



Adding performance counters to htop

Hisham Muhammad

@hisham_hm@mastodon.social

<https://hisham.hm>

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

About me

original author of htop, a project started in 2004

<http://hisham.hm/htop/>

lead dev of LuaRocks, package manager for Lua

<http://luarocks.org/>

co-founder of the GoboLinux distribution

<http://gobolinux.org/>

developer at Kong - open source API gateway

<http://getkong.org/>

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

What is htop

an interactive process manager

intended to be “a better top”

by this all I originally meant was: scrolling!

(versions of top improved a lot since!)

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Hello, htop!

```

./htop
 0 [|||||] 8.4%] Tasks: 94, 42 thr; 1 running
 1 [|||] 3.8%] Load average: 0.47 0.26 0.22
 2 [|] 0.6%] Uptime: 04:28:22
 3 [||] 1.3%] Battery: n/a
Mem[|||||||||||||||||||||||||||||||||||||] 1114/3819MB]
Swp[|] 0/0MB]

IO Priority:
None (based on nice)
Realtme 0 (High)
Realtme 1
Realtme 2
Realtme 3
Realtme 4
Realtme 5
Realtme 6
Realtme 7 (Low)
Best-effort 0 (High)
Best-effort 1
Best-effort 2
Best-effort 3
Best-effort 4
Best-effort 5
Best-effort 6
Best-effort 7 (Low)
Idle

Enter/Set Esc/Cancel

PID USER PRI NI VIRT RES SHR S CPU% MEM% Command
1 root 40 0 1672 572 508 S 0.0 0.0 init [2]
15814 hisham 40 0 16140 6284 3856 S 0.0 0.2 urxvt -cr green -fn *-lode-* -fb *
15815 hisham 40 0 8804 6004 1692 S 0.0 0.2 | zsh
15836 hisham 40 0 45512 30796 13920 S 0.0 0.8 | | gimp
16474 hisham 40 0 26020 12624 8792 S 0.6 0.3 | | | /System/Index/lib/gimp/2.
16047 hisham 40 0 20552 5344 3236 S 0.0 0.1 | | | /System/Index/lib/gimp/2.
15345 hisham 40 0 62852 54784 3920 S 0.0 1.4 | urxvt -cr green -fn *-lode-* -fb *
15346 hisham 40 0 9132 6460 1868 S 0.0 0.2 | | zsh
15357 hisham 40 0 1716 564 468 T 0.0 0.0 | | | cw: wrapping [find] {pid=153
| | | find
15291 hisham 40 0 16332 6492 3920 S 0.0 0.2 | urxvt -cr green -fn *-lode-* -fb *
15292 hisham 40 0 9000 6316 1856 S 0.0 0.2 | | zsh
15340 hisham 40 0 3116 1852 1148 R 4.4 0.0 | | | ./htop
14628 hisham 40 0 16140 6304 3864 S 0.0 0.2 | urxvt -cr green -fn *-lode-* -fb *
14629 hisham 40 0 9240 6496 1868 S 0.0 0.2 | | zsh
14646 hisham 40 0 1712 548 468 S 0.0 0.0 | | | cw: wrapping [env] {pid=1464
| | | /bin/bash /System/Links/E
| | | python /System/Links/E
13475 hisham 40 0 16140 6380 3864 S 0.0 0.2 | urxvt -cr green -fn *-lode-* -fb *
13476 hisham 40 0 8988 6320 1868 S 0.0 0.2 | | zsh
13808 hisham 40 0 3976 1912 1532 S 0.0 0.0 | | | ssh -t loderunner,htop@shell
13384 hisham 40 0 16140 6288 3864 S 0.0 0.2 | urxvt -cr green -fn *-lode-* -fb *
13385 hisham 40 0 9000 6296 1844 S 0.0 0.2 | | zsh
10544 hisham 40 0 26212 16596 3920 S 0.0 0.4 | urxvt -cr green -fn *-lode-* -fb *

