

NAME

systemd-cgtop – Show top control groups by their resource usage

SYNOPSIS

systemd-cgtop [OPTIONS...] [GROUP]

DESCRIPTION

systemd-cgtop shows the top control groups of the local Linux control group hierarchy, ordered by their CPU, memory, or disk I/O load. The display is refreshed in regular intervals (by default every 1s), similar in style to **top**(1). If a control group path is specified, shows only the services of the specified control group.

If **systemd-cgtop** is not connected to a tty, no column headers are printed and the default is to only run one iteration. The **--iterations=** argument, if given, is honored. This mode is suitable for scripting.

Resource usage is only accounted for control groups in the relevant hierarchy, i.e. CPU usage is only accounted for control groups in the "cpuacct" hierarchy, memory usage only for those in "memory" and disk I/O usage for those in "blkio". If resource monitoring for these resources is required, it is recommended to add the *CPUAccounting=1*, *MemoryAccounting=1* and *BlockIOAccounting=1* settings in the unit files in question. See **systemd.resource-control**(5) for details.

The CPU load value can be between 0 and 100 times the number of processors the system has. For example, if the system has 8 processors, the CPU load value is going to be between 0% and 800%. The number of processors can be found in `/proc/cpuinfo`.

To emphasize this: unless "CPUAccounting=1", "MemoryAccounting=1" and "BlockIOAccounting=1" are enabled for the services in question, no resource accounting will be available for system services and the data shown by **systemd-cgtop** will be incomplete.

OPTIONS

The following options are understood:

-p, --order=path

Order by control group path name.

-t, --order=tasks

Order by number of tasks/processes in the control group.

-c, --order=cpu

Order by CPU load.

-m, --order=memory

Order by memory usage.

-i, --order=io

Order by disk I/O load.

-b, --batch

Run in "batch" mode: do not accept input and run until the iteration limit set with **--iterations=** is exhausted or until killed. This mode could be useful for sending output from **systemd-cgtop** to other programs or to a file.

-r, --raw

Format byte counts (as in memory usage and I/O metrics) with raw numeric values rather than human-readable numbers.

--cpu=percentage, --cpu=time

Controls whether the CPU usage is shown as percentage or time. By default, the CPU usage is shown as percentage. This setting may also be toggled at runtime by pressing the % key.

-P

Count only userspace processes instead of all tasks. By default, all tasks are counted: each kernel thread and each userspace thread individually. With this setting, kernel threads are excluded from the counting and each userspace process only counts as one, regardless how many threads it consists of. This setting may also be toggled at runtime by pressing the P key. This option may not be combined

with **-k**.

-k

Count only userspace processes and kernel threads instead of all tasks. By default, all tasks are counted: each kernel thread and each userspace thread individually. With this setting, kernel threads are included in the counting and each userspace process only counts as one, regardless how many threads it consists of. This setting may also be toggled at runtime by pressing the **k** key. This option may not be combined with **-P**.

--recursive=

Controls whether the number of processes shown for a control group shall include all processes that are contained in any of the child control groups as well. Takes a boolean argument, which defaults to "yes". If enabled, the processes in child control groups are included, if disabled, only the processes in the control group itself are counted. This setting may also be toggled at runtime by pressing the **r** key. Note that this setting only applies to process counting, i.e. when the **-P** or **-k** options are used. It has no effect if all tasks are counted, in which case the counting is always recursive.

-n, --iterations=

Perform only this many iterations. A value of 0 indicates that the program should run indefinitely.

-1

A shortcut for **--iterations=1**.

-d, --delay=

Specify refresh delay in seconds (or if one of "ms", "us", "min" is specified as unit in this time unit). This setting may also be increased and decreased at runtime by pressing the **+** and **-** keys.

--depth=

Maximum control group tree traversal depth. Specifies how deep **systemd-cgtop** shall traverse the control group hierarchies. If 0 is specified, only the root group is monitored. For 1, only the first level of control groups is monitored, and so on. Defaults to 3.

-M MACHINE, --machine=MACHINE

Limit control groups shown to the part corresponding to the container *MACHINE*. This option may not be used when a control group path is specified.

-h, --help

Print a short help text and exit.

--version

Print a short version string and exit.

KEYS

systemd-cgtop is an interactive tool and may be controlled via user input using the following keys:

h

Shows a short help text.

Space

Immediately refresh output.

q

Terminate the program.

p, t, c, m, i

Sort the control groups by path, number of tasks, CPU load, memory usage, or I/O load, respectively. This setting may also be controlled using the **--order=** command line switch.

%

Toggle between showing CPU time as time or percentage. This setting may also be controlled using the **--cpu=** command line switch.

+, -

Increase or decrease refresh delay, respectively. This setting may also be controlled using the

--delay= command line switch.

P

Toggle between counting all tasks, or only userspace processes. This setting may also be controlled using the **-P** command line switch (see above).

k

Toggle between counting all tasks, or only userspace processes and kernel threads. This setting may also be controlled using the **-k** command line switch (see above).

r

Toggle between recursively including or excluding processes in child control groups in control group process counts. This setting may also be controlled using the **--recursive=** command line switch. This key is not available if all tasks are counted, it is only available if processes are counted, as enabled with the P or k keys.

EXIT STATUS

On success, 0 is returned, a non-zero failure code otherwise.

SEE ALSO

systemd(1), **systemctl(1)**, **systemd-cgls(1)**, **systemd.resource-control(5)**, **top(1)**