

NAME

top - display and update information about the top cpu processes

SYNOPSIS

top [-abCHlijnPpqSTtuvxz] [-d *count*] [-J *jail*] [-m *mode*] [-o *field*] [-p *pid*] [-s *time*] [-U *uid*] [*number*]

DESCRIPTION

top displays the top processes on the system and periodically updates this information. If standard output is an intelligent terminal (see below) then as many processes as will fit on the terminal screen are displayed by default. Otherwise, a good number of them are shown (around 20). Raw cpu percentage is used to rank the processes. If *number* is given, then the top *number* processes will be displayed instead of the default.

top makes a distinction between terminals that support advanced capabilities and those that do not. This distinction affects the choice of defaults for certain options. In the remainder of this document, an "intelligent" terminal is one that supports cursor addressing, clear screen, and clear to end of line. Conversely, a "dumb" terminal is one that does not support such features. If the output of **top** is redirected to a file, it acts as if it were being run on a dumb terminal.

The options are as follows:

- a Display command names derived from the argv[] vector, rather than real executable name. It is useful when you want to watch applications, that puts their status information there. If the real name differs from argv[0], it will be displayed in parenthesis. Non-printable characters in the command line are encoded in C-style backslash sequences or a three digit octal sequences.
- b Use "batch" mode. In this mode, all input from the terminal is ignored. Interrupt characters (such as ^C and ^\) still have an effect. This is the default on a dumb terminal, or when the output is not a terminal.
- C Toggle CPU display mode. By default top displays the weighted CPU percentage in the WCPU column (this is the same value that ps(1) displays as CPU). Each time -C flag is passed it toggles between "raw cpu" mode and "weighted cpu" mode, showing the "CPU" or the "WCPU" column respectively.
- d *count*
Show only *count* displays, then exit. A display is considered to be one update of the screen. The default is 1 for dumb terminals. Note that for *count* = 1 no information is available about the percentage of time spent by the CPU in every state.

- H** Display each thread for a multithreaded process individually. By default a single summary line is displayed for each process.
- I** Do not display idle processes. By default, top displays both active and idle processes.
- i** Use "interactive" mode. In this mode, any input is immediately read for processing. See the section on "Interactive Mode" for an explanation of which keys perform what functions. After the command is processed, the screen will immediately be updated, even if the command was not understood. This mode is the default when standard output is an intelligent terminal.
- J jail** Show only those processes owned by *jail*. This may be either the *jid* or *name* of the jail. Use 0 to limit to host processes. Using this option implies **-j**.
- j** Display the jail(8) ID.
- m mode**
Display statistics in the specified *mode*. Available modes are **cpu** and **io**. Default is **cpu**.
- n** Use "non-interactive" mode. This is identical to "batch" mode.
- o field**
Sort the process display area on the specified field. The field name is the name of the column as seen in the output, but in lower case: "cpu", "size", "res", "time", "pri", "threads", "total", "read", "write", "fault", "vcsw", "ivcsw", "jid", "swap", or "pid".
- P** Display per-cpu CPU usage statistics.
- p pid** Show only the process *pid*.
- q** Renice **top** to -20 so that it will run faster. This can be used when the system is being very sluggish to improve the possibility of discovering the problem. This option can only be used by root.
- S** Show system processes in the display. Normally, system processes such as the pager and the swapper are not shown. This option makes them visible.
- s time**
Set the delay between screen updates to *time* seconds, which may be fractional. The default delay between updates is 1 second.

- T** Toggle displaying thread ID (tid) instead of process id (pid).
- t** Do not display the **top** process itself.
- U *username***
Show only those processes owned by *username*. This option currently only accepts usernames and will not understand uid numbers.
- u** Do not map uid numbers to usernames. Normally, **top** will read as much of the file */etc/passwd* as is necessary to map all the user id numbers it encounters into login names. This option disables all that, while possibly decreasing execution time. The uid numbers are displayed instead of the names.
- v** Write version number information to stderr then exit immediately.
- w** Display approximate swap usage for each process.
- z** Do not display the system idle process.

Both *count* and *number* fields can be specified as "infinite", indicating that they can stretch as far as possible. This is accomplished by using any proper prefix of the keywords "infinity", "maximum", or "all". Boolean flags are toggles. A second specification of any of these options will negate the first.

INTERACTIVE MODE

When **top** is running in "interactive mode", it reads commands from the terminal and acts upon them accordingly. In this mode, the terminal is put in "CBREAK", so that a character will be processed as soon as it is typed. Almost always, a key will be pressed when **top** is between displays; that is, while it is waiting for *time* seconds to elapse. If this is the case, the command will be processed and the display will be updated immediately thereafter (reflecting any changes that the command may have specified). This happens even if the command was incorrect. If a key is pressed while **top** is in the middle of updating the display, it will finish the update and then process the command. Some commands require additional information, and the user will be prompted accordingly. While typing this information in, the user's erase and kill keys (as set up by the command `stty(1)`) are recognized, and a newline terminates the input.

These commands are currently recognized (^L refers to control-L):

- ^L Redraw the screen.
- h Display a summary of the commands (help screen). Version information is included in this

display.

q **Quit `top`**

d Change the number of displays to show (prompt for new number). Remember that the next display counts as one, so typing 'd1' will make **top** show one final display and then immediately exit.

