

NAME

htop - interactive process viewer

SYNOPSIS

htop [*-dChusv*]

DESCRIPTION

Htop is a free (GPL) ncurses-based process viewer for Linux.

It is similar to top, but allows you to scroll vertically and horizontally, so you can see all the processes running on the system, along with their full command lines.

Tasks related to processes (killing, renicing) can be done without entering their PIDs.

COMMAND-LINE OPTIONS

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

-d --delay=DELAY

Delay between updates, in tenths of seconds

-C --no-color --no-colour

Start htop in monochrome mode

-h --help

Display a help message and exit

-p --pid=PID,PID...

Show only the given PIDs

-s --sort-key COLUMN

Sort by this column (use --sort-key help for a column list)

-u --user=USERNAME

Show only the processes of a given user

-v --version

Output version information and exit

INTERACTIVE COMMANDS

The following commands are supported while in htop:

Arrows, PgUP, PgDn, Home, End

Scroll the process list.

Space

Tag or untag a process. Commands that can operate on multiple processes, like kill, will then apply over the list of tagged processes, instead of the currently highlighted one.

U Untag all processes (remove all tags added with the Space key).

s Trace process system calls: if strace(1) is installed, pressing this key will attach it to the currently selected process, presenting a live update of system calls issued by the process.

l Display open files for a process: if lsof(1) is installed, pressing this key will display the list of file descriptors opened by the process.

L Trace process library calls: if ltrace(1) is installed, pressing this key will attach it to the currently selected process, presenting a live update of library calls issued by the process.

F1, h, ?

Go to the help screen

F2, S

Go to the setup screen, where you can configure the meters displayed at the top of the screen, set various display options, choose among color schemes, and select which columns are displayed, in which order.

- F3, /**
Incrementally search the command lines of all the displayed processes. The currently selected (highlighted) command will update as you type. While in search mode, pressing F3 will cycle through matching occurrences.
- F4,** Incremental process filtering: type in part of a process command line and only processes whose names match will be shown. To cancel filtering, enter the Filter option again and press Esc.
- F5, t**
Tree view: organize processes by parenthood, and layout the relations between them as a tree. Toggling the key will switch between tree and your previously selected sort view. Selecting a sort view will exit tree view.
- F6** On sorted view, select a field for sorting, also accessible through < and >. The current sort field is indicated by a highlight in the header. On tree view, expand or collapse the current subtree. A + indicator in the tree node indicates that it is collapsed.
- F7,]**
Increase the selected process's priority (subtract from 'nice' value). This can only be done by the superuser.
- F8, [**
Decrease the selected process's priority (add to 'nice' value)
- F9, k**
Kill process: sends a signal which is selected in a menu, to one or a group of processes. If processes were tagged, sends the signal to all tagged processes. If none is tagged, sends to the currently selected process.
- F10, q**
Quit
- I** Invert the sort order: if sort order is increasing, switch to decreasing, and vice-versa.
- +, -** When in tree view mode, expand or collapse subtree. When a subtree is collapsed a + sign shows to the left of the process name.
- a (on multiprocessor machines)**
Set CPU affinity: mark which CPUs a process is allowed to use.
- u** Show only processes owned by a specified user.
- M** Sort by memory usage (top compatibility key).
- P** Sort by processor usage (top compatibility key).
- T** Sort by time (top compatibility key).
- F** Follow process: if the sort order causes the currently selected process to move in the list, make the selection bar follow it. This is useful for monitoring a process: this way, you can keep a process always visible on screen. When a movement key is used, follow loses effect.
- K** Hide kernel threads: prevent the threads belonging the kernel to be displayed in the process list. (This is a toggle key.)
- H** Hide user threads: on systems that represent them differently than ordinary processes (such as recent NPTL-based systems), this can hide threads from userspace processes in the process list. (This is a toggle key.)
- Ctrl-L**
Refresh: redraw screen and recalculate values.
- Numbers**
PID search: type in process ID and the selection highlight will be moved to it.

COLUMNS

The following columns can display data about each process. A value of '-' in all the rows indicates that a column is unsupported on your system, or currently unimplemented in htop. The names below are the ones used in the Available Columns section of the setup screen. If a different name is shown in htop's main screen, it is shown below in parenthesis.

