use --output
    and also --tree if necessary.

-l, --list
    Produce output in the form of a list. The output does not provide
    information about relationships between devices and since version
    2.34 every device is printed only once if --pairs or --raw not
    specified (the parsable outputs are maintained in backwardly
    compatible way).

-M, --merge
    Group parents of sub-trees to provide more readable output for
    RAIDs and Multi-path devices. The tree-like output is required.

-m, --perms
    Output info about device owner, group and mode. This option is
    equivalent to -o NAME,SIZE,OWNER,GROUP,MODE.

-N, --nvme
    Output info about NVMe devices only.

-v, --virtio
    Output info about virtio devices only.

-n, --noheadings
    Do not print a header line.

-o, --output list
    Specify which output columns to print. Use --help to get a list of
    all supported columns. The columns may affect tree-like output. The
    default is to use tree for the column 'NAME' (see also --tree).

    The default list of columns may be extended if list is specified in
    the format +list (e.g., lsblk -o +UUID).

-O, --output-all
    Output all available columns.

-P, --pairs
    Produce output in the form of key="value" pairs. The output lines
    are still ordered by dependencies. All potentially unsafe value
    characters are hex-escaped (\x<code>). See also option --shell.

-p, --paths
    Print full device paths.

-r, --raw
    Produce output in raw format. The output lines are still ordered by
    dependencies. All potentially unsafe characters are hex-escaped
```

(\x<code>) in the NAME, KNAME, LABEL, PARTLABEL and MOUNTPOINT
columns.

-S, --scsi
Output info about SCSI devices only. All partitions, slaves and
holder devices are ignored.

-s, --inverse
Print dependencies in inverse order. If the --list output is
requested then the lines are still ordered by dependencies.

-T, --tree[=column]
Force tree-like output format. If column is specified, then a tree
is printed in the column. The default is NAME column.

-t, --topology
Output info about block-device topology. This option is equivalent
to

-o
NAME,ALIGNMENT,MIN-IO,OPT-IO,PHY-SEC,LOG-SEC,ROTA,SCHED,RQ-SIZE,RA,WSAME.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

-w, --width number
Specifies output width as a number of characters. The default is
the number of the terminal columns, and if not executed on a
terminal, then output width is not restricted at all by default.
This option also forces lsblk to assume that terminal control
characters and unsafe characters are not allowed. The expected
use-case is for example when lsblk is used by the watch(1) command.

-x, --sort column
Sort output lines by column. This option enables --list output
format by default. It is possible to use the option --tree to force
tree-like output and than the tree branches are sorted by the
column.

-y, --shell
The column name will be modified to contain only characters allowed
for shell variable identifiers, for example, MIN_IO and FSUSE_PCT
instead of MIN-IO and FSUSE%. This is usable, for example, with
--pairs. Note that this feature has been automatically enabled for
--pairs in version 2.37, but due to compatibility issues, now it’s
necessary to request this behavior by --shell.

-z, --zoned
Print the zone related information for each device.

```
       --sysroot directory
           Gather data for a Linux instance other than the instance from which
           the lsblk command is issued. The specified directory is the system
           root of the Linux instance to be inspected. The real device nodes
           in the target directory can be replaced by text files with udev
           attributes.

EXIT STATUS
       0
           success

       1
           failure

       32
           none of specified devices found

       64
           some specified devices found, some not found

ENVIRONMENT
       LSBLK_DEBUG=all
           enables lsblk debug output.

       LIBBLKID_DEBUG=all
           enables libblkid debug output.

       LIBMOUNT_DEBUG=all
           enables libmount debug output.

       LIBSMARTCOLS_DEBUG=all
           enables libsmartcols debug output.

       LIBSMARTCOLS_DEBUG_PADDING=on
           use visible padding characters.

NOTES
       For partitions, some information (e.g., queue attributes) is inherited
       from the parent device.

       The lsblk command needs to be able to look up each block device by
       major:minor numbers, which is done by using /sys/dev/block. This sysfs
       block directory appeared in kernel 2.6.27 (October 2008). In case of
       problems with a new enough kernel, check that CONFIG_SYSFS was enabled
       at the time of the kernel build.

AUTHORS
       Milan Broz <gmazyland@gmail.com>, Karel Zak <kzak@redhat.com>

SEE ALSO
       ls(1), blkid(8), findmnt(8)

REPORTING BUGS
```

```
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The lsblk command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                          LSBLK(8)
```

# Bash 的 lscpu 命令

```
lscpu [OPTION]... [DEVICE]...
```

**功能**

列出 CPU 信息。

**类型**

可执行文件（`/usr/bin/lscpu`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 显示所有 CPU；必须和 `--extended` 或 `--parse` 一起使用
  - `-B`，`--bytes` - 使用字节作为单位
  - `-b`，`--online` - 仅显示在线的 CPU；必须和 `--extended` 或 `--parse` 一起使用
  - `-C`，`--caches[=list]` - 显示 CPU 缓存信息，可用通过 `list` 指定输出的列名列表（以逗号分隔）
  - `-c`，`--offline` - 仅显示离线的 CPI；必须和 `--extended` 或 `--parse` 一起使用
  - `--hierarchic[=when]` - 在摘要输出中使用子小节
  - `-e`，`--extended[=list]` - 以易于理解的格式输出，可用通过 `list` 指定输出的列名列表（以逗号分隔）
  - `-J`，`--json` - 以 JSON 格式输出
  - `-p`，`--parse[=list]` - 优化命令输出以便于解析，可用通过 `list` 指定输出的列名列表（以逗号分隔）
  - `-s`，`--sysroot directory` - 指定根目录为 `directory`，可用于查看来自其它系统的 udev 信息
  - `-x`，`--hex` - 使用十六进制掩码表示 CPU 集
  - `-y`，`--physical` - 显示所有包含拓扑元素的列的物理 ID
  - `--output-all` - 显示所有可用的列
  - `--help` - 显示当前帮助
  - `--version` - 显示版本
- `DEVICE` - 要查看的设备列表，默认为查看所有设备

## 列名

## `-e` 和 `-p` 选项可用的列名

| 列名 | 说明 |
| --- | --- |
| BOGOMIPS | CPU 速度的粗略策略 |
| CPU | 逻辑 CPU 编号 |
| CORE | 逻辑核心编号 |
| SOCKET | 逻辑插槽编号 |
| CLUSTER | 逻辑簇编号 |
| NODE | 逻辑非一致性内存访问（NUMA）编号 |
| BOOK | 逻辑书本（包含多个 NUMA）编号 |
| DRAWER | 逻辑抽屉（包含多个 Book）编号 |
| CACHE | CPU 之间如何共享缓存 |
| POLARIZATION | 虚拟硬件上的 CPU 调度模式 |
| ADDRESS | CPU 的物理地址 |
| CONFIGURED | 虚拟机管理程序（Hypervisor）是否已经分配 CPU |
| ONLINE | Linux 当前是否使用 CPU |
| MHZ | CPU 当前频率（MHz） |
| SCALMHZ% | CPU 频率缩放百分比 |
| MAXMHZ | CPU 最高频率（MHz） |
| MINMHZ | CPU 最低频率（MHz） |
| MODELNAME | CPU 型号名称 |

## `-C` 选项可用的列名

| 列名 | 说明 |
| --- | --- |
| ALL-SIZE | 所有系统缓存的大小 |
| LEVEL | 缓存级别 |
| NAME | 缓存名称 |
| ONE-SIZE | 一个缓存的大小 |
| TYPE | 缓存类型 |

| 列名 | 说明 |
| --- | --- |
| WAYS | 缓存的关联 |
| ALLOC-POLICY | 缓存的分配策略 |
| WRITE-POLICY | 缓存的写策略 |
| PHY-LINE | 每个缓存标配对应的对应的物理缓存行数 |
| SETS | 缓存中集合的数量 |
| COHERENCY-SIZE | 从内存传输到缓存的最小数据量 |

## 示例

```
$ lscpu
Architecture:             x86_64
  CPU op-mode(s):         32-bit, 64-bit
  Address sizes:          43 bits physical, 48 bits virtual
  Byte Order:             Little Endian
CPU(s):                   8
  On-line CPU(s) list:    0-7
Vendor ID:                AuthenticAMD
  Model name:             AMD Ryzen 5 3400G with Radeon Vega Graphics
    CPU family:           23
    Model:                24
    Thread(s) per core:   2
    Core(s) per socket:   4
    Socket(s):            1
    Stepping:             1
    Frequency boost:      enabled
    CPU(s) scaling MHz:   36%
    CPU max MHz:          3700.0000
    CPU min MHz:          1400.0000
    BogoMIPS:             7385.99
    Flags:                fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov
pat pse36 clflush mmx fxsr sse sse2 ht syscall nx mmxext fxsr_opt pdpe1gb rdtscp lm
constant_tsc rep_good nopl nonstop_tsc cpuid extd_apicid a
                          perfmperf rapl pni pclmulqdq monitor ssse3 fma cx16 sse4_1
sse4_2 movbe popcnt aes xsave avx f16c rdrand lahf_lm cmp_legacy svm extapic cr8_legacy
abm sse4a misalignsse 3dnowprefetch osvw skinit wdt tce t
                          opoext perfctr_core perfctr_nb bpext perfctr_llc mwaitx cpb
hw_pstate ssbd ibpb vmmcall fsgsbase bmi1 avx2 smep bmi2 rdseed adx smap clflushopt sha_ni
xsaveopt xsavec xgetbv1 clzero irperf xsaveerptr arat
                           npt lbrv svm_lock nrip_save tsc_scale vmcb_clean flushbyasid
decodeassists pausefilter pfthreshold avic v_vmsave_vmload vgif overflow_recov succor smca
sev sev_es ibpb_exit_to_user
Virtualization features:
  Virtualization:         AMD-V
Caches (sum of all):
```

```
  L1d:                    128 KiB (4 instances)
  L1i:                    256 KiB (4 instances)
  L2:                     2 MiB (4 instances)
  L3:                     4 MiB (1 instance)
NUMA:
  NUMA node(s):           1
  NUMA node0 CPU(s):      0-7
Vulnerabilities:
  Gather data sampling:   Not affected
  Itlb multihit:          Not affected
  L1tf:                   Not affected
  Mds:                    Not affected
  Meltdown:               Not affected
  Mmio stale data:        Not affected
  Reg file data sampling: Not affected
  Retbleed:               Mitigation; untrained return thunk; SMT vulnerable
  Spec rstack overflow:   Mitigation; Safe RET
  Spec store bypass:      Mitigation; Speculative Store Bypass disabled via prctl
  Spectre v1:             Mitigation; usercopy/swapgs barriers and __user pointer
sanitization
  Spectre v2:             Mitigation; Retpolines; IBPB conditional; STIBP disabled; RSB
filling; PBRSB-eIBRS Not affected; BHI Not affected
  Srbds:                  Not affected
  Tsx async abort:        Not affected
```

## 推荐阅读

- lscpu(1) — Linux manual page

## 手册

```
LSCPU(1)                        User Commands                        LSCPU(1)

NAME
       lscpu - display information about the CPU architecture

SYNOPSIS
       lscpu [options]

DESCRIPTION
       lscpu gathers CPU architecture information from sysfs, /proc/cpuinfo
       and any applicable architecture-specific libraries (e.g. librtas on
       Powerpc). The command output can be optimized for parsing or for easy
       readability by humans. The information includes, for example, the
       number of CPUs, threads, cores, sockets, and Non-Uniform Memory Access
       (NUMA) nodes. There is also information about the CPU caches and cache
       sharing, family, model, bogoMIPS, byte order, and stepping.

       The default output formatting on terminal is subject to change and
```

maybe optimized for better readability. The output for non-terminals (e.g., pipes) is never affected by this optimization and it is always in "Field: data\n" format. Use for example "lscpu | less" to see the default output without optimizations.

In virtualized environments, the CPU architecture information displayed reflects the configuration of the guest operating system which is typically different from the physical (host) system. On architectures that support retrieving physical topology information, lscpu also displays the number of physical sockets, chips, cores in the host system.

Options that result in an output table have a list argument. Use this argument to customize the command output. Specify a comma-separated list of column labels to limit the output table to only the specified columns, arranged in the specified order. See COLUMNS for a list of valid column labels. The column labels are not case sensitive.

Not all columns are supported on all architectures. If an unsupported column is specified, lscpu prints the column but does not provide any data for it.

The cache sizes are reported as summary from all CPUs. The versions before v2.34 reported per-core sizes, but this output was confusing due to complicated CPUs topology and the way how caches are shared between CPUs. For more details about caches see --cache. Since version v2.37 lscpu follows cache IDs as provided by Linux kernel and it does not always start from zero.

OPTIONS

-a, --all
Include lines for online and offline CPUs in the output (default for -e). This option may only be specified together with option -e or -p.

-B, --bytes
Print the sizes in bytes rather than in a human-readable format.

By default, the unit, sizes are expressed in, is byte, and unit prefixes are in power of 2^10 (1024). Abbreviations of symbols are exhibited truncated in order to reach a better readability, by exhibiting alone the first letter of them; examples: "1 KiB" and "1 MiB" are respectively exhibited as "1 K" and "1 M", then omitting on purpose the mention "iB", which is part of these abbreviations.

-b, --online
Limit the output to online CPUs (default for -p). This option may only be specified together with option -e or -p.

-C, --caches[=list]
Display details about CPU caches. For details about available information see --help output.

If the list argument is omitted, all columns for which data is available are included in the command output.

When specifying the list argument, the string of option, equal sign (=), and list must not contain any blanks or other whitespace. Examples: -C=NAME,ONE-SIZE or --caches=NAME,ONE-SIZE.

The default list of columns may be extended if list is specified in the format +list (e.g., lscpu -C=+ALLOC-POLICY).

-c, --offline
Limit the output to offline CPUs. This option may only be specified together with option -e or -p.

--hierarchic[=when]
Use subsections in summary output. For backward compatibility, the default is to use subsections only when output on a terminal and flattened output on a non-terminal. The optional argument when can be never, always or auto. If the when argument is omitted, it defaults to "always".

-e, --extended[=list]
Display the CPU information in human-readable format.

If the list argument is omitted, the default columns are included in the command output. The default output is subject to change.

When specifying the list argument, the string of option, equal sign (=), and list must not contain any blanks or other whitespace. Examples: '-e=cpu,node' or '--extended=cpu,node'.

The default list of columns may be extended if list is specified in the format +list (e.g., lscpu -e=+MHZ).

-J, --json
Use JSON output format for the default summary or extended output (see --extended). For backward compatibility, JSON output follows the default summary behavior for non-terminals (e.g., pipes) where subsections are missing. See also --hierarchic.

-p, --parse[=list]
Optimize the command output for easy parsing.

If the list argument is omitted, the command output is compatible with earlier versions of lscpu. In this compatible format, two commas are used to separate CPU cache columns. If no CPU caches are identified the cache column is omitted. If the list argument is used, cache columns are separated with a colon (:).

When specifying the list argument, the string of option, equal sign (=), and list must not contain any blanks or other whitespace. Examples: '-p=cpu,node' or '--parse=cpu,node'.

The default list of columns may be extended if list is specified in the format +list (e.g., lscpu -p=+MHZ).

-s, --sysroot directory
Gather CPU data for a Linux instance other than the instance from which the lscpu command is issued. The specified directory is the system root of the Linux instance to be inspected.

-x, --hex
Use hexadecimal masks for CPU sets (for example "ff"). The default is to print the sets in list format (for example 0,1). Note that before version 2.30 the mask has been printed with 0x prefix.

-y, --physical
Display physical IDs for all columns with topology elements (core, socket, etc.). Other than logical IDs, which are assigned by lscpu, physical IDs are platform-specific values that are provided by the kernel. Physical IDs are not necessarily unique and they might not be arranged sequentially. If the kernel could not retrieve a physical ID for an element lscpu prints the dash (-) character.

The CPU logical numbers are not affected by this option.

--output-all
Output all available columns. This option must be combined with either --extended, --parse or --caches.

## BUGS

The basic overview of CPU family, model, etc. is always based on the first CPU only.

Sometimes in Xen Dom0 the kernel reports wrong data.

On virtual hardware the number of cores per socket, etc. can be wrong.

## AUTHORS

Cai Qian <qcai@redhat.com>, Karel Zak <kzak@redhat.com>, Heiko Carstens <heiko.carstens@de.ibm.com>

## SEE ALSO

chcpu(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The lscpu command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 lsipc 命令

```
lsipc [OPTION]...
```

**功能**

列出进程间通信（IPC）设施的信息。

**类型**

可执行文件（`/usr/bin/lsipc`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-i`，`--id id` - 要查看的资源 ID
  - `-g`，`--global` - 查看系统范围内的所有 IPC 资源使用情况和限制
  - `-m`，`--shmems` - 仅查看共享内存
  - `-q`，`--queues` - 仅查看消息队列
  - `-s`，`--semaphores` - 仅查看信号量组
  - `-c`，`--creator` - 显示资源的创建者和所有者
  - `-e`，`--export` - 以键值对（key="value"）的格式输出
  - `-J`，`--json` - 以 JSON 格式输出
  - `-l`，`--list` - 以列表格式输出（默认）
  - `-n`，`--newline` - 每条信息单独显示一行
  - `--noheadings` - 不显示标题行（列名）
  - `--notruncate` - 不要截断输出
  - `-o`，`--output list` - 指定要显示的 列名，用逗号分隔
  - `-b`，`--bytes` - 以字节作为单位
  - `-r`，`--raw` - 以原始格式输出
  - `-t`，`--time` - 显示时间信息
  - `--time-format type` - 指定时间格式
  - `-P`，`--numeric-perms` - 显示权限
  - `-y`，`--shell` - 使列名只包含 SHELL 变量标识符允许的字符
  - `--help` - 显示当前帮助
  - `--version` - 显示版本

**返回值**

- 0 - 成功
- 1 - 参数错误
- 2 - 严重错误

# 列名

## 通用列

- `KEY` - 资源的 key
- `ID` - 资源的 ID
- `OWNER` - 所有者的用户名或 UID
- `PERMS` - 权限
- `CUID` - 创建者的 UID
- `CUSER` - 创建者用户
- `CGID` - 创建者的 GID
- `CGROUP` - 创建者组
- `UID` - 用户 ID
- `USER` - 用户名
- `GID` - 组 ID
- `GROUP` - 组名
- `CTIME` - 最后变更时间

## 共享内存（`--shmems`）的列

- `SIZE` - 分段大小
- `NATTCH` - 附加进程数量
- `STATUS` - 状态
- `ATTACH` - 附加时间
- `DETACH` - 分离时间
- `COMMAND` - 创建者命令行
- `CPID` - 创建者的 PID
- `LPID` - 最后一个用户的 PID

## 消息队列（`--queues`）的列

- `USEDBYTES` - 使用的字节数
- `MSGS` - 消息数量
- `SEND` - 最后发送消息的时间
- `RECV` - 最后接收消息的时间

- `LSPID` - 最后发送消息者的 PID
- `LRPID` - 最后接收消息者的 PID

### 信号量（`--semaphores`）的列

- `NSEMS` - 信号量的数量
- `OTIME` - 最后一次操作的时间

### 摘要（`--global`）的列

- `RESOURCE` - 资源名称
- `DESCRIPTION` - 资源描述
- `LIMIT` - 系统范围限制
- `USED` - 当前占用
- `USE%` - 当前占用百分比

## 示例

```
$ lsipc
RESOURCE DESCRIPTION                                    LIMIT USED  USE%
MSGMNI   Number of message queues                       32000    0 0.00%
MSGMAX   Max size of message (bytes)                       8K    -     -
MSGMNB   Default max size of queue (bytes)                16K    -     -
SHMMNI   Shared memory segments                          4096    0 0.00%
SHMALL   Shared memory pages                 18446744073692774399    0 0.00%
SHMMAX   Max size of shared memory segment (bytes)           16E    -     -
SHMMIN   Min size of shared memory segment (bytes)            1B    -     -
SEMMNI   Number of semaphore identifiers                 32000    0 0.00%
SEMMNS   Total number of semaphores                   1024000000    0 0.00%
SEMMSL   Max semaphores per semaphore set.               32000    -     -
SEMOPM   Max number of operations per semop(2)             500    -     -
SEMVMX   Semaphore max value                             32767    -     -
```

## 推荐阅读

- lsipc(1) — Linux manual page

## 手册

```
LSIPC(1)                        User Commands                        LSIPC(1)

NAME
       lsipc - show information on IPC facilities currently employed in the
```

```
       system

SYNOPSIS
       lsipc [options]

DESCRIPTION
       lsipc shows information on the System V inter-process communication
       facilities for which the calling process has read access.

OPTIONS
       -i, --id id
           Show full details on just the one resource element identified by
           id. This option needs to be combined with one of the three resource
           options: -m, -q or -s. It is possible to override the default
           output format for this option with the --list, --raw, --json or
           --export option.

       -g, --global
           Show system-wide usage and limits of IPC resources. This option may
           be combined with one of the three resource options: -m, -q or -s.
           The default is to show information about all resources.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

   Resource options
       -m, --shmems
           Write information about active shared memory segments.

       -q, --queues
           Write information about active message queues.

       -s, --semaphores
           Write information about active semaphore sets.

   Output formatting
       -c, --creator
           Show creator and owner.

       -e, --export
           Produce output in the form of key="value" pairs. All potentially
           unsafe value characters are hex-escaped (\x<code>). See also option
           --shell.

       -J, --json
           Use the JSON output format.

       -l, --list
           Use the list output format. This is the default, except when --id
           is used.
```

```
       -n, --newline
           Display each piece of information on a separate line.

       --noheadings
           Do not print a header line.

       --notruncate
           Don’t truncate output.

       -o, --output list
           Specify which output columns to print. Use --help to get a list of
           all supported columns.

       -b, --bytes
           Print size in bytes rather than in human readable format.

       -r, --raw
           Raw output (no columnation).

       -t, --time
           Write time information. The time of the last control operation that
           changed the access permissions for all facilities, the time of the
           last msgsnd(2) and msgrcv(2) operations on message queues, the time
           of the last shmat(2) and shmdt(2) operations on shared memory, and
           the time of the last semop(2) operation on semaphores.

       --time-format type
           Display dates in short, full or iso format. The default is short,
           this time format is designed to be space efficient and human
           readable.

       -P, --numeric-perms
           Print numeric permissions in PERMS column.

       -y, --shell
           The column name will be modified to contain only characters allowed
           for shell variable identifiers. This is usable, for example, with
           --export. Note that this feature has been automatically enabled for
           --export in version 2.37, but due to compatibility issues, now it’s
           necessary to request this behavior by --shell.

EXIT STATUS
       0
           if OK,

       1
           if incorrect arguments specified,

       2
           if a serious error occurs.

HISTORY
```

The lsipc utility is inspired by the ipcs(1) utility.

## AUTHORS

Ondrej Oprala <ooprala@redhat.com>, Karel Zak <kzak@redhat.com>

## SEE ALSO

ipcmk(1), ipcrm(1), msgrcv(2), msgsnd(2), semget(2), semop(2), shmat(2), shmdt(2), shmget(2), sysvipc(7)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The lsipc command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 lslocks 命令

```
lslocks [OPTION]...
```

**功能**

列出系统中的锁。

**类型**

可执行文件（`/usr/bin/lslocks`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-b`，`--bytes` - 使用字节作为单位
  - `-i`，`--noinaccessible` - 忽略无法访问的锁
  - `-J`，`--json` - 以 JSON 格式输出
  - `-n`，`--noheadings` - 不显示标题行（列名）
  - `-o`，`--output list` - 指定要打印的列，参考 列名
  - `--output-all` - 打印所有可用的列
  - `-p`，`--pid pid` - 只打印 `PID` 指定进程持有的锁
  - `-r`，`--raw` - 以原始格式输出
  - `-u`，`--notruncate` - 不截断列中的文本
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本

## 列名

| 列名 | 说明 |
| --- | --- |
| `COMMAND` | 持有锁的进程命令 |
| `PID` | 持有锁的进程 PID |
| `TYPE` | 锁的类型 |
| `SIZE` | 锁的大小 |
| `INODE` | inode 节点号 |

| 列名 | 说明 |
| --- | --- |
| `MAJ:MIN` | 主要设备号:次要设备号 |
| `MODE` | 锁的访问模式 |
| `M` | 是（1）否（0）为 强制性锁 |
| `START` | 锁的相对字节偏移量 |
| `END` | 锁末端的偏移量 |
| `PATH` | 锁文件的路径 |
| `BLOCKER` | 被锁阻塞进程的 PID |

## 示例

```
$ $ lslocks
COMMAND             PID  TYPE SIZE MODE  M START        END PATH
mysqld           758338 POSIX      WRITE 0     0          0 /...
mysqld           758338 POSIX      WRITE 0     0          0 /...
mysqld           758338 POSIX      WRITE 0     0          0 /...
mysqld          2844183 POSIX      WRITE 0     0          0 /...
mysqld          2844183 POSIX      WRITE 0     0          0 /...
mysqld          2844183 POSIX      WRITE 0     0          0 /...
mysqld          2844183 POSIX      WRITE 0     0          0 /...
mysqld          2844183 POSIX      WRITE 0     0          0 /...
dockerd            1054 FLOCK      WRITE 0     0          0 /...
dockerd            1054 FLOCK      WRITE 0     0          0 /...
dockerd            1054 FLOCK      WRITE 0     0          0 /...
cron            2844018 FLOCK      WRITE 0     0          0 /run...
...
```

## 推荐阅读

- lslocks(8) — Linux manual page

## 手册

```
LSLOCKS(8)                  System Administration                  LSLOCKS(8)

NAME
       lslocks - list local system locks

SYNOPSIS
       lslocks [options]
```

DESCRIPTION

lslocks lists information about all the currently held file locks in a Linux system.

Note that lslocks also lists OFD (Open File Description) locks, these locks are not associated with any process (PID is -1). OFD locks are associated with the open file description on which they are acquired. This lock type is available since Linux 3.15, see fcntl(2) for more details.

OPTIONS

-b, --bytes

Print the sizes in bytes rather than in a human-readable format.

By default, the unit, sizes are expressed in, is byte, and unit prefixes are in power of 2^10 (1024). Abbreviations of symbols are exhibited truncated in order to reach a better readability, by exhibiting alone the first letter of them; examples: "1 KiB" and "1 MiB" are respectively exhibited as "1 K" and "1 M", then omitting on purpose the mention "iB", which is part of these abbreviations.

-i, --noinaccessible

Ignore lock files which are inaccessible for the current user.

-J, --json

Use JSON output format.

-n, --noheadings

Do not print a header line.

-o, --output list

Specify which output columns to print. Use --help to get a list of all supported columns.

The default list of columns may be extended if list is specified in the format +list (e.g., lslocks -o +BLOCKER).

--output-all

Output all available columns.

-p, --pid pid

Display only the locks held by the process with this pid.

-r, --raw

Use the raw output format.

-u, --notruncate

Do not truncate text in columns.

-h, --help

Display help text and exit.

```
       -V, --version
           Print version and exit.

OUTPUT
       COMMAND
           The command name of the process holding the lock.

       PID
           The process ID of the process which holds the lock or -1 for
           OFDLCK.

       TYPE
           The type of lock; can be FLOCK (created with flock(2)), POSIX
           (created with fcntl(2) and lockf(3)) or OFDLCK (created with
           fcntl(2)).

       SIZE
           Size of the locked file.

       MODE
           The lock's access permissions (read, write). If the process is
           blocked and waiting for the lock, then the mode is postfixed with
           an '*' (asterisk).

       M
           Whether the lock is mandatory; 0 means no (meaning the lock is only
           advisory), 1 means yes. (See fcntl(2).)

       START
           Relative byte offset of the lock.

       END
           Ending offset of the lock.

       PATH
           Full path of the lock. If none is found, or there are no
           permissions to read the path, it will fall back to the device's
           mountpoint and "..." is appended to the path. The path might be
           truncated; use --notruncate to get the full path.

       BLOCKER
           The PID of the process which blocks the lock.

NOTES
       The lslocks command is meant to replace the lslk(8) command, originally
       written by Victor A. Abell <abe@purdue.edu> and unmaintained since
       2001.

AUTHORS
       Davidlohr Bueso <dave@gnu.org>

SEE ALSO
       flock(1), fcntl(2), lockf(3)
```

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The lslocks command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                        LSLOCKS(8)
```

# Bash 的 lslogins 命令

```
lslogins [OPTION]... [USERNAME]
```

**功能**

列出系统中的用户信息。

**类型**

可执行文件（`/usr/bin/lslogins`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--acc-expiration` - 显示上次修改密码的时间以及密码过期时间
  - `--btmp-file path` - 指定 btmp 文件的路径
  - `-c`，`--colon-separate` - 使用冒号作为用户之间的分隔符
  - `-e`，`--export` - 以 `NAME=VALUE` 的格式输出
  - `-f`，`--failed` - 显示用户最近一次登录失败的相关信息
  - `-G`，`--supp-groups` - 显示用户的 补充组 信息
  - `-g`，`--groups=groups` - 仅显示属于特定组的用户信息；多个组用逗号分隔
  - `-L`，`--last` - 显示包含用户上次登录会话的信息
  - `-l`，`--logins=logins` - 仅显示特定登录名的用户信息；多个登录名用逗号分隔
  - `-n`，`--newline` - 每条信息单独显示一行
  - `--noheadings` - 不显示标题行（列名）
  - `--notruncate` - 不截断输出
  - `-o`，`--output list` - 只输出指定的列
  - `--output-all` - 输出所有可用的列
  - `-p`，`--pwd` - 显示与密码登录相关的信息
  - `-r`，`--raw` - 以原始格式输出
  - `-s`，`--system-accs` - 显示系统账户
  - `--time-format type` - 指定时间格式：`short`，`full` 或 `iso`；**默认为 `short`**
  - `-u`，`--user-accs` - 显示用户账户
  - `--wtmp-file path` - 指定 wtmp 文件的路径
  - `--lastlog path` - 指定 lastlog 的路径
  - `-y`，`--shell` - 使列名只包含 SHELL 变量标识符允许的字符

    - `-Z`, `--context` - 显示用户的安全上下文
    - `-z`, `--print0` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
    - `-h`, `--help` - 显示帮助
    - `-V`, `--version` - 显示版本
- `USERNAME` - 指定要查看的用户

## 列名

| 列名 | 说明 |
| --- | --- |
| USER | 用户名 |
| UID | 用户 ID |
| GECOS | 完整用户名 |
| HOMEDIR | 用户的主目录 |
| SHELL | 用户的登录 SHELL |
| NOLOGIN | 被 nologin 或 pam_nologin 禁止登录 |
| PWD-LOCK | 密码已设置，但被锁定 |
| PWD-EMPTY | 未设置密码 |
| PWD-DENY | 已禁止密码登录 |
| PWD-METHOD | 密码加密方法 |
| GROUP | 主要组名 |
| GID | 主要组 ID |
| SUPP-GROUPS | 补充组名 |
| SUPP-GIDS | 补充组 ID |
| LAST-LOGIN | 上次登录日期 |
| LAST-TTY | 上次使用的 TTY |
| LAST-HOSTNAME | 上次会话的主机名 |
| FAILED-LOGIN | 上次登录失败的日期 |
| FAILED-TTY | 上次登录失败的 TTY |
| HUSHED | 用户静音设置 |
| PWD-WARN | 密码过期警告剩余天数 |

| 列名 | 说明 |
| --- | --- |
| PWD-CHANGE | 上次更改密码的日期 |
| PWD-MIN | 两次修改密码之间最少需要间隔的天数 |
| PWD-MAX | 密码最多可保持不变的天数 |
| PWD-EXPIR | 密码过期日期 |
| CONTEXT | 用户安全上下文 |
| PROC | 用户运行的进程数 |

## 示例

```
$ lslogins
  UID USER              PROC PWD-LOCK PWD-DENY  LAST-LOGIN GECOS
    0 root               201        0        0             root
    1 daemon               0                                daemon
    2 bin                  0                                bin
    3 sys                  0                                sys
    4 sync                 0                                sync
    5 games                0                                games
    6 man                  0                                man
    7 lp                   0                                lp
    8 mail                 0                                mail
    9 news                 0                                news
...
```

## 推荐阅读

- lslogins(1) — Linux manual page

## 手册

```
LSLOGINS(1)                     User Commands                     LSLOGINS(1)

NAME
       lslogins - display information about known users in the system

SYNOPSIS
       lslogins [options] [-s|-u[=UID]] [-g groups] [-l logins] [username]

DESCRIPTION
       Examine the wtmp and btmp logs, /etc/shadow (if necessary) and /passwd
       and output the desired data.
```

The optional argument username forces lslogins to print all available details about the specified user only. In this case the output format is different than in case of -l or -g and unknown is username reported as an error.

The default action is to list info about all the users in the system.

OPTIONS

Mandatory arguments to long options are mandatory for short options too.

-a, --acc-expiration
Display data about the date of last password change and the account expiration date (see shadow(5) for more info). (Requires root privileges.)

--btmp-file path
Alternate path for btmp.

-c, --colon-separate
Separate info about each user with a colon instead of a newline.

-e, --export
Output data in the format of NAME=VALUE. See also option --shell.

-f, --failed
Display data about the users' last failed login attempts.

-G, --supp-groups
Show information about supplementary groups.

-g, --groups=groups
Only show data of users belonging to groups. More than one group may be specified; the list has to be comma-separated. Unknown group names are ignored.

Note that the relation between user and group may be invisible for the primary group if the user is not explicitly specified as group member (e.g., in /etc/group). If the command lslogins scans for groups then it uses the groups database only, and the user database with primary GID is not used at all.

-L, --last
Display data containing information about the users' last login sessions.

-l, --logins=logins
Only show data of users with a login specified in logins (user names or user IDs). More than one login may be specified; the list has to be comma-separated. Unknown login names are ignored.

-n, --newline

```
    Display each piece of information on a separate line.

--noheadings
    Do not print a header line.

--notruncate
    Don’t truncate output.

-o, --output list
    Specify which output columns to print. The default list of columns
    may be extended if list is specified in the format +list.

--output-all
    Output all available columns. --help to get a list of all supported
    columns.

-p, --pwd
    Display information related to login by password (see also -afL).

-r, --raw
    Raw output (no columnation).

-s, --system-accs
    Show system accounts. These are by default all accounts with a UID
    between 101 and 999 (inclusive), with the exception of either
    nobody or nfsnobody (UID 65534). This hardcoded default may be
    overwritten by parameters SYS_UID_MIN and SYS_UID_MAX in the file
    /etc/login.defs.

--time-format type
    Display dates in short, full or iso format. The default is short,
    this time format is designed to be space efficient and human
    readable.

-u, --user-accs
    Show user accounts. These are by default all accounts with UID
    above 1000 (inclusive), with the exception of either nobody or
    nfsnobody (UID 65534). This hardcoded default maybe overwritten by
    parameters UID_MIN and UID_MAX in the file /etc/login.defs.

-h, --help
    Display help text and exit.

-V, --version
    Print version and exit.

--wtmp-file path
    Alternate path for wtmp.

--lastlog path
    Alternate path for lastlog(8).

-y, --shell
```

```
            The column name will be modified to contain only characters allowed
            for shell variable identifiers. This is usable, for example, with
            --export. Note that this feature has been automatically enabled for
            --export in version 2.37, but due to compatibility issues, now it’s
            necessary to request this behavior by --shell.

        -Z, --context
            Display the users' security context.

        -z, --print0
            Delimit user entries with a nul character, instead of a newline.

EXIT STATUS
        0
            if OK,

        1
            if incorrect arguments specified,

        2
            if a serious error occurs (e.g., a corrupt log).

NOTES
        The default UID thresholds are read from /etc/login.defs.

    Password status
        Multiple fields describe password status.

        "Password is locked"
            The password is prefixed by '!!', and the user cannot login
            although the password is set or empty. This is common for new
            accounts without a set password.

        "Password not required (empty)"
            The password is not set (hash is missing); this is common for
            locked system accounts. Not requiring a password does not mean the
            user can log-in without a password. It depends on the password
            "lock" status.

        "Login by password disabled"
            'yes' means that there is no valid password. The password hash is
            missing, or the hash method is unknown or contains invalid chars.

HISTORY
        The lslogins utility is inspired by the logins utility, which first
        appeared in FreeBSD 4.10.

AUTHORS
        Ondrej Oprala <ooprala@redhat.com>, Karel Zak <kzak@redhat.com>

SEE ALSO
        group(5), passwd(5), shadow(5), utmp(5)
```

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The lslogins command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-30                       LSLOGINS(1)
```

# Bash 的 lsmem 命令

```
lsmem [OPTION]...
```

**功能**

列出可用内存范围及其在线状态。

**类型**

可执行文件（`/usr/bin/lsmem`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 单独列出每个内存块
  - `-b`，`--bytes` - 以字节为单位
  - `-J`，`--json` - 以 JSON 格式输出
  - `-n`，`--noheadings` - 不显示标题行（列名）
  - `-o`，`--output list` - 只输出指定的列
  - `--output-all` - 输出所有列
  - `-P`，`--pairs` - 以键值对（`key="value"`）的格式输出
  - `-r`，`--raw` - 以原始格式输出
  - `-S`，`--split list` - 指定用于分割内存块范围的列
  - `-s`，`--sysroot directory` - 指定根目录，可用于查看来自其它系统的信息
  - `--summary[=when]` - 何时显示摘要行；取值为 `never`，`always` 或 `only`（**默认**）
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本

## 列名

| 列名 | 说明 |
| --- | --- |
| `RANGE` | 内存地址范围 |
| `SIZE` | 内存大小 |
| `STATE` | 内存在线状态 |
| `REMOVABLE` | 内存是否可删除 |

| 列名 | 说明 |
| --- | --- |
| `BLOCK` | 内存块编号或范围 |
| `NODE` | 内存的 NUMA 节点 |
| `ZONES` | 内存范围的有效域 |

## 示例

```
$ lsmem
RANGE                                  SIZE  STATE REMOVABLE  BLOCK
0x0000000000000000-0x000000005fffffff  1.5G online       yes   0-11
0x0000000100000000-0x000000041fffffff 12.5G online       yes 32-131

Memory block size:       128M
Total online memory:      14G
Total offline memory:      0B
$ lsmem --output-all
RANGE                                  SIZE  STATE REMOVABLE  BLOCK NODE  ZONES
0x0000000000000000-0x0000000007ffffff  128M online       yes      0    0   None
0x0000000008000000-0x000000005fffffff  1.4G online       yes   1-11    0  DMA32
0x0000000100000000-0x000000041fffffff 12.5G online       yes 32-131    0 Normal

Memory block size:       128M
Total online memory:      14G
Total offline memory:      0B
```

## 推荐阅读

- lsmem(1) — Linux manual page

## 手册

```
LSMEM(1)                        User Commands                        LSMEM(1)

NAME
       lsmem - list the ranges of available memory with their online status

SYNOPSIS
       lsmem [options]

DESCRIPTION
       The lsmem command lists the ranges of available memory with their
       online status. The listed memory blocks correspond to the memory block
       representation in sysfs. The command also shows the memory block size
       and the amount of memory in online and offline state.
```

The default output is compatible with original implementation from s390-tools, but it’s strongly recommended to avoid using default outputs in your scripts. Always explicitly define expected columns by using the --output option together with a columns list in environments where a stable output is required.

The lsmem command lists a new memory range always when the current memory block distinguish from the previous block by some output column. This default behavior is possible to override by the --split option (e.g., lsmem --split=ZONES). The special word "none" may be used to ignore all differences between memory blocks and to create as large as possible continuous ranges. The opposite semantic is --all to list individual memory blocks.

Note that some output columns may provide inaccurate information if a split policy forces lsmem to ignore differences in some attributes. For example if you merge removable and non-removable memory blocks to the one range than all the range will be marked as non-removable on lsmem output.

Not all columns are supported on all systems. If an unsupported column is specified, lsmem prints the column but does not provide any data for it.

Use the --help option to see the columns description.

OPTIONS

-a, --all
List each individual memory block, instead of combining memory blocks with similar attributes.

-b, --bytes
Print the sizes in bytes rather than in a human-readable format.

By default, the unit, sizes are expressed in, is byte, and unit prefixes are in power of 2^10 (1024). Abbreviations of symbols are exhibited truncated in order to reach a better readability, by exhibiting alone the first letter of them; examples: "1 KiB" and "1 MiB" are respectively exhibited as "1 K" and "1 M", then omitting on purpose the mention "iB", which is part of these abbreviations.

-J, --json
Use JSON output format.

-n, --noheadings
Do not print a header line.

-o, --output list
Specify which output columns to print. Use --help to get a list of all supported columns. The default list of columns may be extended if list is specified in the format +list (e.g., lsmem -o +NODE).

```
       --output-all
           Output all available columns.

       -P, --pairs
           Produce output in the form of key="value" pairs. All potentially
           unsafe value characters are hex-escaped (\x<code>).

       -r, --raw
           Produce output in raw format. All potentially unsafe characters are
           hex-escaped (\x<code>).

       -S, --split list
           Specify which columns (attributes) use to split memory blocks to
           ranges. The supported columns are STATE, REMOVABLE, NODE and ZONES,
           or "none". The other columns are silently ignored. For more details
           see DESCRIPTION above.

       -s, --sysroot directory
           Gather memory data for a Linux instance other than the instance
           from which the lsmem command is issued. The specified directory is
           the system root of the Linux instance to be inspected.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

       --summary[=when]
           This option controls summary lines output. The optional argument
           when can be never, always or only. If the when argument is omitted,
           it defaults to "only". The summary output is suppressed for --raw,
           --pairs and --json.

AUTHORS
       lsmem was originally written by Gerald Schaefer for s390-tools in Perl.
       The C version for util-linux was written by Clemens von Mann, Heiko
       Carstens and Karel Zak.

SEE ALSO
       chmem(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The lsmem command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          LSMEM(1)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 lsns 命令

```
lsns [OPTION]... [NAMESPACE]
```

**功能**

列出系统中的命名空间。

**类型**

可执行文件（`/usr/bin/lsns`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-J`，`--json` - 以 JSON 格式输出
  - `-l`，`--list` - 以列表格式输出
  - `-n`，`--noheadings` - 不显示标题行（列名）
  - `-o`，`--output list` - 仅显示指定的列
  - `--output-all` - 显示所有列
  - `-P`，`--persistent` - 仅显示不包含进程的命名空间
  - `-p`，`--task PID` - 仅显示进程 ID 为 `PID` 的进程持有的命名空间
  - `-r`，`--raw` - 以原始格式输出
  - `-t`，`--type type` - 仅显示指定类型的命名空间
  - `-u`，`--notruncate` - 请勿在列中截断文本
  - `-W`，`--nowrap` - 请勿在列中使用多行文本输出
  - `-T`，`--tree rel` - 以树状格式输出
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `NAMESPACE` - 查看指定的命名空间

## 列名

| 列名 | 说明 |
| --- | --- |
| `NS` | 命名空间标识符 (inode 编号) |
| `TYPE` | 命名空间类型 |

| 列名 | 说明 |
| --- | --- |
| `PATH` | 命名空间的路径 |
| `NPROCS` | 命名空间中的进程数 |
| `PID` | 命名空间中最小的 PID |
| `PPID` | 进程的 PPID（父进程 PID） |
| `COMMAND` | 进程的命令行 |
| `UID` | 进程的 UID |
| `USER` | 进程的用户名 |
| `NETNSID` | 网络子系统使用的命名空间 ID |
| `NSFS` | NSFS 挂载点 |
| `PNS` | 父命名空间标识符 (inode 编号) |
| `ONS` | 所有者命名空间标识符 (inode 编号) |

## 示例

```
$ lsns
        NS TYPE   NPROCS   PID USER     COMMAND
4026531834 time       22 98505 primers  -bash
4026531835 cgroup     22 98505 primers  -bash
4026531836 pid        22 98505 primers  -bash
4026531837 user       22 98505 primers  -bash
4026531838 uts        22 98505 primers  -bash
4026531839 ipc        22 98505 primers  -bash
4026531840 net        22 98505 primers  -bash
4026531841 mnt        22 98505 primers  -bash
```

## 推荐阅读

- lsns(8) — Linux manual page

## 手册

```
LSNS(8)                       System Administration                       LSNS(8)

NAME
       lsns - list namespaces
```

```
SYNOPSIS
       lsns [options] namespace

DESCRIPTION
       lsns lists information about all the currently accessible namespaces or
       about the given namespace. The namespace identifier is an inode number.

       The default output is subject to change. So whenever possible, you
       should avoid using default outputs in your scripts. Always explicitly
       define expected output mode (--tree or --list) and columns by using the
       --output option together with a columns list in environments where a
       stable output is required.

       The NSFS column, printed when net is specified for the --type option,
       is special; it uses multi-line cells. Use the option --nowrap to switch
       to ","-separated single-line representation.

       Note that lsns reads information directly from the /proc filesystem and
       for non-root users it may return incomplete information. The current
       /proc filesystem may be unshared and affected by a PID namespace (see
       unshare --mount-proc for more details). lsns is not able to see
       persistent namespaces without processes where the namespace instance is
       held by a bind mount to /proc/pid/ns/type.

OPTIONS
       -J, --json
           Use JSON output format.

       -l, --list
           Use list output format.

       -n, --noheadings
           Do not print a header line.

       -o, --output list
           Specify which output columns to print. Use --help to get a list of
           all supported columns.

           The default list of columns may be extended if list is specified in
           the format +list (e.g., lsns -o +PATH).

       --output-all
           Output all available columns.

       -P, --persistent
           Display only the namespaces without processes (aka persistent
           namespaces), created by bind mounting /proc/pid/ns/type files to a
           filesystem path.

       -p, --task PID
           Display only the namespaces held by the process with this PID.

       -r, --raw
```

```
           Use the raw output format.

       -t, --type type
           Display the specified type of namespaces only. The supported types
           are mnt, net, ipc, user, pid, uts, cgroup and time. This option may
           be given more than once.

       -u, --notruncate
           Do not truncate text in columns.

       -W, --nowrap
           Do not use multi-line text in columns.

       -T, --tree rel
           Use tree-like output format. If process is given as rel, print
           process tree(s) in each name space. This is default when --tree is
           not specified. If parent is given, print tree(s) constructed by the
           parent/child relationship. If owner is given, print tree(s)
           constructed by the owner/owned relationship. owner is used as
           default when rel is omitted.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       The lsns utility exits with one of the following values:

       0
           Success.

       1
           General error.

       2
           An ioctl was unknown to the kernel.

AUTHORS
       Karel Zak <kzak@redhat.com>

SEE ALSO
       nsenter(1), unshare(1), clone(2), namespaces(7), ioctl_ns(2),
       ip-netns(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The lsns command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
```

<https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 mcookie 命令

```
mcookie [OPTION]...
```

**功能**

生成用作 xauth cookie 的 128 位十六进制随机数。

**类型**

可执行文件（`/usr/bin/mcookie`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-f`，`--file file` - 指定随机数设备文件，默认位 `/dev/urandom`；如果此参数设为 `-` 则读取标准输入
  - `-m`，`--max-size number` - 仅从随机数设备文件中读取 `number` 字节
  - `-v`，`--verbose` - 显示详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本

## 示例

```
$ mcookie                               # 生成随机数
ad149d9f7f46d254bc93b9fc66dc70a4
$ xauth add :0 . $(mcookie)             # 生成 cookie 并添加到 DISPLAY :0 中
$ xauth list                            # 查看 cookie
dev/unix:0  MIT-MAGIC-COOKIE-1  6b662e86481f97c1bf0e6cb3be355ee7
```

## 推荐阅读

- mcookie(1) — Linux manual page

## 手册

```
MCOOKIE(1)                      User Commands                      MCOOKIE(1)

NAME
```

```
       mcookie - generate magic cookies for xauth

SYNOPSIS
       mcookie [options]

DESCRIPTION
       mcookie generates a 128-bit random hexadecimal number for use with the
       X authority system. Typical usage:

          xauth add :0 . mcookie

       The "random" number generated is actually the MD5 message digest of
       random information coming from one of the sources getrandom(2) system
       call, /dev/urandom, /dev/random, or the libc pseudo-random functions,
       in this preference order. See also the option --file.

OPTIONS
       -f, --file file
           Use this file as an additional source of randomness (for example
           /dev/urandom). When file is '-', characters are read from standard
           input.

       -m, --max-size number
           Read from file only this number of bytes. This option is meant to
           be used when reading additional randomness from a file or device.

           The number argument may be followed by the multiplicative suffixes
           KiB=1024, MiB=1024*1024, and so on for GiB, TiB, PiB, EiB, ZiB and
           YiB (the "iB" is optional, e.g., "K" has the same meaning as "KiB")
           or the suffixes KB=1000, MB=1000*1000, and so on for GB, TB, PB,
           EB, ZB and YB.

       -v, --verbose
           Inform where randomness originated, with amount of entropy read
           from each source.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

FILES
       /dev/urandom

       /dev/random

BUGS
       It is assumed that none of the randomness sources will block.

SEE ALSO
       md5sum(1), X(7), xauth(1), rand(3)
```

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mcookie command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                        MCOOKIE(1)
```

# Bash 的 md5sum 命令

```
md5sum [OPTION]... [FILE]...
```

**功能**

计算或校验 MD5 值。

**类型**

可执行文件（`/usr/bin/md5sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 MD5 值进行校验
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 MD5 值**

```
$ md5sum 1.txt                                  # 单个文件
79171af5e65e23a97b58c343c4de7411  1.txt
$ md5sum 1.txt 2.txt 3.txt                      # 多个文件
79171af5e65e23a97b58c343c4de7411  1.txt
d68dae595d597fb67c85a8301521b097  2.txt
72d93867588107cf46b2cc7ea400d0aa  3.txt
$ md5sum --tag 1.txt 2.txt 3.txt                # BSD 风格
```

```
MD5 (1.txt) = 79171af5e65e23a97b58c343c4de7411
MD5 (2.txt) = d68dae595d597fb67c85a8301521b097
MD5 (3.txt) = 72d93867588107cf46b2cc7ea400d0aa
```

**校验 MD5 值**

```
$ md5sum 1.txt 2.txt 3.txt > md5.txt
$ md5sum -c md5.txt # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| b2sum | 计算和校验 BLAKE2 值 |
| sha1sum | 计算和校验 SHA1 值 |
| sha224sum | 计算和校验 SHA224 值 |
| sha256sum | 计算和校验 SHA256 值 |
| sha384sum | 计算和校验 SHA384 值 |
| sha512sum | 计算和校验 SHA512 值 |

## 推荐阅读

- Print or check MD5 digests - GNU Coreutils
- md5sum(1) — Linux manual page

## 手册

```
MD5SUM(1)                        User Commands                        MD5SUM(1)

NAME
       md5sum - compute and check MD5 message digest

SYNOPSIS
       md5sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check MD5 (128-bit) checksums.
```

```
       With no FILE, or when FILE is -, read standard input.

       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in RFC 1321.  When checking, the in-
       put should be a former output of this program.  The default mode is  to
       print a line with: checksum, a space, a character indicating input mode
       ('*'  for  binary,  ' ' for text or where binary is insignificant), and
       name for each FILE.

       Note: There is no difference between binary mode and text mode  on  GNU
       systems.

BUGS
       Do  not  use the MD5 algorithm for security related purposes.  Instead,
       use an SHA-2  algorithm,  implemented  in  the  programs  sha224sum(1),
       sha256sum(1),  sha384sum(1), sha512sum(1), or the BLAKE2 algorithm, im-
       plemented in b2sum(1)
```

AUTHOR
       Written by Ulrich Drepper, Scott Miller, and David Madore.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       cksum(1)

       Full documentation <https://www.gnu.org/software/coreutils/md5sum>
       or available locally via: info '(coreutils) md5sum invocation'

# Bash 的 mesg 命令

```
mesg [OPTION]... [n|y]
```

**功能**

允许（或禁止）显示来自其它用户的消息。

> 通过 `talk` 或 `write` 命令可以向其他用户发送信息。

**类型**

可执行文件（`/usr/bin/mesg`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-v`，`--verbose` - 显示详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `n` - 禁止显示来自其它用户的消息
- `y` - 允许显示来自其它用户的消息

## 示例

```
$ mesg        # 查看是否允许显示来自其它用户的消息
is y
$ mesg n      # 禁止显示来自其它用户的消息
$ mesg
is n
```

## 推荐阅读

- mesg(1) — Linux manual page

## 手册

```
MESG(1)                         User Commands                         MESG(1)
```

NAME

mesg - display (or do not display) messages from other users

SYNOPSIS

mesg [option] [n|y]

DESCRIPTION

The mesg utility is invoked by a user to control write access others have to the terminal device associated with standard error output. If write access is allowed, then programs such as talk(1) and write(1) may display messages on the terminal.

Traditionally, write access is allowed by default. However, as users become more conscious of various security risks, there is a trend to remove write access by default, at least for the primary login shell. To make sure your ttys are set the way you want them to be set, mesg should be executed in your login scripts.

The mesg utility silently exits with error status 2 if not executed on a terminal. In this case executing mesg is pointless. The command line option --verbose forces mesg to print a warning in this situation. This behaviour has been introduced in version 2.33.

ARGUMENTS

n

Disallow messages.

y

Allow messages to be displayed.

If no arguments are given, mesg shows the current message status on standard error output.

OPTIONS

-v, --verbose

Explain what is being done.

-h, --help

Display help text and exit.

-V, --version

Print version and exit.

EXIT STATUS

The mesg utility exits with one of the following values:

0

Messages are allowed.

1

Messages are not allowed.

>1

```
           An error has occurred.

FILES
       /dev/[pt]ty*, /dev/pts/[0-9]*

HISTORY
       mesg (I) appears in the UNIX Programmer’s Manual. It used to invert the
       current state with no argument before Version 7 AT&T UNIX.

SEE ALSO
       login(1), talk(1), write(1), wall(1), xterm(1)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mesg command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                          MESG(1)
```

# Bash 的 mkdir 命令

```
mkdir [OPTION]... DIRECTORY...
```

**功能**

创建目录。

**类型**

可执行文件（`/usr/bin/mkdir`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-m`，`--mode=MODE` - 设置文件的模式（权限），如果不设置此选项则默认使用 `a=rwa` 减去 `umask` 权限
  - `-p`，`--parents` - 自动创建上级目录；如果目录已经存在，不报错
  - `-v`，`--verbose` - 为创建的每一个目录打印信息
  - `-Z` - 将每个新建目录的 SELinux 安全上下文设为默认值
  - `--context[=CTX]` - 将每个新建目录的 SELinux 安全上下文设为 `CTX`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `DIRECTORY` - 要创建的目录

## 示例

**创建一个目录**

```
mkdir project
```

创建名为 `project` 的目录。

**一次创建多个目录**

```
mkdir src include bin
```

同时创建三个目录：`src/`、`include/`、`bin/`。

**创建多级目录**

```
mkdir -p project/src/utils
```

若 `project` 和 `src` 不存在，`-p` 会自动创建它们。

结果：

```
project/
└── src/
    └── utils/
```

**指定权限创建**

```
mkdir -m 700 private
```

权限为：

```
drwx------
```

仅当前用户可访问。

**用符号式权限**

```
mkdir -m u=rwx,g=rx,o=rx public
```

等价于 `chmod 755 public` 即：

```
drwxr-xr-x
```

**结合使用**

```
mkdir -p -m 755 /opt/app/logs
```

自动创建 `/opt/app`（若不存在）并设置 `logs` 目录权限为 755。

**显示创建过程**

```
mkdir -v test
```

输出：

```
mkdir: created directory 'test'
```

**显示层级**

```
mkdir -vp a/b/c
```

输出：

```
mkdir: created directory 'a'
mkdir: created directory 'a/b'
mkdir: created directory 'a/b/c'
```

**安全创建目录（不报错）**

```
mkdir -p ~/Downloads/tmp
```

若目录已存在不会报错，非常适合脚本中使用。

**在脚本中先判断再创建**

```
[ ! -d logs ] && mkdir logs
```

如果 `logs` 目录不存在，则创建它。

**一次创建一组带编号的目录**

```
mkdir dir_{1..5}
```

创建：

```
dir_1/ dir_2/ dir_3/ dir_4/ dir_5/
```

**结合变量使用**

```
name="backup_$(date +%Y%m%d)"
mkdir -p ~/backups/$name
```

创建带日期的备份目录，比如：

```
~/backups/backup_20251025/
```

**批量创建层级结构**

```
mkdir -p project/{src,include,docs,build}
```

结果：

```
project/
├── src/
├── include/
├── docs/
└── build/
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| touch | 创建普通文件或刷新访问时间 |
| mkfifo | 创建命名管道 |
| mknod | 创建设备文件 |
| mktemp | 创建临时文件或目录 |
| link | 创建硬链接 |
| ln | 创建硬链接或符号链接 |

## 推荐阅读

- Make directories - GNU Coreutils
- mkdir(1) — Linux manual page

## 手册

```
MKDIR(1)                          User Commands                          MKDIR(1)

NAME
       mkdir - make directories

SYNOPSIS
       mkdir [OPTION]... DIRECTORY...

DESCRIPTION
       Create the DIRECTORY(ies), if they do not already exist.