/ Display only processes that contain the specified string in their command name. If displaying arguments is enabled, the arguments are searched too. '+' shows all processes.

m Toggle the display between 'cpu' and 'io' modes.

n or # Change the number of processes to display (prompt for new number).

s Change the number of seconds to delay between displays (prompt for new number).

S Toggle the display of system processes.

a Toggle the display of process titles.

k Send a signal (SIGKILL by default) to a list of processes. This acts similarly to the command kill(1).

r Change the priority (the "nice") of a list of processes. This acts similarly to renice(8).

u Display only processes owned by a specific set of usernames (prompt for username). If the username specified is simply "+" or "-", then processes belonging to all users will be displayed. Usernames can be added to and removed from the set by prepending them with "+" and "-", respectively.

o Change the order in which the display is sorted. The sort key names include "cpu", "res", "size", and "time." The default is cpu.

p Display a specific process (prompt for pid). If the pid specified is simply "+", then show all processes.

e Display a list of system errors (if any) generated by the last command.

H Toggle the display of threads.

- i or I Toggle the display of idle processes.
- j Toggle the display of jail(8) ID.
- J Display only processes owned by a specific jail (prompt for jail). If the jail specified is simply "+", then processes belonging to all jails and the host will be displayed. This will also enable the display of JID.
- P Toggle the display of per-CPU statistics.
- T Toggle display of TID and PID
- t Toggle the display of the **top** process.
- w Toggle the display of swap usage.
- z Toggle the display of the system idle process.

THE DISPLAY

The top few lines of the display show general information about the state of the system, including the last process id assigned to a process (on most systems), the three load averages, the current time, the number of existing processes, the number of processes in each state (sleeping, running, starting, zombies, and stopped), and a percentage of time spent in each of the processor states (user, nice, system, and idle). It also includes information about physical and virtual memory allocation.

The remainder of the screen displays information about individual processes. This display is similar in spirit to ps(1) but it is not exactly the same. PID is the process id, JID, when displayed, is the jail(8) ID corresponding to the process, USERNAME is the name of the process's owner (if **-u** is specified, a UID column will be substituted for USERNAME), PRI is the current priority of the process, NICE is the nice(1) amount, SIZE is the total size of the process (text, data, and stack), RES is the current amount of resident memory, SWAP is the approximate amount of swap, if enabled (SIZE, RES and SWAP are given in kilobytes), STATE is the current state (one of "START", "RUN" (shown as "CPU_n" on SMP systems), "SLEEP", "STOP", "ZOMB", "WAIT", "LOCK", or the event on which the process waits), C is the processor number on which the process is executing (visible only on SMP systems), TIME is the number of system and user cpu seconds that the process has used, WCPU, when displayed, is the weighted cpu percentage (this is the same value that ps(1) displays as CPU), CPU is the raw percentage and is the field that is sorted to determine the order of the processes, and COMMAND is the name of the command that the process is currently running (if the process is swapped out, this column is marked "<swapped>").

If a process is in the "SLEEP" or "LOCK" state, the state column will report the name of the event or lock on which the process is waiting. Lock names are prefixed with an asterisk "*" while sleep events are not.

DESCRIPTION OF MEMORY

Mem: 61M Active, 86M Inact, 368K Laundry, 22G Wired, 102G Free
 ARC: 15G Total, 9303M MFU, 6155M MRU, 1464K Anon, 98M Header, 35M Other
 15G Compressed, 27G Uncompressed, 1.75:1 Ratio, 174M Overhead
 Swap: 4096M Total, 532M Free, 13% Inuse, 80K In, 104K Out

Physical Memory Stats

Active number of bytes active
Inact number of clean bytes inactive
Laundry number of dirty bytes queued for laundering
Wired number of bytes wired down, including IO-level cached file data pages
Buf number of bytes used for IO-level disk caching
Free number of bytes free

ZFS ARC Stats

These stats are only displayed when the ARC is in use.

Total number of wired bytes used for the ZFS ARC
MRU number of ARC bytes holding most recently used data
MFU number of ARC bytes holding most frequently used data
Anon number of ARC bytes holding in flight data
Header number of ARC bytes holding headers
Other miscellaneous ARC bytes
Compressed bytes of memory used by ARC caches
Uncompressed bytes of data stored in ARC caches before compression
Ratio compression ratio of data cached in the ARC

Swap Stats

Total total available swap usage
Free total free swap usage
Inuse swap usage
In bytes paged in from swap devices (last interval)
Out bytes paged out to swap devices (last interval)

ENVIRONMENT

TOP Default set of arguments to **top**.

LC_CTYPE The locale to use when displaying the *argv* vector when **-a** flag is specified.

SEE ALSO

kill(1), ps(1), stty(1), getrusage(2), humanize_number(3), mem(4), renice(8)

AUTHORS

William LeFebvre, EECS Department, Northwestern University

BUGS

The command name for swapped processes should be tracked down, but this would make the program run slower.

As with ps(1), things can change while **top** is collecting information for an update. The picture it gives is only a close approximation to reality.