```

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Metrics in htop

I started with the default set of top

Most people don't go beyond the default set

(A lot of people don't even know about tree view!)

```
PID USER      PRI  NI   VIRT   RES   SHR S CPU% MEM%   TIME+  Command
```

htop has a lot more

hit F2 (or uppercase C) to enter Setup and play around:

```

1 [ |           0.7%] Tasks: 43; 1 running
2 [ |           0.7%] Load average: 0.58 0.35 0.18
3 [ ||          2.0%] Uptime: 01:03:03
4 [ |           0.0%]
Mem[|||||      959M/15.5G]
Swp[|           0K/0K]

Setup          Active Columns  Available Columns
Meters         PID              PID - Process/thread ID
Display options USER            STATE - Process state (S sleeping, R run
Colors         PRIORITY        PPID - Parent process ID
Columns        NICE            PGRP - Process group ID
               M_SIZE         SESSION - Process's session ID
               M_RESIDENT    TTY_NR - Controlling terminal
               M_SHARE       TPID - Process ID of the fg process gro
               STATE        MINFLT - Number of minor faults which ha
               PERCENT_CPU  CMINFLT - Children processes' minor faul
               PERCENT_MEM  MAJFLT - Number of major faults which ha
               TIME        CMAJFLT - Children processes' major faul
               Command    UTIME - User CPU time - time the process
               STATE        STIME - System CPU time - time the kerne
               CUTIME - Children processes' user CPU ti

F1 [          F2 [          F3 [          F4 [          F5 [          F6 [          F7 [          F8 [          F9 [          F10 [Don

```

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

I/O metrics

inspired by iotop

these have in fact been available in htop for years!

```
IO_RATE
IO_READ_RATE
IO_WRITE_RATE
IO_PRIORITY
```

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Delay accounting

contributed by André Carvalho

```
PERCENT_CPU_DELAY  
PERCENT_IO_DELAY  
PERCENT_SWAP_DELAY
```

requires libnl for Netlink support

to learn more about delay accounting:

<https://andrestc.com/post/linux-delay-accounting/>

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Hardware Performance Counters

low-overhead performance metrics

managed by the Linux kernel

available in userspace via perf-tools

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Using performance counters

- 1) perform a syscall to request a counter
- 2) syscall returns a file descriptor
- 3) periodically read the file descriptor

be aware: fds and perf counters are a limited resource!

based on tiptop by Erven Rohou and Antoine Naudin

<https://hal.inria.fr/hal-00639173>

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Performance counters in htop

IPC
MCYCLE
MINSTR
PERCENT_MISS
PERCENT_BMISS
L1DREADS
L1DRMISSES
L1DWrites
L1DWMISSES

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Multiple screens

making a larger number of metrics usable

press [Tab] to switch screens

required a larger change to the format of htoprc

Slide: [|||||||]

13] Date: 2018-02-04

Talk: Adding perf counters to htop

Presenter: <https://hisham.hm>

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Demo time

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Screenshot time: I/O

PID	USER	IO	DISK R/W	DISK READ	DISK WRITE	SWD%	IOD%	Command
25829	hisham	B4	33.26 M/s	33.26 M/s	0.00 B/s	0.0	0.0	find /
21709	hisham	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	zsh
15278	hisham	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/Programs/Fir
25800	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	./htop
1	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	init [2]
180	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/sbin/udev
273	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	klogd -n -c 1
274	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	syslogd -n -m
688	messagebu	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/bin/dbus-dae
694	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/usr/bin/acpi
716	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	listener -c /
802	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/bin/login --
803	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/System/Index
804	root	B4	0.00 B/s	0.00 B/s	0.00 B/s	0.0	0.0	/System/Index

F1 Help F2 Setup F3 Search F4 Filter F5 Tree F6 SortBy F7 Nice - F8 Nice + F9 Kill F10 Qui

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Screenshot time: Perf

```

1  [|||||] 35.6%] Tasks: 47; 2 running
2  [|||||] 42.8%] Load average: 1.34 0.63 0.44
3  [|||||] 34.2%] Uptime: 01:18:45
4  [|||||] 30.6%]
Mem[||||] 1.05G/15.5G]
Swp[     ] 0K/0K]

```

PID	USER	CPU%	CPU	Mcycle	Minstr	IPC	MIS%	BrM%	Command
25803	hisham	57.0	2	869.39	1691.64	1.95	0.2	0.1	urxvt
25829	hisham	60.4	1	338.31	562.18	1.66	0.1	0.1	find /
938	root	6.1	2	45.00	36.55	0.81	3.5	0.5	/usr/bin/X :10 vt1
21709	hisham	1.4	3	20.65	22.49	1.09	0.7	0.8	zsh
942	hisham	1.4	4	12.75	21.28	1.67	0.6	0.3	awesome
966	hisham	1.4	3	11.76	3.60	0.31	9.9	1.5	compton -bcf --fad
25800	root	0.7	3	5.17	5.98	1.16	0.5	0.2	./htop
21708	hisham	0.0	1	3.00	1.35	0.45	8.2	1.6	urxvt
21633	hisham	0.7	1	2.67	3.80	1.42	0.6	0.3	/bin/htop
21600	hisham	0.0	3	2.16	2.94	1.36	0.7	0.3	urxvt
21688	hisham	0.0	3	1.82	2.83	1.56	0.8	0.2	urxvt
15278	hisham	0.0	1	1.53	0.33	0.21	23.4	2.6	/Programs/Firefox/
15677	hisham	0.0	3	0.24	0.05	0.20	21.2	3.3	/Programs/Firefox/
15389	hisham	0.0	2	0.21	0.05	0.23	13.7	3.0	/Programs/Firefox/