Command

The full command line of the process (i.e program name and arguments).

PID

The process ID.

PPID

The parent process ID.

PGRP

The process's group ID.

SESSION (SESN)

The process's session ID.

TTY_NR (TTY)

The controlling terminal of the process.

TPGID

The process ID of the foreground process group of the controlling terminal.

STATE (S)

The state of the process: **S** for sleeping (idle) **R** for running **D** for disk sleep (uninterruptible) **Z** for zombie (waiting for parent to read its exit status) **T** for traced or suspended (e.g by SIGTSTP) **W** for paging

PROCESSOR (CPU)

The ID of the CPU the process last executed on.

NLWP

The number of threads in the process.

NICE (NI)

The nice value of a process, from 19 (low priority) to -20 (high priority). A high value means the process is being nice, letting others have a higher relative priority. Only root can lower the value.

PERCENT_CPU (CPU%)

The percentage of the CPU time that the process is currently using.

UTIME (UTIME+)

The user CPU time, which is the amount of time the process has spent executing on the CPU in user mode (i.e everything but system calls), measured in clock ticks.

STIME (STIME+)

The system CPU time, which is the amount of time the kernel has spent executing system calls on behalf of the process, measured in clock ticks.

TIME (TIME+)

The time, measured in clock ticks that the process has spent in user and system time (see UTIME, STIME above).

CUTIME

The children's user CPU time, which is the amount of time the process's waited-for children have spent executing in user mode (see UTIME above).

CSTIME

The children's system CPU time, which is the amount of time the kernel has spent executing system calls on behalf of all the process's waited-for children (see STIME above).

PRIORITY (PRI)

The kernel's internal priority for the process, usually just its nice value plus twenty. Different for real-time processes.

PERCENT_MEM

The percentage of memory the process is currently using (based on the process's resident memory size, see `M_RESIDENT` below).

M_SIZE (VIRT)

Size in memory of the total program size.

M_RESIDENT (RES)

The resident set size, i.e the size of the text and data sections, plus stack usage.

M_SHARE (SHR)

The size of the process's shared pages

M_TRS (CODE)

The size of the text segment of the process (i.e the size of the processes executable instructions).

M_LRS (LIB)

The library size of the process.

M_DRS (DATA)

The size of the data segment plus stack usage of the process.

M_DT (DIRTY)

The size of the dirty pages of the process.

ST_UID (UID)

The user ID of the process owner.

USER

The username of the process owner, or the user ID if the name can't be determined.

STARTTIME

The time the process was started.

RCHAR (RD_CHAR)

The number of bytes the process has read.

WCHAR (WR_CHAR)

The number of bytes the process has written.

SYSCR (RD_SYSC)

The number of [read\(2\)](#) syscalls for the process.

SYSCW (WR_SYSC)

The number of [write\(2\)](#) syscalls for the process.

RBYTES (IO_RBYTES)

Bytes of [read\(2\)](#) I/O for the process.

WBYTES (IO_WBYTES)

Bytes of [write\(2\)](#) I/O for the process.

IO_READ_RATE (IORR)

The I/O rate of [read\(2\)](#) in bytes per second, for the process.

IO_WRITE_RATE (IOWR)

The I/O rate of [write\(2\)](#) in bytes per second, for the process.

IO_RATE (IO)

The I/O rate, `IO_READ_RATE + IO_WRITE_RATE` (see above).

CNCLWB (IO_CANCEL)

Bytes of cancelled [write\(2\)](#) I/O.

CGROUP

Which cgroup the process is in.

CTID

OpenVZ container ID, a.k.a virtual environment ID.

VPID

OpenVZ process ID.

VXID

VServer process ID.

OOM

OOM killer score.

All other flags

Currently unsupported (always displays '-').

SEE ALSO

[proc\(5\)](#), [top\(1\)](#), [free\(1\)](#), [ps\(1\)](#), [uptime\(1\)](#)

AUTHORS

htop is developed by Hisham Muhammad <hisham@gobolinux.org>.

This man page was written by Bartosz Fenski <fenio@o2.pl> for the Debian GNU/Linux distribution (but it may be used by others). It was updated by Hisham Muhammad, and later by Vincent Launchbury, who wrote the 'Columns' section.