
Subject: Re: [RFC][PATCH] allow "unlimited" limit value.
Posted by [Paul Menage](#) on Tue, 25 Sep 2007 20:12:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 9/25/07, David Rientjes <rientjes@google.com> wrote:
> It doesn't matter. When I cat my cgroup's memory.limit (or
> memory.limit_in_bytes), I should see the total number of bytes that my
> applications are allowed. That's not an unrealistic expectation of a
> system that is expressly designed to control my memory. I don't want to
> see a value that is close to what I'm allowed, thanks.

So round up to the nearest page. Then you'll get what you asked for so
you can't get broken by the rounding.

>
> Storing it internally as the number of pages makes the implementation
> simpler since memory controls are only imposed on pages anyway and you get
> the added bonus of integer division truncating in C so that when you cat
> the file it will display the correct number of bytes modulo PAGE_SIZE.

Storing it internally as a number of pages is fine. I'm more concerned
about the userspace API, since that will be very hard to change.

>
> That fundamental unit being charged are pages,

No, that just happens to be the implementation mechanism in this controller.

You could be internally charging in much larger units (e.g. fake numa
nodes) or smaller units (slab object allocations).

Paul

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