       Mandatory  arguments  to  long  options are mandatory for short options
```

too.

-m, --mode=MODE
set file mode (as in chmod), not a=rwx - umask

-p, --parents
no error if existing, make parent directories as needed, with their file modes unaffected by any -m option.

-v, --verbose
print a message for each created directory

-Z set SELinux security context of each created directory to the default type

--context[=CTX]
like -Z, or if CTX is specified then set the SELinux or SMACK security context to CTX

--help display this help and exit

--version
output version information and exit

## AUTHOR

Written by David MacKenzie.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

mkdir(2)

Full documentation <https://www.gnu.org/software/coreutils/mkdir>
or available locally via: info '(coreutils) mkdir invocation'

# Bash 的 mkfifo 命令

```
mkfifo [OPTION]... NAME...
```

**功能**

创建命名管道文件。

**类型**

可执行文件（`/usr/bin/mkfifo`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-m`, `--mode=MODE` - 设置文件的访问权限；默认为 `a=rw - umask`
  - `-Z` - 将 SELinux 安全上下文设为默认类型
  - `--context[=CTX]` - 将 SELinux 安全上下文设为 `CTX`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NAME` - 文件名

## 示例

```
$ mkfifo my_fifo
$ ls -l my_fifo
prw-rw-r-- 1 primers primers 0 Nov 18 13:08 my_fifo
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| touch | 创建普通文件或刷新访问时间 |
| mkdir | 创建目录 |
| mknod | 创建设备文件 |
| mktemp | 创建临时文件或目录 |
| link | 创建硬链接 |

| 命令 | 说明 |
| --- | --- |
| ln | 创建硬链接或符号链接 |

## 推荐阅读

- Make FIFOs (named pipes) - GNU Coreutils
- mkfifo(1) — Linux manual page

## 手册

```
MKFIFO(1)                          User Commands                          MKFIFO(1)

NAME
       mkfifo - make FIFOs (named pipes)

SYNOPSIS
       mkfifo [OPTION]... NAME...

DESCRIPTION
       Create named pipes (FIFOs) with the given NAMEs.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -m, --mode=MODE
              set file permission bits to MODE, not a=rw - umask

       -Z     set the SELinux security context to default type

       --context[=CTX]
              like -Z, or if CTX is specified then set the  SELinux  or  SMACK
              security context to CTX

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
```

```
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       mkfifo(3)

       Full documentation <https://www.gnu.org/software/coreutils/mkfifo>
       or available locally via: info '(coreutils) mkfifo invocation'

GNU coreutils 9.4                April 2024                        MKFIFO(1)
```

# Bash 的 mkfs.bfs 命令

```
mkfs.bfs [OPTION]... DEVICE [BLOCK-COUNT]
```

**功能**

创建一个 Boot File System 文件系统。

**类型**

可执行文件（`/usr/sbin/mkfs.bfs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-N`, `--inodes number` - 指定所需的 inode 数量（最多 512 个）
  - `-V`, `--vname label` - 请指定卷标
  - `-F`, `--fname name` - 请指定文件系统名称
  - `--lock[=mode]` - 是否对操作的设备或文件使用独占 BSD 锁
    - `yes` - 对设备加锁，如果锁被占用，则阻塞等待；省略 `mode` 时默认为 `yes`
    - `no` - 不对设备加锁；没有 `--lock` 选项时默认不加锁
    - `nonblock` - 对设备加锁，如果锁被占用，则失败
  - `-c` - ~~坏块检查~~；废弃选项，将被静默忽略
  - `-l` - ~~列出坏块~~；废弃选项，将被静默忽略
  - `-v`, `--verbose` - 显示详细信息
  - `-h`, `--help` - 显示帮助
  - `-V`, `--version` - 显示版本
- `DEVICE` - 要操作的设备
- `BLOCK-COUNT` - 块数量；默认为整个分区

## 示例

```
$ sudo mkfs.bfs /dev/sda        # 将 /dev/sda 格式化为 bfs 文件系统
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| mkfs | 创建 Linux 文件系统 |

## 推荐阅读

- mkfs.bfs(8) — Linux manual page
- BFS Filesystem for Linux - The Linux Kernel documentation

## 手册

```
MKFS.BFS(8)                 System Administration                MKFS.BFS(8)

NAME
       mkfs.bfs - make an SCO bfs filesystem

SYNOPSIS
       mkfs.bfs [options] device [block-count]

DESCRIPTION
       mkfs.bfs creates an SCO bfs filesystem on a block device (usually a
       disk partition or a file accessed via the loop device).

       The block-count parameter is the desired size of the filesystem, in
       blocks. If nothing is specified, the entire partition will be used.

OPTIONS
       -N, --inodes number
           Specify the desired number of inodes (at most 512). If nothing is
           specified, some default number in the range 48-512 is picked
           depending on the size of the partition.

       -V, --vname label
           Specify the volume label. I have no idea if/where this is used.

       -F, --fname name
           Specify the filesystem name. I have no idea if/where this is used.

       --lock[=mode]
           Use exclusive BSD lock for device or file it operates. The optional
           argument mode can be yes, no (or 1 and 0) or nonblock. If the mode
           argument is omitted, it defaults to yes. This option overwrites
           environment variable $LOCK_BLOCK_DEVICE. The default is not to use
           any lock at all, but it's recommended to avoid collisions with
           systemd-udevd(8) or other tools.

       -v, --verbose
           Explain what is being done.
```

```
       -c
           This option is silently ignored.

       -l
           This option is silently ignored.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit. Option -V only works as --version when it
           is the only option.

EXIT STATUS
       The exit status returned by mkfs.bfs is 0 when all went well, and 1
       when something went wrong.

SEE ALSO
       mkfs(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mkfs.bfs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                      MKFS.BFS(8)
```

# Bash 的 mkfs.cramfs 命令

```
mkfs.cramfs [OPTION]... DIRECTORY FILE
```

**功能**

制作压缩的 ROM 文件系统。

**类型**

可执行文件（`/usr/sbin/mkfs.cramfs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-v` - 打印详细信息
  - `-E` - 将所有警告视为错误
  - `-b blocksize` - 块大小
  - `-e edition` - 在超级块中使用已定义的文件系统版本号
  - `-N big, little, host` - 指定字节序；默认值为主机字节序
  - `-i file` - 将 `file` 文件插入 cramfs 的元数据中
  - `-n name` - 设置文件系统的名称
  - `-p` - 为引导填充 512 字节
  - `-s` - ~~启用目录条目排序~~；废弃选项，被忽略
  - `-z` - 显式空洞
  - `-l[=mode]` - 对操作的设备或文件使用独占 BSD 锁
    - `yes` - 对设备加锁，如果锁被占用，则阻塞等待；省略 `mode` 时默认为 `yes`
    - `no` - 不对设备加锁；没有 `--lock` 选项时默认不加锁
    - `nonblock` - 对设备加锁，如果锁被占用，则失败
  - `-h`, `--help` - 显示帮助
  - `-V`, `--version` - 显示版本
- `DIRECTORY` - 要压缩的目录
- `FILE` - 存储镜像的文件

## 示例

```
$ mkfs.cramfs dir image.img                # 将 dir 目录压缩打包成 image.img
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| mkfs | 创建 Linux 文件系统 |
| fsck.cramfs | 检查并修复压缩 ROM 文件系统 |

## 推荐阅读

- mkfs.cramfs(8) — Linux manual page
- cram a filesystem onto a small ROM - The Linux Kernel documentation

## 手册

```
MKFS.CRAMFS(8)              System Administration              MKFS.CRAMFS(8)

NAME
       mkfs.cramfs - make compressed ROM file system

SYNOPSIS
       mkfs.cramfs [options] directory file

DESCRIPTION
       Files on cramfs file systems are zlib-compressed one page at a time to
       allow random read access. The metadata is not compressed, but is
       expressed in a terse representation that is more space-efficient than
       conventional file systems.

       The file system is intentionally read-only to simplify its design;
       random write access for compressed files is difficult to implement.
       cramfs ships with a utility (mkcramfs(8)) to pack files into new cramfs
       images.

       File sizes are limited to less than 16 MB.

       Maximum file system size is a little under 272 MB. (The last file on
       the file system must begin before the 256 MB block, but can extend past
       it.)

ARGUMENTS
       The directory is simply the root of the directory tree that we want to
       generate a compressed filesystem out of.

       The file will contain the cram file system, which later can be mounted.

OPTIONS
       -v
```

```
            Enable verbose messaging.

        -E
            Treat all warnings as errors, which are reflected as command exit
            status.

        -b blocksize
            Use defined block size, which has to be divisible by page size.

        -e edition
            Use defined file system edition number in superblock.

        -N big, little, host
            Use defined endianness. Value defaults to host.

        -i file
            Insert a file to cramfs file system.

        -n name
            Set name of the cramfs file system.

        -p
            Pad by 512 bytes for boot code.

        -s
            This option is ignored. Originally the -s turned on directory entry
            sorting.

        -z
            Make explicit holes.

        -l[=mode]
            Use exclusive BSD lock for device or file it operates. The optional
            argument mode can be yes, no (or 1 and 0) or nonblock. If the mode
            argument is omitted, it defaults to "yes". This option overwrites
            environment variable $LOCK_BLOCK_DEVICE. The default is not to use
            any lock at all, but it’s recommended to avoid collisions with
            udevd or other tools.

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

EXIT STATUS
        0
            success

        8
            operation error, such as unable to allocate memory

SEE ALSO
```

```
       fsck.cramfs(8), mount(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mkfs.cramfs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-11-21                    MKFS.CRAMFS(8)
```

# Bash 的 mkfs.minix 命令

```
mkfs.minix [OPTION]... DEVICE [SIZE]
```

**功能**

创建 Minix 文件系统。

**类型**

可执行文件（`/usr/sbin/mkfs.cramfs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-c`，`--check` - 创建文件系统之前，请检查设备是否存在坏块
  - `-n`，`--namelength length` - 文件名的最大长度；默认值为 30
  - `--lock[=mode]` - 使用独占的 BSD 锁来控制设备或其操作的文件
    - `yes` - 对设备加锁，如果锁被占用，则阻塞等待；省略 `mode` 时默认为 `yes`
    - `no` - 不对设备加锁；没有 `--lock` 选项时默认不加锁
    - `nonblock` - 对设备加锁，如果锁被占用，则失败
  - `-i`，`--inodes number` - 指定文件系统的 inode 数量
  - `-l`，`--badblocks filename` - 从 `filename` 中读取坏块列表
  - `-1` - 创建 Minix 版本 1 文件系统；**默认**
  - `-2`，`-v` - 创建 Minix 版本 2 文件系统
  - `-3` - 创建 Minix 版本 3 文件系统
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 要操作的设备
- `SIZE` - 文件系统大小（以块为单位）

## 示例

```
$ sudo mkfs.minix /dev/sda      # 将 /dev/sda 格式化为 Minix 文件系统
21856 inodes
65535 blocks
Firstdatazone=696 (696)
Zonesize=1024
```

```
Maxsize=268966912
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| mkfs | 创建 Linux 文件系统 |
| fsck.minix | 检查并修复 Minix 文件系统 |

## 推荐阅读

- mkfs.minix(8) — Linux manual page

## 手册

```
MKFS.MINIX(8)                System Administration                MKFS.MINIX(8)

NAME
       mkfs.minix - make a Minix filesystem

SYNOPSIS
       mkfs.minix [options] device [size-in-blocks]

DESCRIPTION
       mkfs.minix creates a Linux MINIX filesystem on a device (usually a disk
       partition).

       The device is usually of the following form:

           /dev/hda[1-8] (IDE disk 1)
           /dev/hdb[1-8] (IDE disk 2)
           /dev/sda[1-8] (SCSI disk 1)
           /dev/sdb[1-8] (SCSI disk 2)

       The device may be a block device or an image file of one, but this is
       not enforced. Expect not much fun on a character device :-).

       The size-in-blocks parameter is the desired size of the file system, in
       blocks. It is present only for backwards compatibility. If omitted the
       size will be determined automatically. Only block counts strictly
       greater than 10 and strictly less than 65536 are allowed.

OPTIONS
       -c, --check
           Check the device for bad blocks before creating the filesystem. If
```

```
            any are found, the count is printed.

        -n, --namelength length
            Specify the maximum length of filenames. Currently, the only
            allowable values are 14 and 30 for file system versions 1 and 2.
            Version 3 allows only value 60. The default is 30.

        --lock[=mode]
            Use exclusive BSD lock for device or file it operates. The optional
            argument mode can be yes, no (or 1 and 0) or nonblock. If the mode
            argument is omitted, it defaults to yes. This option overwrites
            environment variable $LOCK_BLOCK_DEVICE. The default is not to use
            any lock at all, but it’s recommended to avoid collisions with
            systemd-udevd(8) or other tools.

        -i, --inodes number
            Specify the number of inodes for the filesystem.

        -l, --badblocks filename
            Read the list of bad blocks from filename. The file has one
            bad-block number per line. The count of bad blocks read is printed.

        -1
            Make a Minix version 1 filesystem. This is the default.

        -2, -v
            Make a Minix version 2 filesystem.

        -3
            Make a Minix version 3 filesystem.

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit. The long option cannot be combined with
            other options.

ENVIRONMENT
        LOCK_BLOCK_DEVICE=<mode>
            use exclusive BSD lock. The mode is "1" or "0". See --lock for more
            details.

EXIT STATUS
        The exit status returned by mkfs.minix is one of the following:

        0
            No errors

        8
            Operational error

        16
```

```
           Usage or syntax error

SEE ALSO
       fsck(8), mkfs(8), reboot(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mkfs.minix command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                    MKFS.MINIX(8)
```

# Bash 的 mkfs 命令

```
mkfs [OPTION]... [fs-options] DEVICE [SIZE]
```

**功能**

创建 Linux 文件系统。

**类型**

可执行文件（`/usr/sbin/mkfs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-t`，`--type type` - 文件系统类型，默认为 `ext2`
  - `-v`，`--verbose` - 显示详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `fs-options` - 传递给实际创建文件系统的程序的选项
- `DEVICE` - 要操作的设备
- `SIZE` - 文件系统的大小

## 示例

```
$ sudo mkfs --type fat -I /dev/sda              # 将 /dev/sda 格式化为 FAT32 文件系统
mkfs.fat 4.2 (2021-01-31)
```

```
$ dd if=/dev/zero of=disk.img bs=1M count=16    # 创建一个 16M 的空文件
$ mkfs disk.img                                 # 在 disk.img 上创建文件系统（默认为 ext2）
mke2fs 1.47.0 (5-Feb-2023)
Discarding device blocks: done
Creating filesystem with 4096 4k blocks and 4096 inodes

Allocating group tables: done
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
$ mkdir ./disk                                  # 创建一个文件夹
$ sudo mount -o loop disk.img disk              # 将 disk.img 挂载到 disk
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| fsck | 检查并修复 Linux 文件系统 |
| mkfs.bfs | 创建 bfs 文件系统 |
| mkfs.cramfs | 创建 cramfs 文件系统 |
| mkfs.ext2 | 创建 ext2 文件系统 |
| mkfs.ext3 | 创建 ext3 文件系统 |
| mkfs.ext4 | 创建 ext4 文件系统 |
| mkfs.fat | 创建 fat 文件系统 |
| mkfs.minix | 创建 minix 文件系统 |
| mkfs.msdos | 创建 msdos 文件系统 |
| mkfs.vfat | 创建 vfat 文件系统 |
| mkfs.xfs | 创建 xfs 文件系统 |
| mkswap | 创建 sawp 文件系统 |
| mkisofs | 创建 iso 文件系统 |

## 推荐阅读

- mkfs(8) — Linux manual page

## 手册

```
MKFS(8)                         System Administration                         MKFS(8)

NAME
       mkfs - build a Linux filesystem

SYNOPSIS
       mkfs [options] [-t type] [fs-options] device [size]

DESCRIPTION
       This mkfs frontend is deprecated in favour of filesystem specific
       mkfs.<type> utils.

       mkfs is used to build a Linux filesystem on a device, usually a hard
```

disk partition. The device argument is either the device name (e.g., /dev/hda1, /dev/sdb2), or a regular file that shall contain the filesystem. The size argument is the number of blocks to be used for the filesystem.

The exit status returned by mkfs is 0 on success and 1 on failure.

In actuality, mkfs is simply a front-end for the various filesystem builders (mkfs.fstype) available under Linux. The filesystem-specific builder is searched for via your PATH environment setting only. Please see the filesystem-specific builder manual pages for further details.

## OPTIONS

**-t, --type type**
Specify the type of filesystem to be built. If not specified, the default filesystem type (currently ext2) is used.

**fs-options**
Filesystem-specific options to be passed to the real filesystem builder.

**-V, --verbose**
Produce verbose output, including all filesystem-specific commands that are executed. Specifying this option more than once inhibits execution of any filesystem-specific commands. This is really only useful for testing.

**-h, --help**
Display help text and exit.

**-V, --version**
Print version and exit. (Option -V will display version information only when it is the only parameter, otherwise it will work as --verbose.)

## BUGS

All generic options must precede and not be combined with filesystem-specific options. Some filesystem-specific programs do not automatically detect the device size and require the size parameter to be specified.

## AUTHORS

David Engel <david@ods.com>, Fred N. van Kempen <waltje@uwalt.nl.mugnet.org>, Ron Sommeling <sommel@sci.kun.nl>.

The manual page was shamelessly adapted from Remy Card’s version for the ext2 filesystem.

## SEE ALSO

fs(5), badblocks(8), fsck(8), mkdosfs(8), mke2fs(8), mkfs.bfs(8), mkfs.ext2(8), mkfs.ext3(8), mkfs.ext4(8), mkfs.minix(8), mkfs.msdos(8), mkfs.vfat(8), mkfs.xfs(8)

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mkfs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                          MKFS(8)
```

# Bash 的 mknod 命令

```
mknod [OPTION]... NAME TYPE [MAJOR MINOR]
```

**功能**

创建块设备或字符设备文件。

**类型**

可执行文件（`/usr/bin/mknod`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-m`, `--mode=MODE` - 设置文件的访问权限；默认为 `a=rw - umask`
  - `-Z` - 将 SELinux 安全上下文设为默认类型
  - `--context[=CTX]` - 将 SELinux 安全上下文设为 `CTX`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NAME` - 文件名
- `TYPE` - 类型：
  - `b` - 块设备（有缓冲）
  - `c` - 字符设备（无缓冲）
  - `u` - 字符设备（无缓冲）；同 `c`
  - `p` - 命名管道；同 mkfifo
- `MAJOR` - 主版本号
- `MINOR` - 次版本号

## 示例

```
$ sudo mknod my_block b 1 0     # 创建块设备
brw-r--r-- 1 root root 1, 0 Nov 18 13:18 my_block
$ sudo mknod my_char c 1 0      # 创建字符设备
crw-r--r-- 1 root root 1, 0 Nov 18 13:19 my_char
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| touch | 创建普通文件或刷新访问时间 |
| mkdir | 创建目录 |
| mkfifo | 创建命名管道 |
| mktemp | 创建临时文件或目录 |
| link | 创建硬链接 |
| ln | 创建硬链接或符号链接 |

## 推荐阅读

- Make block or character special files - GNU Coreutils
- mknod(1) — Linux manual page

## 手册

```
MKNOD(1)                          User Commands                          MKNOD(1)

NAME
       mknod - make block or character special files

SYNOPSIS
       mknod [OPTION]... NAME TYPE [MAJOR MINOR]

DESCRIPTION
       Create the special file NAME of the given TYPE.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -m, --mode=MODE
              set file permission bits to MODE, not a=rw - umask

       -Z     set the SELinux security context to default type

       --context[=CTX]
              like -Z, or if CTX is specified then set the  SELinux  or  SMACK
              security context to CTX

       --help display this help and exit

       --version
              output version information and exit
```

Both MAJOR and MINOR must be specified when TYPE is b, c, or u, and they must be omitted when TYPE is p. If MAJOR or MINOR begins with 0x or 0X, it is interpreted as hexadecimal; otherwise, if it begins with 0, as octal; otherwise, as decimal. TYPE may be:

b create a block (buffered) special file

c, u create a character (unbuffered) special file

p create a FIFO

NOTE: your shell may have its own version of mknod, which usually supersedes the version described here. Please refer to your shell's documentation for details about the options it supports.

## AUTHOR

Written by David MacKenzie.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

mknod(2)

Full documentation <https://www.gnu.org/software/coreutils/mknod>
or available locally via: info '(coreutils) mknod invocation'

# Bash 的 mkswap 命令

```
mkswap [OPTION]... DEVICE [SIZE]
```

**功能**

建立 Linux 交换区域。

**类型**

可执行文件（`/usr/sbin/mkswap`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-c`，`--check` - 创建交换区之前，请检查设备是否存在坏块
  - `-f`，`--force` - 强制执行
  - `-q`，`--quiet` - 不显示输出
  - `-L`，`--label label` - 为设备指定标签
  - `--lock[=mode]` - 对操作的设备或文件使用独占 BSD 锁
    - `yes` - 对设备加锁，如果锁被占用，则阻塞等待；省略 `mode` 时默认为 `yes`
    - `no` - 不对设备加锁；没有 `--lock` 选项时默认不加锁
    - `nonblock` - 对设备加锁，如果锁被占用，则失败
  - `-p`，`--pagesize size` - 页大小；单位是字节
  - `-U`，`--uuid UUID` - 指定特定的 UUID，或以下特殊值：
    - `clear` - 清除 UUID
    - `random` - 生成一个新的随机生成的 UUID
    - `time` - 生成一个新的基于时间的 UUID
  - `-e`，`--endianness ENDIANNESS` - 指定字节序：`big`，`little` 或 `native`
  - `-v`，`--swapversion 1` - 指定交换空间版本；只支持 1
  - `--verbose` - 打印详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 要操作的设备
- `SIZE` - ~~交换区域大小~~；被忽略

## 示例

```
$ sudo mkswap /dev/sda
Setting up swapspace version 1, size = 29.3 GiB (31474151424 bytes)
no label, UUID=26c40fb1-e926-4220-a3e2-2e790409a7ab
```

## 推荐阅读

- mkswap(8) — Linux manual page

## 手册

```
MKSWAP(8)                   System Administration                   MKSWAP(8)

NAME
       mkswap - set up a Linux swap area

SYNOPSIS
       mkswap [options] device [size]

DESCRIPTION
       mkswap sets up a Linux swap area on a device or in a file.

       The device argument will usually be a disk partition (something like
       /dev/sdb7) but can also be a file. The Linux kernel does not look at
       partition IDs, but many installation scripts will assume that
       partitions of hex type 82 (LINUX_SWAP) are meant to be swap partitions.
       (Warning: Solaris also uses this type. Be careful not to kill your
       Solaris partitions.)

       The size parameter is superfluous but retained for backwards
       compatibility. (It specifies the desired size of the swap area in
       1024-byte blocks. mkswap will use the entire partition or file if it is
       omitted. Specifying it is unwise - a typo may destroy your disk.)

       After creating the swap area, you need the swapon(8) command to start
       using it. Usually swap areas are listed in /etc/fstab so that they can
       be taken into use at boot time by a swapon -a command in some boot
       script.

WARNING
       The swap header does not touch the first block. A boot loader or disk
       label can be there, but it is not a recommended setup. The recommended
       setup is to use a separate partition for a Linux swap area.

       mkswap, like many others mkfs-like utils, erases the first partition
       block to make any previous filesystem invisible.

       However, mkswap refuses to erase the first block on a device with a
       disk label (SUN, BSD, ...).
```

```
OPTIONS
       -c, --check
           Check the device (if it is a block device) for bad blocks before
           creating the swap area. If any bad blocks are found, the count is
           printed.

       -f, --force
           Go ahead even if the command is stupid. This allows the creation of
           a swap area larger than the file or partition it resides on.

           Also, without this option, mkswap will refuse to erase the first
           block on a device with a partition table.

       -q, --quiet
           Suppress output and warning messages.

       -L, --label label
           Specify a label for the device, to allow swapon(8) by label.

       --lock[=mode]
           Use exclusive BSD lock for device or file it operates. The optional
           argument mode can be yes, no (or 1 and 0) or nonblock. If the mode
           argument is omitted, it defaults to yes. This option overwrites
           environment variable $LOCK_BLOCK_DEVICE. The default is not to use
           any lock at all, but it’s recommended to avoid collisions with
           systemd-udevd(8) or other tools.

       -p, --pagesize size
           Specify the page size (in bytes) to use. This option is usually
           unnecessary; mkswap reads the size from the kernel.

       -U, --uuid UUID
           Specify the UUID to use. The default is to generate a UUID. The
           format of the UUID is a series of hex digits separated by hyphens,
           like this: "c1b9d5a2-f162-11cf-9ece-0020afc76f16". The UUID
           parameter may also be one of the following:

           clear
               clear the filesystem UUID

           random
               generate a new randomly-generated UUID

           time
               generate a new time-based UUID

       -e, --endianness ENDIANNESS
           Specify the ENDIANNESS to use, valid arguments are native, little
           or big. The default is native.

       -v, --swapversion 1
           Specify the swap-space version. (This option is currently
```

pointless, as the old -v 0 option has become obsolete and now only -v 1 is supported. The kernel has not supported v0 swap-space format since 2.5.22 (June 2002). The new version v1 is supported since 2.1.117 (August 1998).)

--verbose
Verbose execution. With this option mkswap will output more details about detected problems during swap area set up.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

ENVIRONMENT

LIBBLKID_DEBUG=all
enables libblkid debug output.

LOCK_BLOCK_DEVICE=<mode>
use exclusive BSD lock. The mode is "1" or "0". See --lock for more details.

NOTES

The maximum useful size of a swap area depends on the architecture and the kernel version.

The maximum number of the pages that is possible to address by swap area header is 4294967295 (32-bit unsigned int). The remaining space on the swap device is ignored.

Presently, Linux allows 32 swap areas. The areas in use can be seen in the file /proc/swaps.

mkswap refuses areas smaller than 10 pages.

If you don’t know the page size that your machine uses, you can look it up with getconf PAGESIZE.

To set up a swap file, it is necessary to create that file before initializing it with mkswap, e.g. using a command like

```
# dd if=/dev/zero of=swapfile bs=1MiB count=$((8*1024))
```

to create 8GiB swapfile.

Please read notes from swapon(8) about the swap file use restrictions (holes, preallocation and copy-on-write issues).

SEE ALSO

fdisk(8), swapon(8)

REPORTING BUGS

```
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mkswap command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                        MKSWAP(8)
```

# Bash 的 mktemp 命令

```
mktemp [OPTION]... [TEMPLATE]
```

**功能**

安全的创建临时文件或目录，并打印名称。

**类型**

可执行文件（`/usr/bin/mknod`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-d`，`--directory` - 创建临时目录，而不是临时文件
  - `-u`，`--dry-run` - 只打印临时文件名，而不创建文件
  - `-q`，`--quiet` - 不打印失败信息
  - `--suffix=SUFF` - 在文件名 `TEMPLATE` 后面添加后缀 `SUFF`
  - `-p DIR`，`--tmpdir[=DIR]` - 相对于 `DIR` 目录创建临时文件
  - `-t` - ~~相对于环境变量 `TMPDIR` 表示的目录创建临时文件；已废弃~~
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `TEMPLATE` - 临时文件名模板，末尾至少 3 个 `X` 表示临时名称；例如 `my_temp.XXX`
  - 不指定 `TEMPLATE` 参数时，默认格式为 `tmp.XXXXXXXXXX`，并默认设置 `--tmpdir=/tmp`
  - 指定 `TEMPLATE` 参数时，不自动设置 `--tmpdir` 选项

## 示例

```
$ mktemp                          # 创建临时文件
/tmp/tmp.MWGmRyLisU
$ mktemp my_temp.XXX              # 在当前目录创建临时文件
my_temp.EQ4
$ mktemp -p /tmp my_temp.XXX      # 在 /tmp 目录创建自定义格式的临时文件
/tmp/my_temp.PZj
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| touch | 创建普通文件或刷新访问时间 |
| mkdir | 创建目录 |
| mkfifo | 创建命名管道 |
| mknod | 创建设备文件 |
| link | 创建硬链接 |
| ln | 创建硬链接或符号链接 |

## 推荐阅读

- Create temporary file or directory - GNU Coreutils
- mktemp(1) — Linux manual page

## 手册

```
MKTEMP(1)                          User Commands                          MKTEMP(1)

NAME
       mktemp - create a temporary file or directory

SYNOPSIS
       mktemp [OPTION]... [TEMPLATE]

DESCRIPTION
       Create a temporary file or directory, safely, and print its name.  TEM-
       PLATE  must  contain at least 3 consecutive 'X's in last component.  If
       TEMPLATE is not specified, use tmp.XXXXXXXXXX, and --tmpdir is implied.
       Files are created u+rw, and directories  u+rwx,  minus  umask  restric-
       tions.

       -d, --directory
              create a directory, not a file

       -u, --dry-run
              do not create anything; merely print a name (unsafe)

       -q, --quiet
              suppress diagnostics about file/dir-creation failure

       --suffix=SUFF
              append  SUFF  to  TEMPLATE; SUFF must not contain a slash.  This
              option is implied if TEMPLATE does not end in X
```

```
       -p DIR, --tmpdir[=DIR]
              interpret TEMPLATE relative to DIR; if DIR is not specified, use
              $TMPDIR if set, else /tmp.  With this option, TEMPLATE must  not
              be  an  absolute  name;  unlike  with  -t,  TEMPLATE may contain
              slashes, but mktemp creates only the final component

       -t     interpret TEMPLATE as a single file name component, relative  to
              a  directory:  $TMPDIR, if set; else the directory specified via
              -p; else /tmp [deprecated]

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Jim Meyering and Eric Blake.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       mkstemp(3), mkdtemp(3), mktemp(3)

       Full documentation <https://www.gnu.org/software/coreutils/mktemp>
       or available locally via: info '(coreutils) mktemp invocation'

GNU coreutils 9.4                April 2024                        MKTEMP(1)
```

# Bash 的 more 命令

```
more [OPTION]... FILE...
```

**功能**

分页查看文件内容。

**类型**

可执行文件（`/usr/bin/more`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - **`-d`**，**`--silent`**
    - 提示：“[按空格继续，按 'q' 退出]”；
    - 当按下非法按键时，不响铃，而是显示：“[按 'h' 查看说明]”。
  - **`-l`**，**`--logical`**
    - 遇到包含 `^L`（换页符）的行时不暂停。
  - **`-e`**，**`--exit-on-eof`**
    - 读到文件末尾（EOF）立即退出。
    - 若未设置 `POSIXLY_CORRECT` 环境变量或未在终端执行时，该行为默认开启。
  - **`-f`**，**`--no-pause`**
    - 按“逻辑行”计数，而不是按屏幕行计数（即长行不会被折行）。
  - **`-p`**，**`--print-over`**
    - 不滚动屏幕，而是清空整个屏幕后再显示文本。
    - 注意：如果可执行文件名为 `page`，此选项会被自动启用。
  - **`-c`**，**`--clean-print`**
    - 不滚动屏幕，而是从屏幕顶部重新绘制，每行显示时清除其余内容。
  - **`-s`**，**`--squeeze`**

        - 将多个连续空行压缩为一个空行。
    - `-u`, `--plain`
        - 取消下划线显示。
        - 为向后兼容，此选项实际上被忽略。
    - `-n`, `--lines number`
        - 指定每屏显示的行数。`number` 必须是正整数。
        - `--lines` 设置的值优先级最高，会覆盖终端报告的行数等其他来源的值。
    - `-number`
        - 与 `--lines number` 等价。
    - `+number`
        - 从第 number 行开始显示文件内容。
    - `+/string`
        - 在开始显示前，先在文件中搜索 string，并从首次出现处开始显示。
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `FILE` 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 相关命令

| 命令 | 说明 |
| --- | --- |
| less | 增强版文件分页阅读命令 |

## 推荐阅读

- more(1) — Linux manual page

## 手册

```
MORE(1)                          User Commands                          MORE(1)

NAME
       more - display the contents of a file in a terminal
```

```
SYNOPSIS
       more [options] file ...

DESCRIPTION
       more is a filter for paging through text one screenful at a time. This
       version is especially primitive. Users should realize that less(1)
       provides more(1) emulation plus extensive enhancements.

OPTIONS
       Options are also taken from the environment variable MORE (make sure to
       precede them with a dash (-)) but command-line options will override
       those.

       -d, --silent
           Prompt with "[Press space to continue, 'q' to quit.]", and display
           "[Press 'h' for instructions.]" instead of ringing the bell when an
           illegal key is pressed.

       -l, --logical
           Do not pause after any line containing a ^L (form feed).

       -e, --exit-on-eof
           Exit on End-Of-File, enabled by default if POSIXLY_CORRECT
           environment variable is not set or if not executed on terminal.

       -f, --no-pause
           Count logical lines, rather than screen lines (i.e., long lines are
           not folded).

       -p, --print-over
           Do not scroll. Instead, clear the whole screen and then display the
           text. Notice that this option is switched on automatically if the
           executable is named page.

       -c, --clean-print
           Do not scroll. Instead, paint each screen from the top, clearing
           the remainder of each line as it is displayed.

       -s, --squeeze
           Squeeze multiple blank lines into one.

       -u, --plain
           Suppress underlining. This option is silently ignored as backwards
           compatibility.

       -n, --lines number
           Specify the number of lines per screenful. The number argument is a
           positive decimal integer. The --lines option shall override any
           values obtained from any other source, such as number of lines
           reported by terminal.

       -number
           A numeric option means the same as --lines option argument.
```

```
       +number
           Start displaying each file at line number.

       +/string
           The string to be searched in each file before starting to display
           it.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

COMMANDS
       Interactive commands for more are based on vi(1). Some commands may be
       preceded by a decimal number, called k in the descriptions below. In
       the following descriptions, ^X means control-X.

       h or ?
           Help; display a summary of these commands. If you forget all other
           commands, remember this one.

       SPACE
           Display next k lines of text. Defaults to current screen size.

       z
           Display next k lines of text. Defaults to current screen size.
           Argument becomes new default.

       RETURN
           Display next k lines of text. Defaults to 1. Argument becomes new
           default.

       d or ^D
           Scroll k lines. Default is current scroll size, initially 11.
           Argument becomes new default.

       q or Q or INTERRUPT
           Exit.

       s
           Skip forward k lines of text. Defaults to 1.

       f
           Skip forward k screenfuls of text. Defaults to 1.

       b or ^B
           Skip backwards k screenfuls of text. Defaults to 1. Only works with
           files, not pipes.

       '
           Go to the place where the last search started.
```

=
    Display current line number.

/pattern
    Search for kth occurrence of regular expression. Defaults to 1.

n
    Search for kth occurrence of last regular expression. Defaults to
    1.

!command or :!command
    Execute command in a subshell.

v
    Start up an editor at current line. The editor is taken from the
    environment variable VISUAL if defined, or EDITOR if VISUAL is not
    defined, or defaults to vi(1) if neither VISUAL nor EDITOR is
    defined.

^L
    Redraw screen.

:n
    Go to kth next file. Defaults to 1.

:p
    Go to kth previous file. Defaults to 1.

:f
    Display current file name and line number.

.
    Repeat previous command.

ENVIRONMENT
The more command respects the following environment variables, if they
exist:

MORE
    This variable may be set with favored options to more.

SHELL
    Current shell in use (normally set by the shell at login time).

TERM
    The terminal type used by more to get the terminal characteristics
    necessary to manipulate the screen.

VISUAL
    The editor the user prefers. Invoked when command key v is pressed.

EDITOR

```
            The editor of choice when VISUAL is not specified.

       POSIXLY_CORRECT
            Disable exit-on-eof (see option -e for more details).

HISTORY
       The more command appeared in 3.0BSD. This man page documents more
       version 5.19 (Berkeley 6/29/88), which is currently in use in the Linux
       community. Documentation was produced using several other versions of
       the man page, and extensive inspection of the source code.

AUTHORS
       Eric Shienbrood, UC Berkeley.

       Modified by Geoff Peck, UCB to add underlining, single spacing.

       Modified by John Foderaro, UCB to add -c and MORE environment variable.

SEE ALSO
       less(1), vi(1)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The more command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-11-21                          MORE(1)
```

# Bash 的 mountpoint 命令

```
mountpoint [OPTION]... FILE
```

**功能**

判断文件是否是挂载点。

**类型**

可执行文件（`/usr/bin/mountpoint`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-d`，`--fs-devno` - 查看挂载点对应设备的 `主设备号:次设备号`
  - `-q`，`--quiet` - 不打印任何内容，仅通过返回值返回
  - `--nofollow` - 检测符号链接本身，而非其链接的文件
  - `-x`，`--devno` - 判断文件是否是块设备，如果是则查看该块设备的 `主设备号:次设备号`
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FILE` - 要判断的文件

**返回值**

- 0 - 是挂载点（或块设备）
- 1 - 调用失败
- 32 - 不是挂载点（或块设备）

## 示例

```
$ mountpoint /                  # 查看 / 是否是挂载点
/ is a mountpoint
$ mountpoint /home              # 查看 /home 是否是挂载点
/home is not a mountpoint
$ mountpoint /run               # 查看 /run 是否是挂载点
/run is a mountpoint
$ mountpoint /dev/nvme0n1p1     # 查看 /dev/nvme0n1p1 是否是挂载点
/dev/nvme0n1p1 is not a mountpoint
$ mountpoint /dev/nvme0n1p1 -x  # 查看 /dev/nvme0n1p1 是否是块设备，打印设备号
259:1
```

## 推荐阅读

- mountpoint(1) — Linux manual page

## 手册

```
MOUNTPOINT(1)                   User Commands                   MOUNTPOINT(1)

NAME
       mountpoint - see if a directory or file is a mountpoint

SYNOPSIS
       mountpoint [-d|-q] directory|file

       mountpoint -x device

DESCRIPTION
       mountpoint checks whether the given directory or file is mentioned in
       the /proc/self/mountinfo file.

OPTIONS
       -d, --fs-devno
           Show the major/minor numbers of the device that is mounted on the
           given directory.

       -q, --quiet
           Be quiet - don't print anything.

       --nofollow
           Do not follow symbolic link if it the last element of the directory
           path.

       -x, --devno
           Show the major/minor numbers of the given blockdevice on standard
           output.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       mountpoint has the following exit status values:

       0
           success; the directory is a mountpoint, or device is block device
           on --devno
```

```
       1
           failure; incorrect invocation, permissions or system error

       32
           failure; the directory is not a mountpoint, or device is not a
           block device on --devno

ENVIRONMENT
       LIBMOUNT_DEBUG=all
           enables libmount debug output.

NOTES
       The util-linux mountpoint implementation was written from scratch for
       libmount. The original version for sysvinit suite was written by Miquel
       van Smoorenburg.

AUTHORS
       Karel Zak <kzak@redhat.com>

SEE ALSO
       mount(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The mountpoint command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                     MOUNTPOINT(1)
```

# Bash 的 mv 命令

```
mv [OPTION]... SOURCE... DIRECTORY
```

**功能**

移动（或重命名）文件。

**类型**

可执行文件（`/usr/bin/mv`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--backup[=CONTROL]` - 选中怎样备份被覆盖的目标文件
  - `-b` - 备份被覆盖的目标文件
  - `--debug` - 输出调试信息（文件的复制方式）
  - `-f`，`--force` - 强制操作，目标存在且无法打开时会尝试将其删除
  - `-i`，`--interactive` - 覆盖前提示
  - `-n`，`--no-clobber` - 不覆盖现有文件，也不失败
  - `--no-copy` - 重命名失败时不复制文件
  - `--strip-trailing-slashes` - 删除每个 `SOURCE` 尾部的 `/`
  - `-S`，`--suffix=SUFFIX` - 设置备份文件的后缀
  - `-t`，`--target-directory=DIRECTORY` - 指定目标目录
  - `-T`，`--no-target-directory` - 将 `DEST` 视为文件（如果 `DEST` 是目录，则进行合并而非移动到它内部）
  - `--update[=UPDATE]` - 控制哪些现有文件需要更新；取值：`all`，`none`，`older`
  - `-u` - 仅在源文件比目标新的时候复制；同 `--update[=older]`
  - `-v`，`--verbose` - 输出详细信息（进行的全部操作）
  - `-Z` - 将目标文件的 SELinux 安全上下文设为默认类型
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `SOURCE` - 源文件列表
- `DEST` - 目标路径

## 示例

### 移动文件

```
$ ls                            # 空目录下操作
$ touch src.txt                 # 创建文件
$ ls
src.txt
$ mv src.txt dest.txt           # 重命名文件
$ ls
dest.txt
$ mkdir dir                     # 创建目录
$ mv src.txt ./dir              # 移动到目录中
$ ls dir
src.txt
```

### 移动目录

```
$ mkdir src dest1 dest2         # 创建目录
$ touch src/1.txt               # 创建文件
$ tree src
src
└── 1.txt
$ mv src dest1                  # 将 src 移动到 dest1 内部
$ tree dest1
dest1
└── src
    └── 1.txt
$ mv -T dest1 dest2             # 移动 dest1 合并到 dest2
$ tree dest2
dest2
└── 1.txt
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| cp | 复制文件 |
| ln | 链接文件 |

## 推荐阅读

- Move (rename) files - GNU Coreutils
- mv(1) — Linux manual page

## 手册

```
MV(1)                           User Commands                           MV(1)

NAME
       mv - move (rename) files

SYNOPSIS
       mv [OPTION]... [-T] SOURCE DEST
       mv [OPTION]... SOURCE... DIRECTORY
       mv [OPTION]... -t DIRECTORY SOURCE...

DESCRIPTION
       Rename SOURCE to DEST, or move SOURCE(s) to DIRECTORY.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       --backup[=CONTROL]
              make a backup of each existing destination file

       -b     like --backup but does not accept an argument

       --debug
              explain how a file is copied.  Implies -v

       -f, --force
              do not prompt before overwriting

       -i, --interactive
              prompt before overwrite

       -n, --no-clobber
              do not overwrite an existing file

       If you specify more than one of -i, -f, -n, only the  final  one  takes
       effect.

       --no-copy
              do not copy if renaming fails

       --strip-trailing-slashes
              remove any trailing slashes from each SOURCE argument

       -S, --suffix=SUFFIX
              override the usual backup suffix

       -t, --target-directory=DIRECTORY
              move all SOURCE arguments into DIRECTORY

       -T, --no-target-directory
              treat DEST as a normal file

       --update[=UPDATE]
```

```
              control     which     existing     files     are     updated;     UP-
              DATE={all,none,older(default)}.  See below

       -u     equivalent to --update[=older]

       -v, --verbose
              explain what is being done

       -Z, --context
              set SELinux security context of destination file to default type

       --help display this help and exit

       --version
              output version information and exit

       UPDATE controls which existing files in the destination  are  replaced.
       'all'  is  the  default operation when an --update option is not speci-
       fied, and results in all existing files in the  destination  being  re-
       placed.  'none' is similar to the --no-clobber option, in that no files
       in the destination are replaced, but also skipped files do not induce a
       failure.   'older' is the default operation when --update is specified,
       and results in files being replaced if they're older  than  the  corre-
       sponding source file.

       The   backup   suffix   is   '~', unless set with --suffix or SIM-
       PLE_BACKUP_SUFFIX.  The version control method may be selected via  the
       --backup  option  or  through the VERSION_CONTROL environment variable.
       Here are the values:

       none, off
              never make backups (even if --backup is given)

       numbered, t
              make numbered backups

       existing, nil
              numbered if numbered backups exist, simple otherwise

       simple, never
              always make simple backups

AUTHOR
       Written by Mike Parker, David MacKenzie, and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
```

```
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       rename(2)

       Full documentation <https://www.gnu.org/software/coreutils/mv>
       or available locally via: info '(coreutils) mv invocation'

GNU coreutils 9.4                April 2024                            MV(1)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 namei 命令

```
namei [OPTION]... PATH
```

**功能**

逐级访问路径并显示其信息。

**类型**

可执行文件（`/usr/bin/namei`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-l`，`--long` - 使用长格式输出；等价于 `-m -o -v`
  - `-m`，`--modes` - 打印文件权限
  - `-n`，`--nosymlinks` - 不跟随符号链接
  - `-o`，`--owners` - 打印文件的所有者
  - `-v`，`--vertical` - 垂直对齐
  - `-x`，`--mountpoints` - 作为挂载点的目录显示为 `D` 而不是 `d`
  - `-Z`，`--context` - 显示安全上下文
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FILE` - 要判断的文件

## 文件类型标识

- `f:` - 正在解析的路径
- `d` - 目录
- `l` - 符号链接
- `s` - socket
- `b` - 块设备
- `c` - 字符设备
- `p` - FIFO (命名管道)
- `-` - 普通文件
- `?` - 错误

# 示例

### 检查路径错误的具体位置

```
$ namei /home/user1/dir/file        # 检查 /home/user1/dir/file
f: /var/www/html/dasdaxasds
 d /
 d home
 d user1
    dir - No such file or directory
```

### 检查权限错误的具体位置

```
$ namei /home/user2/dir/file        # 检查 /home/user2/dir/file
f: /home/user2/dir/file
 d /
 d home
 d user2
    dir - Permission denied
```

### 调试 Web 服务访问失败的原因

```
$ sudo -u www-data namei /var/www/html/index.html   # 以 www-data 用户的权限检查
/var/www/html/index.html
f: /var/www/html/index.html
 d /
 d var
 d www
    html - Permission denied
```

# 推荐阅读

- namei(1) — Linux manual page

# 手册

```
NAMEI(1)                        User Commands                        NAMEI(1)

NAME
       namei - follow a pathname until a terminal point is found

SYNOPSIS
       namei [options] pathname...

DESCRIPTION
```

namei interprets its arguments as pathnames to any type of Unix file (symlinks, files, directories, and so forth). namei then follows each pathname until an endpoint is found (a file, a directory, a device node, etc). If it finds a symbolic link, it shows the link, and starts following it, indenting the output to show the context.

This program is useful for finding "too many levels of symbolic links" problems.

For each line of output, namei uses the following characters to identify the file type found:

```
f: = the pathname currently being resolved
 d = directory
 l = symbolic link (both the link and its contents are output)
 s = socket
 b = block device
 c = character device
 p = FIFO (named pipe)
 - = regular file
 ? = an error of some kind
```

namei prints an informative message when the maximum number of symbolic links this system can have has been exceeded.

OPTIONS

-l, --long
Use the long listing format (same as -m -o -v).

-m, --modes
Show the mode bits of each file type in the style of ls(1), for example 'rwxr-xr-x'.

-n, --nosymlinks
Don’t follow symlinks.

-o, --owners
Show owner and group name of each file.

-v, --vertical
Vertically align the modes and owners.

-x, --mountpoints
Show mountpoint directories with a 'D' rather than a 'd'.

-Z, --context
Show security context of the file or "?" if not available. The support for security contexts is optional and does not have to be compiled to the namei binary.

-h, --help
Display help text and exit.

```
       -V, --version
           Print version and exit.

BUGS
       To be discovered.

AUTHORS
       The original namei program was written by Roger Southwick
       <rogers@amadeus.wr.tek.com>.

       The program was rewritten by Karel Zak Karel Zak <kzak@redhat.com>.

SEE ALSO
       ls(1), stat(1), symlink(7)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The namei command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                          NAMEI(1)
```

# Bash 的 nice 命令

```
nice [OPTION] [COMMAND [ARG]...]
```

**功能**

运行命令时附带调度优先级。

**类型**

可执行文件（`/usr/bin/nice`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-n`, `--adjustment=N` - 将优先级设为 `N`；范围是 `[-20, 19]`，数值越小越优先，默认为 10
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `COMMAND` - 要运行的命令；如果不传递此参数则打印当前优先级
- `ARG` - 传递给 `COMMAND` 的参数列表

## 示例

```
$ nice                          # 查看直接运行命令的优先级（通常是 0）
0
$ nice my_cmd                   # 默认将优先级设为 10，低于直接运行命令的优先级
$ sudo nice -n -20 my_cmd       # 默认将优先级设为 -20，只有 root 用户才能提高优先级（小于 0）
```

## 推荐阅读

- Run a command with modified niceness - GNU Coreutils
- nice(1) — Linux manual page

## 手册

```
NICE(1)                         User Commands                         NICE(1)
```

NAME
       nice - run a program with modified scheduling priority

SYNOPSIS
       nice [OPTION] [COMMAND [ARG]...]

DESCRIPTION
       Run  COMMAND  with an adjusted niceness, which affects process schedul-
       ing.  With no COMMAND, print the  current  niceness.   Niceness  values
       range  from  -20 (most favorable to the process) to 19 (least favorable
       to the process).

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -n, --adjustment=N
              add integer N to the niceness (default 10)

       --help display this help and exit

       --version
              output version information and exit

       NOTE: your shell may have its own version of nice, which usually super-
       sedes  the  version described here.  Please refer to your shell's docu-
       mentation for details about the options it supports.

   Exit status:
       125    if the nice command itself fails

       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found

       -      the exit status of COMMAND otherwise

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       nice(2), renice(1)

       Full documentation <https://www.gnu.org/software/coreutils/nice>

```
       or available locally via: info '(coreutils) nice invocation'

GNU coreutils 9.4                 April 2024                           NICE(1)
```

# Bash 的 nl 命令

```
nl [OPTION]... [FILE]...
```

**功能**

查看文件并添加行号。

**类型**

可执行文件（`/usr/bin/nl`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--body-numbering=STYLE` - 使用 `STYLE` 作为正文编号 风格；默认为 `n` 不编号
  - `-d`，`--section-delimiter=CC` - 使用 `CC` 作为逻辑页的分隔符，三个分隔符表示标题，两个表示正文，一个表示页脚；默认为 `'\:'`
  - `-f`，`--footer-numbering=STYLE` - 使用 `STYLE` 作为页脚编号 风格；默认为 `n` 不编号
  - `-h`，`--header-numbering=STYLE` - 使用 `STYLE` 作为标题编号 风格；默认为 `t` 对非空行编号
  - `-i`，`--line-increment=NUMBER` - 每行编号递增 `NUMBER`
  - `-l`，`--join-blank-lines=NUMBER` - 将 `NUMBER` 行空行视作一行
  - `-n`，`--number-format=FORMAT` - 行号 格式
  - `-p`，`--no-renumber` - 新章节不重置行号
  - `-s`，`--number-separator=STRING` - 在行号后面添加字符串 `STRING`
  - `-v`，`--starting-line-number=NUMBER` - 行号的初始值为 `NUMBER`
  - `-w`，`--number-width=NUMBER` - 行号的宽度为 `NUMBER`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 风格

| 风格 | 说明 |
| --- | --- |
| `a` | 对所有行编号 |
| `t` | 对非空行编号 |

| 风格 | 说明 |
| --- | --- |
| `n` | 不编号 |
| `pBRE` | 对正则表达式 `BRE` 匹配的行编号 |

## 格式

| 格式 | 说明 |
| --- | --- |
| `ln` | 左对齐 |
| `rn` | 右对齐 |
| `rz` | 右对齐，左侧补零 |

## 示例

### 基本使用

```
$ cat 1.txt                             # 查看文件
Primers 编程伙伴
https://xplanc.org/primers/
$ nl 1.txt                              # 附带行号
     1  Primers 编程伙伴
     2  https://xplanc.org/primers/
$ nl -v 100 1.txt                       # 行号从 100 开始
   100  Primers 编程伙伴
   101  https://xplanc.org/primers/
$ nl -n ln 1.txt                        # 左对齐
1       Primers 编程伙伴
2       https://xplanc.org/primers/
```

### 标题和页脚

```
$ cat 1.txt                         # 查看文件
\:\:\:
Primers 编程伙伴
\:\:
Primers 致力于为各类编程学习者提供全面、系统的编程教程和实践资源。

无论你是编程新手，还是有一定基础的开发者，Primers 都提供了适合的学习路径和丰富的实战项目，帮助你从零开始，逐步掌握编程技能，成为一名优秀的开发者。
\:
https://xplanc.org/primers/
$ nl 1.txt                          # 默认只编号正文的非空行

       Primers 编程伙伴
```

```
     1  Primers 致力于为各类编程学习者提供全面、系统的编程教程和实践资源。

     2  无论你是编程新手，还是有一定基础的开发者，Primers 都提供了适合的学习路径和丰富的实战项目，帮助你从零开始，逐步掌握编程技能，成为一名优秀的开发者。

       https://xplanc.org/primers/
$ nl -h a -f a -b a 1.txt       # 标题、页脚、正文都编号所有行

     1  Primers 编程伙伴

     1  Primers 致力于为各类编程学习者提供全面、系统的编程教程和实践资源。
     2
     3  无论你是编程新手，还是有一定基础的开发者，Primers 都提供了适合的学习路径和丰富的实战项目，帮助你从零开始，逐步掌握编程技能，成为一名优秀的开发者。

     1  https://xplanc.org/primers/
$ nl -h a -f a -b a -p 1.txt    # 不重置行号

     1  Primers 编程伙伴

     2  Primers 致力于为各类编程学习者提供全面、系统的编程教程和实践资源。
     3
     4  无论你是编程新手，还是有一定基础的开发者，Primers 都提供了适合的学习路径和丰富的实战项目，帮助你从零开始，逐步掌握编程技能，成为一名优秀的开发者。

     5  https://xplanc.org/primers/
```

- 三个 `\:` 表示标题
- 两个 `\:` 表示正文
- 一个 `\:` 表示页脚

## 推荐阅读

- Number lines and write files - GNU Coreutils
- nl(1) — Linux manual page

## 手册

```
NL(1)                           User Commands                           NL(1)

NAME
       nl - number lines of files

SYNOPSIS
       nl [OPTION]... [FILE]...

DESCRIPTION
```

```
Write each FILE to standard output, with line numbers added.

With no FILE, or when FILE is -, read standard input.

Mandatory  arguments  to  long  options are mandatory for short options
too.

-b, --body-numbering=STYLE
       use STYLE for numbering body lines

-d, --section-delimiter=CC
       use CC for logical page delimiters

-f, --footer-numbering=STYLE
       use STYLE for numbering footer lines

-h, --header-numbering=STYLE
       use STYLE for numbering header lines

-i, --line-increment=NUMBER
       line number increment at each line

-l, --join-blank-lines=NUMBER
       group of NUMBER empty lines counted as one

-n, --number-format=FORMAT
       insert line numbers according to FORMAT

-p, --no-renumber
       do not reset line numbers for each section

-s, --number-separator=STRING
       add STRING after (possible) line number

-v, --starting-line-number=NUMBER
       first line number for each section

-w, --number-width=NUMBER
       use NUMBER columns for line numbers

--help display this help and exit

--version
       output version information and exit

Default options are: -bt -d'\:' -fn -hn -i1 -l1 -n'rn' -s<TAB> -v1 -w6

CC are two delimiter characters used to construct logical  page  delim-
iters;  a missing second character implies ':'.  As a GNU extension one
can specify more than two characters, and  also  specifying  the  empty
string (-d '') disables section matching.

STYLE is one of:
```

```
       a      number all lines

       t      number only nonempty lines

       n      number no lines

       pBRE   number only lines that contain a match for the basic regular ex-
              pression, BRE

       FORMAT is one of:

       ln     left justified, no leading zeros

       rn     right justified, no leading zeros

       rz     right justified, leading zeros

AUTHOR
       Written by Scott Bartram and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/nl>
       or available locally via: info '(coreutils) nl invocation'

GNU coreutils 9.4                April 2024                          NL(1)
```

# Bash 的 nohup 命令

```
nohup [OPTION] [COMMAND [ARG]...]
```

**功能**

运行命令时忽略挂断信号（SIGHUP），以便该命令在用户注销后仍可在后台继续运行。

**类型**

可执行文件（`/usr/bin/nohup`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `COMMAND` - 要运行的命令；如果不传递此参数则打印当前优先级
- `ARG` - 传递给 `COMMAND` 的参数列表

## 示例

```
$ nohup java -jar minecraft_server.jar nogui &      # & 表示后台运行
```

- 在后台运行 `minecraft_server` 并忽略挂断信号（SIGHUP），用户注销后仍可继续运行。

## 推荐阅读

- Run a command immune to hangups - GNU Coreutils
- nohup(1) — Linux manual page

## 手册

```
NOHUP(1)                          User Commands                          NOHUP(1)

NAME
       nohup - run a command immune to hangups, with output to a non-tty

SYNOPSIS
```

```
       nohup COMMAND [ARG]...
       nohup OPTION

DESCRIPTION
       Run COMMAND, ignoring hangup signals.

       --help display this help and exit

       --version
              output version information and exit

       If  standard  input is a terminal, redirect it from an unreadable file.
       If standard output is a terminal, append output to 'nohup.out' if  pos-
       sible,  '$HOME/nohup.out'  otherwise.  If standard error is a terminal,
       redirect it to standard output.  To save output  to  FILE,  use  'nohup
       COMMAND > FILE'.

       NOTE:  your  shell may have its own version of nohup, which usually su-
       persedes the version described here.  Please refer to your shell's doc-
       umentation for details about the options it supports.

   Exit status:
       125    if the nohup command itself fails

       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found

       -      the exit status of COMMAND otherwise

AUTHOR
       Written by Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/nohup>
       or available locally via: info '(coreutils) nohup invocation'

GNU coreutils 9.4                April 2024                          NOHUP(1)
```

# Bash 的 nproc 命令

```
nproc [OPTION]...
```

**功能**

打印中央处理器（CPU）可用的逻辑核心数量。

> 参考 `/sys/devices/system/cpu/` 目录。

**类型**

可执行文件（`/usr/bin/nproc`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--all` - 打印已安装的全部处理器数量
  - `--ignore=N` - 忽略 `N` 个逻辑核心
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ nproc
8
```

## 推荐阅读

- Print the number of available processors - GNU Coreutils
- nproc(1) — Linux manual page

## 手册

```
NPROC(1)                        User Commands                        NPROC(1)

NAME
       nproc - print the number of processing units available
```

```
SYNOPSIS
       nproc [OPTION]...

DESCRIPTION
       Print  the number of processing units available to the current process,
       which may be less than the number of online processors

       --all  print the number of installed processors

       --ignore=N
              if possible, exclude N processing units

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Giuseppe Scrivano.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/nproc>
       or available locally via: info '(coreutils) nproc invocation'

GNU coreutils 9.4                April 2024                          NPROC(1)
```

# Bash 的 nsenter 命令

```
nsenter [OPTION]... [CMD [ARG]...]
```

**功能**

进入指定的命名空间然后运行命令。

**类型**

可执行文件（`/usr/bin/nsenter`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 根据默认的命名空间路径 `/proc/${PID}/ns/*` 进入目标进程的所有命名空间；与 `t` 选项一同使用
  - `-t`，`--target PID` - 通过 `PID` 指定目标进程；进入目标进程的命名空间
  - `-m`，`--mount[=file]` - 进入 `file` 指定的 mount namespace；如果没有 `file`，则进入目标进程的 mount namespace
  - `-u`，`--uts[=file]` - 进入 `file` 指定的 UTS namespace；如果没有 `file`，则进入目标进程的 UTS namespace
  - `-i`，`--ipc[=file]` - 进入 `file` 指定的 IPC namespace；如果没有 `file`，则进入目标进程的 IPC namespace
  - `-n`，`--net[=file]` - 进入 `file` 指定的 network namespace；如果没有 `file`，则进入目标进程的 network namespace
  - `-p`，`--pid[=file]` - 进入 `file` 指定的 PID namespace；如果没有 `file`，则进入目标进程的 PID namespace
  - `-U`，`--user[=file]` - 进入 `file` 指定的 user namespace；如果没有 `file`，则进入目标进程的 user namespace
  - `-C`，`--cgroup[=file]` - 进入 `file` 指定的 cgroup namespace；如果没有 `file`，则进入目标进程的 cgroup namespace
  - `-T`，`--time[=file]` - 进入 `file` 指定的 time namespace；如果没有 `file`，则进入目标进程的 time namespace
  - `-G`，`--setgid gid` - 设置组 ID 并删除补充组 ID；默认为 0
  - `-S`，`--setuid uid` - 设置用户 ID；默认为 0
  - `--preserve-credentials` - 不修改用户 ID 和 组 ID
  - `-r`，`--root[=directory]` - 设置根目录
  - `-w`，`--wd[=directory]` - 设置工作目录；在进入命名空间之前打开

    - `-W`, `--wdns[=directory]` - 设置工作目录；在 chroot 指定打开
    - `-e`, `--env` - 将目标进程的环境变量传递给新创建的进程
    - `-F`, `--no-fork` - 不要在 exec 之前调用 fork 函数
    - `-Z`, `--follow-context` - 将目标进程的安全上下文传递给新创建的进程
    - `-h`, `--help` - 显示帮助
    - `-V`, `--version` - 显示版本
- `CMD` - 要执行的命令；默认执行 `${SHELL}`
- `ARG` - 命令的参数列表

## Linux 的命名空间

| 命名空间 | 中文说明 |
| --- | --- |
| mount namespace | 挂载或卸载文件系统不会影响系统的其他部分，除非该文件系统被显式标记为共享（使用 `mount --make-shared`；共享标志可在 `/proc/self/mountinfo` 中查看）。更多细节参见 `mount_namespaces(7)` 以及 `clone(2)` 中关于 `CLONE_NEWNS` 标志的说明。 |
| UTS namespace | 设置主机名（hostname）或域名（domainname）不会影响系统的其他部分。更多细节参见 `uts_namespaces(7)`。 |
| IPC namespace | 进程将拥有独立的 IPC 命名空间，包括 POSIX 消息队列，以及 System V 消息队列、信号量集合和共享内存段。更多细节参见 `ipc_namespaces(7)`。 |
| network namespace | 进程将拥有独立的 IPv4 和 IPv6 协议栈、IP 路由表、防火墙规则、`/proc/net` 和 `/sys/class/net` 目录树、套接字等资源。更多细节参见 `network_namespaces(7)`。 |
| PID namespace | 子进程将拥有一套与 `nsenter` 进程分离的 PID 到进程映射关系。如果更改 PID 命名空间，`nsenter` 默认会执行 `fork`，以便新程序及其子进程共享同一个 PID 命名空间并彼此可见。如果使用 `--no-fork`，则新程序将在不进行 `fork` 的情况下通过 `exec` 启动。更多细节参见 `pid_namespaces(7)`。 |
| user namespace | 进程将拥有独立的一组 UID、GID 以及能力（capabilities）。更多细节参见 `user_namespaces(7)`。 |
| cgroup namespace | 进程将看到虚拟化后的 `/proc/self/cgroup` 视图，新创建的 cgroup 挂载点将以该命名空间的 cgroup 根目录为起点。更多细节参见 `cgroup_namespaces(7)`。 |
| time namespace | 进程可以拥有独立的 `CLOCK_MONOTONIC` 和/或 `CLOCK_BOOTTIME` 视图，这些时间偏移可通过 `/proc/self/timens_offsets` 进行修改。更多细节参见 `time_namespaces(7)`。 |

## 示例

```
$ nsenter -t 12345 -a ls    # 进入进程 12345 的所有命名空间，执行 ls
$ nsenter -t 12345 -a       # 进入进程 12345 的所有命名空间，执行 /bin/bash
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| unshare | 创建新的命名空间执行命令 |

## 推荐阅读

- nsenter(1) — Linux manual page

## 手册

```
NSENTER(1)                      User Commands                      NSENTER(1)

NAME
       nsenter - run program in different namespaces

SYNOPSIS
       nsenter [options] [program [arguments]]

DESCRIPTION
       The nsenter command executes program in the namespace(s) that are
       specified in the command-line options (described below). If program is
       not given, then "${SHELL}" is run (default: /bin/sh).

