NAME

slabinfo – kernel slab allocator statistics

SYNOPSIS

cat /proc/slabinfo

DESCRIPTION

Frequently used objects in the Linux kernel (buffer heads, inodes, dentries, etc.) have their own cache. The file */proc/slabinfo* gives statistics on these caches. The following (edited) output shows an example of the contents of this file:

```
$ sudo cat /proc/slabinfo
slabinfo - version: 2.1
# name    <active_objs> <num_objs> <objsize> <objperslab> <pagesperslab> ...
sigqueue       100  100  160   25  1 : tunables  0  0  0 : slabdata   4   4  0
sighand_cache 355   405 2112  15  8 : tunables  0  0  0 : slabdata  27  27  0
kmalloc-8192   96   96  8192   4  8 : tunables  0  0  0 : slabdata  24  24  0
...
```

The first line of output includes a version number, which allows an application that is reading the file to handle changes in the file format. (See VERSIONS, below.) The next line lists the names of the columns in the remaining lines.

Each of the remaining lines displays information about a specified cache. Following the cache name, the output shown in each line shows three components for each cache:

* statistics
* tunables
* slabdata

The statistics are as follows:

active_objs
The number of objects that are currently active (i.e., in use).

num_objs
The total number of allocated objects (i.e., objects that are both in use and not in use).

objsize The size of objects in this slab, in bytes.

objperslab
The number of objects stored in each slab.

pagesperslab
The number of pages allocated for each slab.

The *tunables* entries in each line show tunable parameters for the corresponding cache. When using the default SLUB allocator, there are no tunables, the */proc/slabinfo* file is not writable, and the value 0 is shown in these fields. When using the older SLAB allocator, the tunables for a particular cache can be set by writing lines of the following form to */proc/slabinfo*:

```
# echo 'name limit batchcount sharedfactor' > /proc/slabinfo
```

Here, *name* is the cache name, and *limit*, *batchcount*, and *sharedfactor* are integers defining new values for the corresponding tunables. The *limit* value should be a positive value, *batchcount* should be a positive value that is less than or equal to *limit*, and *sharedfactor* should be nonnegative. If any of the specified values is invalid, the cache settings are left unchanged.

The *tunables* entries in each line contain the following fields:

limit The maximum number of objects that will be cached.

batchcount
: On SMP systems, this specifies the number of objects to transfer at one time when refilling the available object list.

sharedfactor
: [To be documented]

The *slabdata* entries in each line contain the following fields:

active_slabs
: The number of active slabs.

nums_slabs
: The total number of slabs.

sharedavail
: [To be documented]

Note that because of object alignment and slab cache overhead, objects are not normally packed tightly into pages. Pages with even one in-use object are considered in-use and cannot be freed.

Kernels configured with **CONFIG_DEBUG_SLAB** will also have additional statistics fields in each line, and the first line of the file will contain the string "(statistics)". The statistics field include : the high water mark of active objects; the number of times objects have been allocated; the number of times the cache has grown (new pages added to this cache); the number of times the cache has been reaped (unused pages removed from this cache); and the number of times there was an error allocating new pages to this cache.

VERSIONS

The */proc/slabinfo* file first appeared in Linux 2.1.23. The file is versioned, and over time there have been a number of versions with different layouts:

1.0 Present throughout the Linux 2.2.x kernel series.

1.1 Present in the Linux 2.4.x kernel series.

1.2 A format that was briefly present in the Linux 2.5 development series.

2.0 Present in Linux 2.6.x kernels up to and including Linux 2.6.9.

2.1 The current format, which first appeared in Linux 2.6.10.

NOTES

Only root can read and (if the kernel was configured with **CONFIG_SLAB**) write the */proc/slabinfo* file.

The total amount of memory allocated to the SLAB/SLUB cache is shown in the *Slab* field of */proc/meminfo*.

SEE ALSO

slabtop(1)

The kernel source file *Documentation/vm/slub.txt* and *tools/vm/slabinfo.c*.

COLOPHON

This page is part of release 5.05 of the Linux *man-pages* project. A description of the project, information about reporting bugs, and the latest version of this page, can be found at https://www.kernel.org/doc/man−pages/.