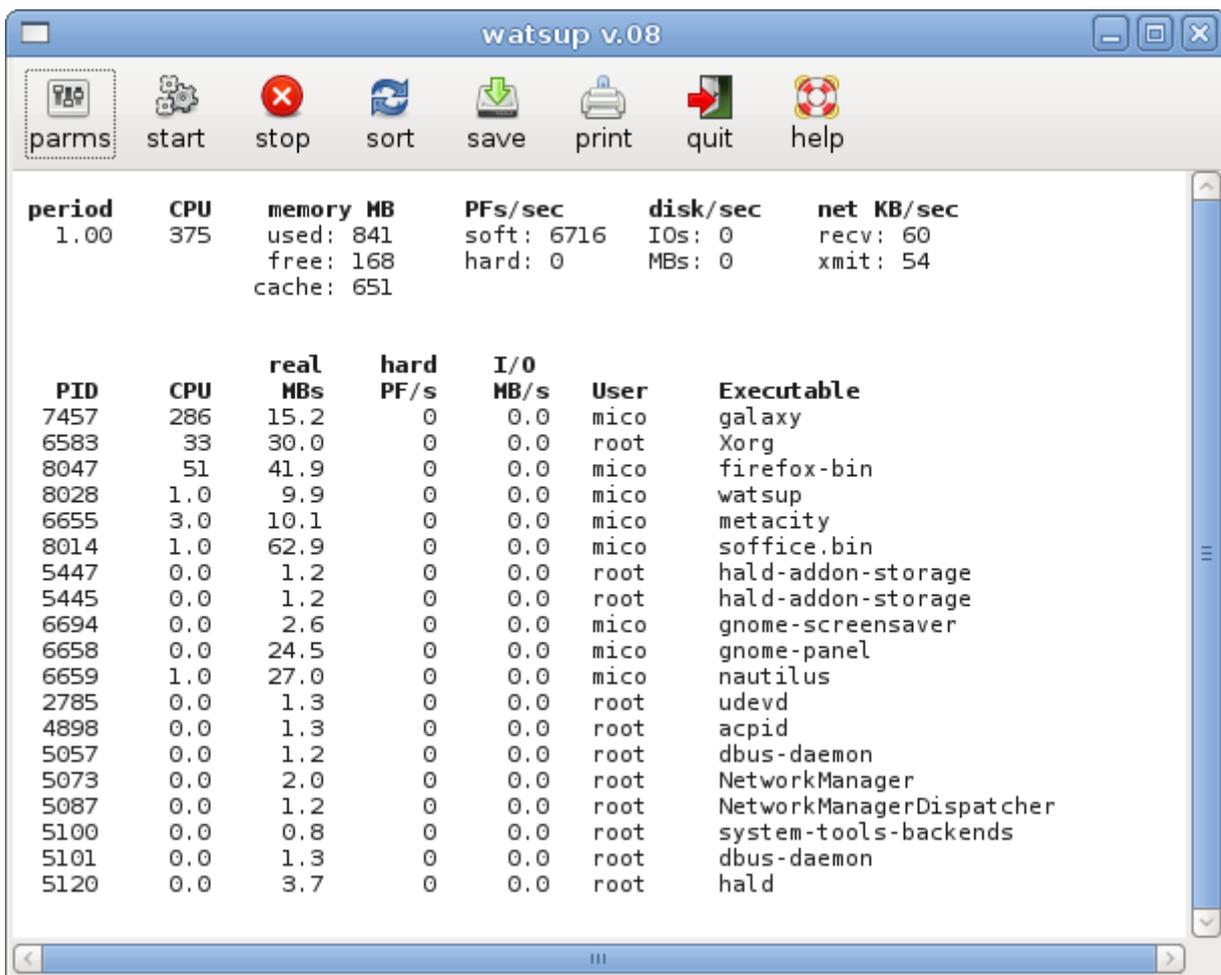


Introduction

Watsup monitors Linux system resources and the processes using those resources. It is different from monitor programs you have seen before:

- Overhead is low, supporting a sample interval down to 0.1 seconds
- Overall system and top process resources are shown on one page
- CPU, memory, disk I/O, network I/O, and page faults are monitored
- Font can be made small (for display in a corner) or large
- The N top-ranked processes fitting in the window are shown with minimal jumping around between samples (easier to watch one or a few processes)
- Process rank is a weighted sum of CPU, hard page faults, and disk I/O over the last several samples (sample weight declines over time)

Screenshot



License and Warranty

Watsup is a free program licensed under the GNU General Public License v.2 (Free Software Foundation). Watsup is not warranted for any purpose, but if you find a bug I will try to fix it.

Origin and Contact

Watsup originates from the author's web site at <http://kornelix.squarespace.com/watsup>

Other web sites may offer it for download. Modifications could have been made.

If you have questions, suggestions or a bug to report: kornelix@yahoo.de

function of toolbar buttons

button	function
parms	dialog for setting the sample period (seconds), and the report font size
start	start sampling and reporting
stop	stop sampling and reporting
sort	re-sort process list in decreasing order of recent average resource usage (weight of old data declines with time - most recent data has highest weight)
save	save current screen to a data file
print	print current screen on default printer
quit	exit the application
help	display help file (this file)

system resources reported

parameter	description	source
period	sample interval, seconds	system clock
CPU	% CPU load (max = no. SMPs x 100)	/proc/stat
memory MB	total real memory used, free, cached	/proc/meminfo
PFs/sec	page faults/sec, hard and soft	/proc/vmstat
disk/sec	disk I/O ops/sec and MB/sec	/proc/diskstats
net KB/sec	network kilobytes/sec received and sent	/proc/net/dev

process resources reported

parameter	description	source
PID	process ID	/proc/<pid>
CPU	% CPU used by process	/proc/<pid>/stat
real MBs	real megabytes used by process	/proc/<pid>/stat
hard PF/s	process hard page faults/sec	/proc/<pid>/stat
I/O MB/s	process disk I/O, MB/sec (see note below)	/proc/<pid>/io
User	login user name	/proc/<pid>/environ
Executable	executable file name	/proc/<pid>/exe

methodology

The files under `/proc` contain current data for memory usage, and counters that continuously increment for other resources (CPU time, disk and network I/O, page faults). Resource usage rates are calculated as $(\text{current value} - \text{prior value}) / (\text{sample interval})$.

CPU

The total CPU load is based on $1 \text{ CPU} = 100\%$. Thus a system having N SMP CPUs (symmetric multi-processors) has a maximum load value of $N \times 100$ percent.

page faults

Hard page faults are those requiring physical I/O, and soft faults are those which use free memory or stale pages from other processes. Modern CPUs can support over 100K soft faults per second. Soft fault overhead is also reflected in CPU consumption.

cached memory

Memory not needed by active processes is used for disk caching. This memory is available for processes when needed.

process CPU

The CPU % used by a process can exceed 100% for multi-threaded processes utilizing multiple SMP CPUs. The top value is $100 \times (\text{SMP count})$. The newest processors have 2 or 4 SMPs.

process disk I/O

This is a new kernel feature (from 2.6.20) that is not enabled in most distros. Watsup tries to read the files `/proc/pid/io` but does not complain if they do not exist. There is a kernel compile time option to turn on these statistics: `CONFIG_TASKSTATS`

running as root

This is necessary to see all processes. Otherwise processes owned by root or other users are not shown. Use a terminal command: `$ sudo /.../watsup` or create a desktop launcher with this command.