       Enterable namespaces are:

       mount namespace
           Mounting and unmounting filesystems will not affect the rest of the
           system, except for filesystems which are explicitly marked as
           shared (with mount --make-shared; see /proc/self/mountinfo for the
           shared flag). For further details, see mount_namespaces(7) and the
           discussion of the CLONE_NEWNS flag in clone(2).

       UTS namespace
           Setting hostname or domainname will not affect the rest of the
           system. For further details, see uts_namespaces(7).

       IPC namespace
           The process will have an independent namespace for POSIX message
           queues as well as System V message queues, semaphore sets and
```

shared memory segments. For further details, see ipc_namespaces(7).

network namespace
The process will have independent IPv4 and IPv6 stacks, IP routing tables, firewall rules, the /proc/net and /sys/class/net directory trees, sockets, etc. For further details, see network_namespaces(7).

PID namespace
Children will have a set of PID to process mappings separate from the nsenter process. nsenter will fork by default if changing the PID namespace, so that the new program and its children share the same PID namespace and are visible to each other. If --no-fork is used, the new program will be exec'ed without forking. For further details, see pid_namespaces(7).

user namespace
The process will have a distinct set of UIDs, GIDs and capabilities. For further details, see user_namespaces(7).

cgroup namespace
The process will have a virtualized view of /proc/self/cgroup, and new cgroup mounts will be rooted at the namespace cgroup root. For further details, see cgroup_namespaces(7).

time namespace
The process can have a distinct view of CLOCK_MONOTONIC and/or CLOCK_BOOTTIME which can be changed using /proc/self/timens_offsets. For further details, see time_namespaces(7).

OPTIONS
Various of the options below that relate to namespaces take an optional file argument. This should be one of the /proc/[pid]/ns/* files described in namespaces(7), or the pathname of a bind mount that was created on one of those files.

-a, --all
Enter all namespaces of the target process by the default /proc/[pid]/ns/* namespace paths. The default paths to the target process namespaces may be overwritten by namespace specific options (e.g., --all --mount=[path]).

The user namespace will be ignored if the same as the caller's current user namespace. It prevents a caller that has dropped capabilities from regaining those capabilities via a call to setns(). See setns(2) for more details.

-t, --target PID
Specify a target process to get contexts from. The paths to the contexts specified by pid are:

/proc/pid/ns/mnt

```
            the mount namespace

        /proc/pid/ns/uts
            the UTS namespace

        /proc/pid/ns/ipc
            the IPC namespace

        /proc/pid/ns/net
            the network namespace

        /proc/pid/ns/pid
            the PID namespace

        /proc/pid/ns/user
            the user namespace

        /proc/pid/ns/cgroup
            the cgroup namespace

        /proc/pid/ns/time
            the time namespace

        /proc/pid/root
            the root directory

        /proc/pid/cwd
            the working directory respectively

    -m, --mount[=file]
        Enter the mount namespace. If no file is specified, enter the mount
        namespace of the target process. If file is specified, enter the
        mount namespace specified by file.

    -u, --uts[=file]
        Enter the UTS namespace. If no file is specified, enter the UTS
        namespace of the target process. If file is specified, enter the
        UTS namespace specified by file.

    -i, --ipc[=file]
        Enter the IPC namespace. If no file is specified, enter the IPC
        namespace of the target process. If file is specified, enter the
        IPC namespace specified by file.

    -n, --net[=file]
        Enter the network namespace. If no file is specified, enter the
        network namespace of the target process. If file is specified,
        enter the network namespace specified by file.

    -p, --pid[=file]
        Enter the PID namespace. If no file is specified, enter the PID
        namespace of the target process. If file is specified, enter the
        PID namespace specified by file.
```

```
-U, --user[=file]
    Enter the user namespace. If no file is specified, enter the user
    namespace of the target process. If file is specified, enter the
    user namespace specified by file. See also the --setuid and
    --setgid options.

-C, --cgroup[=file]
    Enter the cgroup namespace. If no file is specified, enter the
    cgroup namespace of the target process. If file is specified, enter
    the cgroup namespace specified by file.

-T, --time[=file]
    Enter the time namespace. If no file is specified, enter the time
    namespace of the target process. If file is specified, enter the
    time namespace specified by file.

-G, --setgid gid
    Set the group ID which will be used in the entered namespace and
    drop supplementary groups. nsenter always sets GID for user
    namespaces, the default is 0. If the argument "follow" is specified
    the GID of the target process is used.

-S, --setuid uid
    Set the user ID which will be used in the entered namespace.
    nsenter always sets UID for user namespaces, the default is 0. If
    the argument "follow" is specified the UID of the target process is
    used.

--preserve-credentials
    Don’t modify UID and GID when enter user namespace. The default is
    to drops supplementary groups and sets GID and UID to 0.

-r, --root[=directory]
    Set the root directory. If no directory is specified, set the root
    directory to the root directory of the target process. If directory
    is specified, set the root directory to the specified directory.
    The specified directory is open before it switches to the requested
    namespaces.

-w, --wd[=directory]
    Set the working directory. If no directory is specified, set the
    working directory to the working directory of the target process.
    If directory is specified, set the working directory to the
    specified directory. The specified directory is open before it
    switches to the requested namespaces, it means the specified
    directory works as "tunnel" to the current namespace. See also
    --wdns.

-W, --wdns[=directory]
    Set the working directory. The directory is open after switch to
    the requested namespaces and after chroot(2) call. The options --wd
    and --wdns are mutually exclusive.
```

```
       -e, --env
           Pass environment variables from the target process to the new
           process being created. If this option is not provided, the
           environment variables will remain the same as in the current
           namespace..

       -F, --no-fork
           Do not fork before exec’ing the specified program. By default, when
           entering a PID namespace, nsenter calls fork before calling exec so
           that any children will also be in the newly entered PID namespace.

       -Z, --follow-context
           Set the SELinux security context used for executing a new process
           according to already running process specified by --target PID.
           (The util-linux has to be compiled with SELinux support otherwise
           the option is unavailable.)

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

AUTHORS
       Eric Biederman <biederm@xmission.com>, Karel Zak <kzak@redhat.com>

SEE ALSO
       clone(2), setns(2), namespaces(7)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The nsenter command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                        NSENTER(1)
```

# Bash 的 numfmt 命令

```
numfmt [OPTION]... [NUMBER]...
```

**功能**

数值和字符串之间的转换。

**类型**

可执行文件（`/usr/bin/numfmt`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--debug` - 对无效输入打印警告
  - `-d`, `--delimiter=X` - 使用 `X` 代替空格作为分隔符
  - `--field=FIELDS` - 替换 `FIELDS` 指定的 字段（默认为 1）中的数字
  - `--format=FORMAT` - 使用 printf 风格的浮点数格式
  - `--from=UNIT` - 指定输入的 单位格式 为 `UNIT`
  - `--from-unit=N` - 指定输入的单位大小为 `N`；使用此选项但不指定 `N` 时默认为 1
  - `--grouping` - 使用区域定义的数值分组；例如 1,000,000
  - `--header[=N]` - 开头 `N` 行直接打印，不转换；默认为 1
  - `--invalid=MODE` - 无效数值的处理模式：
    - `abort` - 中断（默认）
    - `fail` - 失败
    - `warn` - 警告
    - `ignore` - 忽略
  - `--padding=N` - 将输出对齐到 `N` 个字符
  - `--round=METHOD` - 舍入方式：
    - `up` - 向上取整
    - `down` - 向下取整
    - `from-zero` - 从零远离（默认）
    - `towards-zero` - 向零靠拢
    - `nearest` - 最近整数
  - `--suffix=SUFFIX` - 给输出添加后缀 `SUFFIX`
  - `--to=UNIT` - 指定输出的 单位格式 为 `UNIT`
  - `--to-unit=N` - 指定输入的单位大小为 `N`；默认为 1

    - `-z`, `--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
    - `--ignore=N` - 忽略 `N` 个逻辑核心
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `NUMBER` - 要转换的数值

## 单位格式

| 单位格式 | 说明 |
| --- | --- |
| `none` | 没有单位，后缀会产生错误 |
| `auto` | 1K = 1000, 1Ki = 1024 |
| `si` | 1K = 1000 |
| `iec` | 1K = 1024 |
| `iec-i` | 1Ki = 1024 |

## 字段

| 字段 | 说明 |
| --- | --- |
| `N` | 第 N 个字段 |
| `N-` | 第 N 个字段到末尾 |
| `N-M` | 第 N 个字段到第 M 个字段 |
| `-M` | 从开始到第 M 个字段 |
| `-` | 所有字段 |

- 字段的开始序号是 1

## 示例

```
$ numfmt --to=si 100000000
100M
$ numfmt --to=iec 100000000
96M
$ numfmt --from=si 32G
32000000000
$ $ numfmt --from=si 32G
32000000000
```

```
$ numfmt --to=iec --from=si 1T      # 1T 硬盘的实际容量
932G
```

## 推荐阅读

- Reformat numbers - GNU Coreutils
- numfmt(1) — Linux manual page

## 手册

```
NUMFMT(1)                          User Commands                          NUMFMT(1)

NAME
       numfmt - Convert numbers from/to human-readable strings

SYNOPSIS
       numfmt [OPTION]... [NUMBER]...

DESCRIPTION
       Reformat  NUMBER(s),  or  the  numbers  from standard input if none are
       specified.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       --debug
              print warnings about invalid input

       -d, --delimiter=X
              use X instead of whitespace for field delimiter

       --field=FIELDS
              replace  the  numbers  in  these  input  fields (default=1); see
              FIELDS below

       --format=FORMAT
              use printf style floating-point FORMAT; see FORMAT below for de-
              tails

       --from=UNIT
              auto-scale input numbers to UNITs; default is 'none';  see  UNIT
              below

       --from-unit=N
              specify the input unit size (instead of the default 1)

       --grouping
              use  locale-defined  grouping  of  digits, e.g. 1,000,000 (which
```

```
       means it has no effect in the C/POSIX locale)

   --header[=N]
       print (without converting) the first N header lines; N  defaults
       to 1 if not specified

   --invalid=MODE
       failure  mode for invalid numbers: MODE can be: abort (default),
       fail, warn, ignore

   --padding=N
       pad the output to N characters;  positive  N  will  right-align;
       negative  N will left-align; padding is ignored if the output is
       wider than N; the default is to automatically pad  if  a  white-
       space is found

   --round=METHOD
       use  METHOD  for rounding when scaling; METHOD can be: up, down,
       from-zero (default), towards-zero, nearest

   --suffix=SUFFIX
       add SUFFIX to output numbers, and accept optional SUFFIX in  in-
       put numbers

   --to=UNIT
       auto-scale output numbers to UNITs; see UNIT below

   --to-unit=N
       the output unit size (instead of the default 1)

   -z, --zero-terminated
       line delimiter is NUL, not newline

   --help display this help and exit

   --version
       output version information and exit

UNIT options:
   none   no auto-scaling is done; suffixes will trigger an error

   auto   accept optional single/two letter suffix:

       1K = 1000, 1Ki = 1024, 1M = 1000000, 1Mi = 1048576,

   si     accept optional single letter suffix:

       1K = 1000, 1M = 1000000, ...

   iec    accept optional single letter suffix:

       1K = 1024, 1M = 1048576, ...
```

```
       iec-i  accept optional two-letter suffix:

              1Ki = 1024, 1Mi = 1048576, ...

   FIELDS supports cut(1) style field ranges:
       N      N'th field, counted from 1

       N-     from N'th field, to end of line

       N-M    from N'th to M'th field (inclusive)

       -M     from first to M'th field (inclusive)

       -      all fields

       Multiple fields/ranges can be separated with commas

       FORMAT  must be suitable for printing one floating-point argument '%f'.
       Optional quote (%'f) will enable --grouping (if  supported  by  current
       locale).   Optional  width  value (%10f) will pad output. Optional zero
       (%010f) width will  zero  pad  the  number.  Optional  negative  values
       (%-10f)  will  left align.  Optional precision (%.1f) will override the
       input determined precision.

       Exit status is 0 if all input numbers were successfully converted.   By
       default,  numfmt will stop at the first conversion error with exit sta-
       tus 2.  With --invalid='fail' a warning is printed for each  conversion
       error  and the exit status is 2.  With --invalid='warn' each conversion
       error is diagnosed, but the exit status is 0.  With  --invalid='ignore'
       conversion errors are not diagnosed and the exit status is 0.

EXAMPLES
              $ numfmt --to=si 1000

              -> "1.0K"

              $ numfmt --to=iec 2048

              -> "2.0K"

              $ numfmt --to=iec-i 4096

              -> "4.0Ki"

              $ echo 1K | numfmt --from=si

              -> "1000"

              $ echo 1K | numfmt --from=iec

              -> "1024"

              $ df -B1 | numfmt --header --field 2-4 --to=si
```

```
              $ ls -l  | numfmt --header --field 5 --to=iec
              $ ls -lh | numfmt --header --field 5 --from=iec --padding=10
              $ ls -lh | numfmt --header --field 5 --from=iec --format %10f

AUTHOR
       Written by Assaf Gordon.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/numfmt>
       or available locally via: info '(coreutils) numfmt invocation'

GNU coreutils 9.4                April 2024                        NUMFMT(1)
```

# Bash 的 od 命令

```
od [OPTION]... [FILE]...
```

**功能**

将文件转储为八进制或其它格式。

**类型**

可执行文件（`/usr/bin/od`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-A`，`--address-radix=RADIX` - 文件偏移量的输出格式，`RADIX` 取值：
    - `d` - 十进制
    - `o` - 八进制
    - `x` - 十六进制
    - `n` - None
  - `--endian={big|little}` - 按照此字节序解析输入
  - `-j`，`--skip-bytes=BYTES` - 跳过开头的 `BYTES` 字节
  - `-N`，`--read-bytes=BYTES` - 仅转出 `BYTES` 字节
  - `-S BYTES`，`--strings[=BYTES]` - 仅打印长度不少于 `BYTES` 个可打印字符的以空字符（`\0`）结尾的字符串
  - `-t`，`--format=TYPE` - 输出格式（默认为 `o2`，八进制二字节）：
    - `a` - 命名字符（以字符的名称而非本身显示），忽略高位
    - `c` - 可打印字符或转义字符
    - `d[SIZE]` - `SIZE` 字节的有符号十进制整数
    - `f[SIZE]` - `SIZE` 字节的浮点数
    - `o[SIZE]` - `SIZE` 字节的八进制整数
    - `u[SIZE]` - `SIZE` 字节的无符号十进制整数
    - `x[SIZE]` - `SIZE` 字节的十六进制整数
  - `-v`，`--output-duplicates` - 输出重复项，不用使用 `*` 省略
  - `-w[BYTES]`，`--width[=BYTES]` - 每行输出 `BYTE` 字节；默认为 32
  - `--traditional` - 使用传统模式
  - `--help` - 显示帮助
  - `--version` - 显示版本

- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ echo -e "Primers 编程伙伴 \n https://xplanc.org/primers/" | od
0000000 071120 066551 071145 020163 136347 163626 105650 136344
0000020 162231 132274 005040 064040 072164 071560 027472 074057
0000040 066160 067141 027143 071157 027547 071160 066551 071145
0000060 027563 000012
0000063
$ echo -e "Primers 编程伙伴 \n https://xplanc.org/primers/" | od -t c
0000000   P   r   i   m   e   r   s     347 274 226 347 250 213 344 274
0000020 231 344 274 264      \n       h   t   t   p   s   :   /   /   x
0000040   p   l   a   n   c   .   o   r   g   /   p   r   i   m   e   r
0000060   s   /  \n
0000063
$ echo -e "Primers 编程伙伴 \n https://xplanc.org/primers/" | od -t a
0000000   P   r   i   m   e   r   s  sp   g   <  syn   g   (  vt   d   <
0000020  em   d   <   4  sp  nl  sp   h   t   t   p   s   :   /   /   x
0000040   p   l   a   n   c   .   o   r   g   /   p   r   i   m   e   r
0000060   s   /  nl
0000063
```

## 推荐阅读

- Write files in octal or other formats - GNU Coreutils
- od(1) — Linux manual page

## 手册

```
OD(1)                          User Commands                          OD(1)

NAME
       od - dump files in octal and other formats

SYNOPSIS
       od [OPTION]... [FILE]...
       od [-abcdfilosx]... [FILE] [[+]OFFSET[.][b]]
       od --traditional [OPTION]... [FILE] [[+]OFFSET[.][b] [+][LABEL][.][b]]

DESCRIPTION
       Write an unambiguous representation, octal bytes by default, of FILE to
       standard output.  With more than one FILE argument, concatenate them in
       the listed order to form the input.

       With no FILE, or when FILE is -, read standard input.
```

```
    If  first  and second call formats both apply, the second format is as-
    sumed if the last operand begins with + or (if there are 2 operands)  a
    digit.  An OFFSET operand means -j OFFSET.  LABEL is the pseudo-address
    at  first byte printed, incremented when dump is progressing.  For OFF-
    SET and LABEL, a 0x or 0X prefix indicates hexadecimal; suffixes may be
    . for octal and b for multiply by 512.

    Mandatory arguments to long options are  mandatory  for  short  options
    too.

    -A, --address-radix=RADIX
           output format for file offsets; RADIX is one of [doxn], for Dec-
           imal, Octal, Hex or None

    --endian={big|little}
           swap input bytes according the specified order

    -j, --skip-bytes=BYTES
           skip BYTES input bytes first

    -N, --read-bytes=BYTES
           limit dump to BYTES input bytes

    -S BYTES, --strings[=BYTES]
           show only NUL terminated strings of at least BYTES (3) printable
           characters

    -t, --format=TYPE
           select output format or formats

    -v, --output-duplicates
           do not use * to mark line suppression

    -w[BYTES], --width[=BYTES]
           output  BYTES bytes per output line; 32 is implied when BYTES is
           not specified

    --traditional
           accept arguments in third form above

    --help display this help and exit

    --version
           output version information and exit

Traditional format specifications may be intermixed; they accumulate:
    -a     same as -t a,  select named characters, ignoring high-order bit

    -b     same as -t o1, select octal bytes

    -c     same as -t c,  select printable characters or backslash escapes
```

```
    -d     same as -t u2, select unsigned decimal 2-byte units

    -f     same as -t fF, select floats

    -i     same as -t dI, select decimal ints

    -l     same as -t dL, select decimal longs

    -o     same as -t o2, select octal 2-byte units

    -s     same as -t d2, select decimal 2-byte units

    -x     same as -t x2, select hexadecimal 2-byte units

TYPE is made up of one or more of these specifications:
    a      named character, ignoring high-order bit

    c      printable character or backslash escape

    d[SIZE]
           signed decimal, SIZE bytes per integer

    f[SIZE]
           floating point, SIZE bytes per float

    o[SIZE]
           octal, SIZE bytes per integer

    u[SIZE]
           unsigned decimal, SIZE bytes per integer

    x[SIZE]
           hexadecimal, SIZE bytes per integer

    SIZE is a number.   For  TYPE  in  [doux],  SIZE  may  also  be  C  for
    sizeof(char),   S  for  sizeof(short),  I  for  sizeof(int)  or  L  for
    sizeof(long).  If TYPE is f, SIZE may also be F  for  sizeof(float),  D
    for sizeof(double) or L for sizeof(long double).

    Adding  a z suffix to any type displays printable characters at the end
    of each output line.

BYTES is hex with 0x or 0X prefix, and may have a multiplier suffix:
    b      512

    KB     1000

    K      1024

    MB     1000*1000

    M      1024*1024
```

```
       and so on for G, T, P, E, Z, Y, R, Q.  Binary  prefixes  can  be  used,
       too: KiB=K, MiB=M, and so on.

EXAMPLES
       od -A x -t x1z -v
              Display hexdump format output

       od -A o -t oS -w16
              The default output format used by od

AUTHOR
       Written by Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/od>
       or available locally via: info '(coreutils) od invocation'

GNU coreutils 9.4                April 2024                            OD(1)
```

# Bash 的 partx 命令

```
partx [OPTION]... PARTITION [DISK]
```

**功能**

告知内核磁盘分区的存在及其编号。

**类型**

可执行文件（`/usr/bin/partx`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--add` - 添加指定的分区，或读取磁盘并添加所有分区
  - `-b`，`--bytes` - 以字节为单位
  - `-d`，`--delete` - 删除指定的分区或所有分区
  - `-g`，`--noheadings` - 不打印标题行（列名）
  - `-l`，`--list` - 列出分区
  - `-n`，`--nr M:N` - 指定分区范围
  - `-o`，`--output list` - 指定要输出的列
  - `--output-all` - 输出所有列
  - `-P`，`--pairs` - 以键值对的格式输出
  - `-r`，`--raw` - 以原始格式输出
  - `-s`，`--show` - 显示分区
  - `-t`，`--type type` - 指定分区类型
  - `--list-types` - 列出支持的分区类型
  - `-u`，`--update` - 更新分区
  - `-S`，`--sector-size size` - 指定扇区大小
  - `-v`，`--verbose` - 显示详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本

## 列名

| 列名 | 说明 |
| --- | --- |
| `NR` | 分区编号 |
| `START` | 分区的开始扇区 |
| `END` | 分区的结束扇区 |
| `SECTORS` | 分区的扇区数 |
| `SIZE` | 分区大小 |
| `NAME` | 分区名 |
| `UUID` | 分区的 UUID |
| `TYPE` | 分区类型 |
| `FLAGS` | 分区标志 |
| `SCHEME` | 分区表类型 (dos, gpt, ...) |

## 示例

### 查看分区

```
$ sudo partx --show /dev/nvme0n1    # 查看 nvme0n1 的分区
NR   START        END    SECTORS   SIZE NAME UUID
 1    2048    2203647    2201600     1G      dfedd060-593f-48e6-937d-a573d4a95dc5
 2 2203648 1953521663 1951318016 930.5G      4c3ba215-b476-4c6c-a37b-fb13e02d9d46
```

### 修改分区后更新

```
$ sudo partx -u /dev/nvme0n1        # 更新 nvme0n1 的分区；或
$ sudo partx -a /dev/nvme0n1        # 添加 nvme0n1 的所有分区
```

### 添加分区

```
$ sudo partx -a -n 1:3 /dev/nvme0n1     # 添加第 1 到 3 个分区
```

### 删除分区

```
$ sudo partx -d -n 2:3 /dev/nvme0n1     # 删除第 2 到 3 个分区
```

## 推荐阅读

- partx(8) — Linux manual page

## 手册

```
PARTX(8)                    System Administration                    PARTX(8)

NAME
       partx - tell the kernel about the presence and numbering of on-disk
       partitions

SYNOPSIS
       partx [-a|-d|-P|-r|-s|-u] [-t type] [-n M:_N_] [-] disk

       partx [-a|-d|-P|-r|-s|-u] [-t type] partition [disk]

DESCRIPTION
       Given a device or disk-image, partx tries to parse the partition table
       and list its contents. It can also tell the kernel to add or remove
       partitions from its bookkeeping.

       The disk argument is optional when a partition argument is provided. To
       force scanning a partition as if it were a whole disk (for example to
       list nested subpartitions), use the argument "-" (hyphen-minus). For
       example:

          partx --show - /dev/sda3

       This will see sda3 as a whole-disk rather than as a partition.

       partx is not an fdisk program - adding and removing partitions does not
       change the disk, it just tells the kernel about the presence and
       numbering of on-disk partitions.

OPTIONS
       -a, --add
           Add the specified partitions, or read the disk and add all
           partitions.

       -b, --bytes
           Print the sizes in bytes rather than in a human-readable format.

           By default, the unit, sizes are expressed in, is byte, and unit
           prefixes are in power of 2^10 (1024). Abbreviations of symbols are
           exhibited truncated in order to reach a better readability, by
           exhibiting alone the first letter of them; examples: "1 KiB" and "1
           MiB" are respectively exhibited as "1 K" and "1 M", then omitting
           on purpose the mention "iB", which is part of these abbreviations.

       -d, --delete
           Delete the specified partitions or all partitions. It is not error
```

```
    to remove non-existing partitions, so this option is possible to
    use together with large --nr ranges without care about the current
    partitions set on the device.

-g, --noheadings
    Do not print a header line with --show or --raw.

-l, --list
    List the partitions. Note that all numbers are in 512-byte sectors.
    This output format is DEPRECATED in favour of --show. Do not use it
    in newly written scripts.

-n, --nr M:N
    Specify the range of partitions. For backward compatibility also
    the format M-N is supported. The range may contain negative
    numbers, for example --nr -1:-1 means the last partition, and --nr
    -2:-1 means the last two partitions. Supported range specifications
    are:

    M
        Specifies just one partition (e.g. --nr 3).

    M:
        Specifies the lower limit only (e.g. --nr 2:).

    :N
        Specifies the upper limit only (e.g. --nr :4).

    M:N
        Specifies the lower and upper limits (e.g. --nr 2:4).

-o, --output list
    Define the output columns to use for --show, --pairs and --raw
    output. If no output arrangement is specified, then a default set
    is used. Use --help to get list of all supported columns. This
    option cannot be combined with the --add, --delete, --update or
    --list options.

--output-all
    Output all available columns.

-P, --pairs
    List the partitions using the KEY="value" format.

-r, --raw
    List the partitions using the raw output format.

-s, --show
    List the partitions. The output columns can be selected and
    rearranged with the --output option. All numbers (except SIZE) are
    in 512-byte sectors.

-t, --type type
```

```
            Specify the partition table type.

        --list-types
            List supported partition types and exit.

        -u, --update
            Update the specified partitions.

        -S, --sector-size size
            Overwrite default sector size.

        -v, --verbose
            Verbose mode.

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

ENVIRONMENT
        LIBBLKID_DEBUG=all
            enables libblkid debug output.

EXAMPLE
        partx --show /dev/sdb3, partx --show --nr 3 /dev/sdb, partx --show
        /dev/sdb3 /dev/sdb
            All three commands list partition 3 of /dev/sdb.

        partx --show - /dev/sdb3
            Lists all subpartitions on /dev/sdb3 (the device is used as
            whole-disk).

        partx -o START -g --nr 5 /dev/sdb
            Prints the start sector of partition 5 on /dev/sdb without header.

        partx -o SECTORS,SIZE /dev/sda5 /dev/sda
            Lists the length in sectors and human-readable size of partition 5
            on /dev/sda.

        partx --add --nr 3:5 /dev/sdd
            Adds all available partitions from 3 to 5 (inclusive) on /dev/sdd.

        partx -d --nr :-1 /dev/sdd
            Removes the last partition on /dev/sdd.

AUTHORS
        Davidlohr Bueso <dave@gnu.org>, Karel Zak <kzak@redhat.com>

        The original version was written by Andries E. Brouwer <aeb@cwi.nl>

SEE ALSO
        addpart(8), delpart(8), fdisk(8), parted(8), partprobe(8)
```

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The partx command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          PARTX(8)
```

# Bash 的 paste 命令

```
paste [OPTION]... [FILE]...
```

**功能**

合并文件的行。

**类型**

可执行文件（`/usr/bin/paste`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-d`，`--delimiters=LIST` - 使用 `LIST` 中的字符作为分隔符，而不是制表符
  - `-s`，`--serial` - 按文件顺序串行处理
  - `-z`，`--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ cat 1.txt
A
B
C
$ cat 2.txt
1
2
3
$ cat 3.txt
x
y
z
$ paste 1.txt 2.txt 3.txt       # 合并 1.txt 2.txt 3.txt
A       1       x
B       2       y
C       3       z
$ paste -s 1.txt 2.txt 3.txt    # 串行处理
A       B       C
1       2       3
```

```
x       y       z
$ paste -d .- 1.txt 2.txt 3.txt # 依次使用 .- 作为分隔符
A.1-x
B.2-y
C.3-z
```

## 推荐阅读

- Merge lines of files - GNU Coreutils
- paste(1) — Linux manual page

## 手册

```
PASTE(1)                         User Commands                         PASTE(1)

NAME
       paste - merge lines of files

SYNOPSIS
       paste [OPTION]... [FILE]...

DESCRIPTION
       Write  lines  consisting  of  the sequentially corresponding lines from
       each FILE, separated by TABs, to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -d, --delimiters=LIST
              reuse characters from LIST instead of TABs

       -s, --serial
              paste one file at a time instead of in parallel

       -z, --zero-terminated
              line delimiter is NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David M. Ihnat and David MacKenzie.

REPORTING BUGS
```

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/paste>
or available locally via: info '(coreutils) paste invocation'

# Bash 的 pathchk 命令

```
pathchk [OPTION]... [NAME]...
```

**功能**

检查文件名（路径）是否有效或可移植。

**类型**

可执行文件（`/usr/bin/pathchk`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-p` - 检查文件名是否在大多数 `POSIX` 系统有效
  - `-P` - 空的文件名和以 `-` 开头的文件名无效
  - `--portability` - 检查是否在所有 `POSIX` 系统上有效；相当于 `-p -P`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NAME` - 文件名列表

**返回值**

- 0 - 所有文件名都通过了检查
- 1 - 存在某个文件名未通过检查

## 示例

```
$ pathchk -p filename; echo $?
0
$ pathchk -p -- -filename; echo $?
0
$ pathchk -P -- -filename; echo $?
pathchk: leading '-' in a component of file name '-filename'
1
```

- `$?` 表示上一个命令的返回值
- `--` 表示后面的参数不被当作选项

# 推荐阅读

- Check file name validity and portability - GNU Coreutils
- pathchk(1) — Linux manual page

# 手册

```
PATHCHK(1)                      User Commands                      PATHCHK(1)

NAME
       pathchk - check whether file names are valid or portable

SYNOPSIS
       pathchk [OPTION]... NAME...

DESCRIPTION
       Diagnose invalid or non-portable file names.

       -p     check for most POSIX systems

       -P     check for empty names and leading "-"

       --portability
              check for all POSIX systems (equivalent to -p -P)

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Paul Eggert, David MacKenzie, and Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/pathchk>
       or available locally via: info '(coreutils) pathchk invocation'

GNU coreutils 9.4               April 2024                      PATHCHK(1)
```

# Bash 的 pinky 命令

```
pinky [OPTION]... [USER]...
```

**功能**

打印用户信息。

> 这个命令是 finger 命令的简单实现。

**类型**

可执行文件（`/usr/bin/pinky`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-l` - 为指定的用户生成长格式的输出
  - `-b` - 长格式中省略的用户主目录和 shell
  - `-h` - 长格式中省略 `~/.project` 文件
  - `-p` - 长格式中省略 `~/.plan` 文件
  - `-s` - 使用短格式输出；默认
  - `-f` - 短格式中省略列标题行
  - `-w` - 短格式中省略用户名
  - `-i` - 短格式中省略用户名和远程地址
  - `-q` - 短格式中省略用户名，远程地址和空闲时间
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `USER` - 用户名列表；默认为当前用户

## 示例

```
$ pinky                 # 默认使用短格式
Login       Name                TTY      Idle    When             Where
primers     primers             pts/0    2d      2025-11-19 13:31 192.168.1.100
primers     primers             pts/3            2025-11-19 13:35 192.168.1.100
$ pinky -q              # 省略用户名 远程地址 空闲时间
Login       TTY      When
primers     pts/0    2025-11-19 13:31
primers     pts/3    2025-11-19 13:35
```

```
$ pinky -l planc        # 使用长格式，必须指定用户名
Login name: primers                     In real life:  primers
Directory: /home/primers                Shell:  /bin/bash
```

## 推荐阅读

- Print information about users - GNU Coreutils
- pinky(1) — Linux manual page

## 手册

```
PINKY(1)                        User Commands                        PINKY(1)

NAME
       pinky - lightweight finger

SYNOPSIS
       pinky [OPTION]... [USER]...

DESCRIPTION
       -l     produce long format output for the specified USERs

       -b     omit the user's home directory and shell in long format

       -h     omit the user's project file in long format

       -p     omit the user's plan file in long format

       -s     do short format output, this is the default

       -f     omit the line of column headings in short format

       -w     omit the user's full name in short format

       -i     omit the user's full name and remote host in short format

       -q     omit  the  user's  full name, remote host and idle time in short
              format

       --help display this help and exit

       --version
              output version information and exit

       A lightweight 'finger' program;  print user information.  The utmp file
       will be /var/run/utmp.

AUTHOR
```

Written by Joseph Arceneaux, David MacKenzie, and Kaveh Ghazi.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/pinky>
or available locally via: info '(coreutils) pinky invocation'

# Bash 的 pivot_root 命令

```
pivot_root [OPTION]... NEW_ROOT PUT_OLD
```

### 功能

更改根文件系统。

> `pivot_root` 将当前进程的根文件系统移动到 `PUT_OLD` 目录，并将 `NEW_ROOT` 设置为新的根文件系统。

### 类型

可执行文件（`/usr/sbin/pivot_root`），属于 util-linux。

### 参数

- `OPTIONS` 选项：
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `NEW_ROOT` - 新的根文件系统
- `PUT_OLD` - 将旧的根文件系统挂载到此处

## 示例

```
$ mkdir ./mnt                                  # 创建挂载点
$ sudo mount /dev/sda ./mnt/                   # 挂载文件系统
$ wget https://busybox.net/downloads/binaries/1.35.0-x86_64-linux-musl/busybox -O
./mnt/busybox
$ sudo mkdir ./mnt/old_root                    # 创建目录挂载旧的根文件系统
$ sudo unshare -m /bin/bash                    # 创建新的 mount namespace，在其中运行一个新的
bash
$ sudo pivot_root ./mnt/ ./mnt/old_root        # 将根文件系统切换到 ./mnt/
# cd /                                         # 切换到根目录（原先的工作目录已经失效）
# ./busybox ls                                 # 通过 busybox 执行命令
busybox   old_root
# ./busybox mkdir /proc                        # 挂载基础文件系统
# ./busybox mount -t proc proc /proc
# ./busybox mkdir /sys
# ./busybox mount -t sysfs sysfs /sys
```

> 可以通过 `Ctrl + D` 退出新建的 bash ，从而恢复原先的状态。

> 切换根文件系统后，没有 `/bin` 目录下的各种命令，因此下载 `busybox` 来执行基础命令。

## 相关命令

| 命令 | 说明 |
| --- | --- |
| chroot | 执行命令时指定根路径 |
| switch_root | 切换根文件系统（不可恢复） |

## 推荐阅读

- pivot_root(8) — Linux manual page

## 手册

```
PIVOT_ROOT(8)                System Administration                PIVOT_ROOT(8)

NAME
       pivot_root - change the root filesystem

SYNOPSIS
       pivot_root new_root put_old

DESCRIPTION
       pivot_root moves the root file system of the current process to the
       directory put_old and makes new_root the new root file system. Since
       pivot_root(8) simply calls pivot_root(2), we refer to the man page of
       the latter for further details.

       Note that, depending on the implementation of pivot_root, root and
       current working directory of the caller may or may not change. The
       following is a sequence for invoking pivot_root that works in either
       case, assuming that pivot_root and chroot are in the current PATH:

           cd new_root
           pivot_root . put_old
           exec chroot . command

       Note that chroot must be available under the old root and under the new
       root, because pivot_root may or may not have implicitly changed the
       root directory of the shell.

       Note that exec chroot changes the running executable, which is
       necessary if the old root directory should be unmounted afterwards.
       Also note that standard input, output, and error may still point to a
       device on the old root file system, keeping it busy. They can easily be
       changed when invoking chroot (see below; note the absence of leading
```

```
       slashes to make it work whether pivot_root has changed the shell’s root
       or not).

OPTIONS
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXAMPLE
       Change the root file system to /dev/hda1 from an interactive shell:

           mount /dev/hda1 /new-root
           cd /new-root
           pivot_root . old-root
           exec chroot . sh <dev/console >dev/console 2>&1
           umount /old-root

       Mount the new root file system over NFS from 10.0.0.1:/my_root and run
       init:

           ifconfig lo 127.0.0.1 up   # for portmap
           # configure Ethernet or such
           portmap   # for lockd (implicitly started by mount)
           mount -o ro 10.0.0.1:/my_root /mnt
           killall portmap   # portmap keeps old root busy
           cd /mnt
           pivot_root . old_root
           exec chroot . sh -c 'umount /old_root; exec /sbin/init' \
             <dev/console >dev/console 2>&1

SEE ALSO
       chroot(1), pivot_root(2), mount(8), switch_root(8), umount(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The pivot_root command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                     PIVOT_ROOT(8)
```

# Bash 的 pr 命令

```
pr [OPTION]... [FILE]...
```

**功能**

对文件进行分页或分栏处理以进行打印。

**类型**

可执行文件（`/usr/bin/pr`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `+FIRST_PAGE[:LAST_PAGE], --pages=FIRST_PAGE[:LAST_PAGE]`
    - 从指定的第一页开始打印（可选指定最后一页）。
  - `-COLUMN, --columns=COLUMN`
    - 以 COLUMN 列输出，并按列 **自上而下** 排列（除非使用 `-a`）。
    - 平衡每页各列的行数。
  - `-a, --across`
    - 按 **横向** 排列列，而不是纵向。需配合 `-COLUMN` 使用。
  - `-c, --show-control-chars`
    - 使用帽子符号（^G）和八进制反斜线表示控制字符。
  - `-d, --double-space`
    - 输出结果双倍行距。
  - `-D, --date-format=FORMAT`
    - 使用 FORMAT 指定页眉中的日期格式。
  - `-e[CHAR[WIDTH]], --expand-tabs[=CHAR[WIDTH]]`
    - 将输入中的 CHAR（默认 TAB）扩展为给定宽度（默认 8）。
  - `-F, -f, --form-feed`

    - 用换页符代替换行符分隔页面：
        - 使用 `-F` → 页眉占 3 行
        - 未使用 `-F` → 页眉和页脚共占 5 行
- `-h, --header=HEADER`
    - 在页眉使用居中的 HEADER 替代文件名。
    - 使用 `-h ""` 输出空白行（**不要写成** `-h""`）。
- `-i[CHAR[WIDTH]], --output-tabs[=CHAR[WIDTH]]`
    - 将空格转换为 CHAR（默认 TAB），宽度为 WIDTH（默认 8）。
- `-J, --join-lines`
    - 合并完整行；关闭 `-W` 的行截断；不进行列对齐。
    - 可使用 `--sep-string=STRING` 设置列分隔符。
- `-l, --length=PAGE_LENGTH`
    - 设置页面长度为 PAGE_LENGTH（默认 66 行）。
    - 默认正文为 56 行；使用 `-F` 时正文为 63 行。
    - 如果长度 ≤ 10，则隐含启用 `-t`（不显示页眉页脚）。
- `-m, --merge`
    - 并行打印多个文件，每个文件一列。
    - 截断行；若配合 `-J` 则合并完整行。
- `-n[SEP[DIGITS]], --number-lines[=SEP[DIGITS]]`
    - 为行编号。
    - DIGITS（默认 5）为编号位数；SEP（默认 TAB）为分隔符。
    - 默认从文件第一行开始编号。
- `-N, --first-line-number=NUMBER`
    - 设置输出第一页第一行从 NUMBER 开始编号（参考 `+FIRST_PAGE`）。
- `-o, --indent=MARGIN`
    - 每行前缩进 MARGIN 个空格。
    - 不影响 `-w` 与 `-W`。
    - MARGIN 会加到页面宽度中。
- `-r, --no-file-warnings`

        - 当文件无法打开时不显示警告。
    - `-s[CHAR], --separator[=CHAR]`
        - 用单个 CHAR 分隔列。
        - CHAR 默认是 TAB（无 `-w` 时）或“无字符”（使用 `-w` 时）。
        - 使用 `-s` 会关闭所有列模式（`-COLUMN`、`-a`、`-m`）的行截断，除非设置了 `-w`。
    - `-S[STRING], --sep-string[=STRING]`
        - 使用 STRING 分隔列。
        - 默认分隔符：
            - 使用 `-J` → TAB
            - 未使用 `-J` → 空格（相当于 `-S" "`）
        - 不影响列模式。
    - `-t, --omit-header`
        - 不输出页眉和页脚。
        - 页面长度 ≤ 10 时自动启用。
    - `-T, --omit-pagination`
        - 不输出页眉和页脚，并忽略输入中的分页符。
    - `-v, --show-nonprinting`
        - 使用八进制反斜线表示不可打印字符。
    - `-w, --width=PAGE_WIDTH`
        - 设置页面宽度为 PAGE_WIDTH（默认 72）。
        - 仅用于多列文本输出。
        - 使用 `-s` 会关闭默认页宽 72 的限制。
    - `-W, --page-width=PAGE_WIDTH`
        - 始终设置页面宽度为 PAGE_WIDTH（默认 72）。
        - 截断行（除非使用 `-J`）。
        - 与 `-S` 和 `-s` 不冲突。
    - `--help` - 显示帮助
    - `--version` - 显示版本
- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

# 示例

```
$ man pr > 1.txt            # 创建文件
$ pr -n -l 20 1.txt +3:5    # 打印 1.txt 显示行号 每页 20 行 显示第三页到第五页


2025-11-23 14:09                      1.txt                      Page 3


   21                output COLUMN columns and print columns down, unless -a is used.
   22                Balance number of lines in the columns on each page
   23
   24          -a, --across
   25                print columns across rather than down, used together with  -COL-
   26                UMN
   27
   28          -c, --show-control-chars
   29                use hat notation (^G) and octal backslash notation
   30




2025-11-23 14:09                      1.txt                      Page 4


   31          -d, --double-space
   32                double space the output
   33
   34          -D, --date-format=FORMAT
   35                use FORMAT for the header date
   36
   37          -e[CHAR[WIDTH]], --expand-tabs[=CHAR[WIDTH]]
   38                expand input CHARs (TABs) to tab WIDTH (8)
   39
   40          -F, -f, --form-feed




2025-11-23 14:09                      1.txt                      Page 5


   41                use  form  feeds  instead  of  newlines  to separate pages (by a
   42                3-line page header with -F or a 5-line header and trailer  with-
```

```
   43                out -F)
   44
   45          -h, --header=HEADER
   46                use  a centered HEADER instead of filename in page header, -h ""
   47                prints a blank line, don't use -h""
   48
   49          -i[CHAR[WIDTH]], --output-tabs[=CHAR[WIDTH]]
   50                replace spaces with CHARs (TABs) to tab WIDTH (8)
```

## 推荐阅读

- Paginate or columnate files for printing - GNU Coreutils
- pr(1) — Linux manual page

## 手册

```
PR(1)                            User Commands                           PR(1)

NAME
       pr - convert text files for printing

SYNOPSIS
       pr [OPTION]... [FILE]...

DESCRIPTION
       Paginate or columnate FILE(s) for printing.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       +FIRST_PAGE[:LAST_PAGE], --pages=FIRST_PAGE[:LAST_PAGE]
              begin [stop] printing with page FIRST_[LAST_]PAGE

       -COLUMN, --columns=COLUMN
              output COLUMN columns and print columns down, unless -a is used.
              Balance number of lines in the columns on each page

       -a, --across
              print columns across rather than down, used together with  -COL-
              UMN
```

```
-c, --show-control-chars
       use hat notation (^G) and octal backslash notation

-d, --double-space
       double space the output

-D, --date-format=FORMAT
       use FORMAT for the header date

-e[CHAR[WIDTH]], --expand-tabs[=CHAR[WIDTH]]
       expand input CHARs (TABs) to tab WIDTH (8)

-F, -f, --form-feed
       use  form  feeds  instead  of  newlines  to separate pages (by a
       3-line page header with -F or a 5-line header and trailer  with-
       out -F)

-h, --header=HEADER
       use  a centered HEADER instead of filename in page header, -h ""
       prints a blank line, don't use -h""

-i[CHAR[WIDTH]], --output-tabs[=CHAR[WIDTH]]
       replace spaces with CHARs (TABs) to tab WIDTH (8)

-J, --join-lines
       merge full lines, turns off -W line truncation, no column align-
       ment, --sep-string[=STRING] sets separators

-l, --length=PAGE_LENGTH
       set the page length to PAGE_LENGTH (66) lines (default number of
       lines of text 56, and with -F 63).  implies -t if PAGE_LENGTH <=
       10

-m, --merge
       print all files in parallel, one in each column, truncate lines,
       but join lines of full length with -J

-n[SEP[DIGITS]], --number-lines[=SEP[DIGITS]]
       number lines, use DIGITS (5) digits,  then  SEP  (TAB),  default
       counting starts with 1st line of input file

-N, --first-line-number=NUMBER
       start  counting  with  NUMBER  at 1st line of first page printed
       (see +FIRST_PAGE)

-o, --indent=MARGIN
       offset each line with MARGIN (zero) spaces, do not affect -w  or
       -W, MARGIN will be added to PAGE_WIDTH

-r, --no-file-warnings
       omit warning when a file cannot be opened
```

```
       -s[CHAR], --separator[=CHAR]
              separate  columns by a single character, default for CHAR is the
              <TAB> character without -w and  'no  char'  with  -w.   -s[CHAR]
              turns  off  line  truncation of all 3 column options (-COLUMN|-a
              -COLUMN|-m) except -w is set

       -S[STRING], --sep-string[=STRING]
              separate columns by STRING, without -S: Default separator  <TAB>
              with -J and <space> otherwise (same as -S" "), no effect on col-
              umn options

       -t, --omit-header
              omit page headers and trailers; implied if PAGE_LENGTH <= 10

       -T, --omit-pagination
              omit page headers and trailers, eliminate any pagination by form
              feeds set in input files

       -v, --show-nonprinting
              use octal backslash notation

       -w, --width=PAGE_WIDTH
              set  page  width  to  PAGE_WIDTH  (72)  characters  for multiple
              text-column output only, -s[char] turns off (72)

       -W, --page-width=PAGE_WIDTH
              set page width to PAGE_WIDTH (72)  characters  always,  truncate
              lines, except -J option is set, no interference with -S or -s

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Pete TerMaat and Roland Huebner.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/pr>
       or available locally via: info '(coreutils) pr invocation'

GNU coreutils 9.4                April 2024                           PR(1)
```

# Bash 的 printenv 命令

```
printenv [OPTION]... [VARIABLE]...
```

**功能**

打印全部或部分环境变量。

**类型**

可执行文件（`/usr/bin/printenv`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-0`, `--null` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `VARIABLE` - 环境变量名列表

## 示例

```
$ printenv SHELL
/bin/bash
$ printenv USER
primers
$ printenv TERM
xterm
```

## 推荐阅读

- Print all or some environment variables - GNU Coreutils
- printenv(1) — Linux manual page

## 手册

```
PRINTENV(1)                     User Commands                     PRINTENV(1)

NAME
```

```
       printenv - print all or part of environment

SYNOPSIS
       printenv [OPTION]... [VARIABLE]...

DESCRIPTION
       Print the values of the specified environment VARIABLE(s).  If no VARI-
       ABLE is specified, print name and value pairs for them all.

       -0, --null
              end each output line with NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

       NOTE:  your  shell  may have its own version of printenv, which usually
       supersedes the version described here.  Please refer  to  your  shell's
       documentation for details about the options it supports.

AUTHOR
       Written by David MacKenzie and Richard Mlynarik.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/printenv>
       or available locally via: info '(coreutils) printenv invocation'

GNU coreutils 9.4                April 2024                       PRINTENV(1)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 printf 命令

```
printf [OPTION] FORMAT [ARGUMENT]...
```

**功能**

格式化并打印数据。

**类型**

可执行文件（`/usr/bin/printf`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FORMAT` - 格式字符串
- `ARGUMENT` - 参数列表

## 格式

| 符号 | 说明 |
| --- | --- |
| `\"` | 双引号 |
| `\\` | 反斜杠 |
| `\a` | 警告声（BEL） |
| `\b` | 退格 |
| `\c` | 不再产生任何输出 |
| `\e` | Escape 字符 |
| `\f` | 换页符（form feed） |
| `\n` | 换行符 |
| `\r` | 回车符 |
| `\t` | 水平制表符 |
| `\v` | 垂直制表符 |

| 符号 | 说明 |
| --- | --- |
| `\NNN` | 八进制值为 NNN 的字节（1 到 3 位） |
| `\xHH` | 十六进制值为 HH 的字节（1 到 2 位） |
| `\uHHHH` | Unicode 字符（ISO/IEC 10646），十六进制值 HHHH（4 位） |
| `\UHHHHHHHH` | Unicode 字符，十六进制值 HHHHHHHH（8 位） |
| `%%` | 百分号（`%`）本身 |
| `%b` | 将参数作为带有反斜杠转义的字符串输出；八进制转义格式为 \0 或 \0NNN |
| `%q` | 输出可作为 shell 输入的格式，并使用 POSIX 提议的 $'' 语法对不可打印字符进行转义 |
| `%c` | 单个字符 |
| `%s` | 字符串 |
| `%d` | 有符号十进制整数 |
| `%i` | 有符号十进制整数，同 `%d` |
| `%u` | 无符号十进制整数 |
| `%o` | 无符号八进制整数 |
| `%x` | 无符号十六进制整数，字母小写 |
| `%X` | 无符号十六进制整数，字母大写 |
| `%f` | 浮点数，字母小写 |
| `%F` | 浮点数，同 `%f` |
| `%e` | 浮点数，科学计数法，字母小写 |
| `%E` | 浮点数，科学计数法，字母大写 |
| `%g` | 浮点数，省略小数末尾的 0，值较大时使用科学计数法，字母小写 |
| `%G` | 浮点数，省略小数末尾的 0，值较大时使用科学计数法，字母大写 |

## 示例

```
$ printf "Hello World\n"
Hello World
$ printf "Name: %s, Age: %d\n" "Alice" 18
Name: Alice, Age: 18
```

## 推荐阅读

- Format and print data - GNU Coreutils
- printf(1) — Linux manual page

## 手册

```
PRINTF(1)                        User Commands                        PRINTF(1)

NAME
       printf - format and print data

SYNOPSIS
       printf FORMAT [ARGUMENT]...
       printf OPTION

DESCRIPTION
       Print ARGUMENT(s) according to FORMAT, or execute according to OPTION:

       --help display this help and exit

       --version
              output version information and exit

       FORMAT controls the output as in C printf.  Interpreted sequences are:

       \"     double quote

       \\     backslash

       \a     alert (BEL)

       \b     backspace

       \c     produce no further output

       \e     escape

       \f     form feed

       \n     new line

       \r     carriage return

       \t     horizontal tab

       \v     vertical tab
```

```
       \NNN   byte with octal value NNN (1 to 3 digits)

       \xHH   byte with hexadecimal value HH (1 to 2 digits)

       \uHHHH Unicode (ISO/IEC 10646) character with hex value HHHH (4 digits)

       \UHHHHHHHH
              Unicode character with hex value HHHHHHHH (8 digits)

       %%     a single %

       %b     ARGUMENT  as  a string with '\' escapes interpreted, except that
              octal escapes are of the form \0 or \0NNN

       %q     ARGUMENT is printed in a format that can be reused as shell  in-
              put,  escaping  non-printable characters with the proposed POSIX
              $'' syntax.

       and all C format specifications ending with one of diouxXfeEgGcs,  with
       ARGUMENTs converted to proper type first.  Variable widths are handled.

       NOTE:  your shell may have its own version of printf, which usually su-
       persedes the version described here.  Please refer to your shell's doc-
       umentation for details about the options it supports.

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       printf(3)

       Full documentation <https://www.gnu.org/software/coreutils/printf>
       or available locally via: info '(coreutils) printf invocation'

GNU coreutils 9.4                April 2024                        PRINTF(1)
```

# Bash 的 prlimit 命令

```
prlimit [OPTION]... [CMD [ARG]...]
```

**功能**

查看或设置进程资源限制。

**类型**

可执行文件（`/usr/bin/prlimit`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `--noheadings` - 不显示标题行（列名）
  - `-o`，`--output list` - 打印指定的列
  - `-p`，`--pid` - 指定要操作的进程 ID；如果没有这个选项则运行 `CMD`
  - `--raw` - 以原始格式输出
  - `--verbose` - 打印详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
  - `-c`，`--core[=limits]` - 核心转储文件的最大大小
  - `-d`，`--data[=limits]` - 最大数据大小
  - `-e`，`--nice[=limits]` - 允许提升的最高 nice 优先级
  - `-f`，`--fsize[=limits]` - 最大文件大小
  - `-i`，`--sigpending[=limits]` - 待处理信号的最大数量
  - `-l`，`--memlock[=limits]` - 内存锁定地址空间的最大范围
  - `-m`，`--rss[=limits]` - 最大常驻集大小
  - `-n`，`--nofile[=limits]` - 最大文件打开数量
  - `-q`，`--msgqueue[=limits]` - POSIX 消息队列中的最大字节数
  - `-r`，`--rtprio[=limits]` - 最高实时优先级
  - `-s`，`--stack[=limits]` - 栈的最大尺寸
  - `-t`，`--cpu[=limits]` - CPU 时间（秒）
  - `-u`，`--nproc[=limits]` - 最大进程数
  - `-v`，`--as[=limits]` - 地址空间限制
  - `-x`，`--locks[=limits]` - 可持有文件锁的最大数量

    - `-y`, `--rttime[=limits]` - 实时任务超时时间
- `CMD` - 要执行的命令；
- `ARG` - 命令的参数列表

## 列名

| 列名 | 说明 |
| --- | --- |
| `DESCRIPTION` | 资源描述 |
| `RESOURCE` | 资源名称 |
| `SOFT` | 软限制 |
| `HARD` | 硬限制 |
| `UNITS` | 单位 |

## 示例

### 查看限制

```
$ prlimit       # 查看当前 SHELL 的限制，等价于 prlimit -i $$
RESOURCE   DESCRIPTION                             SOFT       HARD UNITS
AS         address space limit                 unlimited  unlimited bytes
CORE       max core file size                          0  unlimited bytes
CPU        CPU time                            unlimited  unlimited seconds
DATA       max data size                       unlimited  unlimited bytes
FSIZE      max file size                       unlimited  unlimited bytes
LOCKS      max number of file locks held       unlimited  unlimited locks
MEMLOCK    max locked-in-memory address space 1823780864 1823780864 bytes
MSGQUEUE   max bytes in POSIX mqueues             819200     819200 bytes
NICE       max nice prio allowed to raise              0          0
NOFILE     max number of open files                 1024    1048576 files
NPROC      max number of processes                 55346      55346 processes
RSS        max resident set size               unlimited  unlimited bytes
RTPRIO     max real-time priority                      0          0
RTTIME     timeout for real-time tasks         unlimited  unlimited microsecs
SIGPENDING max number of pending signals           55346      55346 signals
STACK      max stack size                        8388608  unlimited bytes
```

### 修改限制

```
$ prlimit -n=10 -p $$               # 当前 SHELL 最多允许打开 10 个文件
$ prlimit -n -p $$                  # 查看当前 SHELL 最多允许打开的文件数量
RESOURCE DESCRIPTION              SOFT HARD UNITS
NOFILE   max number of open files   10   10 files
```

- `$$` 是当前 SHELL 的 PID

**在限制下运行程序**

```
prlimit --v=500M -- ./my_program    # 运行 my_program 时地址空间限制为 500M，即最多只能使用 500M 内存
```

# 推荐阅读

- prlimit(1) — Linux manual page

# 手册

```
PRLIMIT(1)                      User Commands                      PRLIMIT(1)

NAME
       prlimit - get and set process resource limits

SYNOPSIS
       prlimit [options] [--resource[=limits]] [--pid PID]

       prlimit [options] [--resource[=limits]] command [argument...]

DESCRIPTION
       Given a process ID and one or more resources, prlimit tries to retrieve
       and/or modify the limits.

       When command is given, prlimit will run this command with the given
       arguments.

       The limits parameter is composed of a soft and a hard value, separated
       by a colon (:), in order to modify the existing values. If no limits
       are given, prlimit will display the current values. If one of the
       values is not given, then the existing one will be used. To specify the
       unlimited or infinity limit (RLIM_INFINITY), the -1 or 'unlimited'
       string can be passed.

       Because of the nature of limits, the soft limit must be lower or equal
       to the high limit (also called the ceiling). To see all available
       resource limits, refer to the RESOURCE OPTIONS section.

       •   soft:_hard_ Specify both limits.

       •   soft: Specify only the soft limit.

       •   :hard Specify only the hard limit.

       •   value Specify both limits to the same value.
```

```
GENERAL OPTIONS
       --noheadings
           Do not print a header line.

       -o, --output list
           Define the output columns to use. If no output arrangement is
           specified, then a default set is used. Use --help to get a list of
           all supported columns.

       -p, --pid
           Specify the process id; if none is given, the running process will
           be used.

       --raw
           Use the raw output format.

       --verbose
           Verbose mode.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

RESOURCE OPTIONS
       -c, --core[=limits]
           Maximum size of a core file.

       -d, --data[=limits]
           Maximum data size.

       -e, --nice[=limits]
           Maximum nice priority allowed to raise.

       -f, --fsize[=limits]
           Maximum file size.

       -i, --sigpending[=limits]
           Maximum number of pending signals.

       -l, --memlock[=limits]
           Maximum locked-in-memory address space.

       -m, --rss[=limits]
           Maximum Resident Set Size (RSS).

