

NAME

`pidstat` – Report statistics for Linux tasks.

SYNOPSIS

```
pidstat [ -d ] [ -H ] [ -h ] [ -I ] [ -l ] [ -R ] [ -r ] [ -s ] [ -t ] [ -U [ username ] ] [ -u ] [ -V ] [ -v ] [ -w ] [ -C comm ] [ -G process_name ] [ --dec={ 0 | 1 | 2 } ] [ --human ] [ -p { pid [,...] | SELF | ALL } ] [ -T { TASK | CHILD | ALL } ] [ interval [ count ] ] [ -e program args ]
```

DESCRIPTION

The **pidstat** command is used for monitoring individual tasks currently being managed by the Linux kernel. It writes to standard output activities for every task selected with option **-p** or for every task managed by the Linux kernel if option **-p ALL** has been used. Not selecting any tasks is equivalent to specifying **-p ALL** but only active tasks (tasks with non-zero statistics values) will appear in the report.

The **pidstat** command can also be used for monitoring the child processes of selected tasks. Read about option **-T** below.

The *interval* parameter specifies the amount of time in seconds between each report. A value of 0 (or no parameters at all) indicates that tasks statistics are to be reported for the time since system startup (boot). The *count* parameter can be specified in conjunction with the *interval* parameter if this one is not set to zero. The value of *count* determines the number of reports generated at *interval* seconds apart. If the *interval* parameter is specified without the *count* parameter, the **pidstat** command generates reports continuously.

You can select information about specific task activities using flags. Not specifying any flags selects only CPU activity.

OPTIONS

-C *comm*

Display only tasks whose command name includes the string *comm*. This string can be a regular expression.

-d Report I/O statistics (kernels 2.6.20 and later only). The following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

kB_rd/s

Number of kilobytes the task has caused to be read from disk per second.

kB_wr/s

Number of kilobytes the task has caused, or shall cause to be written to disk per second.

kB_ccwr/s

Number of kilobytes whose writing to disk has been cancelled by the task. This may occur when the task truncates some dirty pagecache. In this case, some IO which another task has been accounted for will not be happening.

iodelay

Block I/O delay of the task being monitored, measured in clock ticks. This metric

includes the delays spent waiting for sync block I/O completion and for swpin block I/O completion.

Command

The command name of the task.

- `--dec={ 0 | 1 | 2 }`
Specify the number of decimal places to use (0 to 2, default value is 2).
- `-e program args`
Execute *program* with given arguments *args* and monitor it with **pidstat**. **pidstat** stops when *program* terminates.
- `-G process_name`
Display only processes whose command name includes the string *process_name*. This string can be a regular expression. If option `-t` is used together with option `-G` then the threads belonging to that process are also displayed (even if their command name doesn't include the string *process_name*).
- `-H` Display timestamp in seconds since the epoch.
- `-h` Display all activities horizontally on a single line, with no average statistics at the end of the report. This is intended to make it easier to be parsed by other programs.
- `--human`
Print sizes in human readable format (e.g. 1.0k, 1.2M, etc.) The units displayed with this option supersede any other default units (e.g. kilobytes, sectors...) associated with the metrics.
- `-I` In an SMP environment, indicate that tasks CPU usage (as displayed by option `-u`) should be divided by the total number of processors.
- `-l` Display the process command name and all its arguments.
- `-p { pid [...] | SELF | ALL }`
Select tasks (processes) for which statistics are to be reported. *pid* is the process identification number. The **SELF** keyword indicates that statistics are to be reported for the **pidstat** process itself, whereas the **ALL** keyword indicates that statistics are to be reported for all the tasks managed by the system.
- `-R` Report realtime priority and scheduling policy information. The following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

prio

The realtime priority of the task being monitored.

policy

The scheduling policy of the task being monitored.

Command

The command name of the task.

-r Report page faults and memory utilization.

When reporting statistics for individual tasks, the following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

minflt/s

Total number of minor faults the task has made per second, those which have not required loading a memory page from disk.

majflt/s

Total number of major faults the task has made per second, those which have required loading a memory page from disk.

VSZ

Virtual Size: The virtual memory usage of entire task in kilobytes.

RSS

Resident Set Size: The non-swapped physical memory used by the task in kilobytes.

%MEM

The tasks's currently used share of available physical memory.

Command

The command name of the task.

When reporting global statistics for tasks and all their children, the following values may be displayed:

UID

The real user identification number of the task which is being monitored together with its children.

USER

The name of the real user owning the task which is being monitored together with its children.

PID

The identification number of the task which is being monitored together with its children.

minflt-nr

Total number of minor faults made by the task and all its children, and collected during the interval of time.

majflt-nr

Total number of major faults made by the task and all its children, and collected during the interval of time.

Command

The command name of the task which is being monitored together with its children.

–s Report stack utilization. The following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

StkSize

The amount of memory in kilobytes reserved for the task as stack, but not necessarily used.

StkRef

The amount of memory in kilobytes used as stack, referenced by the task.

Command

The command name of the task.

