
Subject: Re: [PATCH] Cpu statistics accounting based on Paul Menage patches
Posted by [Andrew Morton](#) on Wed, 11 Apr 2007 18:49:27 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, 11 Apr 2007 19:02:27 +0400
Pavel Emelianov <xemul@sw.ru> wrote:

> Provides a per-container statistics concerning the numbers of tasks
> in various states, system and user times, etc. Patch is inspired
> by Paul's example of the used CPU time accounting. Although this
> patch is independent from Paul's example to make it possible playing
> with them separately.

Why is this actually needed? If userspace has a list of the tasks which are in a particular container, it can run around and add up the stats for those tasks without kernel changes?

It's a bit irksome that we have so much accounting of this form in core kernel, yet we have to go and add a completely new implementation to create something which is similar to what we already have. But I don't immediately see a fix for that. Apart from paragraph #1 ;)

Should there be linkage between per-container stats and delivery-via-taskstats? I can't think of one, really.

You have cpu stats. Later, presumably, we'll need IO stats, MM stats, context-switch stats, number-of-syscall stats, etc, etc. Are we going to reimplement all of those things as well? See paragraph #1!

Bottom line: I think we seriously need to find some way of consolidating per-container stats with our present per-task stats. Perhaps we should instead be looking at ways in which we can speed up paragraph #1.