       -n, --nofile[=limits]
           Maximum number of open files.

       -q, --msgqueue[=limits]
           Maximum number of bytes in POSIX message queues.
```

```
       -r, --rtprio[=limits]
           Maximum real-time priority.

       -s, --stack[=limits]
           Maximum size of the stack.

       -t, --cpu[=limits]
           CPU time, in seconds.

       -u, --nproc[=limits]
           Maximum number of processes.

       -v, --as[=limits]
           Address space limit.

       -x, --locks[=limits]
           Maximum number of file locks held.

       -y, --rttime[=limits]
           Timeout for real-time tasks.

NOTES
       The prlimit(2) system call is supported since Linux 2.6.36, older
       kernels will break this program.

EXAMPLES
       prlimit --pid 13134
           Display limit values for all current resources.

       prlimit --pid 13134 --rss --nofile=1024:4095
           Display the limits of the RSS, and set the soft and hard limits for
           the number of open files to 1024 and 4095, respectively.

       prlimit --pid 13134 --nproc=512:
           Modify only the soft limit for the number of processes.

       prlimit --pid $$ --nproc=unlimited
           Set for the current process both the soft and ceiling values for
           the number of processes to unlimited.

       prlimit --cpu=10 sort -u hugefile
           Set both the soft and hard CPU time limit to ten seconds and run
           sort(1).

AUTHORS
       Davidlohr Bueso <dave@gnu.org> - In memory of Dennis M. Ritchie.

SEE ALSO
       ulimit(1p), prlimit(2)

REPORTING BUGS
       For bug reports, use the issue tracker at
```

```
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The prlimit command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-12-01                       PRLIMIT(1)
```

# Bash 的 ptx 命令

```
ptx [OPTION]... [FILE]...
```

**功能**

生成文件内容的关键词排列索引。

**类型**

可执行文件（`/usr/bin/ptx`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-A`，`--auto-reference` - 输出自动生成引用
  - `-G`，`--traditional` - 使用传统（System V）模式
  - `-F`，`--flag-truncation=STRING` - 使用 `STRING` 作为行的截断标记；默认为 `'/'`
  - `-M`，`--macro-name=STRING` - 使用 `STRING` 作为宏名称；默认为 `xx`
  - `-O`，`--format=roff` - 以 roff 命令的格式生成输出
  - `-R`，`--right-side-refs` - 将引用放置在右侧
  - `-S`，`--sentence-regexp=REGEXP` - 使用正则表达式 `REGEXP` 作为行或句子的结束标记
  - `-T`，`--format=tex` - 以 TeX 命令的格式生成输出
  - `-W`，`--word-regexp=REGEXP` - 使用正则表达式 `REGEXP` 匹配关键词
  - `-b`，`--break-file=FILE` - 以 `FILE` 文件中的字符作用单词分隔符
  - `-f`，`--ignore-case` - 排序时忽略大小写（转换为大写进行排序）
  - `-g`，`--gap-size=NUMBER` - 输出字段之间的间隙大写为 `NUMBER`
  - `-i`，`--ignore-file=FILE` - 以 `FILE` 文件中的单词作为忽略词列表
  - `-o`，`--only-file=FILE` - 以 `FILE` 文件中的单词作为只读单词列表
  - `-r`，`--references` - 每行的第一个字段是引用
  - `-t`，`--typeset-mode` - ~~此选项为实现~~
  - `-w`，`--width=NUMBER` - 输出宽度为 `NUMBER`（不含引用）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ man ptx > 1.txt           # 创建文件
$ ptx -W ptx -A -R 1.txt    # 创建关键词 ptx 的索引 生成引用 引用显示在右侧
file contents/      /PTX(1)  NAME   ptx - produce a permuted index of      1.txt:4
  /dex of file contents  SYNOPSIS   ptx [OPTION]... [INPUT]...             1.txt:7
]  DESCRIPTION/      (without -G)   ptx -G [OPTION]... [INPUT [OUTPUT]     1.txt:8
     /behave more like System V '   ptx'         -F, --flag-truncation/    1.txt:23
 /www.gnu.org/software/coreutils/   ptx>        or available locally v/    1.txt:88
  /locally via: info '(coreutils)   ptx invocation'  GNU coreutils 9.4/    1.txt:89
```

## 推荐阅读

- Produce permuted indexes - GNU Coreutils
- ptx(1) — Linux manual page

## 手册

```
PTX(1)                          User Commands                          PTX(1)

NAME
       ptx - produce a permuted index of file contents

SYNOPSIS
       ptx [OPTION]... [INPUT]...   (without -G)
       ptx -G [OPTION]... [INPUT [OUTPUT]]

DESCRIPTION
       Output  a  permuted index, including context, of the words in the input
       files.

       With no FILE, or when FILE is -, read standard input.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -A`, --auto-reference`
              output automatically generated references

       -G, --traditional
              behave more like System V 'ptx'

       -F, --flag-truncation=STRING
              use STRING for flagging line truncations.  The default is '/'

       -M, --macro-name=STRING
              macro name to use instead of 'xx'

       -O, --format=roff
              generate output as roff directives
```

```
       -R, --right-side-refs
              put references at right, not counted in -w

       -S, --sentence-regexp=REGEXP
              for end of lines or end of sentences

       -T, --format=tex
              generate output as TeX directives

       -W, --word-regexp=REGEXP
              use REGEXP to match each keyword

       -b, --break-file=FILE
              word break characters in this FILE

       -f, --ignore-case
              fold lower case to upper case for sorting

       -g, --gap-size=NUMBER
              gap size in columns between output fields

       -i, --ignore-file=FILE
              read ignore word list from FILE

       -o, --only-file=FILE
              read only word list from this FILE

       -r, --references
              first field of each line is a reference

       -t, --typeset-mode               - not implemented -

       -w, --width=NUMBER
              output width in columns, reference excluded

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by F. Pinard.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.
```

```
SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/ptx>
       or available locally via: info '(coreutils) ptx invocation'

GNU coreutils 9.4                April 2024                          PTX(1)
```

# Bash 的 pwd 命令

```
pwd [OPTION]...
```

**功能**

查看当前工作目录。

**类型**

Bash 内置命令。

> `pwd` 也是 coreutils 的一个可执行文件；为了保持 POSIX 标准的兼容性，确保在其他 shell（比如一些极简的或古老的 shell）或者某些特定环境下，命令依然可用。

**参数**

- `OPTION` 选项：
  - `-L`，`--logical` - 查看逻辑路径（符号链接显示自身路径）
  - `-P`，`--physical` - 查看物理路径（符号链接显示源文件路径）
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ cd /lib
$ pwd
/lib
$ pwd -L
/lib
$ pwd -P
/usr/lib
```

## 推荐阅读

- Print working directory - GNU Coreutils
- pwd(1) — Linux manual page

## 手册

NAME

pwd - print name of current/working directory

SYNOPSIS

pwd [OPTION]...

DESCRIPTION

Print the full filename of the current working directory.

-L, --logical

use PWD from environment, even if it contains symlinks

-P, --physical

avoid all symlinks

--help display this help and exit

--version

output version information and exit

If no option is specified, -P is assumed.

NOTE: your shell may have its own version of pwd, which usually supersedes the version described here. Please refer to your shell's documentation for details about the options it supports.

AUTHOR

Written by Jim Meyering.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

getcwd(3)

Full documentation <https://www.gnu.org/software/coreutils/pwd>
or available locally via: info '(coreutils) pwd invocation'

# Bash 的 readlink 命令

```
readlink [OPTION]... FILE...
```

**功能**

获取符号链接的源文件路径。

**类型**

可执行文件（`/usr/bin/readlink`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-f`，`--canonicalize` - 将路径规范化为绝对路径（realpath），路径的所有组成部分（除了最后一部分）都必须存在
  - `-e`，`--canonicalize-existing` - 将路径规范化为绝对路径，路径的所有组成部分都必须存在
  - `-m`，`--canonicalize-missing` - 将路径规范化为绝对路径，路径的所有组成部分都可以不存在
  - `-n`，`--no-newline` - 不要输出尾随的分隔符
  - `-q`，`--quiet` - 不打印错误信息
  - `-s`，`--silent` - 不打印错误信息；**默认**
  - `-v`，`--verbose` - 打印错误信息
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ touch 1.txt           # 创建普通文件
$ readlink 1.txt        # 查看普通文件 没有任何输出
$ readlink -v 1.txt     # 开启错误信息
readlink: 1.txt: Invalid argument
$ ln -s 1.txt 2.txt     # 创建软链接
$ readlink 2.txt        # 查看软链接
1.txt
```

```
$ readlink -f 2.txt     # 规范化为绝对路径
/home/primers/1.txt
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| ln | 创建文件的链接（硬链接或符号链接） |

## 推荐阅读

- Print value of a symlink or canonical file name - GNU Coreutils
- readlink(1) — Linux manual page

## 手册

```
READLINK(1)                     User Commands                     READLINK(1)

NAME
       readlink - print resolved symbolic links or canonical file names

SYNOPSIS
       readlink [OPTION]... FILE...

DESCRIPTION
       Note  realpath(1)  is the preferred command to use for canonicalization
       functionality.

       Print value of a symbolic link or canonical file name

       -f, --canonicalize
              canonicalize by following every symlink in  every  component  of
              the  given name recursively; all but the last component must ex-
              ist

       -e, --canonicalize-existing
              canonicalize by following every symlink in  every  component  of
              the given name recursively, all components must exist

       -m, --canonicalize-missing
              canonicalize  by  following  every symlink in every component of
              the given name recursively, without requirements  on  components
              existence

       -n, --no-newline
              do not output the trailing delimiter
```

```
       -q, --quiet

       -s, --silent
              suppress most error messages (on by default)

       -v, --verbose
              report error messages

       -z, --zero
              end each output line with NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Dmitry V. Levin.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       readlink(2), realpath(1), realpath(3)

       Full documentation <https://www.gnu.org/software/coreutils/readlink>
       or available locally via: info '(coreutils) readlink invocation'

GNU coreutils 9.4                April 2024                       READLINK(1)
```

# Bash 的 readprofile 命令

```
readprofile [OPTION]...
```

**功能**

读取内核分析信息。

> `readprofile` 命令读取 `/proc/profile` 文件中的信息，将 ASCII 数据打印到标准输出。输出分为三列：
>
> 1. 第一列是时钟周期数
> 2. 第二列是内核中出现该时钟周期数的 C 函数名称
> 3. 第三列是过程的归一化“负载”，计算方法是时钟周期数与过程长度的比值

**类型**

可执行文件（`/usr/sbin/readprofile`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-a`，`--all` - 打印映射文件中的所有符号
  - `-b`，`--histbin` - 打印各个直方图区间计数
  - `-i`，`--info` - 只打印内核使用的性能分析步长
  - `-m`，`--mapfile mapfile` - 指定映射文件；默认为 `/usr/src/linux/System.map`
  - `-M`，`--multiplier multiplier` - 指定性能分析的中断频率乘数
  - `-p`，`--profile pro-file` - 指定性能分析文件；默认为 `/proc/profile`
  - `-r`，`--reset` - 重置分析缓冲区
  - `-s`，`--counters` - 打印函数内的各个计数器
  - `-v`，`--verbose` - 打印详细输出
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本

## 推荐阅读

- readprofile(8) — Linux manual page

## 手册

NAME

readprofile - read kernel profiling information

SYNOPSIS

readprofile [options]

VERSION

This manpage documents version 2.0 of the program.

DESCRIPTION

The readprofile command uses the /proc/profile information to print ascii data on standard output. The output is organized in three columns: the first is the number of clock ticks, the second is the name of the C function in the kernel where those many ticks occurred, and the third is the normalized `load' of the procedure, calculated as a ratio between the number of ticks and the length of the procedure. The output is filled with blanks to ease readability.

OPTIONS

-a, --all
Print all symbols in the mapfile. By default the procedures with reported ticks are not printed.

-b, --histbin
Print individual histogram-bin counts.

-i, --info
Info. This makes readprofile only print the profiling step used by the kernel. The profiling step is the resolution of the profiling buffer, and is chosen during kernel configuration (through make config), or in the kernel’s command line. If the -t (terse) switch is used together with -i only the decimal number is printed.

-m, --mapfile mapfile
Specify a mapfile, which by default is /usr/src/linux/System.map. You should specify the map file on cmdline if your current kernel isn’t the last one you compiled, or if you keep System.map elsewhere. If the name of the map file ends with .gz it is decompressed on the fly.

-M, --multiplier multiplier
On some architectures it is possible to alter the frequency at which the kernel delivers profiling interrupts to each CPU. This option allows you to set the frequency, as a multiplier of the system clock frequency, HZ. Linux 2.6.16 dropped multiplier support for most systems. This option also resets the profiling buffer, and requires superuser privileges.

-p, --profile pro-file
Specify a different profiling buffer, which by default is

/proc/profile. Using a different pro-file is useful if you want to `freeze' the kernel profiling at some time and read it later. The /proc/profile file can be copied using cat(1) or cp(1). There is no more support for compressed profile buffers, like in readprofile-1.1, because the program needs to know the size of the buffer in advance.

-r, --reset
Reset the profiling buffer. This can only be invoked by root, because /proc/profile is readable by everybody but writable only by the superuser. However, you can make readprofile set-user-ID 0, in order to reset the buffer without gaining privileges.

-s, --counters
Print individual counters within functions.

-v, --verbose
Verbose. The output is organized in four columns and filled with blanks. The first column is the RAM address of a kernel function, the second is the name of the function, the third is the number of clock ticks and the last is the normalized load.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

FILES

/proc/profile
A binary snapshot of the profiling buffer.

/usr/src/linux/System.map
The symbol table for the kernel.

/usr/src/linux/*
The program being profiled :-)

BUGS

readprofile only works with a 1.3.x or newer kernel, because /proc/profile changed in the step from 1.2 to 1.3.

This program only works with ELF kernels. The change for a.out kernels is trivial, and left as an exercise to the a.out user.

To enable profiling, the kernel must be rebooted, because no profiling module is available, and it wouldn’t be easy to build. To enable profiling, you can specify profile=2 (or another number) on the kernel commandline. The number you specify is the two-exponent used as profiling step.

Profiling is disabled when interrupts are inhibited. This means that many profiling ticks happen when interrupts are re-enabled. Watch out

```
       for misleading information.

EXAMPLE
       Browse the profiling buffer ordering by clock ticks:

              readprofile | sort -nr | less

       Print the 20 most loaded procedures:

              readprofile | sort -nr +2 | head -20

       Print only filesystem profile:

              readprofile | grep _ext2

       Look at all the kernel information, with ram addresses:

              readprofile -av | less

       Browse a 'frozen' profile buffer for a non current kernel:

              readprofile -p ~/profile.freeze -m /zImage.map.gz

       Request profiling at 2kHz per CPU, and reset the profiling buffer:

              sudo readprofile -M 20

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The readprofile command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                    READPROFILE(8)
```

# Bash 的 realpath 命令

```
realpath [OPTION]... FILE...
```

**功能**

解析路径。

**类型**

可执行文件（`/usr/bin/realpath`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-e`，`--canonicalize-existing` - 路径的所有组成部分都必须存在
  - `-m`，`--canonicalize-missing` - 路径的所有组成部分都可以不存在
  - `-L`，`--logical` - 逻辑路径；在跟踪符号链接之前解析 `..`
  - `-P`，`--physical` - 物理路径；在跟踪符号链接之后解析 `..`（默认）
  - `-q`，`--quiet` - 忽略大部分错误信息
  - `--relative-to=DIR` - 打印相对于 `DIR` 的相对路径
  - `--relative-base=DIR` - 如果路径在 `DIR` 内，则相对路径，否则打印绝对路径
  - `-s`，`--strip`，`--no-symlinks` - 不展开符号链接
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件路径

## 示例

**绝对路径**

```
$ pwd                                                    # 查看当前路径
/home/user/primers
$ realpath -m ./file.txt
/home/user/primers/file.txt
$ realpath -m ../data/../config/./settings.yaml
/home/user/config/settings.yaml
```

**相对路径**

```
$ pwd                                                    # 查看当前路径
/home/user/primers
$ realpath -m --relative-to=/home/user ./file.txt        # 打印相对于 /home/user 的路径
primers/file.txt
$ realpath -m --relative-to=/home/user2 ./file.txt       # 打印相对于 /home/user2 的路径
../user/primers/file.txt
$ realpath -m --relative-base=/home/user ./file.txt      # 在 /home/user 内，打印相对于
/home/user 的路径
primers/file.txt
$ realpath -m --relative-base=/home/user2 ./file.txt     # 不在 /home/user 内，打印绝对路径
/home/user/primers/file.txt
```

**符号链接**

```
$ pwd                                                    # 查看当前路径
/home/user/primers
$ ln -s /usr/lib lib                                     # 创建符号链接
$ ls -l lib
lrwxrwxrwx 1 planc planc 8 Oct 31 14:06 lib -> /usr/lib
$ realpath ./lib                                         # 跟踪符号链接（默认）
/usr/lib
$ realpath -s ./lib                                      # 不跟踪符号链接
/home/user/primers/lib
$ realpath ./lib/..                                      # 在跟踪符号链接之后解析 ..（默认）
/usr
$ realpath -L ./lib/..                                   # 在跟踪符号链接之前解析 ..
/home/user/primers
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| basename | 从文件路径中去除目录和后缀 |
| dirname | 从文件路径中提取目录 |

## 推荐阅读

- Print the resolved file name - GNU Coreutils
- realpath(1) — Linux manual page

## 手册

```
REALPATH(1)                      User Commands                      REALPATH(1)
```

```
NAME
       realpath - print the resolved path

SYNOPSIS
       realpath [OPTION]... FILE...

DESCRIPTION
       Print  the resolved absolute file name; all but the last component must
       exist

       -e, --canonicalize-existing
              all components of the path must exist

       -m, --canonicalize-missing
              no path components need exist or be a directory

       -L, --logical
              resolve '..' components before symlinks

       -P, --physical
              resolve symlinks as encountered (default)

       -q, --quiet
              suppress most error messages

       --relative-to=DIR
              print the resolved path relative to DIR

       --relative-base=DIR
              print absolute paths unless paths below DIR

       -s, --strip, --no-symlinks
              don't expand symlinks

       -z, --zero
              end each output line with NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Padraig Brady.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
```

```
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       readlink(1), readlink(2), realpath(3)

       Full documentation <https://www.gnu.org/software/coreutils/realpath>
       or available locally via: info '(coreutils) realpath invocation'

GNU coreutils 9.4                April 2024                      REALPATH(1)
```

# Bash 的 rename.ul 命令

```
rename.ul [OPTION]... EXPRESSION REPLACEMENT FILE...
```

> 现在大部分 Linux 发行版上默认的 `rename` 不是 `rename.ul`，而是功能更全面的 file-rename。
>
> 可以通过 `alternatives` 命令切换 `rename` 的版本：
>
> 1. 注册
>
>    ```
>    $ sudo alternatives --install /usr/bin/rename rename /usr/bin/rename.ul 30
>    ```
>
> 2. 切换
>
>    ```
>    $ sudo alternatives --config rename
>    ```

**功能**

通过表达式重命名文件。

**类型**

可执行文件（`/usr/bin/rename.ul`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-s`，`--symlink` - 不要重命名符号链接，但是修改它的指向
  - `-v`，`--verbose` - 显示被重命名的文件
  - `-n`，`--no-act` - 不实际修改文件名，配合 `--verbose` 查看哪些文件会被修改
  - `-a`，`--all` - 替换表达式的所有匹配位置，而不只是第一个匹配位置
  - `-l`，`--last` - 替换表达式的最后一个匹配位置，而不是第一个匹配位置
  - `-o`，`--no-overwrite` - 不要覆盖现有文件
  - `-i`，`--interactive` - 覆盖现有文件之前进行询问
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `EXPRESSION` - 匹配表达式；普通字符串，不支持通配符和正则表达式
- `REPLACEMENT` - 文件名中匹配补番替换为的内容
- `FILE` - 要重命名的文件列表

# 示例

### 批量重命名文件

```
$ rename.ul --last .txt .md *.txt       # 将 .txt 文件批量重命名为 .md 文件
```

### 批量修改符号链接指向

```
$ mv /path/to/dir1 path/to/dir2                 # 修改了某个目录名
$ rename.ul -s /path/to/dir1 /path/to/dir2 *    # 将符号链接中的 /path/to/dir1 批量修改为
/path/to/dir2
```

# 推荐阅读

- rename(1) — Linux manual page

# 手册

```
RENAME(1)                       User Commands                       RENAME(1)

NAME
       rename - rename files

SYNOPSIS
       rename [options] expression replacement file...

DESCRIPTION
       rename will rename the specified files by replacing the first
       occurrence of expression in their name by replacement.

OPTIONS
       -s, --symlink
           Do not rename a symlink but change where it points.

       -v, --verbose
           Show which files were renamed, if any.

       -n, --no-act
           Do not make any changes; add --verbose to see what would be made.

       -a, --all
           Replace all occurrences of expression rather than only the first
           one.

       -l, --last
           Replace the last occurrence of expression rather than the first
```

```
            one.

       -o, --no-overwrite
            Do not overwrite existing files. When --symlink is active, do not
            overwrite symlinks pointing to existing targets.

       -i, --interactive
            Ask before overwriting existing files.

       -h, --help
            Display help text and exit.

       -V, --version
            Print version and exit.

WARNING
       The renaming has no safeguards by default or without any one of the
       options --no-overwrite, --interactive or --no-act. If the user has
       permission to rewrite file names, the command will perform the action
       without any questions. For example, the result can be quite drastic
       when the command is run as root in the /lib directory. Always make a
       backup before running the command, unless you truly know what you are
       doing.

EDGE CASES
       If the expression is empty, then by default replacement will be added
       to the start of the filename. With --all, replacement will be inserted
       in between every two characters of the filename, as well as at the
       start and end.

       Normally, only the final path component of a filename is updated. (Or
       with --symlink, only the final path component of the link.) But if
       either expression or replacement contains a /, the full path is
       updated. This can cause a file to be moved between folders. Creating
       folders, and moving files between filesystems, is not supported.

INTERACTIVE MODE
       As most standard utilities rename can be used with a terminal device
       (tty in short) in canonical mode, where the line is buffered by the tty
       and you press ENTER to validate the user input. If you put your tty in
       cbreak mode however, rename requires only a single key press to answer
       the prompt. To set cbreak mode, run for example:

            sh -c 'stty -icanon min 1; "$0" "$@"; stty icanon' rename -i from to files

EXIT STATUS
       0
            all requested rename operations were successful

       1
            all rename operations failed

       2
```

```
           some rename operations failed

       4
           nothing was renamed

       64
           unanticipated error occurred

EXAMPLES
       Given the files foo1, ..., foo9, foo10, ..., foo278, the commands

           rename foo foo00 foo?
           rename foo foo0 foo??

       will turn them into foo001, ..., foo009, foo010, ..., foo278. And

           rename .htm .html *.htm

       will fix the extension of your html files. Provide an empty string for
       shortening:

           rename '_with_long_name' '' file_with_long_name.*

       will remove the substring in the filenames.

SEE ALSO
       mv(1)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The rename command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                         RENAME(1)
```

# Bash 的 resizepart 命令

```
resizepart [OPTION]... DEVICE PARTITION LENGTH
```

**功能**

通知内核磁盘分区大小改变。

**类型**

可执行文件（`/usr/bin/resizepart`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 磁盘设备
- `PARTITION` - 分区
- `LENGTH` - 分区的新大小

## 示例

```
$ sudo resizepart /dev/sda 1 209715200
```

## 推荐阅读

- resizepart(8) — Linux manual page

## 手册

```
RESIZEPART(8)                System Administration                RESIZEPART(8)

NAME
       resizepart - tell the kernel about the new size of a partition

SYNOPSIS
       resizepart device partition length
```

```
DESCRIPTION
       resizepart tells the Linux kernel about the new size of the specified
       partition. The command is a simple wrapper around the "resize
       partition" ioctl.

       This command doesn’t manipulate partitions on a block device.

PARAMETERS
       device
           The disk device.

       partition
           The partition number.

       length
           The new length of the partition (in 512-byte sectors).

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

SEE ALSO
       addpart(8), delpart(8), fdisk(8), parted(8), partprobe(8), partx(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The resizepart command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                    RESIZEPART(8)
```

# Bash 的 rev 命令

```
rev [OPTION]... [FILE]...
```

**功能**

反转行里的字符。

**类型**

可执行文件（`/usr/bin/rev`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-0`，`--zero` - 以空字符（`\0`）作为行的分隔符，而不是换行符（`\n`）
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FILE` - 要查看的文件列表；如果没有指定这个参数则读取标准输入

## 示例

```
$ echo -e "hello\nworld" > 1.txt
$ cat 1.txt
hello
world
$ rev 1.txt
olleh
dlrow
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| cat | 查看文件内容 |
| tac | 逆序查看文件内容 |

## 推荐阅读

- rev(1) — Linux manual page

## 手册

```
REV(1)                          User Commands                          REV(1)

NAME
       rev - reverse lines characterwise

SYNOPSIS
       rev [option] [file...]

DESCRIPTION
       The rev utility copies the specified files to standard output,
       reversing the order of characters in every line. If no files are
       specified, standard input is read.

       This utility is a line-oriented tool and it uses in-memory allocated
       buffer for a whole wide-char line. If the input file is huge and
       without line breaks then allocating the memory for the file may be
       unsuccessful.

OPTIONS
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

       -0, --zero
           Zero termination. Use the byte '\0' as line separator.

SEE ALSO
       tac(1)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The rev command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                          REV(1)
```

# Bash 的 rm 命令

```
rm [OPTION]... [FILE]...
```

**功能**

删除文件或目录。

**类型**

可执行文件（`/usr/bin/rm`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-f`，`--force` - 强制删除，不询问，忽略不存在的文件
  - `-i` - 删除每个文件前进行询问；同 `--interactive=always`
  - `-I` - 删除三个以上文件或递归删除前进行询问；同 `--interactive=once`
  - `--interactive[=WHEN]` - 何时进行询问；有效值：`never`，`once`，`always`
  - `--one-file-system` - 递归删除时跳过不同文件系统的目录
  - `--no-preserve-root` - 不保护 `/` 目录
  - `--preserve-root[=all]` - 保护 `/` 目录，并拒绝删除与其父目录不在同一设备的文件
  - `-r`，`-R`，`--recursive` - 递归删除目录
  - `-d`，`--dir` - 删除空目录
  - `-v`，`--verbose` - 打印详细信息
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `-` - 在所有选项之后，同 `-i` 选项
- `NAME=VALUE` - 要设置的环境变量
- `COMMAND` - 要执行的命令，如果为空则打印所有环境变量
- `ARG` - 命令 `COMMAND` 的参数，在 shebang 行中需要使用 `-S` 选项

## 示例

**删除文件**

```
$ rm file.txt                       # 删除单个文件
$ rm file1.txt file2.log file3.tmp  # 删除多个文件
```

```
$ rm *.tmp                          # 使用通配符
```

### 删除目录

```
$ rm -d dir1                        # 删除空目录
$ rm -r dir2                        # 递归删除目录
$ rm -rf dir3                       # 强制递归删除目录
```

### 删除根目录（危险！）

```
$ rm -rf --no-preserve-root /       # 危险！千万不要执行此命令！
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| rmdir | 删除目录 |
| shred | 粉碎文件（避免恢复） |
| unlink | 删除文件 |

## 推荐阅读

- Remove files or directories - GNU Coreutils
- rm(1) — Linux manual page

## 手册

```
RM(1)                            User Commands                           RM(1)

NAME
       rm - remove files or directories

SYNOPSIS
       rm [OPTION]... [FILE]...

DESCRIPTION
       This  manual  page  documents  the  GNU version of rm.  rm removes each
       specified file.  By default, it does not remove directories.

       If the -I or --interactive=once option is given,  and  there  are  more
       than  three  files  or  the  -r,  -R, or --recursive are given, then rm
       prompts the user for whether to proceed with the entire operation.   If
       the response is not affirmative, the entire command is aborted.
```

Otherwise, if a file is unwritable, standard input is a terminal, and the -f or --force option is not given, or the -i or --interactive=always option is given, rm prompts the user for whether to remove the file. If the response is not affirmative, the file is skipped.

OPTIONS

Remove (unlink) the FILE(s).

-f, --force
ignore nonexistent files and arguments, never prompt

-i prompt before every removal

-I prompt once before removing more than three files, or when removing recursively; less intrusive than -i, while still giving protection against most mistakes

--interactive[=WHEN]
prompt according to WHEN: never, once (-I), or always (-i); without WHEN, prompt always

--one-file-system
when removing a hierarchy recursively, skip any directory that is on a file system different from that of the corresponding command line argument

--no-preserve-root
do not treat '/' specially

--preserve-root[=all]
do not remove '/' (default); with 'all', reject any command line argument on a separate device from its parent

-r, -R, --recursive
remove directories and their contents recursively

-d, --dir
remove empty directories

-v, --verbose
explain what is being done

--help display this help and exit

--version
output version information and exit

By default, rm does not remove directories. Use the --recursive (-r or -R) option to remove each listed directory, too, along with all of its contents.

To remove a file whose name starts with a '-', for example '-foo', use

one of these commands:

```
rm -- -foo

rm ./-foo
```

Note that if you use rm to remove a file, it might be possible to recover some of its contents, given sufficient expertise and/or time. For greater assurance that the contents are truly unrecoverable, consider using shred(1).

## AUTHOR

Written by Paul Rubin, David MacKenzie, Richard M. Stallman, and Jim Meyering.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

unlink(1), unlink(2), chattr(1), shred(1)

Full documentation <https://www.gnu.org/software/coreutils/rm>
or available locally via: info '(coreutils) rm invocation'

# Bash 的 rmdir 命令

```
rmdir [OPTION]... DIRECTORY...
```

**功能**

删除空目录。

**类型**

可执行文件（`/usr/bin/rmdir`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--ignore-fail-on-non-empty` - 忽略删除非空目录的失败
  - `-p`，`--parents` - 同时删除祖先目录
  - `-v`，`--verbose` - 打印详细信息
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `DIRECTORY` - 要删除的目录，必须是空目录

## 示例

**删除一个空目录**

```
rmdir folder
```

如果 `folder` 目录里有文件，则会失败并打印错误信息：

```
rmdir: failed to remove 'myfolder': Directory not empty
```

添加 `--ignore-fail-on-non-empty` 则不打印错误信息。

**删除多个空目录**

```
rmdir dir1 dir2 dir3
```

依次尝试删除这些目录。

**递归删除空目录**

```
rmdir -p a/b/c
```

等价于：

```
rmdir a/b/c
rmdir a/b
rmdir a
```

## 推荐阅读

- Remove empty directories - GNU Coreutils
- rmdir(1) — Linux manual page

## 手册

```
RMDIR(1)                         User Commands                         RMDIR(1)

NAME
       rmdir - remove empty directories

SYNOPSIS
       rmdir [OPTION]... DIRECTORY...

DESCRIPTION
       Remove the DIRECTORY(ies), if they are empty.

       --ignore-fail-on-non-empty
              ignore each failure to remove a non-empty directory

       -p, --parents
              remove DIRECTORY and its ancestors; e.g., 'rmdir -p a/b' is sim-
              ilar to 'rmdir a/b a'

       -v, --verbose
              output a diagnostic for every directory processed

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.
```

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

rmdir(2)

Full documentation <https://www.gnu.org/software/coreutils/rmdir>
or available locally via: info '(coreutils) rmdir invocation'

# Bash 的 rtcwake 命令

```
rtcwake [OPTION]...
```

**功能**

系统进入睡眠状态，直到指定的唤醒时间。

**类型**

可执行文件（`/usr/sbin/rtcwake`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-A`，`--adjfile file` - 指定调整文件的路径
  - `-a`，`--auto` - 从 `adjtime` 文件中读取时钟模式（硬件时钟是本地时间还是 UTC 时间）
  - `--date timestamp` - 设置唤醒时间
  - `-d`，`--device device` - 指定实时时钟设备
  - `-l`，`--local` - 硬件时钟是本地时间
  - `--list-modes` - 列出可用的 `--mode` 选项参数
  - `-m`，`--mode mode` - 进入指定的待机状态
    - `standby` - ACPI 状态 S1，CPU 停止执行指令，CPU 和内存均维持供电
    - `freeze` - 进程冻结，所有设备挂起，所有处理器空闲
    - `mem` - ACPI 状态 S3，仅 RAM 维持供电
    - `disk` - ACPI 状态 S4，内存中的数据存储到硬盘中，CPU、内存、硬盘均停止供电
    - `off` - ACPI 状态 S5，软关机，可以通过键盘唤醒
    - `no` - 不进入睡眠，只需设置 RTC 唤醒时间
    - `on` - 不进入睡眠，读取 RTC 设备直到出现闹钟
    - `disable` - 禁用先前设置的闹钟
    - `show` - 查看闹钟信息
  - `-n`，`--dry-run` - 不执行实际操作
  - `-s`，`--seconds seconds` - 未来 `seconds` 秒唤醒
  - `-t`，`--time time_t` - 在绝对时间 `time_t` 唤醒
  - `-u`，`--utc` - 硬件时钟是 UTC 时间
  - `-v`，`--verbose` - 打印详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本

## 示例

```
$ sudo rtcwake -s 10            # 休眠 10s 后唤醒
rtcwake: assuming RTC uses UTC ...
rtcwake: wakeup from "suspend" using /dev/rtc0 at Wed Mar 18 14:51:36 2026
```

## 推荐阅读

- rtcwake(8) — Linux manual page
- ACPI - Wikipedia

## 手册

```
RTCWAKE(8)                  System Administration                  RTCWAKE(8)

NAME
       rtcwake - enter a system sleep state until specified wakeup time

SYNOPSIS
       rtcwake [options] [-d device] [-m standby_mode] {-s seconds|-t time_t}

DESCRIPTION
       This program is used to enter a system sleep state and to automatically
       wake from it at a specified time.

       This uses cross-platform Linux interfaces to enter a system sleep
       state, and leave it no later than a specified time. It uses any RTC
       framework driver that supports standard driver model wakeup flags.

       This is normally used like the old apmsleep utility, to wake from a
       suspend state like ACPI S1 (standby) or S3 (suspend-to-RAM). Most
       platforms can implement those without analogues of BIOS, APM, or ACPI.

       On some systems, this can also be used like nvram-wakeup, waking from
       states like ACPI S4 (suspend to disk). Not all systems have persistent
       media that are appropriate for such suspend modes.

       Note that alarm functionality depends on hardware; not every RTC is
       able to setup an alarm up to 24 hours in the future.

       The suspend setup may be interrupted by active hardware; for example
       wireless USB input devices that continue to send events for some
       fraction of a second after the return key is pressed. rtcwake tries to
       avoid this problem and it waits to the terminal to settle down before
       entering a system sleep.

OPTIONS
```

-A, --adjfile file
    Specify an alternative path to the adjust file.

-a, --auto
    Read the clock mode (whether the hardware clock is set to UTC or
    local time) from the adjtime file, where hwclock(8) stores that
    information. This is the default.

--date timestamp
    Set the wakeup time to the value of the timestamp. Format of the
    timestamp can be any of the following:

| | |
|---|---|
| YYYYMMDDhhmmss | |
| YYYY-MM-DD hh:mm:ss | |
| YYYY-MM-DD hh:mm | (seconds will be set to 00) |
| YYYY-MM-DD | (time will be set to 00:00:00) |
| hh:mm:ss | (date will be set to today) |
| hh:mm | (date will be set to today, seconds to 00) |
| tomorrow | (time is set to 00:00:00) |
| +5min | |

-d, --device device
    Use the specified device instead of rtc0 as realtime clock. This
    option is only relevant if your system has more than one RTC. You
    may specify rtc1, rtc2, ... here.

-l, --local
    Assume that the hardware clock is set to local time, regardless of
    the contents of the adjtime file.

--list-modes
    List available --mode option arguments.

-m, --mode mode
Go into the given standby state. Valid values for mode are:

standby
ACPI state S1. This state offers minimal, though real, power savings, while providing a very low-latency transition back to a working system. This is the default mode.

freeze
The processes are frozen, all the devices are suspended and all the processors idled. This state is a general state that does not need any platform-specific support, but it saves less power than Suspend-to-RAM, because the system is still in a running state. (Available since Linux 3.9.)

mem
ACPI state S3 (Suspend-to-RAM). This state offers significant power savings as everything in the system is put into a low-power state, except for memory, which is placed in self-refresh mode to retain its contents.

disk
ACPI state S4 (Suspend-to-disk). This state offers the greatest power savings, and can be used even in the absence of low-level platform support for power management. This state operates similarly to Suspend-to-RAM, but includes a final step of writing memory contents to disk.

off
ACPI state S5 (Poweroff). This is done by calling '/sbin/shutdown'. Not officially supported by ACPI, but it usually works.

no
Don’t suspend, only set the RTC wakeup time.

on
Don’t suspend, but read the RTC device until an alarm time appears. This mode is useful for debugging.

disable
Disable a previously set alarm.

show
Print alarm information in format: "alarm: off|on <time>". The time is in ctime() output format, e.g., "alarm: on Tue Nov 16 04:48:45 2010".

-n, --dry-run
This option does everything apart from actually setting up the alarm, suspending the system, or waiting for the alarm.

-s, --seconds seconds

```
           Set the wakeup time to seconds in the future from now.

       -t, --time time_t
           Set the wakeup time to the absolute time time_t. time_t is the time
           in seconds since 1970-01-01, 00:00 UTC. Use the date(1) tool to
           convert between human-readable time and time_t.

       -u, --utc
           Assume that the hardware clock is set to UTC (Universal Time
           Coordinated), regardless of the contents of the adjtime file.

       -v, --verbose
           Be verbose.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

NOTES
       Some PC systems can’t currently exit sleep states such as mem using
       only the kernel code accessed by this driver. They need help from
       userspace code to make the framebuffer work again.

FILES
       /etc/adjtime

HISTORY
       The program was posted several times on LKML and other lists before
       appearing in kernel commit message for Linux 2.6 in the GIT commit
       87ac84f42a7a580d0dd72ae31d6a5eb4bfe04c6d.

AUTHORS
       The program was written by David Brownell
       <dbrownell@users.sourceforge.net> and improved by Bernhard Walle
       <bwalle@suse.de>.

COPYRIGHT
       This is free software. You may redistribute copies of it under the
       terms of the GNU General Public License
       <http://www.gnu.org/licenses/gpl.html>. There is NO WARRANTY, to the
       extent permitted by law.

SEE ALSO
       adjtime_config(5), hwclock(8), date(1)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The rtcwake command is part of the util-linux package which can be
```

```
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                          RTCWAKE(8)
```

# Bash 的 runcon 命令

```
runcon [OPTION]... CONTEXT COMMAND [args]...
```

**功能**

使用指定的 SELinux 安全上下文运行命令。

**类型**

可执行文件（`/usr/bin/runcon`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--compute` - 在应用指定的安全上下文之前先计算进程可能的转变安全上下文
  - `-t`，`--type=TYPE` - 指定安全上下文的的 `TYPE`
  - `-u`，`--user=USER` - 指定安全上下文的的 `USER`
  - `-r`，`--role=ROLE` - 指定安全上下文的的 `ROLE`
  - `-l`，`--range=LEVEL` - 指定安全上下文的的 `LEVEL`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CONTEXT` - 安全上下文；通常格式为 `USER:ROLE:TYPE:LEVEL`
- `COMMAND` - 命令
- `args` - 命令的参数列表

## 示例

```
$ runcon system_u:system_r:httpd_t:s0 id -Z
system_u:system_r:httpd_t:s0
```

## 推荐阅读

- Run a command in specified SELinux context - GNU Coreutils
- runcon(1) — Linux manual page

## 手册

RUNCON(1) User Commands RUNCON(1)

NAME

runcon - run command with specified security context

SYNOPSIS

runcon CONTEXT COMMAND [args]

runcon [ -c ] [-u USER] [-r ROLE] [-t TYPE] [-l RANGE] COMMAND [args]

DESCRIPTION

Run COMMAND with completely-specified CONTEXT, or with current or transitioned security context modified by one or more of LEVEL, ROLE, TYPE, and USER.

If none of -c, -t, -u, -r, or -l, is specified, the first argument is used as the complete context. Any additional arguments after COMMAND are interpreted as arguments to the command.

Note that only carefully-chosen contexts are likely to successfully run.

Run a program in a different SELinux security context. With neither CONTEXT nor COMMAND, print the current security context.

Mandatory arguments to long options are mandatory for short options too.

CONTEXT

Complete security context

-c, --compute

compute process transition context before modifying

-t, --type=TYPE

type (for same role as parent)

-u, --user=USER

user identity

-r, --role=ROLE

role

-l, --range=RANGE

levelrange

--help display this help and exit

--version

output version information and exit

Exit status:

125 if the runcon command itself fails

```
       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found

       -      the exit status of COMMAND otherwise

AUTHOR
       Written by Russell Coker.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/runcon>
       or available locally via: info '(coreutils) runcon invocation'

GNU coreutils 9.4                April 2024                         RUNCON(1)
```

# Bash 的 runuser 命令

```
runuser [OPTION]... [USER] [CMD [ARG]...]
```

### 功能

使用指定的用户和组运行程序。

### 类型

可执行文件（`/usr/sbin/runuser`），属于 util-linux。

### 参数

- `OPTIONS` 选项：
  - `-c`，`--command=command` - 要执行的命令
  - `-f`，`--fast` - 向 shell 传递 `-f`
  - `-g`，`--group=group` - 指定主要组
  - `-G`，`--supp-group=group` - 指定补充组
  - `-`，`-l`，`--login` - 启动 login shell
  - `-P`，`--pty` - 为会话创建一个伪终端
  - `-m`，`-p`，`--preserve-environment` - 保留所有环境变量
  - `-s`，`--shell=shell` - 运行指定的 shell
  - `--session-command=command` - 要执行的命令；类似 `-c` 但不创建会话
  - `-u user` - 指定用户名
  - `-w`，`--whitelist-environment=list`
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `USER` - 以此用户的身份启动 shell；执行命令时必须使用 `-u` 选项
- `CMD` - 要执行的命令
- `ARG` - `CMD` 的参数

## 示例

### 执行命令

```
$ sudo runuser -u user ls /home/user    # 以 user 用户的身份运行 ls /home/user
1.txt  2.txt  3.txt  Downloads  Projects
```

**启动 shell**

```
$ sudo runuser user                    # 以 user 用户的身份启动 shell
```

# 推荐阅读

- runuser(8) — Linux manual page

# 手册

```
RUNUSER(1)                      User Commands                      RUNUSER(1)

NAME
       runuser - run a command with substitute user and group ID

SYNOPSIS
       runuser [options] -u user [[--] command [argument...]]

       runuser [options] [-] [user [argument...]]

DESCRIPTION
       runuser can be used to run commands with a substitute user and group
       ID. If the option -u is not given, runuser falls back to su-compatible
       semantics and a shell is executed. The difference between the commands
       runuser and su is that runuser does not ask for a password (because it
       may be executed by the root user only) and it uses a different PAM
       configuration. The command runuser does not have to be installed with
       set-user-ID permissions.

       If the PAM session is not required, then the recommended solution is to
       use the setpriv(1) command.

       When called without arguments, runuser defaults to running an
       interactive shell as root.

       For backward compatibility, runuser defaults to not changing the
       current directory and to setting only the environment variables HOME
       and SHELL (plus USER and LOGNAME if the target user is not root). This
       version of runuser uses PAM for session management.

       Note that runuser in all cases use PAM (pam_getenvlist()) to do the
       final environment modification. Command-line options such as --login
       and --preserve-environment affect the environment before it is modified
       by PAM.

       Since version 2.38 runuser resets process resource limits RLIMIT_NICE,
       RLIMIT_RTPRIO, RLIMIT_FSIZE, RLIMIT_AS and RLIMIT_NOFILE.
```

OPTIONS

-c, --command=command
Pass command to the shell with the -c option.

-f, --fast
Pass -f to the shell, which may or may not be useful, depending on the shell.

-g, --group=group
The primary group to be used. This option is allowed for the root user only.

-G, --supp-group=group
Specify a supplementary group. This option is available to the root user only. The first specified supplementary group is also used as a primary group if the option --group is not specified.

-, -l, --login
Start the shell as a login shell with an environment similar to a real login:

- clears all the environment variables except for TERM and variables specified by --whitelist-environment
- initializes the environment variables HOME, SHELL, USER, LOGNAME, and PATH
- changes to the target user’s home directory
- sets argv[0] of the shell to '-' in order to make the shell a login shell

-P, --pty
Create a pseudo-terminal for the session. The independent terminal provides better security as the user does not share a terminal with the original session. This can be used to avoid TIOCSTI ioctl terminal injection and other security attacks against terminal file descriptors. The entire session can also be moved to the background (e.g., runuser --pty -u username -- command &). If the pseudo-terminal is enabled, then runuser works as a proxy between the sessions (sync stdin and stdout).

This feature is mostly designed for interactive sessions. If the standard input is not a terminal, but for example a pipe (e.g., echo "date" | runuser --pty -u user), then the ECHO flag for the pseudo-terminal is disabled to avoid messy output.

-m, -p, --preserve-environment
Preserve the entire environment, i.e., do not set HOME, SHELL, USER or LOGNAME. The option is ignored if the option --login is specified.

-s, --shell=shell

Run the specified shell instead of the default. The shell to run is selected according to the following rules, in order:

- the shell specified with --shell
- the shell specified in the environment variable SHELL if the --preserve-environment option is used
- the shell listed in the passwd entry of the target user
- /bin/sh

If the target user has a restricted shell (i.e., not listed in /etc/shells), then the --shell option and the SHELL environment variables are ignored unless the calling user is root.

--session-command=command
Same as -c, but do not create a new session. (Discouraged.)

-w, --whitelist-environment=list
Don't reset the environment variables specified in the comma-separated list when clearing the environment for --login. The whitelist is ignored for the environment variables HOME, SHELL, USER, LOGNAME, and PATH.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

CONFIG FILES

runuser reads the /etc/default/runuser and /etc/login.defs configuration files. The following configuration items are relevant for runuser:

ENV_PATH (string)
Defines the PATH environment variable for a regular user. The default value is /usr/local/bin:/bin:/usr/bin.

ENV_ROOTPATH (string), ENV_SUPATH (string)
Defines the PATH environment variable for root. ENV_SUPATH takes precedence. The default value is /usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin.

ALWAYS_SET_PATH (boolean)
If set to yes and --login and --preserve-environment were not specified runuser initializes PATH.

The environment variable PATH may be different on systems where /bin and /sbin are merged into /usr; this variable is also affected by the --login command-line option and the PAM system setting (e.g., pam_env(8)).

```
EXIT STATUS
       runuser normally returns the exit status of the command it executed. If
       the command was killed by a signal, runuser returns the number of the
       signal plus 128.

       Exit status generated by runuser itself:

       1
           Generic error before executing the requested command

       126
           The requested command could not be executed

       127
           The requested command was not found

FILES
       /etc/pam.d/runuser
           default PAM configuration file

       /etc/pam.d/runuser-l
           PAM configuration file if --login is specified

       /etc/default/runuser
           runuser specific logindef config file

       /etc/login.defs
           global logindef config file

HISTORY
       This runuser command was derived from coreutils' su, which was based on
       an implementation by David MacKenzie, and the Fedora runuser command by
       Dan Walsh.

SEE ALSO
       setpriv(1), su(1), login.defs(5), shells(5), pam(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The runuser command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                         RUNUSER(1)
```

# Bash 的 seq 命令

```
seq [OPTION]... FIRST [INCREMENT] LAST
```

**功能**

打印一串数值序列。

**类型**

可执行文件（`/usr/bin/seq`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-f`，`--format=FORMAT` - 使用 printf 风格的浮点的格式
  - `-s`，`--separator=STRING` - 使用 `STRING` 作为分隔符；默认为换行符（`\n`）
  - `-w`，`--equal-width` - 通过填 0 来补充宽度
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FIRST` - 第一个数值
- `INCREMENT` - 增量；默认为 1
- `LAST` - 最后一个数值

## 示例

```
$ seq 1 10                  # 打印 1 到 10
1
2
3
4
5
6
7
8
9
10
$ seq -s , -10 10           # 打印 -10 到 10，使用 `,` 分隔，
-10,-9,-8,-7,-6,-5,-4,-3,-2,-1,0,1,2,3,4,5,6,7,8,9,10
$ seq -s , -w 10 -1 -10     # 打印 -10 到 10，使用 `,` 分隔，补零平衡宽度
010,009,008,007,006,005,004,003,002,001,000,-01,-02,-03,-04,-05,-06,-07,-08,-09,-10
```

```
$ seq -s - 1 2 10          # 打印 1 到 10，增量为 2，使用 `-` 分隔
1-3-5-7-9
```

## 推荐阅读

- Print numeric sequences - GNU Coreutils
- seq(1) — Linux manual page

## 手册

```
SEQ(1)                          User Commands                          SEQ(1)

NAME
       seq - print a sequence of numbers

SYNOPSIS
       seq [OPTION]... LAST
       seq [OPTION]... FIRST LAST
       seq [OPTION]... FIRST INCREMENT LAST

DESCRIPTION
       Print numbers from FIRST to LAST, in steps of INCREMENT.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -f, --format=FORMAT
              use printf style floating-point FORMAT

       -s, --separator=STRING
              use STRING to separate numbers (default: \n)

       -w, --equal-width
              equalize width by padding with leading zeroes

       --help display this help and exit

       --version
              output version information and exit

       If FIRST or INCREMENT is omitted, it defaults to 1.  That is, an  omit-
       ted  INCREMENT defaults to 1 even when LAST is smaller than FIRST.  The
       sequence of numbers ends when the sum of the current number and  INCRE-
       MENT  would  become  greater than LAST.  FIRST, INCREMENT, and LAST are
       interpreted as floating point values.  INCREMENT is usually positive if
       FIRST is smaller than LAST, and INCREMENT is usually negative if  FIRST
       is  greater  than LAST.  INCREMENT must not be 0; none of FIRST, INCRE-
       MENT and LAST may be NaN.  FORMAT must be suitable for printing one ar-
```

gument of type 'double'; it defaults to %.PRECf if FIRST, INCREMENT, and LAST are all fixed point decimal numbers with maximum precision PREC, and to %g otherwise.

AUTHOR

Written by Ulrich Drepper.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/seq>
or available locally via: info '(coreutils) seq invocation'

# Bash 的 setarch 命令

```
setarch [ARCH] [OPTION]... [CMD [ARG]...]
```

**功能**

使用指定的架构或个性化标志运行程序。

**类型**

可执行文件（`/usr/bin/setarch`），属于 util-linux。

**参数**

- `ARCH` - 要设为的架构
- `OPTION` 选项：
  - `--list` - 列出已知架构
  - `--show[=personality]` - 显示当前的个性化标志
  - `--uname-2.6` - 使程序检测到的内核版本号为 `2.6`
  - `-v`，`--verbose` - 显示详细信息
  - `-3`，`--3gb` - 指定程序最多使用 3GB 的地址空间（3GB 用户空间，1GB 内核空间）
  - `--4gb` - ~~指定程序最多使用 4GB 的地址空间（4GB 用户空间，没有内核空间）~~；此选项已废弃不再生效
  - `-B`，`--32bit` - 将地址空间设为 32 位以模拟硬件；支持 ARM 和 Alpha 架构
  - `-F`，`--fdpic-funcptrs` - 将用户空间函数指向信号处理程序的指针视为指向地址描述符的指针；仅支持 ARM, Blackfin, Fujitsu FR-V 和 SuperH 架构
  - `-I`，`--short-inode` - ~~使用 16 位 inode 号~~；此选项已废弃不再生效
  - `-L`，`--addr-compat-layout` - 提供传统虚拟地址空间布局
  - `-R`，`--addr-no-randomize` - 禁用虚拟地址空间的随机化
  - `-S`，`--whole-seconds` - ~~时间值截断为整数秒~~；此选项已废弃不再生效
  - `-T`，`--sticky-timeouts` - 这使得 `select(2)`、`pselect(2)` 和 `ppoll(2)` 系统调用保留超时值
  - `-X`，`--read-implies-exec` - 如果设置此项，则 `mmap(2)` `PROT_READ` 还会添加 `PROT_EXEC`` 位
  - `-Z`，`--mmap-page-zero` - SVr4 错误模拟，会将 `mmap(2)` 第 0 页设置为只读。
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CMD` - 要运行的命令

- `ARG` - 传递给 `CMD` 的参数列表

## 示例

### 查看已知架构

```
$ setarch --list
uname26
linux32
linux64
i386
i486
i586
i686
athlon
x86_64
```

### 查看当前个性化标志

```
$ setarch --show
PER_LINUX
```

### 设置架构

```
$ arch          # 查看架构
x86_64
$ setarch i386 arch
i686
$ setarch linux32 arch
i686
```

## 推荐阅读

- setarch(8) — Linux manual page

## 手册

```
SETARCH(8)                  System Administration                  SETARCH(8)

NAME
       setarch - change reported architecture in new program environment
       and/or set personality flags

SYNOPSIS
```

```
       setarch [arch] [options] [program [argument...]]

       setarch --list|-h|-V

       arch [options] [program [argument...]]

DESCRIPTION
       setarch modifies execution domains and process personality flags.

       The execution domains currently only affects the output of uname -m.
       For example, on an AMD64 system, running setarch i386 program will
       cause program to see i686 instead of x86_64 as the machine type. It can
       also be used to set various personality options. The default program is
       /bin/sh.

       Since version 2.33 the arch command line argument is optional and
       setarch may be used to change personality flags (ADDR_LIMIT_*,
       SHORT_INODE, etc) without modification of the execution domain.

OPTIONS
       --list
           List the architectures that setarch knows about. Whether setarch
           can actually set each of these architectures depends on the running
           kernel.

       --show[=personality]
           Show the currently active personality and flags. If the personality
           argument is provided, it is shown instead of the current one.
           personality is a hexadecimal number with values was described in
           sys/personality.h.

       --uname-2.6
           Causes the program to see a kernel version number beginning with
           2.6. Turns on UNAME26.

       -v, --verbose
           Be verbose.

       -3, --3gb
           Specifies program should use a maximum of 3GB of address space.
           Supported on x86. Turns on ADDR_LIMIT_3GB.

       --4gb
           This option has no effect. It is retained for backward
           compatibility only, and may be removed in future releases.

       -B, --32bit
           Limit the address space to 32 bits to emulate hardware. Supported
           on ARM and Alpha. Turns on ADDR_LIMIT_32BIT.

       -F, --fdpic-funcptrs
           Treat user-space function pointers to signal handlers as pointers
           to address descriptors. This option has no effect on architectures
```

```
            that do not support FDPIC ELF binaries. In kernel v4.14 support is
            limited to ARM, Blackfin, Fujitsu FR-V, and SuperH CPU
            architectures.

        -I, --short-inode
            Obsolete bug emulation flag. Turns on SHORT_INODE.

        -L, --addr-compat-layout
            Provide legacy virtual address space layout. Use when the program
            binary does not have PT_GNU_STACK ELF header. Turns on
            ADDR_COMPAT_LAYOUT.

        -R, --addr-no-randomize
            Disables randomization of the virtual address space. Turns on
            ADDR_NO_RANDOMIZE.

        -S, --whole-seconds
            Obsolete bug emulation flag. Turns on WHOLE_SECONDS.

        -T, --sticky-timeouts
            This makes select(2), pselect(2), and ppoll(2) system calls
            preserve the timeout value instead of modifying it to reflect the
            amount of time not slept when interrupted by a signal handler. Use
            when program depends on this behavior. For more details see the
            timeout description in select(2) manual page. Turns on
            STICKY_TIMEOUTS.

        -X, --read-implies-exec
            If this is set then mmap(2) PROT_READ will also add the PROT_EXEC
            bit - as expected by legacy x86 binaries. Notice that the ELF
            loader will automatically set this bit when it encounters a legacy
            binary. Turns on READ_IMPLIES_EXEC.

        -Z, --mmap-page-zero
            SVr4 bug emulation that will set mmap(2) page zero as read-only.
            Use when program depends on this behavior, and the source code is
            not available to be fixed. Turns on MMAP_PAGE_ZERO.

        -h, --help
            Display help text and exit.

        -V, --version
            Print version and exit.

EXAMPLE
            setarch --addr-no-randomize mytestprog
            setarch ppc32 rpmbuild --target=ppc --rebuild foo.src.rpm
            setarch ppc32 -v -vL3 rpmbuild --target=ppc --rebuild bar.src.rpm
            setarch ppc32 --32bit rpmbuild --target=ppc --rebuild foo.src.rpm

AUTHORS
        Elliot Lee <sopwith@redhat.com>, Jindrich Novy <jnovy@redhat.com>,
        Karel Zak <kzak@redhat.com>
```

```
SEE ALSO
       personality(2), select(2)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The setarch command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-11-21                        SETARCH(8)
```

# Bash 的 setpriv 命令

```
setpriv [OPTION]... [CMD [ARG]...]
```

**功能**

使用指定的 Linux 特权设置运行程序

**类型**

可执行文件（`/usr/bin/setpriv`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `--clear-groups` - 清除补充组
  - `-d`, `--dump` - 转储当前特权状态
  - `--groups group...` - 设置补充组
  - `--inh-caps (+|-)cap...` - 设置可继承能力
  - `--ambient-caps (+|-)cap...` - 环境能力
  - `--bounding-set (+|-)cap...` - 设置能力边界集
  - `--keep-groups` - 保留补充组
  - `--init-groups` - 初始化补充组
  - `--list-caps` - 列出所有已知能力
  - `--no-new-privs` - 限制进程在执行新程序时 不能获得新的权限
  - `--rgid gid` - 设置实际组 ID
  - `--egid gid` - 设置有效组 ID
  - `--regid gid` - 设置实际组 ID 和有效组 ID
  - `--ruid uid` - 设置实际用户 ID
  - `--euid uid` - 设置有效用户 ID
  - `--reuid uid` - 设置实际用户 ID 和有效用户 ID
  - `--securebits (+|-)securebit...` - 设置或清除安全位
  - `--pdeathsig keep|clear|<signal>` - 保留、清除或设置父进程死亡信号
  - `--selinux-label label` - 请求特定的 SELinux 转换
  - `--apparmor-profile profile` - 请求特定的 AppArmor 配置
  - `--reset-env` - 清除除了 `TERM` 以外的所有环境变量
  - `--help` - 显示帮助

    - `--version` - 显示版本
- `CMD` - 要运行的命令
- `ARG` - 传递给 `CMD` 的参数列表

## 示例

### 查看当前特权设置

```
$ setpriv --dump
uid: 1000
euid: 1000
gid: 1000
egid: 1000
Supplementary groups: 4,24,27,30,46,101,1000
no_new_privs: 0
Inheritable capabilities: [none]
Ambient capabilities: [none]
Capability bounding set:
chown,dac_override,dac_read_search,fowner,fsetid,kill,setgid,setuid,setpcap,linux_immutabl
e,net_bind_service,net_broadcast,net_admin,net_raw,ipc_lock,ipc_owner,sys_module,sys_rawio
,sys_chroot,sys_ptrace,sys_pacct,sys_admin,sys_boot,sys_nice,sys_resource,sys_time,sys_tty
_config,mknod,lease,audit_write,audit_control,setfcap,mac_override,mac_admin,syslog,wake_a
larm,block_suspend,audit_read,perfmon,bpf,checkpoint_restore
Securebits: [none]
Parent death signal: [none]
AppArmor profile: unconfined
$ sudo setpriv --dump
uid: 0
euid: 0
gid: 0
egid: 0
Supplementary groups: 0
no_new_privs: 0
Inheritable capabilities: [none]
Ambient capabilities: [none]
Capability bounding set:
chown,dac_override,dac_read_search,fowner,fsetid,kill,setgid,setuid,setpcap,linux_immutabl
e,net_bind_service,net_broadcast,net_admin,net_raw,ipc_lock,ipc_owner,sys_module,sys_rawio
,sys_chroot,sys_ptrace,sys_pacct,sys_admin,sys_boot,sys_nice,sys_resource,sys_time,sys_tty
_config,mknod,lease,audit_write,audit_control,setfcap,mac_override,mac_admin,syslog,wake_a
larm,block_suspend,audit_read,perfmon,bpf,checkpoint_restore
Securebits: [none]
Parent death signal: [none]
AppArmor profile: unconfined
```

### 修改特权设置

```
$ id    # 查看 ID
uid=1000(primers) gid=1000(primers)
```

```
groups=1000(primers),4(adm),24(cdrom),27(sudo),30(dip),46(plugdev),101(lxd)
$ sudo setpriv --reuid=1001 --regid=1001 --clear-groups id  # 指定用户 ID 和组 ID，并清除补充组 ID
uid=1001(user2) gid=1001(user2) groups=1001(user2)
```

# 推荐阅读

- setpriv(1) — Linux manual page

# 手册

```
SETPRIV(1)                      User Commands                      SETPRIV(1)

NAME
       setpriv - run a program with different Linux privilege settings

SYNOPSIS
       setpriv [options] program [arguments]

DESCRIPTION
       Sets or queries various Linux privilege settings that are inherited
       across execve(2).

       In comparison to su(1) and runuser(1), setpriv neither uses PAM, nor
       does it prompt for a password. It is a simple, non-set-user-ID wrapper
       around execve(2), and can be used to drop privileges in the same way as
       setuidgid(8) from daemontools, chpst(8) from runit, or similar tools
       shipped by other service managers.

OPTIONS
       --clear-groups
           Clear supplementary groups.

       -d, --dump
           Dump the current privilege state. This option can be specified more
           than once to show extra, mostly useless, information. Incompatible
           with all other options.

       --groups group...
           Set supplementary groups. The argument is a comma-separated list of
           GIDs or names.

       --inh-caps (+|-)cap..., --ambient-caps (+|-)cap..., --bounding-set
       (+|-)cap...
           Set the inheritable capabilities, ambient capabilities or the
           capability bounding set. See capabilities(7). The argument is a
           comma-separated list of +cap and -cap entries, which add or remove
           an entry respectively. cap can either be a human-readable name as
           seen in capabilities(7) without the cap_ prefix or of the format
```

cap_N, where N is the internal capability index used by Linux. +all and -all can be used to add or remove all caps.

