

NAME

choom – display and adjust OOM-killer score.

SYNOPSIS

choom -p *pid*

choom -p *pid -n number*

choom -n *number command* [*argument...*]

DESCRIPTION

The **choom** command displays and adjusts Out-Of-Memory killer score setting.

OPTIONS

-p, --pid *pid*

Specifies process ID.

-n, --adjust *value*

Specify the adjust score value.

-h, --help

Display help text and exit.

-V, --version

Display version information and exit.

NOTES

Linux kernel uses the badness heuristic to select which process gets killed in out of memory conditions.

The badness heuristic assigns a value to each candidate task ranging from 0 (never kill) to 1000 (always kill) to determine which process is targeted. The units are roughly a proportion along that range of allowed memory the process may allocate from based on an estimation of its current memory and swap use. For example, if a task is using all allowed memory, its badness score will be 1000. If it is using half of its allowed memory, its score will be 500.

There is an additional factor included in the badness score: the current memory and swap usage is discounted by 3% for root processes.

The amount of "allowed" memory depends on the context in which the oom killer was called. If it is due to the memory assigned to the allocating task's cpuset being exhausted, the allowed memory represents the set of mems assigned to that cpuset. If it is due to a mempolicy's node(s) being exhausted, the allowed memory represents the set of mempolicy nodes. If it is due to a memory limit (or swap limit) being reached, the allowed memory is that configured limit. Finally, if it is due to the entire system being out of memory, the allowed memory represents all allocatable resources.

The adjust score value is added to the badness score before it is used to determine which task to kill. Acceptable values range from -1000 to +1000. This allows userspace to polarize the preference for oom killing either by always preferring a certain task or completely disabling it. The lowest possible value, -1000, is equivalent to disabling oom killing entirely for that task since it will always report a badness score of 0.

Setting an adjust score value of +500, for example, is roughly equivalent to allowing the remainder of tasks sharing the same system, cpuset, mempolicy, or memory controller resources to use at least 50% more memory. A value of -500, on the other hand, would be roughly equivalent to discounting 50% of the task's allowed memory from being considered as scoring against the task.

AUTHORS

Karel Zak <kzak@redhat.com>

SEE ALSO

proc(5)

AVAILABILITY

The **choom** command is part of the util-linux package and is available from Linux Kernel Archive (<https://www.kernel.org/pub/linux/utils/util-linux/>).