–T { TASK | CHILD | ALL }

This option specifies what has to be monitored by the **pidstat** command. The **TASK** keyword indicates that statistics are to be reported for individual tasks (this is the default option) whereas the **CHILD** keyword indicates that statistics are to be globally reported for the selected tasks and all their children. The **ALL** keyword indicates that statistics are to be reported for individual tasks and globally for the selected tasks and their children.

Note: Global statistics for tasks and all their children are not available for all options of **pidstat**. Also these statistics are not necessarily relevant to current time interval: The statistics of a child process are collected only when it finishes or it is killed.

–t Also display statistics for threads associated with selected tasks.

This option adds the following values to the reports:

TGID

The identification number of the thread group leader.

TID

The identification number of the thread being monitored.

–U [username]

Display the real user name of the tasks being monitored instead of the UID. If *username* is specified, then only tasks belonging to the specified user are displayed.

–u Report CPU utilization.

When reporting statistics for individual tasks, the following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

%usr

Percentage of CPU used by the task while executing at the user level (application), with or without nice priority. Note that this field does NOT include time spent running a virtual processor.

%system

Percentage of CPU used by the task while executing at the system level (kernel).

%guest

Percentage of CPU spent by the task in virtual machine (running a virtual processor).

%wait

Percentage of CPU spent by the task while waiting to run.

%CPU

Total percentage of CPU time used by the task. In an SMP environment, the task's CPU usage will be divided by the total number of CPU's if option **-I** has been entered on the command line.

CPU

Processor number to which the task is attached.

Command

The command name of the task.

When reporting global statistics for tasks and all their children, the following values may be displayed:

UID

The real user identification number of the task which is being monitored together with its children.

USER

The name of the real user owning the task which is being monitored together with its children.

PID

The identification number of the task which is being monitored together with its children.

usr-ms

Total number of milliseconds spent by the task and all its children while executing at the user level (application), with or without nice priority, and collected during the interval of time. Note that this field does NOT include time spent running a virtual processor.

system-ms

Total number of milliseconds spent by the task and all its children while executing at the system level (kernel), and collected during the interval of time.

guest-ms

Total number of milliseconds spent by the task and all its children in virtual machine (running a virtual processor).

Command

The command name of the task which is being monitored together with its children.

- V Print version number then exit.
- v Report values of some kernel tables. The following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

threads

Number of threads associated with current task.

fd-nr

Number of file descriptors associated with current task.

Command

The command name of the task.

- w Report task switching activity (kernels 2.6.23 and later only). The following values may be displayed:

UID

The real user identification number of the task being monitored.

USER

The name of the real user owning the task being monitored.

PID

The identification number of the task being monitored.

cswch/s

Total number of voluntary context switches the task made per second. A voluntary context switch occurs when a task blocks because it requires a resource that is unavailable.

nvcswch/s

Total number of non voluntary context switches the task made per second. A involuntary context switch takes place when a task executes for the duration of its time slice and then is forced to relinquish the processor.

Command

The command name of the task.

ENVIRONMENT

The **pidstat** command takes into account the following environment variables:

S_COLORS

When this variable is set, display statistics in color on the terminal. Possible values for this variable are *never*, *always* or *auto* (the latter is the default).

Note: On Debian systems the colors are displayed by default when output is connected to the terminal, even if this variable is not set (i.e. unset variable is treated as if it were set to *auto*).

Please note that the color (being red, yellow, or some other color) used to display a value is not indicative of any kind of issue simply because of the color. It only indicates different ranges of values.

S_COLORS_SGR

Specify the colors and other attributes used to display statistics on the terminal. Its value is a colon-separated list of capabilities that defaults to **H=31;1:I=32;22:M=35;1:N=34;1:Z=34;22**. Supported capabilities are:

H= SGR (Select Graphic Rendition) substring for percentage values greater than or equal to 75%.

I= SGR substring for item values like PID, UID or CPU number.

M= SGR substring for percentage values in the range from 50% to 75%.

N= SGR substring for non-zero statistics values and for tasks names.

Z= SGR substring for zero values and for threads names.

S_TIME_FORMAT

If this variable exists and its value is **ISO** then the current locale will be ignored when printing the date in the report header. The **pidstat** command will use the ISO 8601 format (YYYY-MM-DD) instead. The timestamp will also be compliant with ISO 8601 format.

EXAMPLES

pidstat 2 5

Display five reports of CPU statistics for every active task in the system at two second intervals.

pidstat -r -p 1643 2 5

Display five reports of page faults and memory statistics for PID 1643 at two second intervals.

pidstat -C "fox|bird" -r -p ALL

Display global page faults and memory statistics for all the processes whose command name includes the string "fox" or "bird".

pidstat -T CHILD -r 2 5

Display five reports of page faults statistics at two second intervals for the child processes of all tasks in the system. Only child processes with non-zero statistics values are displayed.

BUGS

/proc filesystem must be mounted for the **pidstat** command to work.

FILES

/proc contains various files with system statistics.

AUTHOR

Sebastien Godard (sysstat <at> orange.fr)

SEE ALSO

sar(1), **top(1)**, **ps(1)**, **mpstat(1)**, **iostat(1)**, **vmstat(8)**

<https://github.com/sysstat/sysstat>

<http://pagesperso-orange.fr/sebastien.godard/>