The set of capabilities starts out as the current inheritable set for --inh-caps, the current ambient set for --ambient-caps and the current bounding set for --bounding-set.

Note the following restrictions (detailed in capabilities(7)) regarding modifications to these capability sets:

- A capability can be added to the inheritable set only if it is currently present in the bounding set.
- A capability can be added to the ambient set only if it is currently present in both the permitted and inheritable sets.
- Notwithstanding the syntax offered by setpriv, the kernel does not permit capabilities to be added to the bounding set.

If you drop a capability from the bounding set without also dropping it from the inheritable set, you are likely to become confused. Do not do that.

--keep-groups
Preserve supplementary groups. Only useful in conjunction with --rgid, --egid, or --regid.

--init-groups
Initialize supplementary groups using initgroups3. Only useful in conjunction with --ruid or --reuid.

--list-caps
List all known capabilities. This option must be specified alone.

--no-new-privs
Set the no_new_privs bit. With this bit set, execve(2) will not grant new privileges. For example, the set-user-ID and set-group-ID bits as well as file capabilities will be disabled. (Executing binaries with these bits set will still work, but they will not gain privileges. Certain LSMs, especially AppArmor, may result in failures to execute certain programs.) This bit is inherited by child processes and cannot be unset. See prctl(2) and Documentation/prctl/no_new_privs.txt in the Linux kernel source.

The no_new_privs bit is supported since Linux 3.5.

--rgid gid, --egid gid, --regid gid
Set the real, effective, or both GIDs. The gid argument can be given as a textual group name.

For safety, you must specify one of --clear-groups, --groups, --keep-groups, or --init-groups if you set any primary gid.

```
--ruid uid, --euid uid, --reuid uid
    Set the real, effective, or both UIDs. The uid argument can be
    given as a textual login name.

    Setting a uid or gid does not change capabilities, although the
    exec call at the end might change capabilities. This means that, if
    you are root, you probably want to do something like:

    setpriv --reuid=1000 --regid=1000 --inh-caps=-all

--securebits (+|-)securebit...
    Set or clear securebits. The argument is a comma-separated list.
    The valid securebits are noroot, noroot_locked, no_setuid_fixup,
    no_setuid_fixup_locked, and keep_caps_locked. keep_caps is cleared
    by execve(2) and is therefore not allowed.

--pdeathsig keep|clear|<signal>
    Keep, clear or set the parent death signal. Some LSMs, most notably
    SELinux and AppArmor, clear the signal when the process'
    credentials change. Using --pdeathsig keep will restore the parent
    death signal after changing credentials to remedy that situation.

--selinux-label label
    Request a particular SELinux transition (using a transition on
    exec, not dyntrans). This will fail and cause setpriv to abort if
    SELinux is not in use, and the transition may be ignored or cause
    execve(2) to fail at SELinux's whim. (In particular, this is
    unlikely to work in conjunction with no_new_privs.) This is similar
    to runcon(1).

--apparmor-profile profile
    Request a particular AppArmor profile (using a transition on exec).
    This will fail and cause setpriv to abort if AppArmor is not in
    use, and the transition may be ignored or cause execve(2) to fail
    at AppArmor's whim.

--reset-env
    Clears all the environment variables except TERM; initializes the
    environment variables HOME, SHELL, USER, LOGNAME according to the
    user's passwd entry; sets PATH to /usr/local/bin:/bin:/usr/bin for
    a regular user and to
    /usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin for
    root.

    The environment variable PATH may be different on systems where
    /bin and /sbin are merged into /usr. The environment variable SHELL
    defaults to /bin/sh if none is given in the user's passwd entry.

-h, --help
    Display help text and exit.

-V, --version
    Print version and exit.
```

NOTES

If applying any specified option fails, program will not be run and setpriv will return with exit status 127.

Be careful with this tool — it may have unexpected security consequences. For example, setting no_new_privs and then execing a program that is SELinux-confined (as this tool would do) may prevent the SELinux restrictions from taking effect.

EXAMPLES

If you're looking for behavior similar to su(1)/runuser(1), or sudo(8) (without the -g option), try something like:

```
setpriv --reuid=1000 --regid=1000 --init-groups
```

If you want to mimic daemontools' setuid(8), try:

```
setpriv --reuid=1000 --regid=1000 --clear-groups
```

AUTHORS

Andy Lutomirski <luto@amacapital.net>

SEE ALSO

runuser(1), su(1), prctl(2), capabilities(7)

REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The setpriv command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 setsid 命令

```
setsid [OPTION]... [CMD [ARG]...]
```

**功能**

使用新的会话运行程序。

**类型**

可执行文件（`/usr/bin/setsid`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-c`，`--ctty` - 将控制终端设为当前终端
  - `-f`，`--fork` - 先创建一个子进程，再让子进程调用 `setsid(2)`；`setsid` 命令自身是进程组 leader 时必须使用此选项
  - `-w`，`--wait` - 等待 `CMD` 执行完毕，以 `CMD` 的返回值作为 `setsid` 的返回值返回
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CMD` - 要运行的命令
- `ARG` - 传递给 `CMD` 的参数列表

## 示例

```
$ setsid java -jar minecraft_server.jar nogui &      # & 表示后台运行
```

- 在新的会话中运行 `minecraft_server`，与控制终端解绑，不会收到终端挂断信号，用户注销后仍可继续运行。

## 推荐阅读

- setsid(1) — Linux manual page

## 手册

NAME

setsid - run a program in a new session

SYNOPSIS

setsid [options] program [arguments]

DESCRIPTION

setsid runs a program in a new session. The command calls fork(2) if already a process group leader. Otherwise, it executes a program in the current process. This default behavior is possible to override by the --fork option.

OPTIONS

-c, --ctty
Set the controlling terminal to the current one.

-f, --fork
Always create a new process.

-w, --wait
Wait for the execution of the program to end, and return the exit status of this program as the exit status of setsid.

-V, --version
Display version information and exit.

-h, --help
Display help text and exit.

AUTHORS

Rick Sladkey <jrs@world.std.com>

SEE ALSO

setsid(2)

REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The setsid command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 setterm 命令

```
setterm [OPTION]...
```

**功能**

设置终端属性。

**类型**

可执行文件（`/usr/bin/setterm`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `--appcursorkeys on|off` - 启用或禁用应用程序光标键模式
  - `--append console_number` - 将输出追加到指定控制台
  - `--background 8-color|default` - 设置背景颜色
  - `--blank[=0-60|force|poke]` - 设置屏幕保护空白时间（分钟），force 强制立即空白，poke 重置计时
  - `--bfreq[=number]` - 设置屏幕闪烁频率（Hz）
  - `--blength[=0-2000]` - 设置屏幕闪烁长度（毫秒）
  - `--blink on|off` - 启用或禁用闪烁文本
  - `--bold on|off` - 启用或禁用粗体文本
  - `--clear[=all|rest]` - 清屏，all 清除全部，rest 清除光标之后内容
  - `--clrtabs[=tab1 tab2 tab3 ...]` - 清除指定制表位
  - `--cursor on|off` - 显示或隐藏光标
  - `--default` - 恢复终端默认设置
  - `--dump[=console_number]` - 将控制台内容转储到文件
  - `--file filename` - 指定 dump 文件名或输入文件
  - `--foreground 8-color|default` - 设置前景颜色
  - `--half-bright on|off` - 启用或禁用半亮显示
  - `--hbcolor 16-color` - 设置半亮颜色
  - `--initialize` - 初始化终端（清除屏幕、重置属性）
  - `--inversescreen on|off` - 启用或禁用反显屏幕
  - `--linewrap on|off` - 启用或禁用自动换行
  - `--msg on|off` - 显示或隐藏终端信息消息
  - `--msglevel 0-8` - 设置终端信息消息等级（0 最少，8 最多）

- `--powerdown[=0-60]` - 设置屏幕电源关闭时间（分钟）
- `--powersave mode` - 设置屏幕节能模式（如 on/off/standby）
- `--regtabs[=1-160]` - 设置标准制表位间隔
- `--repeat on|off` - 启用或禁用键盘重复功能
- `--reset` - 重置终端为默认状态
- `--resize` - 调整终端大小以匹配当前窗口
- `--reverse on|off` - 启用或禁用反色显示
- `--store` - 保存当前终端设置以便 later `--restore`
- `--tabs[=tab1 tab2 tab3 ...]` - 设置制表位
- `--term terminal_name` - 指定终端类型
- `--ulcolor 16-color` - 设置下划线文本颜色
- `--underline on|off` - 启用或禁用下划线
- `--help` - 显示帮助
- `--version` - 显示版本

## 示例

```
$ setterm --clear               # 清屏
$ setterm --foreground red      # 设置前景色
$ setterm --background black    # 设计背景色
$ setterm --cursor off          # 不显示光标
$ setterm --bold on             # 开启字体加粗
```

## 推荐阅读

- setterm(1) — Linux manual page

## 手册

```
SETTERM(1)                      User Commands                      SETTERM(1)

NAME
       setterm - set terminal attributes

SYNOPSIS
       setterm [options]

DESCRIPTION
       setterm writes to standard output a character string that will invoke
       the specified terminal capabilities. Where possible terminfo is
       consulted to find the string to use. Some options however (marked
```

"virtual consoles only" below) do not correspond to a terminfo(5) capability. In this case, if the terminal type is "con" or "linux" the string that invokes the specified capabilities on the PC Minix virtual console driver is output. Options that are not implemented by the terminal are ignored.

OPTIONS

For boolean options (on or off), the default is on.

Below, an 8-color can be black, red, green, yellow, blue, magenta, cyan, or white.

A 16-color can be an 8-color, or grey, or bright followed by red, green, yellow, blue, magenta, cyan, or white.

The various color options may be set independently, at least on virtual consoles, though the results of setting multiple modes (for example, --underline and --half-bright) are hardware-dependent.

The optional arguments are recommended with '=' (equals sign) and not space between the option and the argument. For example --option=argument. setterm can interpret the next non-option argument as an optional argument too.

--appcursorkeys on|off
Sets Cursor Key Application Mode on or off. When on, ESC O A, ESC O B, etc. will be sent for the cursor keys instead of ESC [ A, ESC [ B, etc. See the vi and Cursor-Keys section of the Text-Terminal-HOWTO for how this can cause problems for vi users. Virtual consoles only.

--append console_number
Like --dump, but appends to the snapshot file instead of overwriting it. Only works if no --dump options are given.

--background 8-color|default
Sets the background text color.

--blank[=0-60|force|poke]
Sets the interval of inactivity, in minutes, after which the screen will be automatically blanked (using APM if available). Without an argument, it gets the blank status (returns which vt was blanked, or zero for an unblanked vt). Virtual consoles only.

The force argument keeps the screen blank even if a key is pressed.

The poke argument unblanks the screen.

--bfreq[=number]
Sets the bell frequency in Hertz. Without an argument, it defaults to 0. Virtual consoles only.

--blength[=0-2000]

```
    Sets the bell duration in milliseconds. Without an argument, it
    defaults to 0. Virtual consoles only.

--blink on|off
    Turns blink mode on or off. Except on a virtual console, --blink
    off turns off all attributes (bold, half-brightness, blink,
    reverse).

--bold on|off
    urns bold (extra bright) mode on or off. Except on a virtual
    console, --bold off turns off all attributes (bold,
    half-brightness, blink, reverse).

--clear[=all|rest]
    Without an argument or with the argument all, the entire screen is
    cleared and the cursor is set to the home position, just like
    clear(1) does. With the argument rest, the screen is cleared from
    the current cursor position to the end.

--clrtabs[=tab1 tab2 tab3 ...]
    Clears tab stops from the given horizontal cursor positions, in the
    range 1-160. Without arguments, it clears all tab stops. Virtual
    consoles only.

--cursor on|off
    Turns the terminal's cursor on or off.

--default
    Sets the terminal's rendering options to the default values.

--dump[=console_number]
    Writes a snapshot of the virtual console with the given number to
    the file specified with the --file option, overwriting its
    contents; the default is screen.dump. Without an argument, it dumps
    the current virtual console. This overrides --append.

--file filename
    Sets the snapshot file name for any --dump or --append options on
    the same command line. If this option is not present, the default
    is screen.dump in the current directory. A path name that exceeds
    the system maximum will be truncated, see PATH_MAX from
    linux/limits.h for the value.

--foreground 8-color|default
    Sets the foreground text color.

--half-bright on|off
    Turns dim (half-brightness) mode on or off. Except on a virtual
    console, --half-bright off turns off all attributes (bold,
    half-brightness, blink, reverse).

--hbcolor 16-color
    Sets the color for half-bright characters.
```

```
--initialize
    Displays the terminal initialization string, which typically sets
    the terminal’s rendering options, and other attributes to the
    default values.

--inversescreen on|off
    Swaps foreground and background colors for the whole screen.

--linewrap on|off
    Makes the terminal continue on a new line when a line is full.

--msg on|off
    Enables or disables the sending of kernel printk() messages to the
    console. Virtual consoles only.

--msglevel 0-8
    Sets the console logging level for kernel printk() messages. All
    messages strictly more important than this will be printed, so a
    logging level of 0 has the same effect as --msg on and a logging
    level of 8 will print all kernel messages. klogd(8) may be a more
    convenient interface to the logging of kernel messages.

    Virtual consoles only.

--powerdown[=0-60]
    Sets the VESA powerdown interval in minutes. Without an argument,
    it defaults to 0 (disable powerdown). If the console is blanked or
    the monitor is in suspend mode, then the monitor will go into vsync
    suspend mode or powerdown mode respectively after this period of
    time has elapsed.

--powersave mode
    Valid values for mode are:

    vsync|on
        Puts the monitor into VESA vsync suspend mode.

    hsync
        Puts the monitor into VESA hsync suspend mode.

    powerdown
        Puts the monitor into VESA powerdown mode.

    off
        Turns monitor VESA powersaving features.

--regtabs[=1-160]
    Clears all tab stops, then sets a regular tab stop pattern, with
    one tab every specified number of positions. Without an argument,
    it defaults to 8. Virtual consoles only.

--repeat on|off
```

Turns keyboard repeat on or off. Virtual consoles only.

--reset
Displays the terminal reset string, which typically resets the terminal to its power-on state.

--resize
Reset terminal size by assessing maximum row and column. This is useful when actual geometry and kernel terminal driver are not in sync. Most notable use case is with serial consoles, that do not use ioctl(2) but just byte streams and breaks.

--reverse on|off
Turns reverse video mode on or off. Except on a virtual console, --reverse off turns off all attributes (bold, half-brightness, blink, reverse).

--store
Stores the terminal's current rendering options (foreground and background colors) as the values to be used at reset-to-default. Virtual consoles only.

--tabs[=tab1 tab2 tab3 ...]
Sets tab stops at the given horizontal cursor positions, in the range 1-160. Without arguments, it shows the current tab stop settings.

--term terminal_name
Overrides the TERM environment variable.

--ulcolor 16-color
Sets the color for underlined characters. Virtual consoles only.

--underline on|off
Turns underline mode on or off.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

COMPATIBILITY

Since version 2.25 setterm has support for long options with two hyphens, for example --help, beside the historical long options with a single hyphen, for example -help. In scripts it is better to use the backward-compatible single hyphen rather than the double hyphen. Currently there are no plans nor good reasons to discontinue single-hyphen compatibility.

BUGS

Differences between the Minix and Linux versions are not documented.

```
SEE ALSO
       stty(1), tput(1), tty(4), terminfo(5)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The setterm command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                        SETTERM(1)
```

# Bash 的 sha1sum 命令

```
sha1sum [OPTION]... [FILE]...
```

**功能**

计算或校验 SHA1 值。

**类型**

可执行文件（`/usr/bin/sha1sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 SHA1 值进行校验
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 SHA1 值**

```
$ sha1sum 1.txt                                    # 单个文件
59cb445bf92dffdd01cd01154a521d39189dadb3  1.txt
$ sha1sum 1.txt 2.txt 3.txt                        # 多个文件
59cb445bf92dffdd01cd01154a521d39189dadb3  1.txt
668079210932eb19a1b4995a7d95c71ad2e0a83c  2.txt
22596363b3de40b06f981fb85d82312e8c0ed511  3.txt
$ sha1sum --tag 1.txt 2.txt 3.txt                  # BSD 风格
```

```
SHA1 (1.txt) = 59cb445bf92dffdd01cd01154a521d39189dadb3
SHA1 (2.txt) = 668079210932eb19a1b4995a7d95c71ad2e0a83c
SHA1 (3.txt) = 22596363b3de40b06f981fb85d82312e8c0ed511
```

**校验 SHA1 值**

```
$ sha1sum 1.txt 2.txt 3.txt > sha1.txt
$ sha1sum -c sha1.txt                                   # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

# 相关命令

| 命令 | 说明 |
| --- | --- |
| b2sum | 计算和校验 BLAKE2 值 |
| md5sum | 计算和校验 MD5 值 |
| sha224sum | 计算和校验 SHA224 值 |
| sha256sum | 计算和校验 SHA256 值 |
| sha384sum | 计算和校验 SHA384 值 |
| sha512sum | 计算和校验 SHA512 值 |

# 推荐阅读

- Print or check SHA1 digests - GNU Coreutils
- sha1sum(1) — Linux manual page

# 手册

```
SHA1SUM(1)                      User Commands                      SHA1SUM(1)

NAME
       sha1sum - compute and check SHA1 message digest

SYNOPSIS
       sha1sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check SHA1 (160-bit) checksums.

       With no FILE, or when FILE is -, read standard input.
```

```
       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in FIPS-180-1.  When  checking,  the
       input  should  be a former output of this program.  The default mode is
       to print a line with: checksum, a space, a character  indicating  input
       mode  ('*'  for binary, ' ' for text or where binary is insignificant),
       and name for each FILE.

       Note: There is no difference between binary mode and text mode  on  GNU
       systems.

BUGS
       Do not use the SHA-1 algorithm for security related purposes.  Instead,
       use  an  SHA-2  algorithm,  implemented  in  the programs sha224sum(1),
       sha256sum(1), sha384sum(1), sha512sum(1), or the BLAKE2 algorithm,  im-
       plemented in b2sum(1)

AUTHOR
```

Written by Ulrich Drepper, Scott Miller, and David Madore.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

cksum(1)

Full documentation <https://www.gnu.org/software/coreutils/sha1sum>
or available locally via: info '(coreutils) sha1sum invocation'

# Bash 的 sha224sum 命令

```
sha224sum [OPTION]... [FILE]...
```

**功能**

计算或校验 SHA224 值。

**类型**

可执行文件（`/usr/bin/sha224sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 SHA224 值进行校验
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 SHA224 值**

```
$ sha224sum 1.txt                                       # 单个文件
7295844ef5b3a4f51ffe2567d06a0303c5fc5df240ccd2005ebb1a1d  1.txt
$ sha224sum 1.txt 2.txt 3.txt                           # 多个文件
7295844ef5b3a4f51ffe2567d06a0303c5fc5df240ccd2005ebb1a1d  1.txt
438a03517353d6350a834e2f7b9b296e2a942a911af5087f28301a0c  2.txt
95041dd60ab08c0bf5636d50be85fe9790300f39eb84602858a9b430  3.txt
$ sha224sum --tag 1.txt 2.txt 3.txt                     # BSD 风格
```

```
SHA224 (1.txt) = 7295844ef5b3a4f51ffe2567d06a0303c5fc5df240ccd2005ebb1a1d
SHA224 (2.txt) = 438a03517353d6350a834e2f7b9b296e2a942a911af5087f28301a0c
SHA224 (3.txt) = 95041dd60ab08c0bf5636d50be85fe9790300f39eb84602858a9b430
```

**校验 SHA224 值**

```
$ sha224sum 1.txt 2.txt 3.txt > sha1.txt
$ sha224sum -c sha1.txt                                    # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

# 相关命令

| 命令 | 说明 |
| --- | --- |
| b2sum | 计算和校验 BLAKE2 值 |
| md5sum | 计算和校验 MD5 值 |
| sha1sum | 计算和校验 SHA1 值 |
| sha256sum | 计算和校验 SHA256 值 |
| sha384sum | 计算和校验 SHA384 值 |
| sha512sum | 计算和校验 SHA512 值 |

# 推荐阅读

- Print or check SHA-2 digests - GNU Coreutils
- sha224sum(1) — Linux manual page

# 手册

```
SHA224SUM(1)                    User Commands                    SHA224SUM(1)

NAME
       sha224sum - compute and check SHA224 message digest

SYNOPSIS
       sha224sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check SHA224 (224-bit) checksums.

       With no FILE, or when FILE is -, read standard input.
```

```
       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in RFC 3874.  When checking, the in-
       put should be a former output of this program.  The default mode is  to
       print a line with: checksum, a space, a character indicating input mode
       ('*'  for  binary,  ' ' for text or where binary is insignificant), and
       name for each FILE.

       Note: There is no difference between binary mode and text mode  on  GNU
       systems.

AUTHOR
       Written by Ulrich Drepper, Scott Miller, and David Madore.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>
```

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

cksum(1)

Full documentation <https://www.gnu.org/software/coreutils/sha224sum>
or available locally via: info '(coreutils) sha2 utilities'

# Bash 的 sha256sum 命令

```
sha256sum [OPTION]... [FILE]...
```

**功能**

计算或校验 SHA256 值。

**类型**

可执行文件（`/usr/bin/sha256sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 SHA256 值进行校验
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 SHA256 值**

```
$ sha256sum 1.txt                                       # 单个文件
6fc9ba098aa5065ea38f4c9b4eea877685d6150af612b3c1e042dddf9977acda  1.txt
$ sha256sum 1.txt 2.txt 3.txt                           # 多个文件
6fc9ba098aa5065ea38f4c9b4eea877685d6150af612b3c1e042dddf9977acda  1.txt
952be6c3e9b1f3af5d8daf7698bc9b5ba20323738990713839f6f7c236f8fe12  2.txt
a948904f2f0f479b8f8197694b30184b0d2ed1c1cd2a1ec0fb85d299a192a447  3.txt
$ sha256sum --tag 1.txt 2.txt 3.txt                     # BSD 风格
```

```
SHA256 (1.txt) = 6fc9ba098aa5065ea38f4c9b4eea877685d6150af612b3c1e042dddf9977acda
SHA256 (2.txt) = 952be6c3e9b1f3af5d8daf7698bc9b5ba203237389907138 39f6f7c236f8fe12
SHA256 (3.txt) = a948904f2f0f479b8f8197694b30184b0d2ed1c1cd2a1ec0fb85d299a192a447
```

**校验 SHA256 值**

```
$ sha256sum 1.txt 2.txt 3.txt > sha1.txt
$ sha256sum -c sha1.txt                                    # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| b2sum | 计算和校验 BLAKE2 值 |
| md5sum | 计算和校验 MD5 值 |
| sha1sum | 计算和校验 SHA1 值 |
| sha224sum | 计算和校验 SHA224 值 |
| sha384sum | 计算和校验 SHA384 值 |
| sha512sum | 计算和校验 SHA512 值 |

## 推荐阅读

- Print or check SHA-2 digests - GNU Coreutils
- sha256sum(1) — Linux manual page

## 手册

```
SHA256SUM(1)                    User Commands                    SHA256SUM(1)

NAME
       sha256sum - compute and check SHA256 message digest

SYNOPSIS
       sha256sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check SHA256 (256-bit) checksums.

       With no FILE, or when FILE is -, read standard input.
```

```
       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in FIPS-180-2.  When  checking,  the
       input  should  be a former output of this program.  The default mode is
       to print a line with: checksum, a space, a character  indicating  input
       mode  ('*'  for binary, ' ' for text or where binary is insignificant),
       and name for each FILE.

       Note: There is no difference between binary mode and text mode  on  GNU
       systems.

AUTHOR
       Written by Ulrich Drepper, Scott Miller, and David Madore.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>
```

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

cksum(1)

Full documentation <https://www.gnu.org/software/coreutils/sha256sum>
or available locally via: info '(coreutils) sha2 utilities'

# Bash 的 sha384sum 命令

```
sha384sum [OPTION]... [FILE]...
```

**功能**

计算或校验 SHA384 值。

**类型**

可执行文件（`/usr/bin/sha384sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 SHA384 值进行校验
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 SHA384 值**

```
$ sha384sum 1.txt                                       # 单个文件
cfbd4c8d77a8e9158596a1288f55d56876235983e9a1f6d2bfaece85469a953d30abdfd17a359b706ade374d61
361f6b  1.txt
$ sha384sum 1.txt 2.txt 3.txt                           # 多个文件
cfbd4c8d77a8e9158596a1288f55d56876235983e9a1f6d2bfaece85469a953d30abdfd17a359b706ade374d61
361f6b  1.txt
0a18f1948f0c69fa788fa8367fce4df5ef8b7f8fa82581a32722792963e4ab0353f21560359330b4756862a8c4
```

```
ee8dc4  2.txt
6b3b69ff0a404f28d75e98a066d3fc64fffd9940870cc68bece28545b9a75086b343d7a1366838083e4b8f3ca6
fd3c80  3.txt
$ sha384sum --tag 1.txt 2.txt 3.txt                    # BSD 风格
SHA384 (1.txt) =
cfbd4c8d77a8e9158596a1288f55d56876235983e9a1f6d2bfaece85469a953d30abdfd17a359b706ade374d61
361f6b
SHA384 (2.txt) =
0a18f1948f0c69fa788fa8367fce4df5ef8b7f8fa82581a32722792963e4ab0353f21560359330b4756862a8c4
ee8dc4
SHA384 (3.txt) =
6b3b69ff0a404f28d75e98a066d3fc64fffd9940870cc68bece28545b9a75086b343d7a1366838083e4b8f3ca6
fd3c80
```

**校验 SHA384 值**

```
$ sha384sum 1.txt 2.txt 3.txt > sha1.txt
$ sha384sum -c sha1.txt                                # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| b2sum | 计算和校验 BLAKE2 值 |
| md5sum | 计算和校验 MD5 值 |
| sha1sum | 计算和校验 SHA1 值 |
| sha224sum | 计算和校验 SHA224 值 |
| sha256sum | 计算和校验 SHA256 值 |
| sha512sum | 计算和校验 SHA512 值 |

## 推荐阅读

- Print or check SHA-2 digests - GNU Coreutils
- sha384sum(1) — Linux manual page

## 手册

```
SHA384SUM(1)                    User Commands                    SHA384SUM(1)
```

```
NAME
       sha384sum - compute and check SHA384 message digest

SYNOPSIS
       sha384sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check SHA384 (384-bit) checksums.

       With no FILE, or when FILE is -, read standard input.

       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in FIPS-180-2.  When  checking,  the
       input  should  be a former output of this program.  The default mode is
       to print a line with: checksum, a space, a character  indicating  input
       mode  ('*'  for binary, ' ' for text or where binary is insignificant),
       and name for each FILE.
```

Note: There is no difference between binary mode and text mode on GNU systems.

AUTHOR

Written by Ulrich Drepper, Scott Miller, and David Madore.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

cksum(1)

Full documentation <https://www.gnu.org/software/coreutils/sha384sum>
or available locally via: info '(coreutils) sha2 utilities'

# Bash 的 sha512sum 命令

```
sha512sum [OPTION]... [FILE]...
```

**功能**

计算或校验 SHA512 值。

**类型**

可执行文件（`/usr/bin/sha512sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--binary` - 以二进制模式读取文件；类 UNIX 系统下始终是二进制模式
  - `-c`，`--check` - 从文件中读取 SHA512 值进行校验
  - `--tag` - 生成 BSD 风格的输出
  - `-t`，`--text` - 以文本模式读取文件；类 UNIX 系统下不存在文本模式，此选项无效
  - `-z`，`--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--ignore-missing` - 校验时忽略缺失的文件
  - `--quiet` - 校验时不打印成功的消息
  - `--status` - 校验时不打印任何消息；可以通过返回值检查是否成功
  - `--strict` - 校验时严格检查格式
  - `-w`，`--warn` - 校验时遇到格式不正确的行时产生警告
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**计算 SHA512 值**

```
$ sha512sum 1.txt                                          # 单个文件
65ded19fe8945cba7cd781c9765ea47ef12fc381cfb94dc4114d2608d5e74ececcfb992e601090fe6f2e68e782
2b44fcab9ed4d5d33477f0389f067dabcfb6a2  1.txt
$ sha512sum 1.txt 2.txt 3.txt                              # 多个文件
65ded19fe8945cba7cd781c9765ea47ef12fc381cfb94dc4114d2608d5e74ececcfb992e601090fe6f2e68e782
2b44fcab9ed4d5d33477f0389f067dabcfb6a2  1.txt
c9712f9dc674b4c5541dc744277957e4ec2349bc5a997b082ca379efb7dbdb0622a0c9afd033056f8f1706463c
```

```
7fff229a79692e38850fd3e17d3e9d087b1849  2.txt
db3974a97f2407b7cae1ae637c0030687a11913274d578492558e39c16c017de84eacdc8c62fe34ee4e12b4b14
28817f09b6a2760c3f8a664ceae94d2434a593  3.txt
$ sha512sum --tag 1.txt 2.txt 3.txt                    # BSD 风格
SHA512 (1.txt) =
65ded19fe8945cba7cd781c9765ea47ef12fc381cfb94dc4114d2608d5e74ececcfb992e601090fe6f2e68e782
2b44fcab9ed4d5d33477f0389f067dabcfb6a2
SHA512 (2.txt) =
c9712f9dc674b4c5541dc744277957e4ec2349bc5a997b082ca379efb7dbdb0622a0c9afd033056f8f1706463c
7fff229a79692e38850fd3e17d3e9d087b1849
SHA512 (3.txt) =
db3974a97f2407b7cae1ae637c0030687a11913274d578492558e39c16c017de84eacdc8c62fe34ee4e12b4b14
28817f09b6a2760c3f8a664ceae94d2434a593
```

**校验 SHA512 值**

```
$ sha512sum 1.txt 2.txt 3.txt > sha1.txt
$ sha512sum -c sha1.txt                                # 校验
1.txt: OK
2.txt: OK
3.txt: OK
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| b2sum | 计算和校验 BLAKE2 值 |
| md5sum | 计算和校验 MD5 值 |
| sha1sum | 计算和校验 SHA1 值 |
| sha224sum | 计算和校验 SHA224 值 |
| sha256sum | 计算和校验 SHA256 值 |
| sha512sum | 计算和校验 SHA512 值 |

## 推荐阅读

- Print or check SHA-2 digests - GNU Coreutils
- sha512sum(1) — Linux manual page

## 手册

```
SHA512SUM(1)                    User Commands                    SHA512SUM(1)
```

```
NAME
       sha512sum - compute and check SHA512 message digest

SYNOPSIS
       sha512sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check SHA512 (512-bit) checksums.

       With no FILE, or when FILE is -, read standard input.

       -b, --binary
              read in binary mode

       -c, --check
              read checksums from the FILEs and check them

       --tag  create a BSD-style checksum

       -t, --text
              read in text mode (default)

       -z, --zero
              end  each  output  line  with NUL, not newline, and disable file
              name escaping

   The following five options are useful only when verifying checksums:
       --ignore-missing
              don't fail or report status for missing files

       --quiet
              don't print OK for each successfully verified file

       --status
              don't output anything, status code shows success

       --strict
              exit non-zero for improperly formatted checksum lines

       -w, --warn
              warn about improperly formatted checksum lines

       --help display this help and exit

       --version
              output version information and exit

       The sums are computed as described in FIPS-180-2.  When  checking,  the
       input  should  be a former output of this program.  The default mode is
       to print a line with: checksum, a space, a character  indicating  input
       mode  ('*'  for binary, ' ' for text or where binary is insignificant),
       and name for each FILE.
```

```
       Note: There is no difference between binary mode and text mode  on  GNU
       systems.

AUTHOR
       Written by Ulrich Drepper, Scott Miller, and David Madore.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       cksum(1)

       Full documentation <https://www.gnu.org/software/coreutils/sha512sum>
       or available locally via: info '(coreutils) sha2 utilities'

GNU coreutils 9.4                April 2024                      SHA512SUM(1)
```

# Bash 的 shred 命令

```
shred [OPTION]... FILE...
```

**功能**

粉碎文件数据（使用随机数据多次覆盖磁盘数据）以避免被恢复。

**类型**

可执行文件（`/usr/bin/shred`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-f`，`--force` - 强制执行，如有必要，会修改文件的权限
  - `-n`，`--iterations=N` - 覆盖 `N` 次数据；默认为 3
  - `--random-source=FILE` - 从 `FILE` 文件获取随机数据
  - `-s`，`--size=N` - 粉碎 `N` 字节数据；支持 `K` `M` `G` 等单位后缀
  - `-u` - 粉碎后删除文件
  - `--remove[=HOW]` - 粉碎后删除文件，并指定怎样删除目录
    - `unlink` - 使用标准的 `unlink` 调用
    - `wipe` - 并且首先对名称中的字节进行混淆
    - `wipesync` - 并且将混淆后的字节同步到设备（默认）
  - `-v`，`--verbose` - 显示进度
  - `-x`，`--exact` - 不要将文件大小向上取整到一个标准块；对于非普通文件默认开启
  - `-z`，`--zero` - 最后额外使用全零覆盖一次
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

```
$ shred 1.txt               # 覆盖 1.txt 的数据 3 次（默认）
$ shred -n 10 -z 1.txt      # 覆盖 1.txt 的数据 10 次，然后用全零额外覆盖一次
$ shred -u 1.txt            # 覆盖 1.txt 的数据，然后将其删除
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| rm | 删除文件 |
| rmdir | 删除目录 |
| unlink | 删除文件 |

## 推荐阅读

- Remove files more securely - GNU Coreutils
- shred(1) — Linux manual page

## 手册

```
SHRED(1)                        User Commands                        SHRED(1)

NAME
       shred - overwrite a file to hide its contents, and optionally delete it

SYNOPSIS
       shred [OPTION]... FILE...

DESCRIPTION
       Overwrite  the specified FILE(s) repeatedly, in order to make it harder
       for even very expensive hardware probing to recover the data.

       If FILE is -, shred standard output.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -f, --force
              change permissions to allow writing if necessary

       -n, --iterations=N
              overwrite N times instead of the default (3)

       --random-source=FILE
              get random bytes from FILE

       -s, --size=N
              shred this many bytes (suffixes like K, M, G accepted)

       -u     deallocate and remove file after overwriting
```

--remove[=HOW]
       like -u but give control on HOW to delete;  See below

-v, --verbose
       show progress

-x, --exact
       do not round file sizes up to the next full block;

       this is the default for non-regular files

-z, --zero
       add a final overwrite with zeros to hide shredding

--help display this help and exit

--version
       output version information and exit

Delete FILE(s) if --remove (-u) is specified. The default is not to remove the files because it is common to operate on device files like /dev/hda, and those files usually should not be removed. The optional HOW parameter indicates how to remove a directory entry: 'unlink' => use a standard unlink call. 'wipe' => also first obfuscate bytes in the name. 'wipesync' => also sync each obfuscated byte to the device. The default mode is 'wipesync', but note it can be expensive.

CAUTION: shred assumes the file system and hardware overwrite data in place. Although this is common, many platforms operate otherwise. Also, backups and mirrors may contain unremovable copies that will let a shredded file be recovered later. See the GNU coreutils manual for details.

AUTHOR
       Written by Colin Plumb.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it. There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/shred>
       or available locally via: info '(coreutils) shred invocation'

# Bash 的 shuf 命令

```
shuf [OPTION]... [FILE]
```

**功能**

生成随机的排列组合。

**类型**

可执行文件（`/usr/bin/shuf`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-e`，`--echo` - 将每个参数视为一行输入
  - `-i`，`--input-range=LO-HI` - 将数值 `LO` 到 `HI` 随机排列
  - `-n`，`--head-count=COUNT` - 输出保留前 `COUNT` 行
  - `-o`，`--output=FILE` - 将结果输出到 `FILE` 文件；默认为标准输出
  - `--random-source=FILE` - 从 `FILE` 文件获取随机数
  - `-r`，`--repeat` - 输出可以重复
  - `-z`，`--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

**生成随机数列表**

```
$ shuf -i 1-5           # 将 1-5 随机排列
2
5
4
1
3
```

```
$ shuf -i 1-5 -r -n 10  # 生成 10 个 1-5 的随机数
3
5
```

```
2
1
5
4
1
5
3
2
```

**将文件的行随机排列**

```
$ echo -e "A\nB\nC\nD\nE" > 1.txt    # 创建文件
$ shuf 1.txt                         # 将文件的行随机打乱
D
A
C
E
B
```

```
$ echo -e "A\nB\nC\nD\nE" > 1.txt    # 创建文件
$ shuf -r -n 10 1.txt                # 将文件的行随机输出 10 次
D
A
C
E
B
```

# 推荐阅读

- Shuffling text - GNU Coreutils
- shuf(1) — Linux manual page

# 手册

```
SHUF(1)                          User Commands                          SHUF(1)

NAME
       shuf - generate random permutations

SYNOPSIS
       shuf [OPTION]... [FILE]
       shuf -e [OPTION]... [ARG]...
       shuf -i LO-HI [OPTION]...

DESCRIPTION
       Write a random permutation of the input lines to standard output.
```

With no FILE, or when FILE is -, read standard input.

Mandatory arguments to long options are mandatory for short options too.

-e, --echo
treat each ARG as an input line

-i, --input-range=LO-HI
treat each number LO through HI as an input line

-n, --head-count=COUNT
output at most COUNT lines

-o, --output=FILE
write result to FILE instead of standard output

--random-source=FILE
get random bytes from FILE

-r, --repeat
output lines can be repeated

-z, --zero-terminated
line delimiter is NUL, not newline

--help display this help and exit

--version
output version information and exit

AUTHOR

Written by Paul Eggert.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/shuf>
or available locally via: info '(coreutils) shuf invocation'

# Bash 的 sleep 命令

```
sleep [OPTION]... NUMBER[SUFFIX]...
```

**功能**

延时一段时间。

**类型**

可执行文件（`/usr/bin/sleep`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `NUMBER` - 要延时的时间
- `SUFFIX` - 单位后缀：
  - `s` - 秒（默认）
  - `m` - 分钟
  - `h` - 小时
  - `d` - 天

## 示例

```
$ sleep 1       # 延时 1 秒
$ sleep 1 2 3   # 延时 6 秒
$ sleep 3m      # 延时 3 分钟
$ sleep 3m 20   # 延时 3 分 20 秒
```

## 推荐阅读

- Delay for a specified time - GNU Coreutils
- sleep(1) — Linux manual page

## 手册

SLEEP(1) User Commands SLEEP(1)

NAME
sleep - delay for a specified amount of time

SYNOPSIS
sleep NUMBER[SUFFIX]...
sleep OPTION

DESCRIPTION
Pause for NUMBER seconds. SUFFIX may be 's' for seconds (the default), 'm' for minutes, 'h' for hours or 'd' for days. NUMBER need not be an integer. Given two or more arguments, pause for the amount of time specified by the sum of their values.

--help display this help and exit

--version
output version information and exit

AUTHOR
Written by Jim Meyering and Paul Eggert.

REPORTING BUGS
GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
sleep(3)

Full documentation <https://www.gnu.org/software/coreutils/sleep>
or available locally via: info '(coreutils) sleep invocation'

# Bash 的 sort 命令

```
sort [OPTION]... [FILE]...
```

**功能**

将文本文件的行进行排序。

**类型**

可执行文件（`/usr/bin/sort`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--ignore-leading-blanks` - 忽略开头的空白
  - `-d`，`--dictionary-order` - 使用字典序，忽略字母、数字、空格以外的字符（默认）
  - `-f`，`--ignore-case` - 忽略大小写，小写字母当作大写字母
  - `-g`，`--general-numeric-sort` - 比较广义数值大小（`'NaN'` < `'-inf'` < `'2'` < `'3.14'` < `'10'` < `'1e4'` < `'inf'`）
  - `-i`，`--ignore-nonprinting` - 忽略不可打印的字符（参考 isprint）
  - `-M`，`--month-sort` - 比较月份大小（其它 < `'JAN'` < `'FEB'` < ... < `'DEC'`）
  - `-h`，`--human-numeric-sort` - 比较含单位的数值大小（`2K` < `1G`）
  - `-n`，`--numeric-sort` - 比较数值大小（`'2'` < `'3.14'` < `'10'`），非数值视为 0
  - `-R`，`--random-sort` - 随机排序（但是相同的行会在一起）
  - `--random-source=FILE` - 通过 `FILE` 获取随机数
  - `-r`，`--reverse` - 逆序输出；默认为升序，此选项为降序
  - `--sort=WORD` - 按 `WORD` 排序：
    - `general-numeric` - 同 `-g`，比较广义数值大小（`'NaN'` < `'-inf'` < `'2'` < `'3.14'` < `'10'` < `'1e4'` < `'inf'`）
    - `human-numeric` - 同 `-h`，比较含单位的数值大小（`2K` < `1G`）
    - `month` - 同 `-M`，比较月份大小（其它 < `'JAN'` < `'FEB'` < ... < `'DEC'`）
    - `numeric` - 同 `-n`，比较数值大小（`'2'` < `'3.14'` < `'10'`），非数值视为 0
    - `random` - 同 `-R`，随机排序（但是相同的行会在一起）
    - `version` - 同 `-V`，比较版本号大小（`'v3.2'` < `'v10.1'`）
  - `-V`，`--version-sort` - 比较版本号大小（`'v3.2'` < `'v10.1'`）
  - `--batch-size=NMERGE` - 一个合并最多 `NMERGE` 个输入
  - `-c`，`--check`，`--check=diagnose-first` - 检查输入是否已经排序，而不是进行排序

  - `-C`, `--check=quiet`, `--check=silent` - 检查输入是否已经排序，而不是进行排序，并且不报告错序行
  - `--compress-program=PROG` - 使用 `PROG` 命令临时进行压缩，之后使用 `PROG -d` 命令解压
  - `--debug` - 打印调试信息；对用于排序的行进行注释并将可疑用法输出到标准错误
  - `--files0-from=F` - 从文件 `F` 中读取文件列表作为输入，每个文件名以空字符（`\0`）结尾；如果 `F` 是 `-` 则从标准输入读取文件列表
  - `-k`, `--key=KEYDEF` - 根据键进行排序
  - `-m`, `--merge` - 合并已经排序的文件，而不是进行排序
  - `-o`, `--output=FILE` - 将结果输出到 `FILE` 文件；默认为标准输出
  - `-s`, `--stable` - 稳定排序
  - `-S`, `--buffer-size=SIZE` - 使用 `SIZE` 主内存缓冲
  - `-t`, `--field-separator=SEP` - 使用 `SEP` 作为分隔，而不是非空白到空白的过渡作为分隔
  - `-T`, `--temporary-directory=DIR` - 使用 `DIR` 作为临时目录，而不是 `$TMPDIR` 或 `/tmp`
  - `--parallel=N` - 并发运行数量设为 `N`
  - `-u`, `--unique` - 排序时相同的行只输出一次；检查排序（与 `-c` 一同使用）时严格检查顺序 （不能有重复行）
  - `-z`, `--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

示例文件内容：

```
apple
Apple
!banana
10
2
50K
3M
Jan
Feb
v1
v2
v10
3.5
1e2
NaN
```

**默认排序（字典序）**

```
$ sort 1.txt
10
1e2
2
3.5
3M
50K
apple
Apple
!banana
Feb
Jan
NaN
v1
v10
v2
```

### 按数值排序

```
$ sort -n 1.txt
apple
Apple
!banana
Feb
Jan
NaN
v1
v10
v2
1e2
2
3M
3.5
10
50K
```

### 按广义数值排序

```
$ sort -g 1.txt
apple
Apple
!banana
Feb
Jan
v1
v10
v2
NaN
2
3M
```

```
3.5
10
50K
1e2
```

**含单位的数值排序**

```
$ sort  -h 1.txt
apple
Apple
!banana
Feb
Jan
NaN
v1
v10
v2
1e2
2
3.5
10
50K
3M
```

**版本号排序**

```
$ sort -V 1.txt
1e2
2
3M
3.5
10
50K
Apple
Feb
Jan
NaN
apple
v1
v2
v10
!banana
```

**月份排序**

```
$ sort -M 1.txt
10
1e2
2
3.5
```

```
3M
50K
apple
Apple
!banana
NaN
v1
v10
v2
Jan
Feb
```

## 推荐阅读

- Sort text files - GNU Coreutils
- sort(1) — Linux manual page

## 手册

```
SORT(1)                          User Commands                          SORT(1)

NAME
       sort - sort lines of text files

SYNOPSIS
       sort [OPTION]... [FILE]...
       sort [OPTION]... --files0-from=F

DESCRIPTION
       Write sorted concatenation of all FILE(s) to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.  Ordering options:

       -b, --ignore-leading-blanks
              ignore leading blanks

       -d, --dictionary-order
              consider only blanks and alphanumeric characters

       -f, --ignore-case
              fold lower case to upper case characters

       -g, --general-numeric-sort
              compare according to general numerical value
```

```
-i, --ignore-nonprinting
       consider only printable characters

-M, --month-sort
       compare (unknown) < 'JAN' < ... < 'DEC'

-h, --human-numeric-sort
       compare human readable numbers (e.g., 2K 1G)

-n, --numeric-sort
       compare according to string numerical value

-R, --random-sort
       shuffle, but group identical keys.  See shuf(1)

--random-source=FILE
       get random bytes from FILE

-r, --reverse
       reverse the result of comparisons

--sort=WORD
       sort according to WORD: general-numeric  -g,  human-numeric  -h,
       month -M, numeric -n, random -R, version -V

-V, --version-sort
       natural sort of (version) numbers within text

Other options:

--batch-size=NMERGE
       merge at most NMERGE inputs at once; for more use temp files

-c, --check, --check=diagnose-first
       check for sorted input; do not sort

-C, --check=quiet, --check=silent
       like -c, but do not report first bad line

--compress-program=PROG
       compress temporaries with PROG; decompress them with PROG -d

--debug
       annotate the part of the line used to sort, and warn about ques-
       tionable usage to stderr

--files0-from=F
       read  input  from the files specified by NUL-terminated names in
       file F; If F is - then read names from standard input

-k, --key=KEYDEF
       sort via a key; KEYDEF gives location and type
```

-m, --merge
merge already sorted files; do not sort

-o, --output=FILE
write result to FILE instead of standard output

-s, --stable
stabilize sort by disabling last-resort comparison

-S, --buffer-size=SIZE
use SIZE for main memory buffer

-t, --field-separator=SEP
use SEP instead of non-blank to blank transition

-T, --temporary-directory=DIR
use DIR for temporaries, not $TMPDIR or /tmp; multiple options specify multiple directories

--parallel=N
change the number of sorts run concurrently to N

-u, --unique
with -c, check for strict ordering; without -c, output only the first of an equal run

-z, --zero-terminated
line delimiter is NUL, not newline

--help display this help and exit

--version
output version information and exit

KEYDEF is F[.C][OPTS][,F[.C][OPTS]] for start and stop position, where F is a field number and C a character position in the field; both are origin 1, and the stop position defaults to the line's end. If neither -t nor -b is in effect, characters in a field are counted from the beginning of the preceding whitespace. OPTS is one or more single-letter ordering options [bdfgiMhnRrV], which override global ordering options for that key. If no key is given, use the entire line as the key. Use --debug to diagnose incorrect key usage.

SIZE may be followed by the following multiplicative suffixes: % 1% of memory, b 1, K 1024 (default), and so on for M, G, T, P, E, Z, Y, R, Q.

*** WARNING *** The locale specified by the environment affects sort order. Set LC_ALL=C to get the traditional sort order that uses native byte values.

AUTHOR
Written by Mike Haertel and Paul Eggert.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

shuf(1), uniq(1)

Full documentation <https://www.gnu.org/software/coreutils/sort>
or available locally via: info '(coreutils) sort invocation'

# Bash 的 split 命令

```
split [OPTION]... [FILE [PREFIX]]
```

**功能**

将文件拆分为多个文件。

**类型**

可执行文件（`/usr/bin/split`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--suffix-length=N` - 输出文件名后缀长度设为 `N`；默认为 2
  - `--additional-suffix=SUFFIX` - 输出文件名额外添加后缀 `SUFFIX`
  - `-b`，`--bytes=SIZE` - 每个输出文件的大小为 `SIZE` 字节
  - `-C`，`--line-bytes=SIZE` - 每个输出文件的大小不超过 `SIZE` 字节，保持行的完整性
  - `-d` - 输出文件名后缀使用从 0 开始的数值；默认为 `aa` `ab` 格式
  - `--numeric-suffixes[=FROM]` - 输出文件名后缀使用从 `FROM` 开始的数值
  - `-x` - 输出文件名后缀使用从 0 开始的十六进行数值
  - `--hex-suffixes[=FROM]` - 输出文件名后缀使用从 `FROM` 开始的十六进行数值
  - `-e`，`--elide-empty-files` - 使用 `-n` 选项时，不生成空文件
  - `--filter=COMMAND` - 将输出写入 `COMMAND` 命令处理，并设置 `$FILE` 为文件名
  - `-l`，`--lines=NUMBER` - 每个输出文件有 `NUMBER` 行；默认为 1000 行
  - `-n`，`--number=CHUNKS` - 基于 `CHUNK` 生成文件：
    - `N` - 平均拆分成 `N` 个文件
    - `K/N` - 平均拆分成 `N` 份，将第 `K` 份打印到标准输出
    - `l/N` - 平均拆分成 `N` 个文件，保持行的完整性
    - `l/K/N` - 平均拆分成 `N` 份，将第 `K` 份打印到标准输出，保持行的完整性
    - `r/N` - 类似 `l/N` 但是使用循环分配
    - `r/K/N` - 类似 `l/K/N` 但是使用循环分配
  - `-t`，`--separator=SEP` - 使用 `SEP` 而不是换行符（`\n`）作为行的分隔符
  - `-u`，`--unbuffered` - 输出时不使用缓冲
  - `--verbose` - 在打开每个输出文件前打印诊断信息
  - `--help` - 显示帮助
  - `--version` - 显示版本

- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入
- `PREFIX` - 输出文件名前缀；默认为 `x`

## 示例

### 字母后缀

```
$ split 1.txt                  # 将文件 1.txt 拆分，每个文件 1000 行
$ ls
1.txt xaa
$ split -n 3 1.txt             # 将文件 1.txt 拆分为 3 个文件
$ ls
1.txt xaa  xab  xac
$ split -n 3 1.txt my_file_    # 输出文件名前缀为 my_file_
$ ls
1.txt xaa  xab  xac  my_file_aa  my_file_ab  my_file_ac
```

### 数值后缀

```
$ split -n 3 -d 1.txt my_file_            # 前缀为 my_file_，使用从 0 开始的数值后缀
$ ls
1.txt my_file_00  my_file_01  my_file_02
$ split -n 3 --numeric-suffixes=10 1.txt  # 前缀为 my_file_，使用从 10 开始的数值后缀
$ ls
1.txt my_file_00  my_file_01  my_file_02  my_file_10  my_file_11  my_file_12
```

## 推荐阅读

- Split a file into pieces - GNU Coreutils
- split(1) — Linux manual page

## 手册

```
SPLIT(1)                        User Commands                        SPLIT(1)

NAME
       split - split a file into pieces

SYNOPSIS
       split [OPTION]... [FILE [PREFIX]]

DESCRIPTION
       Output  pieces of FILE to PREFIXaa, PREFIXab, ...; default size is 1000
       lines, and default PREFIX is 'x'.
```

```
With no FILE, or when FILE is -, read standard input.

Mandatory arguments to long options are  mandatory  for  short  options
too.

-a, --suffix-length=N
       generate suffixes of length N (default 2)

--additional-suffix=SUFFIX
       append an additional SUFFIX to file names

-b, --bytes=SIZE
       put SIZE bytes per output file

-C, --line-bytes=SIZE
       put at most SIZE bytes of records per output file

-d     use numeric suffixes starting at 0, not alphabetic

--numeric-suffixes[=FROM]
       same as -d, but allow setting the start value

-x     use hex suffixes starting at 0, not alphabetic

--hex-suffixes[=FROM]
       same as -x, but allow setting the start value

-e, --elide-empty-files
       do not generate empty output files with '-n'

--filter=COMMAND
       write to shell COMMAND; file name is $FILE

-l, --lines=NUMBER
       put NUMBER lines/records per output file

-n, --number=CHUNKS
       generate CHUNKS output files; see explanation below

-t, --separator=SEP
       use  SEP instead of newline as the record separator; '\0' (zero)
       specifies the NUL character

-u, --unbuffered
       immediately copy input to output with '-n r/...'

--verbose
       print a diagnostic just before each output file is opened

--help display this help and exit

--version
       output version information and exit
```

The SIZE argument is an integer and optional unit (example: 10K is 10*1024). Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or KB,MB,... (powers of 1000). Binary prefixes can be used, too: KiB=K, MiB=M, and so on.