F1 Help F2 Setup F3 Search F4 Filter F5 Tree F6 SortBy F7 Nice - F8 Nice + F9 Kill F10 Qui

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Screenshot time: L1

```

1  [|||||] 20.5%] Tasks: 47; 2 running
2  [|||||] 14.3%] Load average: 1.67 0.72 0.47
3  [|||||] 17.7%] Uptime: 01:18:57
4  [|||||] 18.8%]
Mem[|||||] 1.05G/15.5G]
Swp[|] 0K/0K]

```

PID	USER	CPU%	L1Dread	R miss	L1Dwrite	W miss	Command
25803	hisham	29.2	370064.8	19203.4	176724.1	N/A	urxvt
25829	hisham	30.6	79529.7	3790.3	31105.7	N/A	find /
938	root	4.1	11127.4	568.9	4815.2	N/A	/usr/bin/X :10 vt
21709	hisham	1.4	3862.8	151.5	2522.1	N/A	zsh
25800	root	0.0	1279.1	87.5	932.9	N/A	./htop
966	hisham	1.4	1006.0	168.3	547.0	N/A	compton -bcf --fa
21688	hisham	0.0	947.9	54.4	502.3	N/A	urxvt
942	hisham	0.0	933.9	31.4	478.7	N/A	awesome
21600	hisham	0.0	874.1	44.8	461.2	N/A	urxvt
21633	hisham	1.4	790.9	29.1	500.2	N/A	/bin/htop
21489	hisham	0.0	282.3	25.3	108.7	N/A	/Programs/LibreOf
21708	hisham	0.0	245.8	24.1	101.8	N/A	urxvt
15278	hisham	0.0	55.0	9.3	32.8	N/A	/Programs/Firefox
15155	hisham	0.0	24.6	3.5	14.0	N/A	/Programs/Firefox

F1 Help F2 Setup F3 Search F4 Filter F5 Tree F6 SortBy F7 Nice - F8 Nice + F9 Kill F10 Qui


```
PID USER      PRI  NI   VIRT   RES   SHR S CPU% MEM%   TIME+  Command
```

Screenshot time: Screens

```
 1 [|||]          1.3%]   Tasks: 46; 2 running
 2 [|||]          1.3%]   Load average: 0.78 0.75 0.53
 3 [|||]          0.7%]   Uptime: 01:22:37
 4 [|||]          0.7%]
Mem[|||||]      1.07G/15.5G]
Swp[|||||]      0K/0K]
```

Setup	Screens	Active Columns	Available Columns
Meters	Default	PID	PID - Process/thread
Display options	I/O	USER	STATE - Process state
Colors	Perf Counters	PERCENT_CPU	PPID - Parent process ID
Screens	L1 Data Cache	PROCESSOR	PGRP - Process group ID
		MCYCLE	SESSION - Process's session ID
		MINSTR	TTY_NR - Controlling terminal
		IPC	TPGID - Process ID of parent
		PERCENT_MISS	MINFLT - Number of minor faults
		PERCENT_BMISS	CMINFLT - Children minor faults
		Command	MAJFLT - Number of major faults
			CMAJFLT - Children major faults
			UTIME - User CPU time
			STIME - System CPU time
			CUTIME - Children user CPU time

F1 [] F2 Rename F3 [] F4 [] F5 New F6 [] F7 MoveUp F8 MoveDn F9 Remove F10 Don

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

Availability

I/O metrics and delay accounting:

htop 2.1.0, out today

performance counters and multiple screens:

htop 3.0beta1

Linux-only! (patches welcome!)

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
-----	------	-----	----	------	-----	-----	---	------	------	-------	---------

What's next

New header meters? Per-screen headers?

More perf counters?

PS: check out the [lua](#) branch