

Re: PROBLEM: ramfs reports 0 free

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-04/msg06622.html>

- *From:* Willy Tarreau <willy@xxxxxxxxxx>
 - *Date:* Sat, 29 Apr 2006 22:50:45 +0200
-

On Sat, Apr 29, 2006 at 07:52:16PM +0000, Senior Goat wrote:

[1.] One line summary of the problem:

ramfs erroneously reports 0 bytes free, which confuses some programs.

[2.] Full description of the problem/report:

The system call on a mounted ramfs, as indicated by df, reports 0 bytes total, used, and free. If ramfs is expected to perform like a filesystem, it should not do this.

(...)

[X.] Other notes, patches, fixes, workarounds:

I read that Linus himself wrote this module. For some reason he decided to report 0, but I can't figure why. Perhaps the overhead for finding information was too great.

Couldn't you add up the amount of filesystem cache with the free memory and get a crude, but quick estimate of the amount of free space available for any given ramfs.

I'm not sure how to handle the total space, since you probably don't want that fluctuating too much, except that you might just report the total amount of ram(which won't fluctuate), and then report the used ram. This is all information that /free/ reports from system calls with little delay.

The only other place I can think that ramfs might get memory is when the kernel swaps out other processes, but you can't count on that.

To sum it up, the best way to get a semi-valid report would be:

total: total ram

used: used ram (which takes into account memory used in the ramfs, coincidentally)

Re: PROBLEM: ramfs reports 0 free

free: total – used (which ignores disk cache, since that should be freed when needed)

To achieve this, you have to set an arbitrary limit on the maximal FS size. I have updated a patch originally from David Gibson to enforce limits on RAMFS. His original patch for 2.4.19-rc1-ac2 as well as my update for 2.4.32 is available here, in case you're interested in porting it to 2.6 :

<http://w.ods.org/linux/kernel/2.4/lkup/ramfs-limit.html>

The other option would be to make all the little programs (like Debian's package manager) check if the filesystem it wants to write to is a ramfs before reporting an error, but this is a blatant hack.

This is very hard to maintain, that's why I include the patch above in my kernels ;-)

Cheers,
Willy

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>