CHUNKS may be:

N split into N files based on size of input

K/N output Kth of N to stdout

l/N split into N files without splitting lines/records

l/K/N output Kth of N to stdout without splitting lines/records

r/N like 'l' but use round robin distribution

r/K/N likewise but only output Kth of N to stdout

## AUTHOR

Written by Torbjorn Granlund and Richard M. Stallman.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/split>
or available locally via: info '(coreutils) split invocation'

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 stat 命令

```
stat [OPTION]... FILE...
```

**功能**

查看文件的状态。

**类型**

可执行文件（`/usr/bin/sleep`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-L`，`--dereference` - 查看符号链接的源文件，而非符号链接自身
  - `-f`，`--file-system` - 查看文件系统的状态
  - `--cached=MODE` - 指定怎样使用缓存属性（常用于远程文件系统）
  - `-c`，`--format=FORMAT` - 使用指定的格式，并且每个格式后输出一个换行符
  - `--printf=FORMAT` - 类似 `--format=FORMAT` 且会解析转移字符，但不会自动换行
  - `-t`，`--terse` - 使用简明的格式
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 格式说明

**用于文件的有效格式**

| 格式符 | 含义 |
| --- | --- |
| `%a` | 权限位的八进制表示（注意 `#` 和 `0` printf 标志） |
| `%A` | 权限位和文件类型的可读形式 |
| `%b` | 已分配的块数（参见 `%B`） |
| `%B` | `%b` 所报告的每个块的字节数 |
| `%C` | SELinux 安全上下文字符串 |
| `%d` | 设备号（十进制，`st_dev`） |

| 格式符 | 含义 |
| --- | --- |
| `%D` | 设备号（十六进制，`st_dev`） |
| `%Hd` | 主要设备号（十进制） |
| `%Ld` | 次要设备号（十进制） |
| `%f` | 原始模式（十六进制） |
| `%F` | 文件类型 |
| `%g` | 所有者的组 ID |
| `%G` | 所有者的组名 |
| `%h` | 硬链接数 |
| `%i` | inode 号 |
| `%m` | 挂载点 |
| `%n` | 文件名 |
| `%N` | 文件名（带引号，如果是符号链接则解除引用） |
| `%o` | 最优 I/O 传输大小提示 |
| `%s` | 总大小（字节） |
| `%r` | 设备类型（十进制，`st_rdev`） |
| `%R` | 设备类型（十六进制，`st_rdev`） |
| `%Hr` | 字符/块设备特殊文件的主要设备类型（十进制） |
| `%Lr` | 字符/块设备特殊文件的次要设备类型（十进制） |
| `%t` | 字符/块设备特殊文件的主要设备类型（十六进制） |
| `%T` | 字符/块设备特殊文件的次要设备类型（十六进制） |
| `%u` | 所有者的用户 ID |
| `%U` | 所有者的用户名 |
| `%w` | 文件创建时间（可读）；未知则为 `-` |
| `%W` | 文件创建时间，自 Epoch 起的秒数；未知则为 `0` |
| `%x` | 最后访问时间（可读） |
| `%X` | 最后访问时间，自 Epoch 起的秒数 |
| `%y` | 最后数据修改时间（可读） |

| 格式符 | 含义 |
| --- | --- |
| `%Y` | 最后数据修改时间，自 Epoch 起的秒数 |
| `%z` | 最后状态改变时间（可读） |
| `%Z` | 最后状态改变时间，自 Epoch 起的秒数 |

**用于文件系统的有效格式**

| 格式符 | 含义 |
| --- | --- |
| `%a` | 非超级用户可用的空闲块数 |
| `%b` | 文件系统中的数据块总数 |
| `%c` | 文件系统中的文件节点总数 |
| `%d` | 文件系统中空闲的文件节点数 |
| `%f` | 文件系统中的空闲块数 |
| `%i` | 文件系统 ID（十六进制） |
| `%l` | 文件名的最大长度 |
| `%n` | 文件名 |
| `%s` | 块大小（用于加快传输速度） |
| `%S` | 基本块大小（用于块计数） |
| `%t` | 文件系统类型（十六进制） |
| `%T` | 文件系统类型（可读形式） |

## 示例

```
$ stat 1.txt        # 默认格式
  File: 1.txt
  Size: 62              Blocks: 8          IO Block: 4096   regular file
Device: 259,2   Inode: 60559096    Links: 1
Access: (0664/-rw-rw-r--)  Uid: ( 1000/   planc)   Gid: ( 1000/   planc)
Access: 2025-12-04 14:08:58.127960041 +0000
Modify: 2025-12-04 14:08:57.155939186 +0000
Change: 2025-12-04 14:08:57.160939293 +0000
 Birth: 2025-12-04 14:08:57.155939186 +0000
$ stat -t 1.txt     # 简明格式
1.txt 62 8 81b4 1000 1000 10302 60559096 1 0 0 1764857338 1764857337 1764857337 1764857337
4096
```

# 推荐阅读

- Report file or file system status - GNU Coreutils
- stat(1) — Linux manual page

# 手册

```
STAT(1)                          User Commands                          STAT(1)

NAME
       stat - display file or file system status

SYNOPSIS
       stat [OPTION]... FILE...

DESCRIPTION
       Display file or file system status.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -L, --dereference
              follow links

       -f, --file-system
              display file system status instead of file status

       --cached=MODE
              specify how to use cached attributes; useful on remote file sys-
              tems. See MODE below

       -c  --format=FORMAT
              use the specified FORMAT instead of the default; output  a  new-
              line after each use of FORMAT

       --printf=FORMAT
              like  --format, but interpret backslash escapes, and do not out-
              put a mandatory trailing newline; if you want a newline, include
              \n in FORMAT

       -t, --terse
              print the information in terse form

       --help display this help and exit

       --version
              output version information and exit

       The MODE argument of --cached can be: always, never, or default.   'al-
```

```
ways'  will  use cached attributes if available, while 'never' will try
to synchronize with the latest attributes, and 'default' will leave  it
up to the underlying file system.

The valid format sequences for files (without --file-system):

%a     permission bits in octal (note '#' and '0' printf flags)

%A     permission bits and file type in human readable form

%b     number of blocks allocated (see %B)

%B     the size in bytes of each block reported by %b

%C     SELinux security context string

%d     device number in decimal (st_dev)

%D     device number in hex (st_dev)

%Hd    major device number in decimal

%Ld    minor device number in decimal

%f     raw mode in hex

%F     file type

%g     group ID of owner

%G     group name of owner

%h     number of hard links

%i     inode number

%m     mount point

%n     file name

%N     quoted file name with dereference if symbolic link

%o     optimal I/O transfer size hint

%s     total size, in bytes

%r     device type in decimal (st_rdev)

%R     device type in hex (st_rdev)

%Hr    major device type in decimal, for character/block device special
       files
```

```
       %Lr    minor device type in decimal, for character/block device special
              files

       %t     major  device  type  in  hex, for character/block device special
              files

       %T     minor device type in hex,  for  character/block  device  special
              files

       %u     user ID of owner

       %U     user name of owner

       %w     time of file birth, human-readable; - if unknown

       %W     time of file birth, seconds since Epoch; 0 if unknown

       %x     time of last access, human-readable

       %X     time of last access, seconds since Epoch

       %y     time of last data modification, human-readable

       %Y     time of last data modification, seconds since Epoch

       %z     time of last status change, human-readable

       %Z     time of last status change, seconds since Epoch

       Valid format sequences for file systems:

       %a     free blocks available to non-superuser

       %b     total data blocks in file system

       %c     total file nodes in file system

       %d     free file nodes in file system

       %f     free blocks in file system

       %i     file system ID in hex

       %l     maximum length of filenames

       %n     file name

       %s     block size (for faster transfers)

       %S     fundamental block size (for block counts)

       %t     file system type in hex
```

```
       %T     file system type in human readable form

   --terse is equivalent to the following FORMAT:
              %n %s %b %f %u %g %D %i %h %t %T %X %Y %Z %W %o %C

   --terse --file-system is equivalent to the following FORMAT:
              %n %i %l %t %s %S %b %f %a %c %d

       NOTE: your shell may have its own version of stat, which usually super-
       sedes  the  version described here.  Please refer to your shell's docu-
       mentation for details about the options it supports.

AUTHOR
       Written by Michael Meskes.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       stat(2), statfs(2), statx(2)

       Full documentation <https://www.gnu.org/software/coreutils/stat>
       or available locally via: info '(coreutils) stat invocation'

GNU coreutils 9.4                April 2024                          STAT(1)
```

# Bash 的 stdbuf 命令

```
stdbuf OPTION... COMMAND
```

**功能**

使用指定的缓冲模式运行命令。

**类型**

可执行文件（`/usr/bin/stdbuf`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-i`，`--input=MODE` - 指定标准输入的缓冲模式
  - `-o`，`--output=MODE` - 指定标准输出的缓冲模式
  - `-e`，`--error=MODE` - 指定标准错误的缓冲模式
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `COMMAND` - 要执行的命令

## 缓冲模式

- `0` - 无缓冲
- `L` - 行缓冲（不能用于标准输入）
- `N` - `N` 字节的全缓冲，支持 `K`（1024）、`KB`（1000）等后缀

## 示例

```
$ stdbuf -o 0 cat /vat/log/xxx.log        # 使用无缓冲运行命令
```

## 推荐阅读

- Run a command with modified I/O stream buffering - GNU Coreutils
- stdbuf(1) — Linux manual page

# 手册

```
STDBUF(1)                        User Commands                        STDBUF(1)

NAME
       stdbuf  - Run COMMAND, with modified buffering operations for its stan-
       dard streams.

SYNOPSIS
       stdbuf OPTION... COMMAND

DESCRIPTION
       Run COMMAND,  with  modified  buffering  operations  for  its  standard
       streams.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -i, --input=MODE
              adjust standard input stream buffering

       -o, --output=MODE
              adjust standard output stream buffering

       -e, --error=MODE
              adjust standard error stream buffering

       --help display this help and exit

       --version
              output version information and exit

       If MODE is 'L' the corresponding stream will be  line  buffered.   This
       option is invalid with standard input.

       If MODE is '0' the corresponding stream will be unbuffered.

       Otherwise  MODE is a number which may be followed by one of the follow-
       ing: KB 1000, K  1024,  MB  1000*1000,  M  1024*1024,  and  so  on  for
       G,T,P,E,Z,Y,R,Q.   Binary  prefixes can be used, too: KiB=K, MiB=M, and
       so on.  In this case the corresponding stream will  be  fully  buffered
       with the buffer size set to MODE bytes.

       NOTE:  If  COMMAND adjusts the buffering of its standard streams ('tee'
       does for example) then that  will  override  corresponding  changes  by
       'stdbuf'.   Also  some  filters  (like  'dd'  and 'cat' etc.) don't use
       streams for I/O, and are thus unaffected by 'stdbuf' settings.

   Exit status:
       125    if the stdbuf command itself fails
```

```
       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found

       -      the exit status of COMMAND otherwise

EXAMPLES
       tail -f access.log | stdbuf -oL cut -d ' ' -f1 | uniq
       This will immediately display unique entries from access.log

BUGS
       On GLIBC platforms, specifying a buffer size, i.e., using fully
       buffered mode will result in undefined operation.

AUTHOR
       Written by Padraig Brady.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/stdbuf>
       or available locally via: info '(coreutils) stdbuf invocation'

GNU coreutils 9.4                April 2024                        STDBUF(1)
```

# Bash 的 stty 命令

```
stty [OPTION]... [SETTING]...
```

**功能**

查看或修改终端配置。

**类型**

可执行文件（`/usr/bin/stty`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 以易于阅读的形式打印所有当前终端设置
  - `-g`，`--save` - 以紧凑格式打印所有当前终端设置
  - `-F`，`--file=DEVICE` - 指定终端设备
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `SETTING` - 设置项列表

## 设置项说明

| 特殊字符 | 说明 | 备注 |
| --- | --- | --- |
| discard CHAR | CHAR 将切换（开启/关闭）输出丢弃功能 | 非 POSIX 标准 |
| eof CHAR | CHAR 将发送文件结束符（终止输入） | - |
| eol CHAR | CHAR 将结束当前行 | - |
| eol2 CHAR | 用于结束行的备用 CHAR | 非 POSIX 标准 |
| erase CHAR | CHAR 将删除最后输入的一个字符 | - |
| intr CHAR | CHAR 将发送中断信号 | - |
| kill CHAR | CHAR 将删除当前整行 | - |
| lnext CHAR | CHAR 将以字面值方式输入下一字符（引用输入） | 非 POSIX 标准 |
| quit CHAR | CHAR 将发送退出（quit）信号 | - |

| 特殊字符 | 说明 | 备注 |
| --- | --- | --- |
| rprnt CHAR | CHAR 将重新绘制当前行 | 非 POSIX 标准 |
| start CHAR | CHAR 将在停止后恢复输出 | - |
| stop CHAR | CHAR 将停止输出 | - |
| susp CHAR | CHAR 将发送终端挂起（stop）信号 | - |
| swtch CHAR | CHAR 将切换到另一个 shell 层 | 非 POSIX 标准 |
| werase CHAR | CHAR 将删除上一个单词 | 非 POSIX 标准 |

| 特殊设置 | 说明 | 备注 |
| --- | --- | --- |
| N | 将输入和输出速度设置为 N 波特率 | - |
| cols N | 告诉内核终端宽度为 N 列 | 非 POSIX 标准 |
| columns N | 等同于 `cols N` | 非 POSIX 标准 |
| [-]drain | 在应用新设置前等待传输完成（默认启用） | 非 POSIX 标准 |
| ispeed N | 将输入速度设置为 N | - |
| line N | 使用编号为 N 的 line discipline（行规程） | 非 POSIX 标准 |
| min N | 在使用 `-icanon` 时，设置读取完成所需的最少字符数 | - |
| ospeed N | 将输出速度设置为 N | - |
| rows N | 告诉内核终端高度为 N 行 | 非 POSIX 标准 |
| size | 打印内核记录的终端行数与列数 | 非 POSIX 标准 |
| speed | 打印终端当前速度 | - |
| time N | 在使用 `-icanon` 时，设置读取超时时间（N × 0.1 秒） | - |

| 控制设置 | 说明 | 备注 |
| --- | --- | --- |
| [-]clocal | 禁用或启用调制解调器控制信号（`-` 为禁用，默认依系统而定） | - |
| [-]cread | 允许或禁止接收输入 | - |
| [-]crtscts | 启用或关闭 RTS/CTS 硬件握手流控 | 非 POSIX 标准 |
| csN | 设置字符位宽为 N 位（N 取值范围 5～8） | - |
| [-]cstopb | 每个字符使用两个停止位（加 `-` 时为使用一个停止位） | - |

| 控制设置 | 说明 | 备注 |
| --- | --- | --- |
| [-]hup | 当最后一个进程关闭 TTY 时发送挂断信号 | - |
| [-]hupcl | 等同于 `[-]hup` | - |
| [-]parenb | 输出中生成奇偶校验位，并在输入时要求奇偶校验位 | - |
| [-]parodd | 使用奇校验（加 `-` 时使用偶校验） | - |
| [-]cmspar | 使用 "stick" 校验方式（mark/space parity） | 非 POSIX 标准 |

| 输入设置 | 说明 | 备注 |
| --- | --- | --- |
| [-]brkint | break 字符会触发中断信号 | - |
| [-]icrnl | 将回车符（CR）转换为换行符（NL） | - |
| [-]ignbrk | 忽略 break 字符 | - |
| [-]igncr | 忽略回车符（CR） | - |
| [-]ignpar | 忽略带奇偶校验错误的字符 | - |
| [-]imaxbel | 当输入缓冲区满时发出响铃（不清空缓冲区） | 非 POSIX 标准 |
| [-]inlcr | 将换行符（NL）转换为回车符（CR） | - |
| [-]inpck | 启用输入奇偶校验检查 | - |
| [-]istrip | 清除输入字符的第 8 位（只保留低 7 位） | - |
| [-]iutf8 | 假定输入字符是 UTF-8 编码 | 非 POSIX 标准 |
| [-]iuclc | 将大写字符转换为小写 | 非 POSIX 标准 |
| [-]ixany | 输入任意字符都可重新启动输出（不只限 start 字符） | 非 POSIX 标准 |
| [-]ixoff | 启用发送 XOFF/XON 的控制字符（软件流控） | - |
| [-]ixon | 启用接收 XON/XOFF 流控 | - |
| [-]parmrk | 用 255-0-字符 的序列标记奇偶校验错误 | - |
| [-]tandem | 等同于 `[-]ixoff` | - |

| 输出设置 | 说明 | 备注 |
| --- | --- | --- |
| bsN | 退格（Backspace）延时样式，N 取值范围 0～1 | 非 POSIX 标准 |
| crN | 回车（Carriage Return）延时样式，N 取值范围 0～3 | 非 POSIX 标准 |
| ffN | 换页（Form Feed）延时样式，N 取值范围 0～1 | 非 POSIX 标准 |

| 输出设置 | 说明 | 备注 |
| --- | --- | --- |
| nlN | 换行（Newline）延时样式，N 取值范围 0～1 | 非 POSIX 标准 |
| [-]ocrnl | 将回车符（CR）转换为换行符（NL） | 非 POSIX 标准 |
| [-]ofdel | 填充字符使用 DEL（0x7F），而不是 NUL（0x00） | 非 POSIX 标准 |
| [-]ofill | 使用填充字符（padding）而不是基于时间的延迟 | 非 POSIX 标准 |
| [-]olcuc | 将小写字母转换为大写字母 | 非 POSIX 标准 |
| [-]onlcr | 将换行符（NL）转换成回车-换行（CR-NL） | 非 POSIX 标准 |
| [-]onlret | 换行符（NL）同时执行回车（CR）动作 | 非 POSIX 标准 |
| [-]onocr | 在当前列为第一列时不输出回车符（CR） | 非 POSIX 标准 |
| [-]opost | 启用或禁用输出后处理 | - |
| tabN | 水平制表符（tab）延时样式，N 取值范围 0～3 | 非 POSIX 标准 |
| tabs | 等同于 `tab0` | 非 POSIX 标准 |
| -tabs | 等同于 `tab3` | 非 POSIX 标准 |
| vtN | 垂直制表符（vertical tab）延时样式，N 取值范围 0～1 | 非 POSIX 标准 |

| 本地设置 | 说明 | 备注 |
| --- | --- | --- |
| [-]crterase | 将删除字符以“退格-空格-退格”的方式回显 | - |
| crtkill | 按照 `echoprt` 和 `echoe` 的设置删除（kill）整行 | 非 POSIX 标准 |
| -crtkill | 按照 `echoctl` 和 `echok` 的设置删除（kill）整行 | 非 POSIX 标准 |
| [-]ctlecho | 以帽号表示法（如 `^C` ）回显控制字符 | 非 POSIX 标准 |
| [-]echo | 回显输入字符 | - |
| [-]echoctl | 等同于 `[-]ctlecho` | 非 POSIX 标准 |
| [-]echoe | 等同于 `[-]crterase` | - |
| [-]echok | 在 kill 字符后回显一个换行符 | - |
| [-]echoke | 等同于 `[-]crtkill` | 非 POSIX 标准 |
| [-]echonl | 即使不回显其他字符，也回显换行符 | - |
| [-]echoprt | 以反向方式回显被删除的字符，并包围在 `\` 和 `/` 之间 | 非 POSIX 标准 |
| [-]extproc | 启用 LINEMODE（适用于高延迟链接） | 非 POSIX 标准 |
| [-]flusho | 丢弃正在输出的数据 | 非 POSIX 标准 |

| 本地设置 | 说明 | 备注 |
| --- | --- | --- |
| [-]icanon | 启用特殊字符：erase,kill,werase,rprnt | - |
| [-]iexten | 启用非 POSIX 的扩展特殊字符 | - |
| [-]isig | 启用中断,退出与挂起（suspend）特殊字符 | - |
| [-]noflsh | 禁止在中断或退出字符后刷新缓冲区 | - |
| [-]prterase | 等同于 `[-]echoprt` | 非 POSIX 标准 |
| [-]tostop | 阻止后台任务向终端写入（会发送 SIGTTOU） | 非 POSIX 标准 |
| [-]xcase | 在 canonical 模式下，对大写字母加 `\` 进行转义 | 非 POSIX 标准 |

| 组合设置 | 说明 | 备注 |
| --- | --- | --- |
| [-]LCASE | 等同于 `[-]lcase` | 非 POSIX 标准 |
| cbreak | 等同于 `-icanon`（关闭规范模式，单字符输入） | - |
| -cbreak | 等同于 `icanon`（开启规范模式） | - |
| cooked | 将 brkint ignpar istrip icrnl ixon opost isig icanon eof 和 eol 字符恢复默认值 | - |
| -cooked | 等同于 raw（关闭规范模式并应用 raw 设置） | - |
| crt | 等同于 echoe echoctl echoke | - |
| dec | 等同于 echoe echoctl echoke -ixany intr ^C erase 0177 kill ^U | - |
| [-]decctlq | 等同于 `[-]ixany` | 非 POSIX 标准 |
| ek | 将 erase 和 kill 字符恢复默认值 | - |
| evenp | 等同于 parenb -parodd cs7（偶校验，7 位字符） | - |
| -evenp | 等同于 -parenb cs8（关闭奇偶校验，8 位字符） | - |
| [-]lcase | 等同于 xcase iuclc olcuc（小写字母处理相关） | 非 POSIX 标准 |
| litout | 等同于 -parenb -istrip -opost cs8 | - |
| -litout | 等同于 parenb istrip opost cs7 | - |

| 组合设置 | 说明 | 备注 |
| --- | --- | --- |
| nl | 等同于 -icrnl -onlcr | - |
| -nl | 等同于 icrnl -inlcr -igncr onlcr -ocrnl -onlret | - |
| oddp | 等同于 parenb parodd cs7（奇校验，7 位字符） | - |
| -oddp | 等同于 -parenb cs8（关闭奇偶校验，8 位字符） | - |
| [-]parity | 等同于 `[-]evenp` | - |
| pass8 | 等同于 -parenb -istrip cs8（8 位字符，无奇偶校验） | - |
| -pass8 | 等同于 parenb istrip cs7 | - |
| raw | 等同于关闭多种选项：-ignbrk -brkint -ignpar -parmrk -inpck -istrip -inlcr -igncr -icrnl -ixon -ixoff -icanon -opost -isig -iuclc -ixany -imaxbel -xcase min 1 time 0 | - |
| -raw | 等同于 cooked | - |
| sane | 等同于开启 cread -ignbrk brkint -inlcr -igncr icrnl icanon iexten echo echoe echok -echonl -noflsh -ixoff -iutf8 -iuclc -ixany imaxbel -xcase -olcuc -ocrnl opost -ofill onlcr -onocr -onlret nl0 cr0 tab0 bs0 vt0 ff0 isig -tostop -ofdel -echoprt echoctl echoke -extproc -flusho，以及所有特殊字符恢复默认值 | - |

## 示例

```
$ stty -a       # 易于阅读的终端设置
speed 38400 baud; rows 55; columns 230; line = 0;
intr = ^C; quit = ^\; erase = ^?; kill = ^U; eof = ^D; eol = <undef>; eol2 = <undef>;
swtch = <undef>; start = ^Q; stop = ^S; susp = ^Z; rprnt = ^R; werase = ^W; lnext = ^V;
discard = ^O; min = 1; time = 0;
-parenb -parodd -cmspar cs8 -hupcl -cstopb cread -clocal -crtscts
-ignbrk -brkint -ignpar -parmrk -inpck -istrip -inlcr -igncr icrnl ixon -ixoff -iuclc -
ixany -imaxbel -iutf8
opost -olcuc -ocrnl onlcr -onocr -onlret -ofill -ofdel nl0 cr0 tab0 bs0 vt0 ff0
isig icanon iexten echo echoe echok -echonl -noflsh -xcase -tostop -echoprt echoctl echoke
-flusho -extproc
```

## 推荐阅读

- Print or change terminal characteristics - GNU Coreutils
- stty(1) — Linux manual page

# 手册

```
STTY(1)                          User Commands                          STTY(1)

NAME
       stty - change and print terminal line settings

SYNOPSIS
       stty [-F DEVICE | --file=DEVICE] [SETTING]...
       stty [-F DEVICE | --file=DEVICE] [-a|--all]
       stty [-F DEVICE | --file=DEVICE] [-g|--save]

DESCRIPTION
       Print or change terminal characteristics.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -a, --all
              print all current settings in human-readable form

       -g, --save
              print all current settings in a stty-readable form

       -F, --file=DEVICE
              open and use the specified DEVICE instead of stdin

       --help display this help and exit

       --version
              output version information and exit

       Optional - before SETTING indicates negation.   An  *  marks  non-POSIX
       settings.  The underlying system defines which settings are available.

   Special characters:
       * discard CHAR
              CHAR will toggle discarding of output

       eof CHAR
              CHAR will send an end of file (terminate the input)

       eol CHAR
              CHAR will end the line

       * eol2 CHAR
              alternate CHAR for ending the line

       erase CHAR
              CHAR will erase the last character typed
```

```
        intr CHAR
               CHAR will send an interrupt signal

        kill CHAR
               CHAR will erase the current line

        * lnext CHAR
               CHAR will enter the next character quoted

        quit CHAR
               CHAR will send a quit signal

        * rprnt CHAR
               CHAR will redraw the current line

        start CHAR
               CHAR will restart the output after stopping it

        stop CHAR
               CHAR will stop the output

        susp CHAR
               CHAR will send a terminal stop signal

        * swtch CHAR
               CHAR will switch to a different shell layer

        * werase CHAR
               CHAR will erase the last word typed

    Special settings:
        N      set the input and output speeds to N bauds

        * cols N
               tell the kernel that the terminal has N columns

        * columns N
               same as cols N

        * [-]drain
               wait for transmission before applying settings (on by default)

        ispeed N
               set the input speed to N

        * line N
               use line discipline N

        min N  with -icanon, set N characters minimum for a completed read

        ospeed N
               set the output speed to N
```

```
    * rows N
           tell the kernel that the terminal has N rows

    * size print the number of rows and columns according to the kernel

    speed  print the terminal speed

    time N with -icanon, set read timeout of N tenths of a second

Control settings:
    [-]clocal
           disable modem control signals

    [-]cread
           allow input to be received

    * [-]crtscts
           enable RTS/CTS handshaking

    csN    set character size to N bits, N in [5..8]

    [-]cstopb
           use two stop bits per character (one with '-')

    [-]hup send a hangup signal when the last process closes the tty

    [-]hupcl
           same as [-]hup

    [-]parenb
           generate parity bit in output and expect parity bit in input

    [-]parodd
           set odd parity (or even parity with '-')

    * [-]cmspar
           use "stick" (mark/space) parity

Input settings:
    [-]brkint
           breaks cause an interrupt signal

    [-]icrnl
           translate carriage return to newline

    [-]ignbrk
           ignore break characters

    [-]igncr
           ignore carriage return

    [-]ignpar
           ignore characters with parity errors
```

```
    * [-]imaxbel
           beep and do not flush a full input buffer on a character

    [-]inlcr
           translate newline to carriage return

    [-]inpck
           enable input parity checking

    [-]istrip
           clear high (8th) bit of input characters

    * [-]iutf8
           assume input characters are UTF-8 encoded

    * [-]iuclc
           translate uppercase characters to lowercase

    * [-]ixany
           let any character restart output, not only start character

    [-]ixoff
           enable sending of start/stop characters

    [-]ixon
           enable XON/XOFF flow control

    [-]parmrk
           mark parity errors (with a 255-0-character sequence)

    [-]tandem
           same as [-]ixoff

Output settings:
    * bsN  backspace delay style, N in [0..1]

    * crN  carriage return delay style, N in [0..3]

    * ffN  form feed delay style, N in [0..1]

    * nlN  newline delay style, N in [0..1]

    * [-]ocrnl
           translate carriage return to newline

    * [-]ofdel
           use delete characters for fill instead of NUL characters

    * [-]ofill
           use fill (padding) characters instead of timing for delays

    * [-]olcuc
```

```
              translate lowercase characters to uppercase

       * [-]onlcr
              translate newline to carriage return-newline

       * [-]onlret
              newline performs a carriage return

       * [-]onocr
              do not print carriage returns in the first column

       [-]opost
              postprocess output

       * tabN horizontal tab delay style, N in [0..3]

       * tabs same as tab0

       * -tabs
              same as tab3

       * vtN  vertical tab delay style, N in [0..1]

   Local settings:
       [-]crterase
              echo erase characters as backspace-space-backspace

       * crtkill
              kill all line by obeying the echoprt and echoe settings

       * -crtkill
              kill all line by obeying the echoctl and echok settings

       * [-]ctlecho
              echo control characters in hat notation ('^c')

       [-]echo
              echo input characters

       * [-]echoctl
              same as [-]ctlecho

       [-]echoe
              same as [-]crterase

       [-]echok
              echo a newline after a kill character

       * [-]echoke
              same as [-]crtkill

       [-]echonl
              echo newline even if not echoing other characters
```

```
    * [-]echoprt
           echo erased characters backward, between '\' and '/'

    * [-]extproc
           enable "LINEMODE"; useful with high latency links

    * [-]flusho
           discard output

    [-]icanon
           enable special characters: erase, kill, werase, rprnt

    [-]iexten
           enable non-POSIX special characters

    [-]isig
           enable interrupt, quit, and suspend special characters

    [-]noflsh
           disable flushing after interrupt and quit special characters

    * [-]prterase
           same as [-]echoprt

    * [-]tostop
           stop background jobs that try to write to the terminal

    * [-]xcase
           with icanon, escape with '\' for uppercase characters

Combination settings:
    * [-]LCASE
           same as [-]lcase

    cbreak same as -icanon

    -cbreak
           same as icanon

    cooked same  as  brkint ignpar istrip icrnl ixon opost isig icanon, eof
           and eol characters to their default values

    -cooked
           same as raw

    crt    same as echoe echoctl echoke

    dec    same as echoe echoctl echoke -ixany intr ^c erase 0177 kill ^u

    * [-]decctlq
           same as [-]ixany
```

```
       ek     erase and kill characters to their default values

       evenp  same as parenb -parodd cs7

       -evenp same as -parenb cs8

       * [-]lcase
              same as xcase iuclc olcuc

       litout same as -parenb -istrip -opost cs8

       -litout
              same as parenb istrip opost cs7

       nl     same as -icrnl -onlcr

       -nl    same as icrnl -inlcr -igncr onlcr -ocrnl -onlret

       oddp   same as parenb parodd cs7

       -oddp  same as -parenb cs8

       [-]parity
              same as [-]evenp

       pass8  same as -parenb -istrip cs8

       -pass8 same as parenb istrip cs7

       raw    same as -ignbrk -brkint -ignpar -parmrk  -inpck  -istrip  -inlcr
              -igncr  -icrnl  -ixon  -ixoff -icanon -opost -isig -iuclc -ixany
              -imaxbel -xcase min 1 time 0

       -raw   same as cooked

       sane   same as cread -ignbrk brkint -inlcr -igncr icrnl  icanon  iexten
              echo  echoe  echok  -echonl  -noflsh -ixoff -iutf8 -iuclc -ixany
              imaxbel -xcase -olcuc -ocrnl opost -ofill onlcr  -onocr  -onlret
              nl0  cr0  tab0  bs0 vt0 ff0 isig -tostop -ofdel -echoprt echoctl
              echoke -extproc -flusho, all special characters to their default
              values

       Handle the tty line connected to standard  input.   Without  arguments,
       prints  baud  rate, line discipline, and deviations from stty sane.  In
       settings, CHAR is taken literally, or coded as in  ^c,  0x37,  0177  or
       127; special values ^- or undef used to disable special characters.

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>
```

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/stty>
or available locally via: info '(coreutils) stty invocation'

# Bash 的 su 命令

```
su [OPTION]... [USER [ARGS]...]
```

**功能**

切换用户。

**类型**

可执行文件（`/usr/bin/su`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-c`，`--command=command` - 向 SEHLL 传递 `-c command`
  - `-f`，`--fast` - 向 SEHLL 传递 `-f`
  - `-g`，`--group=group` - 指定主要组；仅对 root 用户有效
  - `-G`，`--supp-group=group` - 指定附加组；仅对 root 用户有效
  - `-`，`-l`，`--login` - 以登录 shell 的方式启动 shell
  - `-m`，`-p`，`--preserve-environment` - 保留环境变量；如果指定了 `--login` 则此选项被忽略
  - `-P`，`--pty` - 为会话创建伪终端
  - `-s`，`--shell=shell` - 指定要运行的 shell
  - `--session-command=command` - 向 SEHLL 传递 `-c command`；但不创建新的会话
  - `-w`，`--whitelist-environment=list` - 保留部分环境变量；以逗号分隔
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `USER` - 要切换到的用户；**默认为 `root`**
- `ARGS` - 切换后传递给 shell 的参数

## 示例

```
user$ su user2      # 切换到 user2 用户
Password:
user2$
```

```
user$ su            # 切换到 root 用户
Password:
```

```
root#
```

> 一些发行版（例如 Ubuntu）默认未激活 root 用户，无法直接使用 `su` 切换到 root 用户。
> 如果当前用户属于 `sudo` 组，则可以通过 `sudo su` 切换到 `root` 用户。

## 推荐阅读

- su(1) — Linux manual page

## 手册

```
SU(1)                           User Commands                           SU(1)

NAME
       su - run a command with substitute user and group ID

SYNOPSIS
       su [options] [-] [user [argument...]]

DESCRIPTION
       su allows commands to be run with a substitute user and group ID.

       When called with no user specified, su defaults to running an
       interactive shell as root. When user is specified, additional arguments
       can be supplied, in which case they are passed to the shell.

       For backward compatibility, su defaults to not change the current
       directory and to only set the environment variables HOME and SHELL
       (plus USER and LOGNAME if the target user is not root). It is
       recommended to always use the --login option (instead of its shortcut
       -) to avoid side effects caused by mixing environments.

       This version of su uses PAM for authentication, account and session
       management. Some configuration options found in other su
       implementations, such as support for a wheel group, have to be
       configured via PAM.

       su is mostly designed for unprivileged users, the recommended solution
       for privileged users (e.g., scripts executed by root) is to use
       non-set-user-ID command runuser(1) that does not require authentication
       and provides separate PAM configuration. If the PAM session is not
       required at all then the recommended solution is to use command
       setpriv(1).

       Note that su in all cases uses PAM (pam_getenvlist(3)) to do the final
       environment modification. Command-line options such as --login and
       --preserve-environment affect the environment before it is modified by
       PAM.
```

Since version 2.38 su resets process resource limits RLIMIT_NICE, RLIMIT_RTPRIO, RLIMIT_FSIZE, RLIMIT_AS and RLIMIT_NOFILE.

OPTIONS

-c, --command=command
Pass command to the shell with the -c option.

-f, --fast
Pass -f to the shell, which may or may not be useful, depending on the shell.

-g, --group=group
Specify the primary group. This option is available to the root user only.

-G, --supp-group=group
Specify a supplementary group. This option is available to the root user only. The first specified supplementary group is also used as a primary group if the option --group is not specified.

-, -l, --login
Start the shell as a login shell with an environment similar to a real login:

- clears all the environment variables except TERM and variables specified by --whitelist-environment
- initializes the environment variables HOME, SHELL, USER, LOGNAME, and PATH
- changes to the target user’s home directory
- sets argv[0] of the shell to '-' in order to make the shell a login shell

-m, -p, --preserve-environment
Preserve the entire environment, i.e., do not set HOME, SHELL, USER or LOGNAME. This option is ignored if the option --login is specified.

-P, --pty
Create a pseudo-terminal for the session. The independent terminal provides better security as the user does not share a terminal with the original session. This can be used to avoid TIOCSTI ioctl terminal injection and other security attacks against terminal file descriptors. The entire session can also be moved to the background (e.g., su --pty - username -c application &). If the pseudo-terminal is enabled, then su works as a proxy between the sessions (sync stdin and stdout).

This feature is mostly designed for interactive sessions. If the standard input is not a terminal, but for example a pipe (e.g., echo "date" | su --pty), then the ECHO flag for the pseudo-terminal

```
           is disabled to avoid messy output.

       -s, --shell=shell
           Run the specified shell instead of the default. The shell to run is
           selected according to the following rules, in order:

           •   the shell specified with --shell

           •   the shell specified in the environment variable SHELL, if the
               --preserve-environment option is used

           •   the shell listed in the passwd entry of the target user

           •   /bin/sh

       If the target user has a restricted shell (i.e., not listed in
       /etc/shells), the --shell option and the SHELL environment variables
       are ignored unless the calling user is root.

       --session-command=command
           Same as -c, but do not create a new session. (Discouraged.)

       -w, --whitelist-environment=list
           Don't reset the environment variables specified in the
           comma-separated list when clearing the environment for --login. The
           whitelist is ignored for the environment variables HOME, SHELL,
           USER, LOGNAME, and PATH.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

SIGNALS
       Upon receiving either SIGINT, SIGQUIT or SIGTERM, su terminates its
       child and afterwards terminates itself with the received signal. The
       child is terminated by SIGTERM, after unsuccessful attempt and 2
       seconds of delay the child is killed by SIGKILL.

CONFIG FILES
       su reads the /etc/default/su and /etc/login.defs configuration files.
       The following configuration items are relevant for su:

       FAIL_DELAY (number)
           Delay in seconds in case of an authentication failure. The number
           must be a non-negative integer.

       ENV_PATH (string)
           Defines the PATH environment variable for a regular user. The
           default value is /usr/local/bin:/bin:/usr/bin.

       ENV_ROOTPATH (string), ENV_SUPATH (string)
```

```
            Defines the PATH environment variable for root. ENV_SUPATH takes
            precedence. The default value is
            /usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin.

        ALWAYS_SET_PATH (boolean)
            If set to yes and --login and --preserve-environment were not
            specified su initializes PATH.

            The environment variable PATH may be different on systems where
            /bin and /sbin are merged into /usr; this variable is also affected
            by the --login command-line option and the PAM system setting
            (e.g., pam_env(8)).

EXIT STATUS
        su normally returns the exit status of the command it executed. If the
        command was killed by a signal, su returns the number of the signal
        plus 128.

        Exit status generated by su itself:

        1
            Generic error before executing the requested command

        126
            The requested command could not be executed

        127
            The requested command was not found

FILES
        /etc/pam.d/su
            default PAM configuration file

        /etc/pam.d/su-l
            PAM configuration file if --login is specified

        /etc/default/su
            command specific logindef config file

        /etc/login.defs
            global logindef config file

NOTES
        For security reasons, su always logs failed log-in attempts to the btmp
        file, but it does not write to the lastlog file at all. This solution
        can be used to control su behavior by PAM configuration. If you want to
        use the pam_lastlog(8) module to print warning message about failed
        log-in attempts then pam_lastlog(8) has to be configured to update the
        lastlog file as well. For example by:

           session required pam_lastlog.so nowtmp

HISTORY
```

```
       This su command was derived from coreutils' su, which was based on an
       implementation by David MacKenzie. The util-linux version has been
       refactored by Karel Zak.

SEE ALSO
       setpriv(1), login.defs(5), shells(5), pam(8), runuser(1)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The su command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                             SU(1)
```

# Bash 的 sulogin 命令

```
sulogin [OPTION]... TTY
```

**功能**

单用户模式（救援模式）登录。

**类型**

可执行文件（`/usr/sbin/sulogin`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-e`，`--force` - 强制启动 root shell
  - `-p`，`--login-shell` - 启动 login shell
  - `-t`，`--timeout seconds`
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `TTY` - 此命令连接的终端设备；默认为 `dev/console`

## 示例

```
$ sudo sulogin
Press Enter for maintenance
(or press Control-D to continue):
```

## 推荐阅读

- sulogin(8) — Linux manual page

## 手册

```
SULOGIN(8)                  System Administration                  SULOGIN(8)

NAME
       sulogin - single-user login
```

```
SYNOPSIS
       sulogin [options] [tty]

DESCRIPTION
       sulogin is invoked by init when the system goes into single-user mode.

       The user is prompted:

       Give root password for system maintenance (or type Control-D for normal
       startup):

       If the root account is locked and --force is specified, no password is
       required.

       sulogin will be connected to the current terminal, or to the optional
       tty device that can be specified on the command line (typically
       /dev/console).

       When the user exits from the single-user shell, or presses control-D at
       the prompt, the system will continue to boot.

OPTIONS
       -e, --force
           If the default method of obtaining the root password from the
           system via getpwnam(3) fails, then examine /etc/passwd and
           /etc/shadow to get the password. If these files are damaged or
           nonexistent, or when root account is locked by '!' or '*' at the
           begin of the password then sulogin will start a root shell without
           asking for a password.

           Only use the -e option if you are sure the console is physically
           protected against unauthorized access.

       -p, --login-shell
           Specifying this option causes sulogin to start the shell process as
           a login shell.

       -t, --timeout seconds
           Specify the maximum amount of time to wait for user input. By
           default, sulogin will wait forever.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

ENVIRONMENT
       sulogin looks for the environment variable SUSHELL or sushell to
       determine what shell to start. If the environment variable is not set,
       it will try to execute root’s shell from /etc/passwd. If that fails, it
       will fall back to /bin/sh.
```

AUTHORS

sulogin was written by Miquel van Smoorenburg for sysvinit and later ported to util-linux by Dave Reisner and Karel Zak.

REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The sulogin command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 sum 命令

```
sum [OPTION]... [FILE]...
```

**功能**

查看文件的校验和和块数。

**类型**

可执行文件（`/usr/bin/sum`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-r` - 使用 BSD 求和算法（默认）；块大小为 1K
  - `-s`，`--sysv` - 使用 System V 求和算法；块大小为 512 字节
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ sum 1.txt             # 使用 BSD 算法
08227     1 1.txt
$ sum -s 1.txt          # 使用 System V 算法
3977      1 1.txt
```

## 推荐阅读

- Print checksum and block counts - GNU Coreutils
- sum(1) — Linux manual page

## 手册

```
SUM(1)                          User Commands                          SUM(1)

NAME
       sum - checksum and count the blocks in a file
```

```
SYNOPSIS
       sum [OPTION]... [FILE]...

DESCRIPTION
       Print or check BSD (16-bit) checksums.

       With no FILE, or when FILE is -, read standard input.

       -r     use BSD sum algorithm (the default), use 1K blocks

       -s, --sysv
              use System V sum algorithm, use 512 bytes blocks

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Kayvan Aghaiepour and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/sum>
       or available locally via: info '(coreutils) sum invocation'

GNU coreutils 9.4                April 2024                          SUM(1)
```

# Bash 的 swaplabel 命令

```
swaplabel [OPTION]... DEVICE
```

**功能**

打印或更改交换区域的标签或 UUID。

**类型**

可执行文件（`/usr/sbin/swaplabel`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-L`，`--label label` - 为设备指定一个新的标签；最多 16 个字符
  - `-U`，`--uuid UUID` - 为设备指定一个新的 UUID；必须采用标准的 `8-4-4-4-12` 字符格式
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 要操作的设备

## 示例

```
$ sudo swaplabel /swap.img
UUID:  2a30e3bb-c87e-4d82-87b8-b5dd660cb034
```

## 推荐阅读

- swaplabel(8) — Linux manual page

## 手册

```
SWAPLABEL(8)                System Administration               SWAPLABEL(8)

NAME
       swaplabel - print or change the label or UUID of a swap area

SYNOPSIS
       swaplabel [-L label] [-U UUID] device
```

DESCRIPTION

swaplabel will display or change the label or UUID of a swap partition located on device (or regular file).

If the optional arguments -L and -U are not given, swaplabel will simply display the current swap-area label and UUID of device.

If an optional argument is present, then swaplabel will change the appropriate value on device. These values can also be set during swap creation using mkswap(8). The swaplabel utility allows changing the label or UUID on an actively used swap device.

OPTIONS

-h, --help

Display help text and exit.

-V, --version

Print version and exit.

-L, --label label

Specify a new label for the device. Swap partition labels can be at most 16 characters long. If label is longer than 16 characters, swaplabel will truncate it and print a warning message.

-U, --uuid UUID

Specify a new UUID for the device. The UUID must be in the standard 8-4-4-4-12 character format, such as is output by uuidgen(1).

ENVIRONMENT

LIBBLKID_DEBUG=all

enables libblkid debug output.

AUTHORS

swaplabel was written by Jason Borden <jborden@bluehost.com> and Karel Zak <kzak@redhat.com>.

SEE ALSO

uuidgen(1), mkswap(8), swapon(8)

REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The swaplabel command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 switch_root 命令

```
switch_root [OPTION]... NEWROOT INIT [ARG...]
```

**功能**

切换到另一个文件系统作为挂载树的根目录。

> `switch_root` 将已挂载的 `/proc`、`/dev`、`/sys` 和 `/run` 移动到 `NEWROOT`，并将 `NEWROOT` 设置为新的根文件系统，然后启动 `INIT` 进程

**类型**

可执行文件（`/usr/sbin/switch_root`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `NEWROOT` - 新的根文件系统
- `INIT` - init 进程
- `ARG` - `INIT` 的参数

## 示例

```
$ sudo switch_root new_root /usr/lib/systemd/systemd --system --deserialize=78
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| chroot | 执行命令时指定根路径 |
| pivot_root | 切换根文件系统（可恢复） |

## 推荐阅读

- switch_root(8) — Linux manual page

# 手册

```
SWITCH_ROOT(8)               System Administration              SWITCH_ROOT(8)

NAME
       switch_root - switch to another filesystem as the root of the mount
       tree

SYNOPSIS
       switch_root [-hV]

       switch_root newroot init [arg...]

DESCRIPTION
       switch_root moves already mounted /proc, /dev, /sys and /run to newroot
       and makes newroot the new root filesystem and starts init process.

       WARNING: switch_root removes recursively all files and directories on
       the current root filesystem.

OPTIONS
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       switch_root returns 1 on failure, it never returns on success.

NOTES
       switch_root will fail to function if newroot is not the root of a
       mount. If you want to switch root into a directory that does not meet
       this requirement then you can first use a bind-mounting trick to turn
       any directory into a mount point:

           mount --bind $DIR $DIR

AUTHORS
       Peter Jones <pjones@redhat.com>, Jeremy Katz <katzj@redhat.com>, Karel
       Zak <kzak@redhat.com>

SEE ALSO
       chroot(2), init(8), mkinitrd(8), mount(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The switch_root command is part of the util-linux package which can be
```

```
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                 2023-10-23                    SWITCH_ROOT(8)
```

# Bash 的 sync 命令

```
sync [OPTION] [FILE]...
```

**功能**

将缓存写入存储器。

**类型**

可执行文件（`/usr/bin/sync`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-d`，`--data` - 仅同步文件内容数据，不同步非必要的元数据
  - `-f`，`--file-system` - 同步包含 `FILE` 参数的文件系统
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 推荐阅读

- Synchronize cached writes to persistent storage - GNU Coreutils
- sync(1) — Linux manual page

## 手册

```
SYNC(1)                          User Commands                          SYNC(1)

NAME
       sync - Synchronize cached writes to persistent storage

SYNOPSIS
       sync [OPTION] [FILE]...

DESCRIPTION
       Synchronize cached writes to persistent storage

       If one or more files are specified, sync only them, or their containing
       file systems.
```

```
       -d, --data
              sync only file data, no unneeded metadata

       -f, --file-system
              sync the file systems that contain the files

       --help display this help and exit

       --version
              output version information and exit

BUGS
       Persistence guarantees vary per system.  See the system calls below for
       more details.

AUTHOR
       Written by Jim Meyering and Giuseppe Scrivano.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       fdatasync(2), fsync(2), sync(2), syncfs(2)

       Full documentation <https://www.gnu.org/software/coreutils/sync>
       or available locally via: info '(coreutils) sync invocation'

GNU coreutils 9.4                 April 2024                          SYNC(1)
```

# Bash 的 tac 命令

```
tac [OPTION]... [FILE]...
```

**功能**

将文件的内容顺序逆转，拼接多个文件并打印到标准输出。

> 文件内顺序逆转，文件之间顺序不变。

**类型**

可执行文件（`/usr/bin/tac`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-b`，`--before` - 将分隔符放在前面而不是后面
  - `-r`，`--regex` - 将分隔符当作正则表达式解析
  - `-s`，`--separator=STRING` - 使用 `STRING` 作为分隔符；默认为换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ cat 1.txt 2.txt       # 正序查看文件内容
1
2
3
A
B
C
$ tac 1.txt 2.txt       # 逆序查看文件内容
3
2
1
C
B
A
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| cat | 正序拼接文件内容并打印到标准输出 |
| rev | 反转行中的字符 |

## 推荐阅读

- Concatenate and write files in reverse - GNU Coreutils
- tac(1) — Linux manual page

## 手册

```
TAC(1)                          User Commands                          TAC(1)

NAME
       tac - concatenate and print files in reverse

SYNOPSIS
       tac [OPTION]... [FILE]...

DESCRIPTION
       Write each FILE to standard output, last line first.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -b, --before
              attach the separator before instead of after

       -r, --regex
              interpret the separator as a regular expression

       -s, --separator=STRING
              use STRING as the separator instead of newline

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Jay Lepreau and David MacKenzie.
```

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

cat(1), rev(1)

Full documentation <https://www.gnu.org/software/coreutils/tac>
or available locally via: info '(coreutils) tac invocation'

# Bash 的 tail 命令

```
tail [OPTION]... [FILE]...
```

**功能**

查看文件的末尾。

**类型**

可执行文件（`/usr/bin/tail`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--bytes=[+]NUM` - 查看文件末尾的 `NUM` 字节，如果使用正号，则表示查看第 `NUM` 字节之后的内容
  - `-f`，`--follow[={name|descriptor}]` - 跟随文件，当文件内容增加时打印追加的内容
  - `-F` - 等价于 `--follow=name --retry`
  - `-n`，`--lines=[+]NUM` - 查看文件末尾的 `NUM` 行，如果使用正号，则表示查看第 `NUM` 行之后的内容
  - `--max-unchanged-stats=N` - 使用 `-f` 选项选项时，如果连续 `N` 次读取文件不变，则尝试重新打开文件
  - `--pid=PID` - 使用 `-f` 选项选项时，当 `PID` 进程结束时停止此命令
  - `-q`，`--quiet`，`--silent` - 不打印包含文件名在内的头部信息
  - `--retry` - 如果文件无法打开，则重新尝试打开
  - `-v`，`--verbose` - 打印包含文件名在内的头部信息
  - `-z`，`--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表，如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

**基本功能**

```
$ man tail | tail -n 5                  # 查看末尾 5 行

       Full documentation <https://www.gnu.org/software/coreutils/tail>
```

```
       or available locally via: info '(coreutils) tail invocation'

GNU coreutils 9.4                 April 2024                           TAIL(1)
$ man tail | tail -n +90                # 显示第 90 行之后的内容
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       head(1)

       Full documentation <https://www.gnu.org/software/coreutils/tail>
       or available locally via: info '(coreutils) tail invocation'

GNU coreutils 9.4                 April 2024                           TAIL(1)
```

**监听日志**

```
$ tail -f /var/log/syslog       # 监听该文件，文件内容更新时打印追加的内容
2025-11-15T08:30:00.190259+00:00 dev systemd[1]: Finished sysstat-collect.service - system
activity accounting tool.
2025-11-15T08:35:01.817178+00:00 dev CRON[3239701]: (root) CMD (command -v debian-sa1 >
/dev/null && debian-sa1 1 1)
2025-11-15T08:40:00.422316+00:00 dev systemd[1]: Starting sysstat-collect.service - system
activity accounting tool...
2025-11-15T08:40:00.435212+00:00 dev systemd[1]: sysstat-collect.service: Deactivated
successfully.
2025-11-15T08:40:00.435761+00:00 dev systemd[1]: Finished sysstat-collect.service - system
activity accounting tool.
2025-11-15T08:45:01.832707+00:00 dev CRON[3241555]: (root) CMD (command -v debian-sa1 >
/dev/null && debian-sa1 1 1)
2025-11-15T08:50:00.422442+00:00 dev systemd[1]: Starting sysstat-collect.service - system
activity accounting tool...
2025-11-15T08:50:00.434650+00:00 dev systemd[1]: sysstat-collect.service: Deactivated
successfully.
2025-11-15T08:50:00.435219+00:00 dev systemd[1]: Finished sysstat-collect.service - system
activity accounting tool.
2025-11-15T08:55:01.848254+00:00 dev CRON[3243092]: (root) CMD (command -v debian-sa1 >
/dev/null && debian-sa1 1 1)
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| head | 查看文件开头 |
| more | 分页查看文件内容 |
| less | 增强版文件分页阅读命令 |

## 推荐阅读

- Output the last part of files - GNU Coreutils
- tail(1) — Linux manual page

## 手册

```
TAIL(1)                          User Commands                          TAIL(1)

NAME
       tail - output the last part of files

SYNOPSIS
       tail [OPTION]... [FILE]...

DESCRIPTION
       Print  the  last  10  lines of each FILE to standard output.  With more
       than one FILE, precede each with a header giving the file name.

       With no FILE, or when FILE is -, read standard input.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -c, --bytes=[+]NUM
              output  the  last  NUM  bytes; or use -c +NUM to output starting
              with byte NUM of each file

       -f, --follow[={name|descriptor}]
              output appended data as the file grows;

              an absent option argument means 'descriptor'

       -F     same as --follow=name --retry

       -n, --lines=[+]NUM
              output the last NUM lines, instead of the last  10;  or  use  -n
              +NUM to skip NUM-1 lines at the start

       --max-unchanged-stats=N
              with --follow=name, reopen a FILE which has not

              changed  size  after  N  (default 5) iterations to see if it has
              been unlinked or renamed (this is the usual case of rotated  log
              files); with inotify, this option is rarely useful

       --pid=PID
              with -f, terminate after process ID, PID dies
```

```
       -q, --quiet, --silent
              never output headers giving file names

       --retry
              keep trying to open a file if it is inaccessible

       -s, --sleep-interval=N
              with -f, sleep for approximately N seconds (default 1.0) between
              iterations;  with  inotify and --pid=P, check process P at least
              once every N seconds

       -v, --verbose
              always output headers giving file names

       -z, --zero-terminated
              line delimiter is NUL, not newline

       --help display this help and exit

       --version
              output version information and exit

       NUM may have a multiplier suffix: b 512, kB 1000, K 1024, MB 1000*1000,
       M 1024*1024, GB 1000*1000*1000, G 1024*1024*1024, and so on for  T,  P,
       E,  Z, Y, R, Q.  Binary prefixes can be used, too: KiB=K, MiB=M, and so
       on.

       With --follow (-f), tail defaults to  following  the  file  descriptor,
       which  means that even if a tail'ed file is renamed, tail will continue
       to track its end.  This default behavior is not desirable when you  re-
       ally want to track the actual name of the file, not the file descriptor
       (e.g.,  log  rotation).   Use  --follow=name in that case.  That causes
       tail to track the named file in a way that accommodates  renaming,  re-
       moval and creation.

AUTHOR
       Written  by Paul Rubin, David MacKenzie, Ian Lance Taylor, and Jim Mey-
       ering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       head(1)

       Full documentation <https://www.gnu.org/software/coreutils/tail>
```

```
       or available locally via: info '(coreutils) tail invocation'

GNU coreutils 9.4                April 2024                          TAIL(1)
```

# Bash 的 taskset 命令

```
taskset [OPTION]... MASK CMD [ARGS]...
```

**功能**

设置（或获取）进程的 CPU 亲和性。

**类型**

可执行文件（`/usr/bin/su`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-a`，`--all-tasks` - 设置（或获取）指定进程的所有任务（线程）的 CPU 亲和性
  - `-c`，`--cpu-list` - 将掩码解释为处理器数值列表，用逗号分隔
  - `-p`，`--pid` - 设置（或获取）指定进程的亲和性
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `MASK` - 掩码
- `CMD` - 要执行的命令；使用 `-p` 选项时表示 PID
- `ARGS` - 传递给 `CMD` 的参数

## 示例

**使用指定的 CPU 核心运行程序**

```
$ taskset 0x1 cmd                   # 使用 0 号 CPU 核心运行 cmd
$ taskset 0x3 cmd                   # 使用 0 号 和 1 号 CPU 核心运行 cmd
$ taskset 0xffff cmd                # 使用所有 CPU 核心运行 cmd
$ taskset --cpu-list 0-2,6 cmd      # 使用 #0 #1 #2 #6 核心运行 cmd
$ taskset --cpu-list 0-8:2 cmd      # 使用 #0 #2 #4 #6 #8 核心运行 cmd
```

**查看指定进程的 CPU 亲和性**

```
$ taskset -p $$                     # 查看当前 shell 进程的 CPU 亲和性
pid 694907's current affinity mask: ff
```

- `$$` 是当前 bash 的进 ID

**修改指定进程的 CPU 亲和性**

```
$ sudo taskset -p 1 $$              # 将当前 shell 的 CPU 亲和性设为 1，只在 #0 CPU 核心上运行
pid 694907's current affinity mask: ff
pid 694907's new affinity mask: 1
```

- `$$` 是当前 bash 的进 ID

## 推荐阅读

- taskset(1) — Linux manual page
- Processor affinity - Wikipedia

## 手册

```
TASKSET(1)                      User Commands                      TASKSET(1)

NAME
       taskset - set or retrieve a process's CPU affinity

SYNOPSIS
       taskset [options] mask command [argument...]

       taskset [options] -p [mask] pid

DESCRIPTION
       The taskset command is used to set or retrieve the CPU affinity of a
       running process given its pid, or to launch a new command with a given
       CPU affinity. CPU affinity is a scheduler property that "bonds" a
       process to a given set of CPUs on the system. The Linux scheduler will
       honor the given CPU affinity and the process will not run on any other
       CPUs. Note that the Linux scheduler also supports natural CPU affinity:
       the scheduler attempts to keep processes on the same CPU as long as
       practical for performance reasons. Therefore, forcing a specific CPU
       affinity is useful only in certain applications. The affinity of some
       processes like kernel per-CPU threads cannot be set.

       The CPU affinity is represented as a bitmask, with the lowest order bit
       corresponding to the first logical CPU and the highest order bit
       corresponding to the last logical CPU. Not all CPUs may exist on a
       given system but a mask may specify more CPUs than are present. A
       retrieved mask will reflect only the bits that correspond to CPUs
       physically on the system. If an invalid mask is given (i.e., one that
       corresponds to no valid CPUs on the current system) an error is
       returned. The masks may be specified in hexadecimal (with or without a
       leading "0x"), or as a CPU list with the --cpu-list option. For
       example,
```

```
       0x00000001
           is processor #0,

       0x00000003
           is processors #0 and #1,

       FFFFFFFF
           is processors #0 through #31,

       0x32
           is processors #1, #4, and #5,

       --cpu-list 0-2,6
           is processors #0, #1, #2, and #6.

       --cpu-list 0-10:2
           is processors #0, #2, #4, #6, #8 and #10. The suffix ":N" specifies
           stride in the range, for example 0-10:3 is interpreted as 0,3,6,9
           list.

       When taskset returns, it is guaranteed that the given program has been
       scheduled to a legal CPU.

OPTIONS
       -a, --all-tasks
           Set or retrieve the CPU affinity of all the tasks (threads) for a
           given PID.

       -c, --cpu-list
           Interpret mask as numerical list of processors instead of a
           bitmask. Numbers are separated by commas and may include ranges.
           For example: 0,5,8-11.

       -p, --pid
           Operate on an existing PID and do not launch a new task.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

USAGE
       The default behavior is to run a new command with a given affinity
       mask:
           taskset mask command [arguments]

       You can also retrieve the CPU affinity of an existing task:
           taskset -p pid

       Or set it:
           taskset -p mask pid
```

```
       When a cpu-list is specified for an existing process, the -p and -c
       options must be grouped together:
           taskset -pc cpu-list pid

       The --cpu-list form is applicable only for launching new commands:
           taskset --cpu-list cpu-list command

PERMISSIONS
       A user can change the CPU affinity of a process belonging to the same
       user. A user must possess CAP_SYS_NICE to change the CPU affinity of a
       process belonging to another user. A user can retrieve the affinity
       mask of any process.

RETURN VALUE
       taskset returns 0 in its affinity-getting mode as long as the provided
       PID exists.

       taskset returns 0 in its affinity-setting mode as long as the
       underlying sched_setaffinity(2) system call does. The success of the
       command does not guarantee that the specified thread has actually
       migrated to the indicated CPU(s), but only that the thread will not
       migrate to a CPU outside the new affinity mask. For example, the
       affinity of the kernel thread kswapd can be set, but the thread may not
       immediately migrate and is not guaranteed to ever do so:

       $ ps ax -o comm,psr,pid | grep kswapd
       kswapd0           4      82
       $ sudo taskset -p 1 82
       pid 82's current affinity mask: 1
       pid 82's new affinity mask: 1
       $ echo $?
       0
       $ ps ax -o comm,psr,pid | grep kswapd
       kswapd0           4      82
       $ taskset -p 82
       pid 82's current affinity mask: 1

       In contrast, when the user specifies an illegal affinity, taskset will
       print an error and return 1:

       $ ps ax -o comm,psr,pid | grep ksoftirqd/0
       ksoftirqd/0       0      14
       $ sudo taskset -p 1 14
       pid 14's current affinity mask: 1
       taskset: failed to set pid 14's affinity: Invalid argument
       $ echo $?
       1

AUTHORS
       Written by Robert M. Love.

COPYRIGHT
```

Copyright © 2004 Robert M. Love. This is free software; see the source
for copying conditions. There is NO warranty; not even for
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SEE ALSO

chrt(1), nice(1), renice(1), sched_getaffinity(2), sched_setaffinity(2)

See sched(7) for a description of the Linux scheduling scheme.

REPORTING BUGS

For bug reports, use the issue tracker at
https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The taskset command is part of the util-linux package which can be
downloaded from Linux Kernel Archive
<https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 tee 命令

```
tee [OPTION]... [FILE]...
```

### 功能

读取标准输入，然后将内容打印到标准输出以及每个文件。

### 类型

可执行文件（`/usr/bin/tee`），属于 coreutils。

### 参数

- `OPTION` 选项：
  - `-a`，`--append` - 追加写入文件，而不是覆盖文件
  - `-i`，`--ignore-interrupts` - 忽略中断信号
  - `-p` - 以更合适的模式操作管道
  - `--output-error[=MODE]` - 写入错误时的行为：
    - `warn` - 写入错误时发出警告
    - `warn-nopipe` - 写入非管道输出发生错误时，发出警告
    - `exit` - 写入错误时退出
    - `exit-nopipe` - 写入非管道输出发生错误时，退出
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

### 创建文件

```
$ tee 1.txt << EOF          # 将接下来的输入写入 1.txt，遇到 EOF 时结束
> Primers 编程伙伴
> https://xplanc.org/primers/
> EOF
Primers 编程伙伴
https://xplanc.org/primers/
$ cat 1.txt
Primers 编程伙伴
https://xplanc.org/primers/
```

> `cmd << XXX` 语法表示读取输入直到 `XXX` 结束，然后将这段内容（不含 `XXX`）作为标准输入运行 `cmd`。

**重定向**

```
$ my_cmd | tee -a 1.txt     # 将 my_cmd 命令的输出打印出来的同时追加写入 1.txt
```

# 推荐阅读

- Redirect output to multiple files or processes - GNU Coreutils
- tee(1) — Linux manual page

# 手册

```
TEE(1)                          User Commands                          TEE(1)

NAME
       tee - read from standard input and write to standard output and files

SYNOPSIS
       tee [OPTION]... [FILE]...

DESCRIPTION
       Copy standard input to each FILE, and also to standard output.

       -a, --append
              append to the given FILEs, do not overwrite

       -i, --ignore-interrupts
              ignore interrupt signals

       -p     operate in a more appropriate MODE with pipes.

       --output-error[=MODE]
              set behavior on write error.  See MODE below

       --help display this help and exit

       --version
              output version information and exit

   MODE determines behavior with write errors on the outputs:
       warn   diagnose errors writing to any output

       warn-nopipe
              diagnose errors writing to any output not a pipe
```

```
       exit   exit on error writing to any output

       exit-nopipe
              exit on error writing to any output not a pipe

       The  default  MODE  for  the -p option is 'warn-nopipe'.  With "nopipe"
       MODEs, exit immediately if all outputs become broken  pipes.   The  de-
       fault  operation when --output-error is not specified, is to exit imme-
       diately on error writing to a pipe, and diagnose errors writing to  non
       pipe outputs.

AUTHOR
       Written by Mike Parker, Richard M. Stallman, and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/tee>
       or available locally via: info '(coreutils) tee invocation'

GNU coreutils 9.4                April 2024                          TEE(1)
```

# Bash 的 test 命令

```
test EXPRESSION

# 或

[ EXPRESSION ]
```

### 功能

评估表达式的布尔值；可以检查文件类型、检查字符串、比较整数大小。

### 类型

Bash 内置命令。

> `test` 和 `[` 也是 coreutils 的一个可执行文件；为了保持 POSIX 标准的兼容性，确保在其他 shell（比如一些极简的或古老的 shell）或者某些特定环境下，test 命令依然可用。

### 说明

- 表达式
  - `( EXPRESSION )` - EXPRESSION 是 true
  - `! EXPRESSION` - EXPRESSION 是 false
  - `EXPRESSION1 -a EXPRESSION2` - EXPRESSION1 和 EXPRESSION2 都是 true
  - `EXPRESSION1 -o EXPRESSION2` - EXPRESSION1 或 EXPRESSION2 是 true
- 字符串判断
  - `-n STRING` - STRING 长度不为 0
  - `-z STRING` - STRING 长度为 0
  - `STRING1 = STRING2` - STRING1 等于 STRING2
  - `STRING1 != STRING2` - STRING1 不等于 STRING2
- 整数判断
  - `INTEGER1 -eq INTEGER2` - INTEGER1 等于 INTEGER2
  - `INTEGER1 -ne INTEGER2` - INTEGER1 不等于 INTEGER2
  - `INTEGER1 -lt INTEGER2` - INTEGER1 小于 INTEGER2
  - `INTEGER1 -gt INTEGER2` - INTEGER1 大于 INTEGER2
  - `INTEGER1 -le INTEGER2` - INTEGER1 小于或等于 INTEGER2
  - `INTEGER1 -ge INTEGER2` - INTEGER1 大于或等于 INTEGER2
- 文件判断
  - `-e FILE` - 文件存在

- `-f FILE` - 文件存在且是普通文件（regular file）
- `-d FILE` - 文件存在且是目录（directory）
- `-b FILE` - 文件存在且是块设备文件（block special）
- `-c FILE` - 文件存在且是字符设备文件（character special）
- `-p FILE` - 文件存在且是命名管道（named pipe）
- `-S FILE` - 文件存在且是套接字文件（socket）
- `-h FILE` - 文件存在且是符号链接（symbolic link），同 `-L`
- `-L FILE` - 文件存在且是符号链接（symbolic link），同 `-h`
- `-r FILE` - 文件存在且当前用户有读权限
- `-w FILE` - 文件存在且当前用户有写权限
- `-x FILE` - 文件存在且当前用户有执行权限
- `-s FILE` - 文件存在且大小大于 0
- `-k FILE` - 文件存在且设置了 sticky 位
- `-u FILE` - 文件存在且设置了 set-user-ID
- `-g FILE` - 文件存在且设置了 set-group-ID
- `-O FILE` - 文件存在且所属用户 ID 是当前进程的有效用户 ID
- `-G FILE` - 文件存在且所属组 ID 是当前进程的有效组 ID
- `-N FILE` - 文件存在且自上次读取后已被修改
- `-t FD` - 文件描述符 FD 在终端中被打开
- `FILE1 -ef FILE2` - FILE1 和 FILE2 相同（具有相同的设备和 inode 编号）
- `FILE1 -nt FILE2` - FILE1 比 FILE2 新
- `FILE1 -ot FILE2` - FILE1 比 FILE2 旧

## 示例

```
if [ -d "/tmp" ]; then
    echo "/tmp 是目录"
fi

if [ -f "/tmp/code" ]; then
    echo "/tmp/code 是普通文件"
fi

if [ -c "/dev/null" ]; then
    echo "/dev/null 是字符设备"
fi

if [ 10 -gt 5 ]; then
    echo "10 大于 5"
fi

if [ 6 -lt 8 ]; then
```

```
    echo "6 小于 8"
fi
```

## 推荐阅读

- Check file types and compare values - GNU Coreutils
- test(1) — Linux manual page

## 手册

```
TEST(1)                          User Commands                          TEST(1)

NAME
       test - check file types and compare values

SYNOPSIS
       test EXPRESSION
       test
       [ EXPRESSION ]
       [ ]
       [ OPTION

DESCRIPTION
       Exit with the status determined by EXPRESSION.

       --help display this help and exit

       --version
              output version information and exit

       An omitted EXPRESSION defaults to false.  Otherwise, EXPRESSION is true
       or false and sets exit status.  It is one of:

       ( EXPRESSION )
              EXPRESSION is true

       ! EXPRESSION
              EXPRESSION is false

       EXPRESSION1 -a EXPRESSION2
              both EXPRESSION1 and EXPRESSION2 are true

       EXPRESSION1 -o EXPRESSION2
              either EXPRESSION1 or EXPRESSION2 is true

       -n STRING
              the length of STRING is nonzero
```

```
       STRING equivalent to -n STRING

       -z STRING
              the length of STRING is zero

       STRING1 = STRING2
              the strings are equal

       STRING1 != STRING2
              the strings are not equal

       INTEGER1 -eq INTEGER2
              INTEGER1 is equal to INTEGER2

       INTEGER1 -ge INTEGER2
              INTEGER1 is greater than or equal to INTEGER2

       INTEGER1 -gt INTEGER2
              INTEGER1 is greater than INTEGER2

       INTEGER1 -le INTEGER2
              INTEGER1 is less than or equal to INTEGER2

       INTEGER1 -lt INTEGER2
              INTEGER1 is less than INTEGER2

       INTEGER1 -ne INTEGER2
              INTEGER1 is not equal to INTEGER2

       FILE1 -ef FILE2
              FILE1 and FILE2 have the same device and inode numbers

       FILE1 -nt FILE2
              FILE1 is newer (modification date) than FILE2

       FILE1 -ot FILE2
              FILE1 is older than FILE2

       -b FILE
              FILE exists and is block special

       -c FILE
              FILE exists and is character special

       -d FILE
              FILE exists and is a directory

       -e FILE
              FILE exists

       -f FILE
              FILE exists and is a regular file
```

```
-g FILE
       FILE exists and is set-group-ID

-G FILE
       FILE exists and is owned by the effective group ID

-h FILE
       FILE exists and is a symbolic link (same as -L)

-k FILE
       FILE exists and has its sticky bit set

-L FILE
       FILE exists and is a symbolic link (same as -h)

-N FILE
       FILE exists and has been modified since it was last read

-O FILE
       FILE exists and is owned by the effective user ID

-p FILE
       FILE exists and is a named pipe

-r FILE
       FILE exists and the user has read access

-s FILE
       FILE exists and has a size greater than zero

-S FILE
       FILE exists and is a socket

-t FD  file descriptor FD is opened on a terminal

-u FILE
       FILE exists and its set-user-ID bit is set

-w FILE
       FILE exists and the user has write access

-x FILE
       FILE exists and the user has execute (or search) access

Except  for  -h  and  -L,  all  FILE-related tests dereference symbolic
links.  Beware that parentheses need to  be  escaped  (e.g.,  by  back-
slashes) for shells.  INTEGER may also be -l STRING, which evaluates to
the length of STRING.

NOTE:  Binary  -a  and -o are inherently ambiguous.  Use 'test EXPR1 &&
test EXPR2' or 'test EXPR1 || test EXPR2' instead.

NOTE: [ honors the --help and --version options,  but  test  does  not.
```

test treats each of those as it treats any other nonempty STRING.

NOTE: your shell may have its own version of test and/or [, which usually supersedes the version described here. Please refer to your shell's documentation for details about the options it supports.

## AUTHOR

Written by Kevin Braunsdorf and Matthew Bradburn.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

access(2)

Full documentation <https://www.gnu.org/software/coreutils/test>
or available locally via: info '(coreutils) test invocation'

# Bash 的 timeout 命令

```
timeout [OPTION] DURATION COMMAND [ARG]...
```

**功能**

执行命令时附带超时限制。

**类型**

可执行文件（`/usr/bin/timeout`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--preserve-status` - 超时时仍返回 `COMMAND` 的返回值
  - `--foreground` - 当不从 shell 运行命令时，运行 `COMMAND` 从 tty 读取输入和接收信号；此模式下 `COMMAND` 创建的子进程不会超时
  - `-k`，`--kill-after=DURATION` - 发送信号后如果经过 `DURATION` 时间仍未退出，则发送 `SIGKILL`
  - `-s`，`--signal=SIGNAL` - 指定超时时发送的信号
  - `-v`，`--verbose` - 打印详细信息
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `DURATION` - 超时时间
- `COMMAND` - 执行的命令
- `ARG` - 命令的参数列表

**返回值**

- 124 - 超时（且未启用 `--preserve-status` 选项）
- 125 - `timeout` 自身失败
- 126 - `COMMAND` 不能运行
- 127 - `COMMAND` 不存在
- 137 - `SIGKILL`
- 其它情况返回 `COMMAND` 的返回值

## 示例

```
$ timeout -v 3s sleep 10s
timeout: sending signal TERM to command ‘sleep’
```

## 推荐阅读

- Run a command with a time limit - GNU Coreutils
- timeout(1) — Linux manual page

## 手册

```
TIMEOUT(1)                      User Commands                     TIMEOUT(1)

NAME
       timeout - run a command with a time limit

SYNOPSIS
       timeout [OPTION] DURATION COMMAND [ARG]...
       timeout [OPTION]

DESCRIPTION
       Start COMMAND, and kill it if still running after DURATION.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       --preserve-status

              exit with the same status as COMMAND, even when the

              command times out

       --foreground

              when not running timeout directly from a shell prompt,

              allow COMMAND to read from the TTY and get TTY signals; in  this
              mode, children of COMMAND will not be timed out

       -k, --kill-after=DURATION

              also send a KILL signal if COMMAND is still running

              this long after the initial signal was sent

       -s, --signal=SIGNAL

              specify the signal to be sent on timeout;
```

```
              SIGNAL may be a name like 'HUP' or a number; see 'kill -l' for a
              list of signals

       -v, --verbose
              diagnose to stderr any signal sent upon timeout

       --help display this help and exit

       --version
              output version information and exit

       DURATION  is  a  floating point number with an optional suffix: 's' for
       seconds (the default), 'm' for minutes, 'h' for hours or 'd' for  days.
       A duration of 0 disables the associated timeout.

       Upon timeout, send the TERM signal to COMMAND, if no other SIGNAL spec-
       ified.   The TERM signal kills any process that does not block or catch
       that signal.  It may be necessary to use the KILL  signal,  since  this
       signal can't be caught.

   Exit status:
       124    if COMMAND times out, and --preserve-status is not specified

       125    if the timeout command itself fails

       126    if COMMAND is found but cannot be invoked

       127    if COMMAND cannot be found

       137    if  COMMAND  (or  timeout  itself)  is  sent the KILL (9) signal
              (128+9)

       -      the exit status of COMMAND otherwise

BUGS
       Some platforms don't currently support timeouts beyond the year 2038.

AUTHOR
       Written by Padraig Brady.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       kill(1)
```

```
       Full documentation <https://www.gnu.org/software/coreutils/timeout>
       or available locally via: info '(coreutils) timeout invocation'

GNU coreutils 9.4                April 2024                        TIMEOUT(1)
```

# Bash 的 touch 命令

```
touch [OPTION]... FILE...
```

**功能**

将文件的时间戳修改为当前时间，如果文件不存在则会创建一个空文件。

> 查看文件时间戳：
>
> ```
> $ stat file.txt
> Access: 2025-10-27 20:35:00
> Modify: 2025-10-27 20:35:00
> Change: 2025-10-27 20:35:00
> ```

**类型**

可执行文件（`/usr/bin/rm`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a` - 只修改文件的访问时间（access time）
  - `-c`，`--no-create` - 不创建文件
  - `-d`，`--date=STRING` - 指定时间（格式为 `YYYY-MM-DD HH:MM:SS`），而不是当前时间
  - `-f` -
  - `-h`，`--no-dereference` - 如果文件是符号链接，修改符号链接本身，而不是源文件
  - `-m` - 只修改文件的修改时间（modification time）
  - `-r`，`--reference=FILE` - 使用 `FILE` 文件的时间
  - `-t STAMP` - 指定时间（格式为 `[[CC]YY]MMDDhhmm[.ss]`），而不是当前时间
  - `--time=WORD` - 修改指定的时间类型；有效值：`atime`，`mtime`
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

**更新已有文件的时间戳**

```
$ touch existing.txt
```

- 默认更新 **访问时间** + **修改时间** 为当前时间。

**创建新文件（如果不存在）**

```
$ touch file.txt
```

- 如果 `file.txt` 不存在，则创建空文件。

**不创建新文件（仅修改已有文件）**

```
$ touch -c file.txt
```

- 文件不存在时不创建

**只更新访问时间（atime）**

```
$ touch -a file.txt
```

**只更新修改时间（mtime）**

```
$ touch -m file.txt
```

**设置指定时间**

```
$ touch -t 202510271200 file.txt
```

- 修改时间为 2025-10-27 12:00

```
$ touch --date="2025-10-27 12:00" file.txt
```

- 使用可读时间字符串（GNU 扩展）

**参考另一个文件的时间**

```
$ touch -r file1 file2
```

- 让 `file2` 的时间和 `file1` 一样

```
$ touch -r file1 --time=atime file2   # 只修改访问时间
```

```
$ touch -r file1 --time=mtime file2   # 只修改修改时间
```

- 可以指定修改哪个时间戳

**修改多个文件的时间**

```
$ touch file1.txt file2.txt file3.txt
```

# 推荐阅读

- Remove files or directories - GNU Coreutils
- rm(1) — Linux manual page

# 手册

```
TOUCH(1)                         User Commands                         TOUCH(1)

NAME
       touch - change file timestamps

SYNOPSIS
       touch [OPTION]... FILE...

DESCRIPTION
       Update  the  access  and modification times of each FILE to the current
       time.

       A FILE argument that does not exist is created empty, unless -c  or  -h
       is supplied.

       A  FILE  argument  string of - is handled specially and causes touch to
       change the times of the file associated with standard output.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -a     change only the access time

       -c, --no-create
              do not create any files

       -d, --date=STRING
              parse STRING and use it instead of current time

       -f     (ignored)

       -h, --no-dereference
              affect each symbolic link instead of any referenced file (useful
```

```
              only on systems that can change the timestamps of a symlink)

       -m     change only the modification time

       -r, --reference=FILE
              use this file's times instead of current time

       -t STAMP
              use [[CC]YY]MMDDhhmm[.ss] instead of current time

       --time=WORD
              change the specified time: WORD is access, atime, or use: equiv-
              alent to -a WORD is modify or mtime: equivalent to -m

       --help display this help and exit

       --version
              output version information and exit

       Note that the -d and -t options accept different time-date formats.

DATE STRING
       The  --date=STRING  is  a mostly free format human readable date string
       such as "Sun, 29 Feb 2004 16:21:42 -0800" or "2004-02-29  16:21:42"  or
       even  "next Thursday".  A date string may contain items indicating cal-
       endar date, time of day, time zone, day of week, relative  time,  rela-
       tive date, and numbers.  An empty string indicates the beginning of the
       day.   The date string format is more complex than is easily documented
       here but is fully described in the info documentation.

AUTHOR
       Written by Paul Rubin, Arnold Robbins, Jim  Kingdon,  David  MacKenzie,
       and Randy Smith.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/touch>
       or available locally via: info '(coreutils) touch invocation'

GNU coreutils 9.4                April 2024                          TOUCH(1)
```

# Bash 的 tr 命令

```
tr [OPTION]... STRING1 [STRING2]
```

**功能**

对标准输入中的字符进行翻译、压缩、删除，然后写入标准输出。

**类型**

可执行文件（`/usr/bin/tr`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`-C`，`--complement` - 转换 `STRING1` 中不包含的字符
  - `-d`，`--delete` - 删除 `STRING1` 中的字符
  - `-s`，`--squeeze-repeats` - 将最后一个参数指定的字符去重
  - `-t`，`--truncate-set1` - 首先将 `STRING1` 的长度截断至和 `STRING2` 一样
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `STRING1` - 控制字符串列表1
- `STRING2` - 控制字符串列表2

## 特殊字符

| 转义或表示法 | 含义 |
| --- | --- |
| `\NNN` | 八进制值为 NNN 的字符（1 到 3 个八进制数字） |
| `\\` | 反斜杠 |
| `\a` | 可听的 BEL（响铃） |
| `\b` | 退格符 |
| `\f` | 换页符 |
| `\n` | 换行符 |
| `\r` | 回车符 |
| `\t` | 水平制表符 |

| 转义或表示法 | 含义 |
| --- | --- |
| `\v` | 垂直制表符 |
| `CHAR1-CHAR2` | 从 CHAR1 到 CHAR2 的所有字符（按升序） |
| `[CHAR*]` | 在 ARRAY2 中复制 CHAR，直到 ARRAY1 的长度 |
| `[CHAR*REPEAT]` | 复制 CHAR REPEAT 次，如果 REPEAT 以 0 开头，则为八进制 |
| `[:alnum:]` | 所有字母和数字 |
| `[:alpha:]` | 所有字母 |
| `[:blank:]` | 所有水平空白字符 |
| `[:cntrl:]` | 所有控制字符 |
| `[:digit:]` | 所有数字 |
| `[:graph:]` | 所有可打印字符（不包括空格） |
| `[:lower:]` | 所有小写字母 |
| `[:print:]` | 所有可打印字符（包括空格） |
| `[:punct:]` | 所有标点符号 |
| `[:space:]` | 所有水平或垂直空白字符 |
| `[:upper:]` | 所有大写字母 |
| `[:xdigit:]` | 所有十六进制数字 |
| `[=CHAR=]` | 与 CHAR 等价的所有字符 |

## 示例

```
$ echo "a b c" | tr ' ' '\n'        # 将空格替换为换行
a
b
c
$ echo "a b c" | tr 'abc' 'efg'     # 将 'abc' 依次替换为 'efg'
e f g
$ echo "a b c" | tr -c 'ac' '-'     # 将 'ac' 以外的字符替换为 '-'
a---c-
$ echo "a b c" | tr -d 'a'          # 删除 'a'
 b c
$ echo "aaa bbb ccc" | tr -d 'b'    # 去重 'b'
aaa  ccc
```

## 推荐阅读

- Translate, squeeze, and/or delete characters - GNU Coreutils
- tr(1) — Linux manual page

## 手册

```
TR(1)                           User Commands                           TR(1)

NAME
       tr - translate or delete characters

SYNOPSIS
       tr [OPTION]... STRING1 [STRING2]

DESCRIPTION
       Translate, squeeze, and/or delete characters from standard input, writ-
       ing  to standard output.  STRING1 and STRING2 specify arrays of charac-
       ters ARRAY1 and ARRAY2 that control the action.

       -c, -C, --complement
              use the complement of ARRAY1

       -d, --delete
              delete characters in ARRAY1, do not translate

       -s, --squeeze-repeats
              replace each sequence of a repeated character that is listed  in
              the last specified ARRAY, with a single occurrence of that char-
              acter

       -t, --truncate-set1
              first truncate ARRAY1 to length of ARRAY2

       --help display this help and exit

       --version
              output version information and exit

       ARRAYs  are  specified  as strings of characters.  Most represent them-
       selves.  Interpreted sequences are:

       \NNN   character with octal value NNN (1 to 3 octal digits)

       \\     backslash

       \a     audible BEL

       \b     backspace
```

```
\f     form feed

\n     new line

\r     return

\t     horizontal tab

\v     vertical tab

CHAR1-CHAR2
       all characters from CHAR1 to CHAR2 in ascending order

[CHAR*]
       in ARRAY2, copies of CHAR until length of ARRAY1

[CHAR*REPEAT]
       REPEAT copies of CHAR, REPEAT octal if starting with 0

[:alnum:]
       all letters and digits

[:alpha:]
       all letters

[:blank:]
       all horizontal whitespace

[:cntrl:]
       all control characters

[:digit:]
       all digits

[:graph:]
       all printable characters, not including space

[:lower:]
       all lower case letters

[:print:]
       all printable characters, including space

[:punct:]
       all punctuation characters

[:space:]
       all horizontal or vertical whitespace

[:upper:]
       all upper case letters
```

[:xdigit:]
all hexadecimal digits

[=CHAR=]
all characters which are equivalent to CHAR

Translation occurs if -d is not given and both STRING1 and STRING2 appear. -t is only significant when translating. ARRAY2 is extended to length of ARRAY1 by repeating its last character as necessary. Excess characters of ARRAY2 are ignored. Character classes expand in unspecified order; while translating, [:lower:] and [:upper:] may be used in pairs to specify case conversion. Squeezing occurs after translation or deletion.

## BUGS

Full support is available only for safe single-byte locales, in which every possible input byte represents a single character. The C locale is safe in GNU systems, so you can avoid this issue in the shell by running LC_ALL=C tr instead of plain tr.

## AUTHOR

Written by Jim Meyering.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/tr>
or available locally via: info '(coreutils) tr invocation'

# Bash 的 true 命令

```
true OPTION
```

**功能**

什么也不做，直接返回成功（0），通常用于 `if` 语句。

**类型**

Bash 内置命令。

> `true` 也是 coreutils 的一个可执行文件；为了保持 POSIX 标准的兼容性，确保在其他 shell（比如一些极简的或古老的 shell）或者某些特定环境下，命令依然可用。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
if [ true ]; then
    echo 吃饭
else
    echo 睡觉
fi
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| false | 直接返回失败 |
| test | 评估表达式的布尔值 |

## 推荐阅读

- true(1) — Linux manual page

# 手册

```
TRUE(1)                          User Commands                          TRUE(1)

NAME
       true - do nothing, successfully

SYNOPSIS
       true [ignored command line arguments]
       true OPTION

DESCRIPTION
       Exit with a status code indicating success.

       --help display this help and exit

       --version
              output version information and exit

       NOTE: your shell may have its own version of true, which usually super-
       sedes  the  version described here.  Please refer to your shell's docu-
       mentation for details about the options it supports.

AUTHOR
       Written by Jim Meyering.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/true>
       or available locally via: info '(coreutils) true invocation'

GNU coreutils 9.4                 April 2024                          TRUE(1)
```

# Bash 的 truncate 命令

```
truncate OPTION... FILE...
```

### 功能

截断（或增大）文件。

### 类型

可执行文件（`/usr/bin/truncate`），属于 coreutils。

### 参数

- `OPTION` 选项：
  - `-c`，`--no-create` - 不要创建文件
  - `-o`，`--io-blocks` - 将 `SIZE` 视作 I/O 块的个数，而不是字节；通常一个 I/O 块是 4096 字节
  - `-r`，`--reference=RFILE` - 基于 `RFILE` 文件的大小
  - `-s`，`--size=SIZE` - 调整后的文件大小，默认为字节
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表

## 示例

### 绝对大小

```
$ truncate -s 10k 1.txt             # 将 1.txt 的大小设为 10k，文件不存在时自动创建
$ ls -l 1.txt
-rw-rw-r-- 1 planc planc 10240 Dec  6 12:27 1.txt
$ truncate -s 1 -o 1.txt            # 将 1.txt 的大小设为一个 I/O 块
$ ls -l 1.txt
-rw-rw-r-- 1 planc planc 4096 Dec  6 12:28 1.txt
```

### 相对大小

```
$ truncate -s 1k 2.txt              # 创建一个 1k 大小的文件
$ truncate -r 2.txt 1.txt           # 将 1.txt 的大小设为和 2.txt 一样
$ ls -l 1.txt
-rw-rw-r-- 1 planc planc 1024 Dec  6 12:30 1.txt
```

```
$ truncate -r 2.txt -s +1k 1.txt    # 将 1.txt 的大小设为比 2.txt 大 1k
$ ls -l 1.txt
-rw-rw-r-- 1 planc planc 2048 Dec  6 12:32 1.txt
```

## 推荐阅读

- Shrink or extend the size of a file - GNU Coreutils
- truncate(1) — Linux manual page

## 手册

```
TRUNCATE(1)                      User Commands                      TRUNCATE(1)

NAME
       truncate - shrink or extend the size of a file to the specified size

SYNOPSIS
       truncate OPTION... FILE...

DESCRIPTION
       Shrink or extend the size of each FILE to the specified size

       A FILE argument that does not exist is created.

       If  a  FILE  is larger than the specified size, the extra data is lost.
       If a FILE is shorter, it is  extended  and  the  sparse  extended  part
       (hole) reads as zero bytes.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -c, --no-create
              do not create any files

       -o, --io-blocks
              treat SIZE as number of IO blocks instead of bytes

       -r, --reference=RFILE
              base size on RFILE

       -s, --size=SIZE
              set or adjust the file size by SIZE bytes

       --help display this help and exit

       --version
              output version information and exit
```

The SIZE argument is an integer and optional unit (example: 10K is 10*1024). Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or KB,MB,... (powers of 1000). Binary prefixes can be used, too: KiB=K, MiB=M, and so on.

SIZE may also be prefixed by one of the following modifying characters: '+' extend by, '-' reduce by, '<' at most, '>' at least, '/' round down to multiple of, '%' round up to multiple of.

## AUTHOR

Written by Padraig Brady.

## REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

## COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

## SEE ALSO

dd(1), truncate(2), ftruncate(2)

Full documentation <https://www.gnu.org/software/coreutils/truncate>
or available locally via: info '(coreutils) truncate invocation'

# Bash 的 tsort 命令

```
tsort [OPTION] [FILE]
```

**功能**

对文件内容进行拓扑排序。

**类型**

可执行文件（`/usr/bin/tsort`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ cat 1.txt          # 查看文件
F G
D E
A B
E F
C D
D E
$ tsort 1.txt        # 拓扑排序
A
C
B
D
E
F
G
```

源文件表示 `F -> G`，`D -> E`，`A -> B`，`E -> F`，`C -> D`，`D -> E`，拓扑排序后得到 `A -> B -> C -> D -> E -> F -> G`。

## 推荐阅读

- Topological sort - GNU Coreutils
- tsort(1) — Linux manual page

## 手册

```
TSORT(1)                        User Commands                        TSORT(1)

NAME
       tsort - perform topological sort

SYNOPSIS
       tsort [OPTION] [FILE]

DESCRIPTION
       Write  totally  ordered  list  consistent  with the partial ordering in
       FILE.

       With no FILE, or when FILE is -, read standard input.

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Mark Kettenis.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/tsort>
       or available locally via: info '(coreutils) tsort invocation'

GNU coreutils 9.4                April 2024                          TSORT(1)
```

# Bash 的 tty 命令

```
tty [OPTION]...
```

**功能**

查看连接到标准输入的终端文件名。

**类型**

可执行文件（`/usr/bin/tty`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-s`，`--silent`，`--quiet` - 不打印任何内容，仅通过返回值返回状态
  - `--help` - 显示帮助
  - `--version` - 显示版本

**返回值**

- 0 - 标准输入是终端
- 1 - 标准输入是非终端文件
- 2 - 参数错误
- 3 - 发生写错误
- 5 - 无法确认终端名称

## 示例

```
$ tty
/dev/pts/7
```

## 推荐阅读

- Print file name of terminal on standard input - GNU Coreutils
- tty(1) — Linux manual page

## 手册

NAME

tty - print the file name of the terminal connected to standard input

SYNOPSIS

tty [OPTION]...

DESCRIPTION

Print the file name of the terminal connected to standard input.

-s, --silent, --quiet
print nothing, only return an exit status

--help display this help and exit

--version
output version information and exit

AUTHOR

Written by David MacKenzie.

REPORTING BUGS

GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

Full documentation <https://www.gnu.org/software/coreutils/tty>
or available locally via: info '(coreutils) tty invocation'

# Bash 的 uclampset 命令

```
uclampset [OPTION]... [CMD [ARGS]...]
```

**功能**

设置系统或进程的 CPU 利用率限制属性。

**类型**

可执行文件（`/usr/bin/uclampset`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-m uclamp_min` - 设置最小 CPU 利用率（0–1024）；并非强制利用率不小于该值，而是向调度器表达性能需求
  - `-M uclamp_max` - 设置最大 CPU 利用率（0–1024）
  - `-a`，`--all-tasks` - 设置（或获取）指定进程的所有任务（线程）的利用率限制
  - `-p`，`--pid` - 设置（或获取）指定进程的利用率限制
  - `-s`，`--system` - 设置（或获取）系统的利用率限制
  - `-R`，`--reset-on-fork` - 子进程不继承父进程的调度策略
  - `-v`，`--verbose` - 显示详细信息
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `CMD` - 要执行的命令；使用 `-p` 选项时表示 PID
- `ARGS` - 传递给 `CMD` 的参数

## 示例

**限制 CPU 利用率**

```
$ uclampset -M 512 cmd      # 设置最大 CPU 利用率为 50%
```

- `cmd` 的 CPU 利用率不会超过 50%

**提升性能**

```
$ uclampset -m 512 cmd      # 设置最小 CPU 利用率为 50%
```

- CPU 的实际利用率为 10% 时，调度器会认为利用率是 50%，从而避免 CPU 降频

**操作指定进程**

```
$ uclampset -M 512 -p $$    # 设置当前 bash 的最大 CPU 利用率为 50%
```

- `$$` 是当前 bash 的进 ID

# 推荐阅读

- uclampset(1) — Linux manual page

# 手册

```
UCLAMPSET(1)                    User Commands                    UCLAMPSET(1)

NAME
       uclampset - manipulate the utilization clamping attributes of the
       system or a process

SYNOPSIS
       uclampset [options] [-m uclamp_min] [-M uclamp_max] _command argument

       uclampset [options] [-m uclamp_min] [-M uclamp_max] -p PID

DESCRIPTION
       uclampset sets or retrieves the utilization clamping attributes of an
       existing PID, or runs command with the given attributes.

       Utilization clamping is a new feature added in v5.3. It gives a hint to
       the scheduler about the allowed range of utilization the task should be
       operating at.

       The utilization of the task affects frequency selection and task
       placement. Only schedutil cpufreq governor understands handling util
       clamp hints at the time of writing. Consult your kernel docs for
       further info about other cpufreq governors support.

       If you’re running on asymmetric heterogeneous system like Arm’s
       big.LITTLE. Utilization clamping can help bias task placement. If the
       task is boosted such that util_min value is higher than the little
       cores' capacity, then the scheduler will do its best to place it on a
       big core.

       Similarly, if util_max is smaller than or equal the capacity of the
       little cores, then the scheduler can still choose to place it there
       even if the actual utilization of the task is at max.
```

```
       Setting a task's uclamp_min to a none zero value will effectively boost
       the task as when it runs it'll always start from this utilization
       value.

       By setting a task's uclamp_max below 1024, this will effectively cap
       the task as when it runs it'll never be able to go above this
       utilization value.

       The full utilization range is: [0:1024]. The special value -1 is used
       to reset to system's default.

OPTIONS
       -m
           Set util_min value.

       -M
           Set util_max value.

       -a, --all-tasks
           Set or retrieve the utilization clamping attributes of all the
           tasks (threads) for a given PID.

       -p, --pid
           Operate on an existing PID and do not launch a new task.

       -s, --system
           Set or retrieve the system-wide utilization clamping attributes.

       -R, --reset-on-fork
           Set SCHED_FLAG_RESET_ON_FORK flag.

       -v, --verbose
           Show status information.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

USAGE
       The default behavior is to run a new command:
           uclampset [-m uclamp_min] [-M uclamp_max] command [arguments]

       You can also retrieve the utilization clamping attributes of an
       existing task:
           uclampset -p PID

       Or set them:
           uclampset -p PID [-m uclamp_min] [-M uclamp_max]

       Or control the system-wide attributes:
           uclampset -s [-m uclamp_min] [-M uclamp_max]
```

PERMISSIONS

A user must possess CAP_SYS_NICE to change the scheduling attributes of a process. Any user can retrieve the scheduling information.

NOTES

The system wide utilization clamp attributes are there to control the allowed range the tasks can use. By default both uclamp_min and uclamp_max are set to 1024. This means users can set the utilization clamp values for their task across the full range [0:1024].

For example:

```
uclampset -s -m 512 -M 700
```

will prevent any task from being boosted higher than 512. And all tasks in the systems are capped to a utilization of 700. Effectively rendering the maximum performance of the system to 700.

Consult your kernel docs for the exact expected behavior on that kernel.

AUTHORS

Qais Yousef <qais.yousef@arm.com>

SEE ALSO

nice(1), renice(1), taskset(1), sched(7)

See sched_setscheduler(2) and sched_setattr(2) for a description of the Linux scheduling scheme.

REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The uclampset command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 uname 命令

```
uname [OPTION]...
```

**功能**

查看系统信息。

**类型**

可执行文件（`/usr/bin/uname`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 打印所有信息
  - `-s`，`--kernel-name` - 打印内核名称
  - `-n`，`--nodename` - 打印网络节点主机名
  - `-r`，`--kernel-release` - 打印内核发布版本
  - `-v`，`--kernel-version` - 打印内核版本
  - `-m`，`--machine` - 打印机器硬件名称
  - `-p`，`--processor` - 打印处理器类型
  - `-i`，`--hardware-platform` - 打印硬件架构
  - `-o`，`--operating-system` - 打印操作系统
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$uname -a
Linux dev 6.8.0-87-generic #88-Ubuntu SMP PREEMPT_DYNAMIC Sat Oct 11 09:28:41 UTC 2025
x86_64 x86_64 x86_64 GNU/Linux
$uname -s
Linux
$uname -n
dev
$uname -r
6.8.0-87-generic
$uname -v
#88-Ubuntu SMP PREEMPT_DYNAMIC Sat Oct 11 09:28:41 UTC 2025
$uname -m
```

```
x86_64
$uname -p
x86_64
$uname -i
x86_64
$uname -o
GNU/Linux
```

## 推荐阅读

- Print system information - GNU Coreutils
- uname(1) — Linux manual page

## 手册

```
UNAME(1)                        User Commands                        UNAME(1)

NAME
       uname - print system information

SYNOPSIS
       uname [OPTION]...

DESCRIPTION
       Print certain system information.  With no OPTION, same as -s.

       -a, --all
              print  all  information,  in the following order, except omit -p
              and -i if unknown:

       -s, --kernel-name
              print the kernel name

       -n, --nodename
              print the network node hostname

       -r, --kernel-release
              print the kernel release

       -v, --kernel-version
              print the kernel version

       -m, --machine
              print the machine hardware name

       -p, --processor
              print the processor type (non-portable)
```

```
       -i, --hardware-platform
              print the hardware platform (non-portable)

       -o, --operating-system
              print the operating system

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       arch(1), uname(2)

       Full documentation <https://www.gnu.org/software/coreutils/uname>
       or available locally via: info '(coreutils) uname invocation'

GNU coreutils 9.4                April 2024                          UNAME(1)
```

# Bash 的 unexpand 命令

```
unexpand [OPTION]... [FILE]...
```

**功能**

将空格转换为制表符（`\t`）。

**类型**

可执行文件（`/usr/bin/unexpand`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 转换搜索空格，而不仅是前导空格
  - `--first-only` - 仅转换前导空格
  - `-t`，`--tabs=N` - 制表位宽度为 `N` 个字符；默认为 8
  - `-t`，`--tabs=LIST` - 使用 LIST 指定制表位的位置
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ echo -e "        A        B        C" > 1.txt
$ unexpand 1.txt | cat -t
^IA        B        C
$ unexpand -a 1.txt | cat -t
^IA^I B^I  C
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| expand | 将制表符（`\t`）转换为空格 |

## 推荐阅读

- Convert spaces to tabs - GNU Coreutils
- unexpand(1) — Linux manual page

## 手册

```
UNEXPAND(1)                     User Commands                     UNEXPAND(1)

NAME
       unexpand - convert spaces to tabs

SYNOPSIS
       unexpand [OPTION]... [FILE]...

DESCRIPTION
       Convert blanks in each FILE to tabs, writing to standard output.

       With no FILE, or when FILE is -, read standard input.

       Mandatory  arguments  to  long  options are mandatory for short options
       too.

       -a, --all
              convert all blanks, instead of just initial blanks

       --first-only
              convert only leading sequences of blanks (overrides -a)

       -t, --tabs=N
              have tabs N characters apart instead of 8 (enables -a)

       -t, --tabs=LIST
              use comma separated list of tab positions.  The  last  specified
              position  can  be prefixed with '/' to specify a tab size to use
              after the last explicitly specified tab stop.  Also a prefix  of
              '+'  can  be  used  to align remaining tab stops relative to the
              last specified tab stop instead of the first column

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
```

Copyright © 2023 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

SEE ALSO

expand(1)

Full documentation <https://www.gnu.org/software/coreutils/unexpand>
or available locally via: info '(coreutils) unexpand invocation'

# Bash 的 uniq 命令

```
uniq [OPTION]... [INPUT [OUTPUT]]
```

**功能**

过滤文件中连续的重复行。

**类型**

可执行文件（`/usr/bin/uniq`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--count` - 按出现的次数给行添加前缀
  - `-d`，`--repeated` - 仅打印重复的行一次
  - `-D` - 打印所有重复行
  - `--all-repeated[=METHOD]` - 打印所有重复行，并使用空行分隔分组：
    - `none` - 不使用空行分隔分组；同 `-D`（默认）
    - `prepend` - 每组之前空行
    - `separate` - 两组之间空行
  - `-f`，`--skip-fields=N` - 避免比较前 `N` 个字段
  - `--group[=METHOD]` - 打印所有行，并使用空行分组
    - `separate` - 两组之间空行
    - `prepend` - 每组之前空行
    - `append` - 每组之后空行
    - `both` - 每组前后空行
  - `-i`，`--ignore-case` - 忽略大小写
  - `-s`，`--skip-chars=N` - 不比较行的前 `N` 个字符
  - `-u`，`--unique` - 仅打印唯一的行
  - `-z`，`--zero-terminated` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `-w`，`--check-chars=N` - 只比较行的前 `N` 个字符
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `INPUT` - 输入文件；如果没有这个参数或指定为 `-`，则从标准输入读取
- `OUTPUT` - 输出文件；如果没有这个参数或指定为 `-`，则打印到标准输出

## 示例

```
$ cat 1.txt
Primers 编程伙伴
Primers 编程伙伴
https://xplanc.org/primers/
https://xplanc.org/primers/
Primers 编程伙伴
Primers 编程伙伴
https://xplanc.org/primers/
https://xplanc.org/primers/
Primers 编程伙伴 https://xplanc.org/primers/
$ uniq 1.txt                                    # 过滤连续的重复行
Primers 编程伙伴
https://xplanc.org/primers/
Primers 编程伙伴
https://xplanc.org/primers/
Primers 编程伙伴 https://xplanc.org/primers/
$ uniq -c 1.txt                                 # 打印重复次数
      2 Primers 编程伙伴
      2 https://xplanc.org/primers/
      2 Primers 编程伙伴
      2 https://xplanc.org/primers/
      1 Primers 编程伙伴 https://xplanc.org/primers/
$ uniq -d 1.txt                                 # 仅打印连续的重复行一次
Primers 编程伙伴
https://xplanc.org/primers/
Primers 编程伙伴
https://xplanc.org/primers/
$ uniq -D 1.txt                                 # 仅打印所有重复行
Primers 编程伙伴
Primers 编程伙伴
https://xplanc.org/primers/
https://xplanc.org/primers/
Primers 编程伙伴
Primers 编程伙伴
https://xplanc.org/primers/
https://xplanc.org/primers/
$ uniq -u 1.txt                                 # 仅打印唯一行
Primers 编程伙伴 https://xplanc.org/primers/
```

## 推荐阅读

- Uniquify files - GNU Coreutils
- uniq(1) — Linux manual page

## 手册

```
UNIQ(1)                          User Commands                          UNIQ(1)

NAME
       uniq - report or omit repeated lines

SYNOPSIS
       uniq [OPTION]... [INPUT [OUTPUT]]

DESCRIPTION
       Filter  adjacent matching lines from INPUT (or standard input), writing
       to OUTPUT (or standard output).

       With no options, matching lines are merged to the first occurrence.

       Mandatory arguments to long options are  mandatory  for  short  options
       too.

       -c, --count
              prefix lines by the number of occurrences

       -d, --repeated
              only print duplicate lines, one for each group

       -D     print all duplicate lines

       --all-repeated[=METHOD]
              like  -D,  but  allow  separating  groups  with  an  empty line;
              METHOD={none(default),prepend,separate}

       -f, --skip-fields=N
              avoid comparing the first N fields

       --group[=METHOD]
              show  all  items,  separating  groups  with   an   empty   line;
              METHOD={separate(default),prepend,append,both}

       -i, --ignore-case
              ignore differences in case when comparing

       -s, --skip-chars=N
              avoid comparing the first N characters

       -u, --unique
              only print unique lines

       -z, --zero-terminated
              line delimiter is NUL, not newline

       -w, --check-chars=N
              compare no more than N characters in lines

       --help display this help and exit
```

```
       --version
              output version information and exit

       A field is a run of blanks (usually spaces and/or TABs), then non-blank
       characters.  Fields are skipped before chars.

       Note:  'uniq'  does not detect repeated lines unless they are adjacent.
       You may want to sort the input first, or use 'sort -u' without 'uniq'.

AUTHOR
       Written by Richard M. Stallman and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       comm(1), join(1), sort(1)

       Full documentation <https://www.gnu.org/software/coreutils/uniq>
       or available locally via: info '(coreutils) uniq invocation'

GNU coreutils 9.4                April 2024                          UNIQ(1)
```

# Bash 的 unlink 命令

```
unlink [OPTION]... FILE
```

**功能**

删除文件。

**类型**

可执行文件（`/usr/bin/unlink`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件

## 示例

```
$ unlink 1.txt              # 删除 1.txt
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| rm | 删除文件 |
| rmdir | 删除目录 |
| shred | 粉碎文件（避免恢复） |

## 推荐阅读

- Remove files via the unlink syscall - GNU Coreutils
- unlink(1) — Linux manual page

## 手册

```
UNLINK(1)                       User Commands                       UNLINK(1)

NAME
       unlink - call the unlink function to remove the specified file

SYNOPSIS
       unlink FILE
       unlink OPTION

DESCRIPTION
       Call the unlink function to remove the specified FILE.

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Michael Stone.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       unlink(2)

       Full documentation <https://www.gnu.org/software/coreutils/unlink>
       or available locally via: info '(coreutils) unlink invocation'

GNU coreutils 9.4                April 2024                         UNLINK(1)
```

# Bash 的 unshare 命令

```
unshare [OPTION]... [CMD [ARG]...]
```

**功能**

创建新的命名空间并执行命令。

**类型**

可执行文件（`/usr/bin/unshare`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-i`，`--ipc[=file]` - 创建 IPC namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-m`，`--mount[=file]` - 创建 mount namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-n`，`--net[=file]` - 创建 network namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-p`，`--pid[=file]` - 创建 PID namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-u`，`--uts[=file]` - 创建 UTS namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-U`，`--user[=file]` - 创建 user namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-C`，`--cgroup[=file]` - 创建 cgroup namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-T`，`--time[=file]` - 创建 time namespace，如果指定了 `file` 则在该文件上创建挂载点
  - `-f`，`--fork` - 作为 `unshare` 的子进程运行命令，而非直接运行
  - `--keep-caps` - 创建用户命名空间时，在用户命名空间中授予进程的能力（capabilities）在子进程中保留
  - `--kill-child[=signame]` - unshare 退出时向子进程发送 `signame` 信号；**默认为** `SIGKILL`
  - `--mount-proc[=mountpoint]` - 在运行程序直接挂载 proc 文件系统（默认 `/proc`）
  - `--map-user=uid|name` - 仅在当前有效用户 ID 映射为 `uid` 后运行程序
  - `--map-users=inneruid:outeruid:count|auto` - 仅在从 `outeruid` 开始大小为 `count` 的用户 ID 块映射为 `inneruid` 开始的用户 ID 块后运行程序
  - `--map-group=gid|name` - 仅在当前有效组 ID 映射为 `gid` 后运行程序
  - `--map-groups=innergid:outergid:count|auto` - 仅在从 `outergid` 开始大小为 `count` 的组 ID 块映射为 `innergid` 开始的组 ID 块后运行程序

- `--map-auto` - 将 `/etc/subuid` 中当前用户拥有的第一个用户 ID 块映射到以用户 ID 0 开始的块；将 `/etc/subgid` 中当前用户拥有的第一个组 ID 块映射到以组 ID 0 开始的块
- `-r`, `--map-root-user` - 仅在当前有效用户 ID 和组 ID 映射为 root 用户的 UID 和 GID 后运行程序
- `-c`, `--map-current-user` - 仅在当前有效用户 ID 和组 ID 映射为当前用户的 UID 和 GID 后运行程序
- `--propagation private|shared|slave|unchanged` - 递归地在新挂载命名空间中设置挂载传播标志
- `--setgroups allow|deny` - 运行或禁止在 user namespace 中使用 `setgroups` 系统调用
- `-R`, `--root=dir` - 设置根目录
- `-w`, `--wd=dir` - 设置工作目录
- `-S`, `--setuid uid` - 设置 UID
- `-G`, `--setgid gid` - 设置 GID
- `--monotonic offset` - 设置 `CLOCK_MONOTONIC` 在 time namespace 中的偏移量
- `--boottime offset` - 设置 `CLOCK_BOOTTIME` 在 time namespace 中的偏移量
- `-h`, `--help` - 显示帮助
- `-V`, `--version` - 显示版本

- `CMD` - 要执行的命令；默认执行 `${SHELL}`
- `ARG` - 命令的参数列表

## Linux 的命名空间

Linux 命名空间（namespace）是一种内核级隔离机制，用于让不同进程看到不同的系统资源视图，从而实现资源隔离。是容器技术（如 Docker 与 Kubernetes）的核心基础之一。

| 命名空间 | 说明 |
| --- | --- |
| **mount namespace** | 挂载或卸载文件系统不会影响系统的其他部分，除非该文件系统被显式标记为共享（使用 `mount --make-shared`；共享标志可在 `/proc/self/mountinfo` 中查看）。更多细节参见 `mount_namespaces(7)` 以及 `clone(2)` 中关于 `CLONE_NEWNS` 标志的说明。 |
| **UTS namespace** | 设置主机名（hostname）或域名（domainname）不会影响系统的其他部分。更多细节参见 `uts_namespaces(7)`。 |
| **IPC namespace** | 进程将拥有独立的 IPC 命名空间，包括 POSIX 消息队列，以及 System V 消息队列、信号量集合和共享内存段。更多细节参见 `ipc_namespaces(7)`。 |
| **network namespace** | 进程将拥有独立的 IPv4 和 IPv6 协议栈、IP 路由表、防火墙规则、`/proc/net` 和 `/sys/class/net` 目录树、套接字等资源。更多细节参见 `network_namespaces(7)`。 |

| 命名空间 | 说明 |
| --- | --- |
| PID namespace | 子进程将拥有一套与 `nsenter` 进程分离的 PID 到进程映射关系。如果更改 PID 命名空间，`nsenter` 默认会执行 `fork`，以便新程序及其子进程共享同一个 PID 命名空间并彼此可见。如果使用 `--no-fork`，则新程序将在不进行 `fork` 的情况下通过 `exec` 启动。更多细节参见 `pid_namespaces(7)`。 |
| user namespace | 进程将拥有独立的一组 UID、GID 以及能力（capabilities）。更多细节参见 `user_namespaces(7)`。 |
| cgroup namespace | 进程将看到虚拟化后的 `/proc/self/cgroup` 视图，新创建的 cgroup 挂载点将以该命名空间的 cgroup 根目录为起点。更多细节参见 `cgroup_namespaces(7)`。 |
| time namespace | 进程可以拥有独立的 `CLOCK_MONOTONIC` 和/或 `CLOCK_BOOTTIME` 视图，这些时间偏移可通过 `/proc/self/timens_offsets` 进行修改。更多细节参见 `time_namespaces(7)`。 |

## 示例

### 隔离进程环境

```
$ sudo unshare -p -f --mount-proc       # 创建新的 PID 命名空间，运行 /bin/bash
# ps -afx                               # 查看进程树
    PID TTY      STAT   TIME COMMAND
      1 pts/4    S      0:00 -bash
     16 pts/4    R+     0:00 ps -afx
```

- 由于创建了 PID 命名空间，隔离了进程环境，因此看不到其它进程

### 隔离网络环境

```
$ sudo unshare -n bash
# ip addr
1: lo: <LOOPBACK> mtu 65536 qdisc noop state DOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
```

- 只剩下 loopback 网卡

## 相关命令

| 命令 | 说明 |
| --- | --- |
| nsenter | 进入指定的命名空间执行命令 |

# 推荐阅读

- unshare(1) — Linux manual page
- Linux namespaces - Wikipedia

# 手册

```
UNSHARE(1)                      User Commands                      UNSHARE(1)

NAME
       unshare - run program in new namespaces

SYNOPSIS
       unshare [options] [program [arguments]]

DESCRIPTION
       The unshare command creates new namespaces (as specified by the
       command-line options described below) and then executes the specified
       program. If program is not given, then "${SHELL}" is run (default:
       /bin/sh).

       By default, a new namespace persists only as long as it has member
       processes. A new namespace can be made persistent even when it has no
       member processes by bind mounting /proc/pid/ns/type files to a
       filesystem path. A namespace that has been made persistent in this way
       can subsequently be entered with nsenter(1) even after the program
       terminates (except PID namespaces where a permanently running init
       process is required). Once a persistent namespace is no longer needed,
       it can be unpersisted by using umount(8) to remove the bind mount. See
       the EXAMPLES section for more details.

       unshare since util-linux version 2.36 uses
       /proc/[pid]/ns/pid_for_children and /proc/[pid]/ns/time_for_children
       files for persistent PID and TIME namespaces. This change requires
       Linux kernel 4.17 or newer.

       The following types of namespaces can be created with unshare:

       mount namespace
           Mounting and unmounting filesystems will not affect the rest of the
           system, except for filesystems which are explicitly marked as
           shared (with mount --make-shared; see /proc/self/mountinfo or
           findmnt -o+PROPAGATION for the shared flags). For further details,
           see mount_namespaces(7).

           unshare since util-linux version 2.27 automatically sets
           propagation to private in a new mount namespace to make sure that
           the new namespace is really unshared. It’s possible to disable this
           feature with option --propagation unchanged. Note that private is
```

the kernel default.

UTS namespace
Setting hostname or domainname will not affect the rest of the system. For further details, see uts_namespaces(7).

IPC namespace
The process will have an independent namespace for POSIX message queues as well as System V message queues, semaphore sets and shared memory segments. For further details, see ipc_namespaces(7).

network namespace
The process will have independent IPv4 and IPv6 stacks, IP routing tables, firewall rules, the /proc/net and /sys/class/net directory trees, sockets, etc. For further details, see network_namespaces(7).

PID namespace
Children will have a distinct set of PID-to-process mappings from their parent. For further details, see pid_namespaces(7).

cgroup namespace
The process will have a virtualized view of /proc/self/cgroup, and new cgroup mounts will be rooted at the namespace cgroup root. For further details, see cgroup_namespaces(7).

user namespace
The process will have a distinct set of UIDs, GIDs and capabilities. For further details, see user_namespaces(7).

time namespace
The process can have a distinct view of CLOCK_MONOTONIC and/or CLOCK_BOOTTIME which can be changed using /proc/self/timens_offsets. For further details, see time_namespaces(7).

OPTIONS
-i, --ipc[=file]
Create a new IPC namespace. If file is specified, then the namespace is made persistent by creating a bind mount at file.

-m, --mount[=file]
Create a new mount namespace. If file is specified, then the namespace is made persistent by creating a bind mount at file. Note that file must be located on a mount whose propagation type is not shared (or an error results). Use the command findmnt -o+PROPAGATION when not sure about the current setting. See also the examples below.

-n, --net[=file]
Create a new network namespace. If file is specified, then the namespace is made persistent by creating a bind mount at file.

```
-p, --pid[=file]
    Create a new PID namespace. If file is specified, then the
    namespace is made persistent by creating a bind mount at file.
    (Creation of a persistent PID namespace will fail if the --fork
    option is not also specified.)

    See also the --fork and --mount-proc options.

-u, --uts[=file]
    Create a new UTS namespace. If file is specified, then the
    namespace is made persistent by creating a bind mount at file.

-U, --user[=file]
    Create a new user namespace. If file is specified, then the
    namespace is made persistent by creating a bind mount at file.

-C, --cgroup[=file]
    Create a new cgroup namespace. If file is specified, then the
    namespace is made persistent by creating a bind mount at file.

-T, --time[=file]
    Create a new time namespace. If file is specified, then the
    namespace is made persistent by creating a bind mount at file. The
    --monotonic and --boottime options can be used to specify the
    corresponding offset in the time namespace.

-f, --fork
    Fork the specified program as a child process of unshare rather
    than running it directly. This is useful when creating a new PID
    namespace. Note that when unshare is waiting for the child process,
    then it ignores SIGINT and SIGTERM and does not forward any signals
    to the child. It is necessary to send signals to the child process.

--keep-caps
    When the --user option is given, ensure that capabilities granted
    in the user namespace are preserved in the child process.

--kill-child[=signame]
    When unshare terminates, have signame be sent to the forked child
    process. Combined with --pid this allows for an easy and reliable
    killing of the entire process tree below unshare. If not given,
    signame defaults to SIGKILL. This option implies --fork.

--mount-proc[=mountpoint]
    Just before running the program, mount the proc filesystem at
    mountpoint (default is /proc). This is useful when creating a new
    PID namespace. It also implies creating a new mount namespace since
    the /proc mount would otherwise mess up existing programs on the
    system. The new proc filesystem is explicitly mounted as private
    (with MS_PRIVATE|MS_REC).

--map-user=uid|name
    Run the program only after the current effective user ID has been
```

mapped to uid. If this option is specified multiple times, the last occurrence takes precedence. This option implies --user.

--map-users=inneruid:outeruid:count|auto
Run the program only after the block of user IDs of size count beginning at outeruid has been mapped to the block of user IDs beginning at inneruid. This mapping is created with newuidmap(1). If the range of user IDs overlaps with the mapping specified by --map-user, then a "hole" will be removed from the mapping. This may result in the highest user ID of the mapping not being mapped. The special value auto will map the first block of user IDs owned by the effective user from /etc/subuid to a block starting at user ID 0. If this option is specified multiple times, the last occurrence takes precedence. This option implies --user.

Before util-linux version 2.39, this option expected a comma-separated argument of the form outeruid,inneruid,count but that format is now deprecated for consistency with the ordering used in /proc/[pid]/uid_map and the X-mount.idmap mount option.

--map-group=gid|name
Run the program only after the current effective group ID has been mapped to gid. If this option is specified multiple times, the last occurrence takes precedence. This option implies --setgroups=deny and --user.

--map-groups=innergid:outergid:count|auto
Run the program only after the block of group IDs of size count beginning at outergid has been mapped to the block of group IDs beginning at innergid. This mapping is created with newgidmap(1). If the range of group IDs overlaps with the mapping specified by --map-group, then a "hole" will be removed from the mapping. This may result in the highest group ID of the mapping not being mapped. The special value auto will map the first block of user IDs owned by the effective user from /etc/subgid to a block starting at group ID 0. If this option is specified multiple times, the last occurrence takes precedence. This option implies --user.

Before util-linux version 2.39, this option expected a comma-separated argument of the form outergid,innergid,count but that format is now deprecated for consistency with the ordering used in /proc/[pid]/gid_map and the X-mount.idmap mount option.

--map-auto
Map the first block of user IDs owned by the effective user from /etc/subuid to a block starting at user ID 0. In the same manner, also map the first block of group IDs owned by the effective group from /etc/subgid to a block starting at group ID 0. This option is intended to handle the common case where the first block of subordinate user and group IDs can map the whole user and group ID space. This option is equivalent to specifying --map-users=auto and --map-groups=auto.

```
-r, --map-root-user
    Run the program only after the current effective user and group IDs
    have been mapped to the superuser UID and GID in the newly created
    user namespace. This makes it possible to conveniently gain
    capabilities needed to manage various aspects of the newly created
    namespaces (such as configuring interfaces in the network namespace
    or mounting filesystems in the mount namespace) even when run
    unprivileged. As a mere convenience feature, it does not support
    more sophisticated use cases, such as mapping multiple ranges of
    UIDs and GIDs. This option implies --setgroups=deny and --user.
    This option is equivalent to --map-user=0 --map-group=0.

-c, --map-current-user
    Run the program only after the current effective user and group IDs
    have been mapped to the same UID and GID in the newly created user
    namespace. This option implies --setgroups=deny and --user. This
    option is equivalent to --map-user=$(id -ru) --map-group=$(id -rg).

--propagation private|shared|slave|unchanged
    Recursively set the mount propagation flag in the new mount
    namespace. The default is to set the propagation to private. It is
    possible to disable this feature with the argument unchanged. The
    option is silently ignored when the mount namespace (--mount) is
    not requested.

--setgroups allow|deny
    Allow or deny the setgroups(2) system call in a user namespace.

    To be able to call setgroups(2), the calling process must at least
    have CAP_SETGID. But since Linux 3.19 a further restriction
    applies: the kernel gives permission to call setgroups(2) only
    after the GID map (/proc/pid*/gid_map*) has been set. The GID map
    is writable by root when setgroups(2) is enabled (i.e., allow, the
    default), and the GID map becomes writable by unprivileged
    processes when setgroups(2) is permanently disabled (with deny).

-R, --root=dir
    run the command with root directory set to dir.

-w, --wd=dir
    change working directory to dir.

-S, --setuid uid
    Set the user ID which will be used in the entered namespace.

-G, --setgid gid
    Set the group ID which will be used in the entered namespace and
    drop supplementary groups.

--monotonic offset
    Set the offset of CLOCK_MONOTONIC which will be used in the entered
    time namespace. This option requires unsharing a time namespace
    with --time.
```

```
       --boottime offset
           Set the offset of CLOCK_BOOTTIME which will be used in the entered
           time namespace. This option requires unsharing a time namespace
           with --time.

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

NOTES
       The proc and sysfs filesystems mounting as root in a user namespace
       have to be restricted so that a less privileged user cannot get more
       access to sensitive files that a more privileged user made unavailable.
       In short the rule for proc and sysfs is as close to a bind mount as
       possible.

EXAMPLES
       The following command creates a PID namespace, using --fork to ensure
       that the executed command is performed in a child process that (being
       the first process in the namespace) has PID 1. The --mount-proc option
       ensures that a new mount namespace is also simultaneously created and
       that a new proc(5) filesystem is mounted that contains information
       corresponding to the new PID namespace. When the readlink(1) command
       terminates, the new namespaces are automatically torn down.

           # unshare --fork --pid --mount-proc readlink /proc/self
           1

       As an unprivileged user, create a new user namespace where the user's
       credentials are mapped to the root IDs inside the namespace:

           $ id -u; id -g
           1000
           1000
           $ unshare --user --map-root-user \
                   sh -c 'whoami; cat /proc/self/uid_map /proc/self/gid_map'
           root
                    0       1000          1
                    0       1000          1

       As an unprivileged user, create a user namespace where the first 65536
       IDs are all mapped, and the user's credentials are mapped to the root
       IDs inside the namespace. The map is determined by the subordinate IDs
       assigned in subuid(5) and subgid(5). Demonstrate this mapping by
       creating a file with user ID 1 and group ID 1. For brevity, only the
       user ID mappings are shown:

           $ id -u
           1000
           $ cat /etc/subuid
```

```
1000:100000:65536
$ unshare --user --map-auto --map-root-user
# id -u
0
# cat /proc/self/uid_map
         0       1000          1
         1     100000      65535
# touch file; chown 1:1 file
# ls -ln --time-style=+ file
-rw-r--r-- 1 1 1 0  file
# exit
$ ls -ln --time-style=+ file
-rw-r--r-- 1 100000 100000 0  file
```

The first of the following commands creates a new persistent UTS namespace and modifies the hostname as seen in that namespace. The namespace is then entered with nsenter(1) in order to display the modified hostname; this step demonstrates that the UTS namespace continues to exist even though the namespace had no member processes after the unshare command terminated. The namespace is then destroyed by removing the bind mount.

```
# touch /root/uts-ns
# unshare --uts=/root/uts-ns hostname FOO
# nsenter --uts=/root/uts-ns hostname
FOO
# umount /root/uts-ns
```

The following commands establish a persistent mount namespace referenced by the bind mount /root/namespaces/mnt. In order to ensure that the creation of that bind mount succeeds, the parent directory (/root/namespaces) is made a bind mount whose propagation type is not shared.

```
# mount --bind /root/namespaces /root/namespaces
# mount --make-private /root/namespaces
# touch /root/namespaces/mnt
# unshare --mount=/root/namespaces/mnt
```

The following commands demonstrate the use of the --kill-child option when creating a PID namespace, in order to ensure that when unshare is killed, all of the processes within the PID namespace are killed.

```
# set +m                # Don't print job status messages

# unshare --pid --fork --mount-proc --kill-child -- \
       bash --norc -c '(sleep 555 &) && (ps a &) && sleep 999' &
[1] 53456
#     PID TTY      STAT   TIME COMMAND
      1 pts/3    S+     0:00 sleep 999
      3 pts/3    S+     0:00 sleep 555
      5 pts/3    R+     0:00 ps a
```

```
            # ps h -o 'comm' $! # Show that background job is unshare(1)
            unshare
            # kill $! # Kill unshare(1)
            # pidof sleep
```

The pidof(1) command prints no output, because the sleep processes have been killed. More precisely, when the sleep process that has PID 1 in the namespace (i.e., the namespace's init process) was killed, this caused all other processes in the namespace to be killed. By contrast, a similar series of commands where the --kill-child option is not used shows that when unshare terminates, the processes in the PID namespace are not killed:

```
            # unshare --pid --fork --mount-proc -- \
                   bash --norc -c '(sleep 555 &) && (ps a &) && sleep 999' &
            [1] 53479
            #     PID TTY      STAT   TIME COMMAND
                  1 pts/3    S+     0:00 sleep 999
                  3 pts/3    S+     0:00 sleep 555
                  5 pts/3    R+     0:00 ps a

            # kill $!
            # pidof sleep
            53482 53480
```

The following example demonstrates the creation of a time namespace where the boottime clock is set to a point several years in the past:

```
            # uptime -p             # Show uptime in initial time namespace
            up 21 hours, 30 minutes
            # unshare --time --fork --boottime 300000000 uptime -p
            up 9 years, 28 weeks, 1 day, 2 hours, 50 minutes
```

## AUTHORS

Mikhail Gusarov <dottedmag@dottedmag.net>, Karel Zak <kzak@redhat.com>

## SEE ALSO

newuidmap(1), newgidmap(1), clone(2), unshare(2), namespaces(7), mount(8)

## REPORTING BUGS

For bug reports, use the issue tracker at https://github.com/util-linux/util-linux/issues.

## AVAILABILITY

The unshare command is part of the util-linux package which can be downloaded from Linux Kernel Archive <https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 users 命令

```
users [OPTION]... [FILE]
```

**功能**

查看设备上登录的所有用户。

**类型**

可执行文件（`/usr/bin/users`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 通过此文件获取登录的用户，默认为 `/var/run/utmp`；此参数常设为 `/var/log/wtmp`

## 示例

```
$ users
primers
```

## 推荐阅读

- Print login names of users currently logged in - GNU Coreutils
- users(1) — Linux manual page

## 手册

```
USERS(1)                        User Commands                        USERS(1)

NAME
       users  -  print the user names of users currently logged in to the cur-
       rent host

SYNOPSIS
       users [OPTION]... [FILE]
```

```
DESCRIPTION
       Output who is currently logged in according to FILE.  If  FILE  is  not
       specified, use /var/run/utmp.  /var/log/wtmp as FILE is common.

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Joseph Arceneaux and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       getent(1), who(1)

       Full documentation <https://www.gnu.org/software/coreutils/users>
       or available locally via: info '(coreutils) users invocation'

GNU coreutils 9.4                April 2024                          USERS(1)
```

# Bash 的 utmpdump 命令

```
utmpdump [OPTION]... FILENAME
```

**功能**

转储 UTMP 或 WTMP 文件（用户登录日志）。

**类型**

可执行文件（`/usr/bin/utmpdump`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-f`，`--follow` - 持续转储文件的追加内容
  - `-o`，`--output file` - 转储到 `file` 文件中，而不是打印到标准输出
  - `-r`，`--reverse` - 方向操作，将 `file` 反向写回 UTMP 或 WTMP 文件
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `FILENAME` - UTMP 或 WTMP 文件

## 示例

```
$ utmpdump -f /var/log/wtmp     # 持续监控 wtmp
Utmp dump of /var/log/wtmp
[7] [2760293] [ts/5] [primers ] [pts/5       ] [192.168.1.100       ] [192.168.1.100  ]
[2026-03-02T18:10:57,185538+00:00]
[8] [2347153] [    ] [        ] [pts/1       ] [                    ] [0.0.0.0        ]
[2026-03-03T15:14:54,064908+00:00]
[7] [3040115] [ts/1] [primers ] [pts/1       ] [192.168.1.100       ] [192.168.1.100  ]
[2026-03-03T15:14:57,605562+00:00]
```

## 推荐阅读

- utmpdump(1) — Linux manual page

## 手册

NAME

utmpdump - dump UTMP and WTMP files in raw format

SYNOPSIS

utmpdump [options] filename

DESCRIPTION

utmpdump is a simple program to dump UTMP and WTMP files in raw format, so they can be examined. utmpdump reads from stdin unless a filename is passed.

OPTIONS

-f, --follow
Output appended data as the file grows.

-o, --output file
Write command output to file instead of standard output.

-r, --reverse
Undump, write back edited login information into the utmp or wtmp files.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

NOTES

utmpdump can be useful in cases of corrupted utmp or wtmp entries. It can dump out utmp/wtmp to an ASCII file, which can then be edited to remove bogus entries, and reintegrated using:

```
utmpdump -r < ascii_file > wtmp
```

But be warned, utmpdump was written for debugging purposes only.

File formats
Only the binary version of the utmp(5) is standardised. Textual dumps may become incompatible in future.

The version 2.28 was the last one that printed text output using ctime(3) timestamp format. Newer dumps use millisecond precision ISO-8601 timestamp format in UTC-0 timezone. Conversion from former timestamp format can be made to binary, although attempt to do so can lead the timestamps to drift amount of timezone offset.

BUGS

You may not use the -r option, as the format for the utmp/wtmp files strongly depends on the input format. This tool was not written for

```
       normal use, but for debugging only.

AUTHORS
       Michael Krapp

SEE ALSO
       last(1), w(1), who(1), utmp(5)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The utmpdump command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                       UTMPDUMP(1)
```

# Bash 的 vdir 命令

```
vdir [OPTION]... [FILE]...
```

**功能**

详细列出目录内容。

> 和 ls 命令等价，区别在于默认开启 `-l` 选项。

**类型**

可执行文件（`/usr/bin/vdir`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 显示隐藏文件（文件名以 `.` 开头）
  - `-A`，`--almost-all` - 显示隐藏文件，但不显示 `.` 和 `..`
  - `--author` - 显示文件的作者；和 `-l` 选项一起使用
  - `-b`，`--escape` - 非 图形字符 显示为 C 风格的 转义字符
  - `--block-size=SIZE` - 文件大小的单位；和 `-l` 选项一起使用
  - `-B`，`--ignore-backups` 不显示备份文件（文件名以 `~` 结束）
  - `-c` - 按 ctime（文件状态信息最后更新的时间）排序；和 `-lt` 选项一起使用
  - `-C` - 竖向排列
  - `--color[=WHEN]` - 何时使用彩色，有效值：`always`，`auto`，`never`
  - `-d`，`--directory` - 查看目录本身，而不是目录的内容
  - `-D`，`--dired` - 生成为 Emacs 的 dired 模式设计的输出
  - `-f` - 显示所有文件且不进行排序（按照目录顺序显示）
  - `-F`，`--classify[=WHEN]` - 何时添加 文件类型指示器，有效值：`always`，`auto`，`never`
  - `--file-type` - 显示文件类型指示器；同 `--classify` 但是不添加 `*`
  - `--format=WORD` - 格式，有效值：`across`，`commas`，`horizontal`，`long`，`single-column`，`verbose`，`vertical`
  - `--full-time` - 显示完整的时间；同 `-l --time-style=full-iso`
  - `-g` - 显示文件详情；同 `-l` 但不显示所有者
  - `--group-directories-first` - 分组时将目录排在文件前面
  - `-G`，`--no-group` - 不显示分组名
  - `-h`，`--human-readable` - 让详细信息更适合人类阅读；和 `-l` 或 `-s` 选择一起使用

- `--si` - 让详细信息更适合人类阅读；同 `-h` 但是文件大小以 1000 为单位，而不是 1024
- `-H`，`--dereference-command-line` - 参数是符号链接时，查看源文件
- `--dereference-command-line-symlink-to-dir` - 参数是目录的符号链接时，查看源目录
- `--hide=PATTERN` - 隐藏匹配的文件
- `--hyperlink[=WHEN]` - 何时将文件名显示为超链接（现代终端上点击可以打开文件管理器），有效值：`always`，`auto`，`never`
- `--indicator-style=WORD` - 添加 文件类型指示器 ，有效值：`none`，`slash`，`file-type`，`classify`
- `-i`，`--inode` - 显示文件的 inode 编号
- `-I`，`--ignore=PATTERN` - 不显示匹配的文件
- `-k`，`--kibibytes` - 文件大小以 1024 字节为单位
- `-l` - 显示详细信息
- `-L`，`--dereference` - 符号链接显示源文件的信息，而非符号链接本身的信息
- `-m` - 用逗号分隔条目
- `-n`，`--numeric-uid-gid` - 显示详细信息，但是所有者和所属组显示为 ID 而非名称
- `-N`，`--literal` - 含特殊字符的文件名不添加单引号
- `-o` - 显示详细信息，但不显示所属组信息
- `-p`，`--indicator-style=slash` - 添加 `/` 指示器
- `-q`，`--hide-control-chars` - 非 图形字符 显示为 `?`
- `--show-control-chars` - 直接显示非 图形字符
- `-Q`，`--quote-name` - 用双引号包裹文件名
- `--quoting-style=WORD` - 引号风格，有效值：`literal`，`locale`，`shell`，`shell-always`，`shell-escape`，`shell-escape-always`，`c`，`escape`
- `-r`，`--reverse` - 逆序
- `-R`，`--recursive` - 递归查看子目录
- `-s`，`--size` - 显示文件分配的块数
- `-S` - 按文件大小排序，最大的优先
- `--sort=WORD` - 排序规则，有效值：`none`，`size`，`time`，`version`，`extension`，`width`
- `--time=WORD` - 显示或排序的时间，有效值：`atime`，`ctime`，`mtime`，`btime`
- `--time-style=TIME_STYLE` - 日期和时间的格式
- `-t` - 按时间排序
- `-T`，`--tabsize=COLS` - 制表符的宽度
- `-u` - 显示 atime 或按 atime 排序
- `-U` - 不排序
- `-v` - 自然排序
- `-w`，`--width=COLS` - 一行的宽度，0 表示无限
- `-x` - 横向排列
- `-X` - 按扩展名排序

  - `-Z`, `--context` - 显示安全上下文
  - `--zero` - 以空字符（`\0`）作为行的结尾，而不是换行符（`\n`）
  - `-1` - 每行只显示一个文件
  - `--help` - 显示当前帮助
  - `--version` - 显示版本
- `FILE` 文件，可以省略，默认为当前工作目录

**返回值**

- `0`：成功
- `1`：失败，小问题
- `2`：失败，严重问题

## 文件类型指示器

| 符号 | 含义 | 示例 |
| --- | --- | --- |
| `/` | 目录（directory） | `src/` |
| `@` | 符号链接（symbolic link） | `readme@` |
| `*` | 可执行文件（executable file） | `run.sh*` |
| `=` | UNIX 域套接字（socket） | `mysock=` |
| `\|` | 命名管道（FIFO） | `myfifo\|` |
| `>` | Solaris Door File | `door>` |

## 推荐阅读

- Verbosely list directory contents - GNU Coreutils
- vdir(1) — Linux manual page

## 手册

```
VDIR(1)                          User Commands                          VDIR(1)

NAME
       vdir - list directory contents

SYNOPSIS
       vdir [OPTION]... [FILE]...

DESCRIPTION
```

```
List  information  about  the FILEs (the current directory by default).
Sort entries alphabetically if none of -cftuvSUX nor --sort  is  speci-
fied.

Mandatory  arguments  to  long  options are mandatory for short options
too.

-a, --all
       do not ignore entries starting with .

-A, --almost-all
       do not list implied . and ..

--author
       with -l, print the author of each file

-b, --escape
       print C-style escapes for nongraphic characters

--block-size=SIZE
       with  -l,  scale  sizes  by  SIZE  when  printing  them;  e.g.,
       '--block-size=M'; see SIZE format below

-B, --ignore-backups
       do not list implied entries ending with ~

-c     with  -lt: sort by, and show, ctime (time of last change of file
       status information); with -l: show ctime and sort by name;  oth-
       erwise: sort by ctime, newest first

-C     list entries by columns

--color[=WHEN]
       color the output WHEN; more info below

-d, --directory
       list directories themselves, not their contents

-D, --dired
       generate output designed for Emacs' dired mode

-f     list all entries in directory order

-F, --classify[=WHEN]
       append indicator (one of */=>@|) to entries WHEN

--file-type
       likewise, except do not append '*'

--format=WORD
       across  -x, commas -m, horizontal -x, long -l, single-column -1,
       verbose -l, vertical -C
```

```
--full-time
       like -l --time-style=full-iso

-g     like -l, but do not list owner

--group-directories-first
       group directories before files; can be augmented with  a  --sort
       option, but any use of --sort=none (-U) disables grouping

-G, --no-group
       in a long listing, don't print group names

-h, --human-readable
       with -l and -s, print sizes like 1K 234M 2G etc.

--si   likewise, but use powers of 1000 not 1024

-H, --dereference-command-line
       follow symbolic links listed on the command line

--dereference-command-line-symlink-to-dir
       follow  each  command line symbolic link that points to a direc-
       tory

--hide=PATTERN
       do not list implied entries matching shell  PATTERN  (overridden
       by -a or -A)

--hyperlink[=WHEN]
       hyperlink file names WHEN

--indicator-style=WORD
       append indicator with style WORD to entry names: none (default),
       slash (-p), file-type (--file-type), classify (-F)

-i, --inode
       print the index number of each file

-I, --ignore=PATTERN
       do not list implied entries matching shell PATTERN

-k, --kibibytes
       default  to  1024-byte  blocks  for file system usage; used only
       with -s and per directory totals

-l     use a long listing format

-L, --dereference
       when showing file information for a symbolic link, show informa-
       tion for the file the link references rather than for  the  link
       itself

-m     fill width with a comma separated list of entries
```

```
-n, --numeric-uid-gid
       like -l, but list numeric user and group IDs

-N, --literal
       print entry names without quoting

-o     like -l, but do not list group information

-p, --indicator-style=slash
       append / indicator to directories

-q, --hide-control-chars
       print ? instead of nongraphic characters

--show-control-chars
       show nongraphic characters as-is (the default, unless program is
       'ls' and output is a terminal)

-Q, --quote-name
       enclose entry names in double quotes

--quoting-style=WORD
       use  quoting style WORD for entry names: literal, locale, shell,
       shell-always,  shell-escape,  shell-escape-always,   c,   escape
       (overrides QUOTING_STYLE environment variable)

-r, --reverse
       reverse order while sorting

-R, --recursive
       list subdirectories recursively

-s, --size
       print the allocated size of each file, in blocks

-S     sort by file size, largest first

--sort=WORD
       sort  by  WORD instead of name: none (-U), size (-S), time (-t),
       version (-v), extension (-X), width

--time=WORD
       select which timestamp used to  display  or  sort;  access  time
       (-u): atime, access, use; metadata change time (-c): ctime, sta-
       tus;  modified  time (default): mtime, modification; birth time:
       birth, creation;

       with -l, WORD determines which time to show;  with  --sort=time,
       sort by WORD (newest first)

--time-style=TIME_STYLE
       time/date format with -l; see TIME_STYLE below
```

```
-t     sort by time, newest first; see --time

-T, --tabsize=COLS
       assume tab stops at each COLS instead of 8

-u     with  -lt:  sort by, and show, access time; with -l: show access
       time and sort by name; otherwise: sort by  access  time,  newest
       first

-U     do not sort; list entries in directory order

-v     natural sort of (version) numbers within text

-w, --width=COLS
       set output width to COLS.  0 means no limit

-x     list entries by lines instead of by columns

-X     sort alphabetically by entry extension

-Z, --context
       print any security context of each file

--zero end each output line with NUL, not newline

-1     list one file per line

--help display this help and exit

--version
       output version information and exit

The  SIZE  argument  is  an  integer and optional unit (example: 10K is
10*1024).  Units are K,M,G,T,P,E,Z,Y,R,Q (powers of 1024) or  KB,MB,...
(powers  of 1000).  Binary prefixes can be used, too: KiB=K, MiB=M, and
so on.

The TIME_STYLE argument can be  full-iso,  long-iso,  iso,  locale,  or
+FORMAT.   FORMAT  is  interpreted  like in date(1).  If FORMAT is FOR-
MAT1<newline>FORMAT2, then FORMAT1 applies to non-recent files and FOR-
MAT2 to recent files.  TIME_STYLE prefixed with 'posix-'  takes  effect
only  outside  the POSIX locale.  Also the TIME_STYLE environment vari-
able sets the default style to use.

The WHEN argument defaults to  'always'  and  can  also  be  'auto'  or
'never'.

Using  color  to distinguish file types is disabled both by default and
with --color=never.  With --color=auto, ls emits color codes only  when
standard  output is connected to a terminal.  The LS_COLORS environment
variable can change the settings.  Use the dircolors(1) command to  set
it.
```

```
   Exit status:
       0      if OK,

       1      if minor problems (e.g., cannot access subdirectory),

       2      if serious trouble (e.g., cannot access command-line argument).

AUTHOR
       Written by Richard M. Stallman and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/vdir>
       or available locally via: info '(coreutils) vdir invocation'

GNU coreutils 9.4                April 2024                          VDIR(1)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 wc 命令

```
wc [OPTION]... [FILE]...
```

**功能**

计算文件的行数、单词数和字节数。

**类型**

可执行文件（`/usr/bin/wc`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-c`，`--bytes` - 打印字节数
  - `-m`，`--chars` - 打印字符数
  - `-l`，`--lines` - 打印行数
  - `--files0-from=F` - 从文件 `F` 中读取文件列表作为输入，，每个文件名以空字符（`\0`）结尾；如果 `F` 是 `-` 则从标准输入读取文件列表
  - `-L`，`--max-line-length` - 打印的最大宽度
  - `-w`，`--words` - 打印单词数
  - `--total=WHEN` - 何时打印总数行：
    - `auto` - 自动
    - `always` - 总是打印总数行
    - `only` - 仅打印总数行，不打印各文件行
    - `never` - 从不打印总数行
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `FILE` - 文件列表；如果没有这个参数或指定为 `-`，则读取标准输入

## 示例

```
$ cat 1.txt
Primers 编程伙伴
https://xplanc.org/primers/
$ wc 1.txt
 2  3 49 1.txt
$ wc 1.txt 2.txt --total=always     # 打印总数行
 2  3 49 1.txt
```

```
 3  3  6 2.txt
 5  6 55 total
$ wc 1.txt 2.txt --total=only       # 仅打印总数行
 5  6 55 total
```

## 推荐阅读

- Print newline, word, and byte counts - GNU Coreutils
- wc(1) — Linux manual page

## 手册

```
WC(1)                            User Commands                            WC(1)

NAME
       wc - print newline, word, and byte counts for each file

SYNOPSIS
       wc [OPTION]... [FILE]...
       wc [OPTION]... --files0-from=F

DESCRIPTION
       Print newline, word, and byte counts for each FILE, and a total line if
       more  than one FILE is specified.  A word is a non-zero-length sequence
       of printable characters delimited by white space.

       With no FILE, or when FILE is -, read standard input.

       The options below may be used to select which counts are  printed,  al-
       ways  in  the  following order: newline, word, character, byte, maximum
       line length.

       -c, --bytes
              print the byte counts

       -m, --chars
              print the character counts

       -l, --lines
              print the newline counts

       --files0-from=F
              read input from the files specified by NUL-terminated  names  in
              file F; If F is - then read names from standard input

       -L, --max-line-length
              print the maximum display width
```

```
       -w, --words
              print the word counts

       --total=WHEN
              when  to  print a line with total counts; WHEN can be: auto, al-
              ways, only, never

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Paul Rubin and David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/wc>
       or available locally via: info '(coreutils) wc invocation'

GNU coreutils 9.4                April 2024                              WC(1)
```

# Bash 的 wdctl 命令

```
wdctl [OPTION]... [DEVICE]...
```

**功能**

查看硬件看门狗的状态。

**类型**

可执行文件（`/usr/bin/wdctl`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-f`，`--flags list` - 仅显示指定的看门狗标志
  - `-F`，`--noflags` - 不显示看门狗标志信息
  - `-I`，`--noident` - 不显示看门狗身份信息
  - `-n`，`--noheadings` - 不显示标题行（列名）
  - `-o`，`--output list` - 只显示指定的列
  - `-O`，`--oneline` - 以单行键值对格式输出所有信息
  - `-p`，`--setpretimeout seconds` - 设置看门狗预超时时间
  - `-g`，`--setpregovernor governor` - 设置预超时调节器
  - `-r`，`--raw` - 以原始格式输出
  - `-s`，`--settimeout seconds` - 设置看门狗超时时间
  - `-T`，`--notimeouts` - 不显示看门狗超时信息
  - `-x`，`--flags-only` - 只显示标志信息；等同于 `-I -T`
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 设备列表

## 列名

| 列名 | 说明 |
| --- | --- |
| `FLAG` | 标志名 |
| `DESCRIPTION` | 标志的描述 |

| 列名 | 说明 |
| --- | --- |
| `STATUS` | 标志状态 |
| `BOOT-STATUS` | 重启标志状态 |
| `DEVICE` | 看门狗设备名 |

## 示例

```
$ wdctl
Device:        /dev/watchdog
Identity:      iTCO_wdt [version 0]
Timeout:       30 seconds
Timeleft:      29 seconds
FLAG           DESCRIPTION               STATUS BOOT-STATUS
KEEPALIVEPING  Keep alive ping reply          1           0
MAGICCLOSE     Supports magic close char      0           0
```

## 推荐阅读

- wdctl(8) — Linux manual page

## 手册

```
WDCTL(8)                    System Administration                    WDCTL(8)

NAME
       wdctl - show hardware watchdog status

SYNOPSIS
       wdctl [options] [device...]

DESCRIPTION
       Show hardware watchdog status. The default device is /dev/watchdog. If
       more than one device is specified then the output is separated by one
       blank line.

       If the device is already used or user has no permissions to read from
       the device, then wdctl reads data from sysfs. In this case information
       about supported features (flags) might be missing.

       Note that the number of supported watchdog features is hardware
       specific.

OPTIONS
       -f, --flags list
```

Print only the specified flags.

-F, --noflags
Do not print information about flags.

-I, --noident
Do not print watchdog identity information.

-n, --noheadings
Do not print a header line for flags table.

-o, --output list
Define the output columns to use in table of watchdog flags. If no output arrangement is specified, then a default set is used. Use --help to get list of all supported columns.

-O, --oneline
Print all wanted information on one line in key="value" output format.

-p, --setpretimeout seconds
Set the watchdog pre-timeout in seconds. A watchdog pre-timeout is a notification generated by the watchdog before the watchdog reset might occur in the event the watchdog has not been serviced. This notification is handled by the kernel and can be configured to take an action using sysfs or by --setpregovernor.

-g, --setpregovernor governor
Set pre-timeout governor name. For available governors see default wdctl output.

-r, --raw
Use the raw output format.

-s, --settimeout seconds
Set the watchdog timeout in seconds.

-T, --notimeouts
Do not print watchdog timeouts.

-x, --flags-only
Same as -I -T.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

AUTHORS
Karel Zak <kzak@redhat.com>, Lennart Poettering
<lennart@poettering.net>

```
REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The wdctl command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          WDCTL(8)
```

# Bash 的 whereis 命令

```
whereis [OPTION]... NAME...
```

**功能**

定位二进制文件、源码文件和手册文件的路径。

**类型**

可执行文件（`/usr/bin/whereis`），属于 util-linux。

**参数**

- `OPTION` 选项：
  - `-b` - 查找二进制文件
  - `-m` - 查找手册文件
  - `-s` - 查找源码文件
  - `-u` - 仅显示居然异常条目的命令名称
  - `-B list` - 指定搜索二进制文件的目录列表；以空格分隔条目
  - `-M list` - 指定搜索手册文件的目录列表；以空格分隔条目
  - `-S list` - 指定搜索源码文件的目录列表；以空格分隔条目
  - `-f` - 终止目录列表；在 `-B` `-M` `-S` 选项后必须使用此选项结束
  - `-l` - 显示此命令的搜索目录列表
  - `-g` - 将接下来的名称解释为 glob(7) 模式
  - `-h`, `--help` - 显示帮助
  - `-V`, `--version` - 显示版本
- `NAME` - 要定位的文件列表

## 示例

**定位文件**

```
$ whereis whereis
whereis: /usr/bin/whereis /usr/share/man/man1/whereis.1.gz
```

**查看搜索路径**

```
$ whereis -l        # 查看 whereis 的默认搜索路径
bin: /usr/bin
bin: /usr/sbin
bin: /usr/lib/x86_64-linux-gnu
bin: /usr/lib
bin: /usr/lib32
bin: /usr/lib64
bin: /etc
bin: /usr/games
bin: /usr/local/bin
bin: /usr/local/sbin
bin: /usr/local/etc
bin: /usr/local/lib
bin: /usr/local/games
bin: /usr/include
bin: /usr/local
bin: /usr/libexec
bin: /usr/share
bin: /opt/qt-6.7.1/bin
bin: /opt/qt-wasm-6.7.1/bin
bin: /home/planc/.cargo/bin
man: /usr/share/man/zh_CN
man: /usr/share/man/man1
man: /usr/share/man/hu
man: /usr/share/man/zh_TW
man: /usr/share/man/id
man: /usr/share/man/man3
man: /usr/share/man/da
man: /usr/share/man/vi
man: /usr/share/man/sr
man: /usr/share/man/man8
man: /usr/share/man/man4
man: /usr/share/man/ru
man: /usr/share/man/man6
man: /usr/share/man/sv
man: /usr/share/man/man5
man: /usr/share/man/man2
man: /usr/share/man/pt
man: /usr/share/man/man9
man: /usr/share/man/sl
man: /usr/share/man/pt_BR
man: /usr/share/man/cs
man: /usr/share/man/ko
man: /usr/share/man/uk
man: /usr/share/man/tr
man: /usr/share/man/ja
man: /usr/share/man/fr
man: /usr/share/man/es
man: /usr/share/man/man7
man: /usr/share/man/nl
man: /usr/share/man/de
man: /usr/share/man/pl
```

```
man: /usr/share/man/fi
man: /usr/share/man/ro
man: /usr/share/man/hr
man: /usr/share/man/it
man: /usr/share/info
src: /usr/src/linux-headers-6.8.0-101-generic
src: /usr/src/linux-headers-6.8.0-87-generic
src: /usr/src/linux-headers-6.8.0-87
src: /usr/src/linux-headers-6.8.0-101
src: /usr/src/python3.12
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| which | 定位 PATH 环境变量中命令的路径 |

## 推荐阅读

- whereis(1) — Linux manual page

## 手册

```
WHEREIS(1)                       User Commands                       WHEREIS(1)

NAME
       whereis - locate the binary, source, and manual page files for a
       command

SYNOPSIS
       whereis [options] [-BMS directory... -f] name...

DESCRIPTION
       whereis locates the binary, source and manual files for the specified
       command names. The supplied names are first stripped of leading
       pathname components. Prefixes of s. resulting from use of source code
       control are also dealt with. whereis then attempts to locate the
       desired program in the standard Linux places, and in the places
       specified by $PATH and $MANPATH.

       The search restrictions (options -b, -m and -s) are cumulative and
       apply to the subsequent name patterns on the command line. Any new
       search restriction resets the search mask. For example,

          whereis -bm ls tr -m gcc

       searches for "ls" and "tr" binaries and man pages, and for "gcc" man
```

```
       pages only.

       The options -B, -M and -S reset search paths for the subsequent name
       patterns. For example,

          whereis -m ls -M /usr/share/man/man1 -f cal

       searches for "ls" man pages in all default paths, but for "cal" in the
       /usr/share/man/man1 directory only.

OPTIONS
       -b
           Search for binaries.

       -m
           Search for manuals.

       -s
           Search for sources.

       -u
           Only show the command names that have unusual entries. A command is
           said to be unusual if it does not have just one entry of each
           explicitly requested type. Thus 'whereis -m -u *' asks for those
           files in the current directory which have no documentation file, or
           more than one.

       -B list
           Limit the places where whereis searches for binaries, by a
           whitespace-separated list of directories.

       -M list
           Limit the places where whereis searches for manuals and
           documentation in Info format, by a whitespace-separated list of
           directories.

       -S list
           Limit the places where whereis searches for sources, by a
           whitespace-separated list of directories.

       -f
           Terminates the directory list and signals the start of filenames.
           It must be used when any of the -B, -M, or -S options is used.

       -l
           Output the list of effective lookup paths that whereis is using.
           When none of -B, -M, or -S is specified, the option will output the
           hard-coded paths that the command was able to find on the system.

       -g
           Interpret the next names as a glob(7) patterns. whereis always
           compares only filenames (aka basename) and never complete path.
           Using directory names in the pattern has no effect. Don’t forget
```

```
           that the shell interprets the pattern when specified on the command
           line without quotes. It’s necessary to use quotes for the name, for
           example:

              whereis -g 'find*'

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

FILE SEARCH PATHS
       By default whereis tries to find files from hard-coded paths, which are
       defined with glob patterns. The command attempts to use the contents of
       $PATH and $MANPATH environment variables as default search path. The
       easiest way to know what paths are in use is to add the -l listing
       option. Effects of the -B, -M, and -S are displayed with -l.

ENVIRONMENT
       WHEREIS_DEBUG=all
           enables debug output.

EXAMPLES
       To find all files in /usr/bin which are not documented in /usr/man/man1
       or have no source in /usr/src:

          cd /usr/bin
          whereis -u -ms -M /usr/man/man1 -S /usr/src -f *

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The whereis command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                        WHEREIS(1)
```

Primers 编程伙伴 : https://xplanc.org/primers

# Bash 的 who 命令

```
who [OPTION]... [ FILE | ARG1 ARG2 ]
```

**功能**

查看当前登录的用户。

**类型**

可执行文件（`/usr/bin/who`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `-a`，`--all` - 相当于 `-b -d --login -p -r -t -T -u`
  - `-b`，`--boot` - 打印登录以来的时间
  - `-d`，`--dead` - 打印死亡进程
  - `-H`，`--heading` - 打印列标题行
  - `-l`，`--login` - 打印系统登录进程
  - `--lookup` - 尝试通过 DNS 规范化主机名
  - `-m` - 仅查看与标准输入关联的主机名和用户
  - `-p`，`--process` - 打印由 init 进程启动的活动进程
  - `-q`，`--count` - 打印登录名和登录用户的数量
  - `-r`，`--runlevel` - 打印当前的 运行级别
  - `-s`，`--short` - 打印短格式（默认）
  - `-t`，`--time` - 打印上次系统时钟的变化
  - `-T`，`-w`，`--mesg` - 添加用户消息状态
  - `-u`，`--users` - 列出已登录用户
  - `--message` - 同 `-T`
  - `--writable` - 同 `-T`
  - `--help` - 显示帮助
  - `--version` - 显示版本

## SysV init 运行级别

| 运行级别 | 名称 | 描述 |
| --- | --- | --- |
| 0 | halt | 关机（系统关闭电源） |
| 1 | single-user | 单用户模式（维护模式，无网络，多为 root） |
| 2 | multi-user | 多用户模式，不带网络服务（Debian/Ubuntu 部分配置） |
| 3 | multi-user | 多用户模式，带网络，纯命令行模式（文本模式） |
| 4 | unused | 用户自定义，可自由配置 |
| 5 | multi-user | 多用户模式，带图形界面（GUI） |
| 6 | reboot | 重启系统 |

## 示例

```
$ who
primers    pts/6        2025-11-29 08:35 (192.168.1.100)
$ who -a
           system boot  2025-11-18 07:14
LOGIN      tty1         2025-11-18 07:15              1086 id=tty1
           run-level 5  2025-11-18 07:15
           pts/0        2025-11-23 15:11            129498 id=ts/0  term=0 exit=0
           pts/8        2025-11-18 13:19              8743 id=ts/8  term=0 exit=0
           pts/3        2025-11-22 11:57             56288 id=ts/3  term=0 exit=0
           pts/4        2025-11-19 14:38             59093 id=ts/4  term=0 exit=0
           pts/11       2025-11-26 13:59            210827 id=s/11  term=0 exit=0
           pts/9        2025-11-23 15:11            186011 id=ts/9  term=0 exit=0
           pts/10       2025-11-26 13:59            210640 id=s/10  term=0 exit=0
           pts/12       2025-11-24 15:27            211297 id=s/12  term=0 exit=125
primers  + pts/6        2025-11-29 08:35 00:37      283090 (192.168.1.100)
```

## 推荐阅读

- Print who is currently logged in - GNU Coreutils
- who(1) — Linux manual page

## 手册

```
WHO(1)                          User Commands                          WHO(1)

NAME
       who - show who is logged on

SYNOPSIS
```

```
       who [OPTION]... [ FILE | ARG1 ARG2 ]

DESCRIPTION
       Print information about users who are currently logged in.

       -a, --all
              same as -b -d --login -p -r -t -T -u

       -b, --boot
              time of last system boot

       -d, --dead
              print dead processes

       -H, --heading
              print line of column headings

       -l, --login
              print system login processes

       --lookup
              attempt to canonicalize hostnames via DNS

       -m     only hostname and user associated with stdin

       -p, --process
              print active processes spawned by init

       -q, --count
              all login names and number of users logged on

       -r, --runlevel
              print current runlevel

       -s, --short
              print only name, line, and time (default)

       -t, --time
              print last system clock change

       -T, -w, --mesg
              add user's message status as +, - or ?

       -u, --users
              list users logged in

       --message
              same as -T

       --writable
              same as -T

       --help display this help and exit
```

```
       --version
              output version information and exit

       If  FILE is not specified, use /var/run/utmp.  /var/log/wtmp as FILE is
       common.  If ARG1 ARG2 given, -m presumed: 'am i'  or  'mom  likes'  are
       usual.

AUTHOR
       Written by Joseph Arceneaux, David MacKenzie, and Michael Stone.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/who>
       or available locally via: info '(coreutils) who invocation'

GNU coreutils 9.4                April 2024                           WHO(1)
```

# Bash 的 whoami 命令

```
whoami [OPTION]...
```

**功能**

查看有效用户名。

**类型**

可执行文件（`/usr/bin/whoami`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本

## 示例

```
$ whoami
primers
```

## 推荐阅读

- Print effective user name - GNU Coreutils
- whoami(1) — Linux manual page

## 手册

```
WHOAMI(1)                          User Commands                          WHOAMI(1)

NAME
       whoami - print effective user name

SYNOPSIS
       whoami [OPTION]...

DESCRIPTION
       Print  the  user  name  associated  with the current effective user ID.
```

```
       Same as id -un.

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by Richard Mlynarik.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright © 2023 Free Software Foundation, Inc.   License  GPLv3+:  GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This  is  free  software:  you  are free to change and redistribute it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/whoami>
       or available locally via: info '(coreutils) whoami invocation'

GNU coreutils 9.4                April 2024                        WHOAMI(1)
```

# Bash 的 wipefs 命令

```
wipefs [OPTION]... DEVICE
```

**功能**

擦除设备的签名。

**类型**

可执行文件（`/usr/sbin/wipefs`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-a`，`--all` - 擦除所有签名
  - `-b`，`--backup` - 签名备份到 `$HOME/wipefs--.bak`
  - `-f`，`--force` - 强制擦除
  - `-J`，`--json` - 以 JSON 格式输出
  - `--lock[=mode]` - 是否对操作的设备或文件使用独占 BSD 锁
    - `yes` - 对设备加锁，如果锁被占用，则阻塞等待；省略 `mode` 时默认为 `yes`
    - `no` - 不对设备加锁；没有 `--lock` 选项时默认不加锁
    - `nonblock` - 对设备加锁，如果锁被占用，则失败
  - `-i`，`--noheadings` - 不打印标题行（列名）
  - `-O`，`--output list` - 只打印指定的列；多个列名用逗号分隔
  - `-n`，`--no-act` - 不实际执行擦除操作
  - `-o`，`--offset offset` - 指定要从设备中删除的签名的位置
  - `-p`，`--parsable` - 以可解析格式输出
  - `-q`，`--quiet` - 不打印输出
  - `-t`，`--types list` - 限制打印或擦除签名的种类
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 要操作的设备

## 列名

| 列名 | 说明 |
| --- | --- |
| `UUID` | 文件系统的 UUID |
| `LABEL` | 文件系统的标签 |
| `LENGTH` | 魔法字符串的长度 |
| `TYPE` | 超级块的类型 |
| `OFFSET` | 魔法字符串的偏移量 |
| `USAGE` | 类型描述 |
| `DEVICE` | 块设备名 |

## 示例

### 查看签名

```
$ sudo wipefs -n /dev/sda
DEVICE OFFSET TYPE UUID      LABEL
sda    0x52   vfat FA6B-635A
sda    0x0    vfat FA6B-635A
sda    0x1fe  vfat FA6B-635A
```

### 擦除签名

```
$ sudo wipefs -a /dev/sda
/dev/sda: 8 bytes were erased at offset 0x00000052 (vfat): 46 41 54 33 32 20 20 20
/dev/sda: 1 byte was erased at offset 0x00000000 (vfat): eb
/dev/sda: 2 bytes were erased at offset 0x000001fe (vfat): 55 aa
```

## 推荐阅读

- wipefs(8) — Linux manual page

## 手册

```
WIPEFS(8)                   System Administration                   WIPEFS(8)

NAME
       wipefs - wipe a signature from a device

SYNOPSIS
       wipefs [options] device...
```

```
       wipefs [--backup] -o offset device...

       wipefs [--backup] -a device...

DESCRIPTION
       wipefs can erase filesystem, raid or partition-table signatures (magic
       strings) from the specified device to make the signatures invisible for
       libblkid. wipefs does not erase the filesystem itself nor any other
       data from the device.

       When used without any options, wipefs lists all visible filesystems and
       the offsets of their basic signatures. The default output is subject to
       change. So whenever possible, you should avoid using default outputs in
       your scripts. Always explicitly define expected columns by using
       --output columns-list in environments where a stable output is
       required.

       wipefs calls the BLKRRPART ioctl when it has erased a partition-table
       signature to inform the kernel about the change. The ioctl is called as
       the last step and when all specified signatures from all specified
       devices are already erased. This feature can be used to wipe content on
       partitions devices as well as partition table on a disk device, for
       example by wipefs -a /dev/sdc1 /dev/sdc2 /dev/sdc.

       Note that some filesystems and some partition tables store more magic
       strings on the device (e.g., FAT, ZFS, GPT). The wipefs command (since
       v2.31) lists all the offset where a magic strings have been detected.

       When option -a is used, all magic strings that are visible for
       libblkid(3) are erased. In this case the wipefs scans the device again
       after each modification (erase) until no magic string is found.

       Note that by default wipefs does not erase nested partition tables on
       non-whole disk devices. For this the option --force is required.

OPTIONS
       -a, --all
           Erase all available signatures. The set of erased signatures can be
           restricted with the -t option.

       -b, --backup
           Create a signature backup to the file
           $HOME/wipefs-<devname>-<offset>.bak. For more details see the
           EXAMPLE section.

       -f, --force
           Force erasure, even if the filesystem is mounted. This is required
           in order to erase a partition-table signature on a block device.

       -J, --json
           Use JSON output format.
```

```
       --lock[=mode]
           Use exclusive BSD lock for device or file it operates. The optional
           argument mode can be yes, no (or 1 and 0) or nonblock. If the mode
           argument is omitted, it defaults to "yes". This option overwrites
           environment variable $LOCK_BLOCK_DEVICE. The default is not to use
           any lock at all, but it’s recommended to avoid collisions with
           udevd or other tools.

       -i, --noheadings
           Do not print a header line.

       -O, --output list
           Specify which output columns to print. Use --help to get a list of
           all supported columns.

       -n, --no-act
           Causes everything to be done except for the write(2) call.

       -o, --offset offset
           Specify the location (in bytes) of the signature which should be
           erased from the device. The offset number may include a "0x"
           prefix; then the number will be interpreted as a hex value. It is
           possible to specify multiple -o options.

           The offset argument may be followed by the multiplicative suffixes
           KiB (=1024), MiB (=1024*1024), and so on for GiB, TiB, PiB, EiB,
           ZiB and YiB (the "iB" is optional, e.g., "K" has the same meaning
           as "KiB"), or the suffixes KB (=1000), MB (=1000*1000), and so on
           for GB, TB, PB, EB, ZB and YB.

       -p, --parsable
           Print out in parsable instead of printable format. Encode all
           potentially unsafe characters of a string to the corresponding hex
           value prefixed by '\x'.

       -q, --quiet
           Suppress any messages after a successful signature wipe.

       -t, --types list
           Limit the set of printed or erased signatures. More than one type
           may be specified in a comma-separated list. The list or individual
           types can be prefixed with 'no' to specify the types on which no
           action should be taken. For more details see mount(8).

       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

ENVIRONMENT
       LIBBLKID_DEBUG=all
           enables libblkid(3) debug output.
```

```
       LOCK_BLOCK_DEVICE=<mode>
           use exclusive BSD lock. The mode is "1" or "0". See --lock for more
           details.

EXAMPLES
       wipefs /dev/sda*
           Prints information about sda and all partitions on sda.

       wipefs --all --backup /dev/sdb
           Erases all signatures from the device /dev/sdb and creates a
           signature backup file ~/wipefs-sdb-<offset>.bak for each signature.

       dd if=~/wipefs-sdb-0x00000438.bak of=/dev/sdb seek=$((0x00000438)) bs=1
       conv=notrunc
           Restores an ext2 signature from the backup file
           ~/wipefs-sdb-0x00000438.bak.

AUTHORS
       Karel Zak <kzak@redhat.com>

SEE ALSO
       blkid(8), findfs(8)

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The wipefs command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-10-23                          WIPEFS(8)
```

# Bash 的 x86_64 命令

```
x86_64 [OPTION]... [CMD [ARG]...]
```

**功能**

**类型**

可执行文件（`/usr/bin/x86_64`），属于 util-linux。

> 实际上是 `/usr/bin/setarch` 的符号链接，并默认将架构设为 `x86_64`。

**参数**

- `OPTION` 选项：
  - `--uname-2.6` - 使程序检测到的内核版本号为 `2.6`
  - `-v`，`--verbose` - 显示详细信息
  - `-3`，`--3gb` - 指定程序最多使用 3GB 的地址空间（3GB 用户空间，1GB 内核空间）
  - `--4gb` - ~~指定程序最多使用 4GB 的地址空间（4GB 用户空间，没有内核空间）~~；此选项已废弃不再生效
  - `-B`，`--32bit` - 将地址空间设为 32 位以模拟硬件；支持 ARM 和 Alpha 架构
  - `-F`，`--fdpic-funcptrs` - 将用户空间函数指向信号处理程序的指针视为指向地址描述符的指针；仅支持 ARM, Blackfin, Fujitsu FR-V 和 SuperH 架构
  - `-I`，`--short-inode` - ~~使用 16 位 inode 号~~；此选项已废弃不再生效
  - `-L`，`--addr-compat-layout` - 提供传统虚拟地址空间布局
  - `-R`，`--addr-no-randomize` - 禁用虚拟地址空间的随机化
  - `-S`，`--whole-seconds` - ~~时间值截断为整数秒~~；此选项已废弃不再生效
  - `-T`，`--sticky-timeouts` - 这使得 `select(2)`、`pselect(2)` 和 `ppoll(2)` 系统调用保留超时值
  - `-X`，`--read-implies-exec` - 如果设置此项，则 `mmap(2)` `PROT_READ` 还会添加 `PROT_EXEC`` 位
  - `-Z`，`--mmap-page-zero` - SVr4 错误模拟，会将 `mmap(2)` 第 0 页设置为只读。
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `CMD` - 要运行的命令
- `ARG` - 传递给 `CMD` 的参数列表

## 示例

```
$ x86_64 arch
x86_64
```

## 相关命令

| 命令 | 说明 |
| --- | --- |
| setarch | 使用指定的架构或个性化标志运行程序 |

## 推荐阅读

- setarch(8) — Linux manual page

## 手册

```
SETARCH(8)                  System Administration                  SETARCH(8)

NAME
       setarch - change reported architecture in new program environment
       and/or set personality flags

SYNOPSIS
       setarch [arch] [options] [program [argument...]]

       setarch --list|-h|-V

       arch [options] [program [argument...]]

DESCRIPTION
       setarch modifies execution domains and process personality flags.

       The execution domains currently only affects the output of uname -m.
       For example, on an AMD64 system, running setarch i386 program will
       cause program to see i686 instead of x86_64 as the machine type. It can
       also be used to set various personality options. The default program is
       /bin/sh.

       Since version 2.33 the arch command line argument is optional and
       setarch may be used to change personality flags (ADDR_LIMIT_*,
       SHORT_INODE, etc) without modification of the execution domain.

OPTIONS
       --list
           List the architectures that setarch knows about. Whether setarch
           can actually set each of these architectures depends on the running
           kernel.
```

```
--show[=personality]
    Show the currently active personality and flags. If the personality
    argument is provided, it is shown instead of the current one.
    personality is a hexadecimal number with values was described in
    sys/personality.h.

--uname-2.6
    Causes the program to see a kernel version number beginning with
    2.6. Turns on UNAME26.

-v, --verbose
    Be verbose.

-3, --3gb
    Specifies program should use a maximum of 3GB of address space.
    Supported on x86. Turns on ADDR_LIMIT_3GB.

--4gb
    This option has no effect. It is retained for backward
    compatibility only, and may be removed in future releases.

-B, --32bit
    Limit the address space to 32 bits to emulate hardware. Supported
    on ARM and Alpha. Turns on ADDR_LIMIT_32BIT.

-F, --fdpic-funcptrs
    Treat user-space function pointers to signal handlers as pointers
    to address descriptors. This option has no effect on architectures
    that do not support FDPIC ELF binaries. In kernel v4.14 support is
    limited to ARM, Blackfin, Fujitsu FR-V, and SuperH CPU
    architectures.

-I, --short-inode
    Obsolete bug emulation flag. Turns on SHORT_INODE.

-L, --addr-compat-layout
    Provide legacy virtual address space layout. Use when the program
    binary does not have PT_GNU_STACK ELF header. Turns on
    ADDR_COMPAT_LAYOUT.

-R, --addr-no-randomize
    Disables randomization of the virtual address space. Turns on
    ADDR_NO_RANDOMIZE.

-S, --whole-seconds
    Obsolete bug emulation flag. Turns on WHOLE_SECONDS.

-T, --sticky-timeouts
    This makes select(2), pselect(2), and ppoll(2) system calls
    preserve the timeout value instead of modifying it to reflect the
    amount of time not slept when interrupted by a signal handler. Use
    when program depends on this behavior. For more details see the
```

timeout description in select(2) manual page. Turns on
STICKY_TIMEOUTS.

-X, --read-implies-exec
If this is set then mmap(2) PROT_READ will also add the PROT_EXEC
bit - as expected by legacy x86 binaries. Notice that the ELF
loader will automatically set this bit when it encounters a legacy
binary. Turns on READ_IMPLIES_EXEC.

-Z, --mmap-page-zero
SVr4 bug emulation that will set mmap(2) page zero as read-only.
Use when program depends on this behavior, and the source code is
not available to be fixed. Turns on MMAP_PAGE_ZERO.

-h, --help
Display help text and exit.

-V, --version
Print version and exit.

EXAMPLE

```
setarch --addr-no-randomize mytestprog
setarch ppc32 rpmbuild --target=ppc --rebuild foo.src.rpm
setarch ppc32 -v -vL3 rpmbuild --target=ppc --rebuild bar.src.rpm
setarch ppc32 --32bit rpmbuild --target=ppc --rebuild foo.src.rpm
```

AUTHORS

Elliot Lee <sopwith@redhat.com>, Jindrich Novy <jnovy@redhat.com>,
Karel Zak <kzak@redhat.com>

SEE ALSO

personality(2), select(2)

REPORTING BUGS

For bug reports, use the issue tracker at
https://github.com/util-linux/util-linux/issues.

AVAILABILITY

The setarch command is part of the util-linux package which can be
downloaded from Linux Kernel Archive
<https://www.kernel.org/pub/linux/utils/util-linux/>.

# Bash 的 yes 命令

```
yes [STRING]...
```

**功能**

持续打印字符串 `STRING` 或 `"y"`，常用于自动确认。

**类型**

可执行文件（`/usr/bin/yes`），属于 coreutils。

**参数**

- `OPTION` 选项：
  - `--help` - 显示帮助
  - `--version` - 显示版本
- `STRING` - 要重复打印的字符串，默认为 `y`

## 示例

```
$ yes | apt install gcc g++ make cmake  # 自动确认
```

## 推荐阅读

- Print a string until interrupted - GNU Coreutils
- yes(1) — Linux manual page

## 手册

```
YES(1)                          User Commands                          YES(1)

NAME
       yes - output a string repeatedly until killed

SYNOPSIS
       yes [STRING]...
       yes OPTION

DESCRIPTION
```

```
       Repeatedly output a line with all specified STRING(s), or 'y'.

       --help display this help and exit

       --version
              output version information and exit

AUTHOR
       Written by David MacKenzie.

REPORTING BUGS
       GNU coreutils online help: <https://www.gnu.org/software/coreutils/>
       Report any translation bugs to <https://translationproject.org/team/>

COPYRIGHT
       Copyright  ©  2023  Free Software Foundation, Inc.  License GPLv3+: GNU
       GPL version 3 or later <https://gnu.org/licenses/gpl.html>.
       This is free software: you are free  to  change  and  redistribute  it.
       There is NO WARRANTY, to the extent permitted by law.

SEE ALSO
       Full documentation <https://www.gnu.org/software/coreutils/yes>
       or available locally via: info '(coreutils) yes invocation'

GNU coreutils 9.4                April 2024                           YES(1)
```

# Bash 的 zramctl 命令

```
zramctl [OPTION]... DEVICE
```

**功能**

设置和控制 ZRAM （压缩内存）设备。

**类型**

可执行文件（`/usr/sbin/zramctl`），属于 util-linux。

**参数**

- `OPTIONS` 选项：
  - `-a`，`--algorithm lzo|lz4|lz4hc|deflate|842|zstd` - 设置压缩算法
  - `-f`，`--find` - 找到第一个未使用的 ZRAM 设备
  - `-n`，`--noheadings` - 不打印标题行（列名）
  - `-o`，`--output list` - 只打印指定的列；多个列名用逗号分隔
  - `--output-all` - 打印所有列
  - `--raw` - 以原始格式输出
  - `-r`，`--reset` - 重置指定 ZRAM 设备的选项
  - `-s`，`--size size` - 创建指定大小的 ZRAM 设备
  - `-t`，`--streams number` - 设置设备可使用的最大压缩流数量
  - `-h`，`--help` - 显示帮助
  - `-V`，`--version` - 显示版本
- `DEVICE` - 要操作的设备

## 列名

| 列名 | 说明 |
| --- | --- |
| `NAME` | ZRAM 设备名 |
| `DISKSIZE` | 未压缩数据量的限制 |
| `DATA` | 存储数据的未压缩大小 |
| `COMPR` | 存储数据的压缩大小 |
| `ALGORITHM` | 压缩算法 |

| 列名 | 说明 |
| --- | --- |
| `STREAMS` | 并发压缩操作的数量 |
| `ZERO-PAGES` | 没有分配内存的空白页 |
| `TOTAL` | 总内存开销 |
| `MEM-LIMIT` | 用于存储压缩数据的内存限制 |
| `MEM-USED` | ZRAM 已被用于存储压缩数据的内存 |
| `MIGRATED` | 通过压缩迁移的对象数量 |
| `MOUNTPOINT` | 挂载点 |

## 推荐阅读

- zramctl(8) — Linux manual page

## 手册

```
ZRAMCTL(8)                  System Administration                  ZRAMCTL(8)

NAME
       zramctl - set up and control zram devices

SYNOPSIS
       Get info:
           zramctl [options]

       Reset zram:
           zramctl -r zramdev...

       Print name of first unused zram device:
           zramctl -f

       Set up a zram device:
           zramctl [-f | zramdev] [-s size] [-t number] [-a algorithm]

DESCRIPTION
       zramctl is used to quickly set up zram device parameters, to reset zram
       devices, and to query the status of used zram devices.

       If no option is given, all non-zero size zram devices are shown.

       Note that zramdev node specified on command line has to already exist.
       The command zramctl creates a new /dev/zram<N> nodes only when --find
       option specified. It’s possible (and common) that after system boot
```

```
       /dev/zram<N> nodes are not created yet.

OPTIONS
       -a, --algorithm lzo|lz4|lz4hc|deflate|842|zstd
           Set the compression algorithm to be used for compressing data in
           the zram device.

           The list of supported algorithms could be inaccurate as it depends
           on the current kernel configuration. A basic overview can be
           obtained by using the command "cat
           /sys/block/zram0/comp_algorithm"; however, please note that this
           list might also be incomplete. This is due to the fact that ZRAM
           utilizes the Crypto API, and if certain algorithms were built as
           modules, it becomes impossible to enumerate all of them.

       -f, --find
           Find the first unused zram device. If a --size argument is present,
           then initialize the device.

       -n, --noheadings
           Do not print a header line in status output.

       -o, --output list
           Define the status output columns to be used. If no output
           arrangement is specified, then a default set is used. Use --help to
           get a list of all supported columns.

       --output-all
           Output all available columns.

       --raw
           Use the raw format for status output.

       -r, --reset
           Reset the options of the specified zram device(s). Zram device
           settings can be changed only after a reset.

       -s, --size size
           Create a zram device of the specified size. Zram devices are
           aligned to memory pages; when the requested size is not a multiple
           of the page size, it will be rounded up to the next multiple. When
           not otherwise specified, the unit of the size parameter is bytes.

           The size argument may be followed by the multiplicative suffixes
           KiB (=1024), MiB (=1024*1024), and so on for GiB, TiB, PiB, EiB,
           ZiB and YiB (the "iB" is optional, e.g., "K" has the same meaning
           as "KiB") or the suffixes KB (=1000), MB (=1000*1000), and so on
           for GB, TB, PB, EB, ZB and YB.

       -t, --streams number
           Set the maximum number of compression streams that can be used for
           the device. The default is use all CPUs and one stream for kernels
           older than 4.6.
```

```
       -h, --help
           Display help text and exit.

       -V, --version
           Print version and exit.

EXIT STATUS
       zramctl returns 0 on success, nonzero on failure.

FILES
       /dev/zram[0..N]
           zram block devices

EXAMPLE
       The following commands set up a zram device with a size of one gigabyte
       and use it as swap device.

            # zramctl --find --size 1024M
            /dev/zram0
            # mkswap /dev/zram0
            # swapon /dev/zram0
            ...
            # swapoff /dev/zram0
            # zramctl --reset /dev/zram0

AUTHORS
       Timofey Titovets <nefelim4ag@gmail.com>, Karel Zak <kzak@redhat.com>

SEE ALSO
       Linux kernel documentation
       <https://docs.kernel.org/admin-guide/blockdev/zram.html>

REPORTING BUGS
       For bug reports, use the issue tracker at
       https://github.com/util-linux/util-linux/issues.

AVAILABILITY
       The zramctl command is part of the util-linux package which can be
       downloaded from Linux Kernel Archive
       <https://www.kernel.org/pub/linux/utils/util-linux/>.

util-linux 2.39.3                2023-11-21                        ZRAMCTL(8)
```