
Subject: Re: [RFC][PATCH] allow "unlimited" limit value.
Posted by [David Rientjes](#) on Tue, 25 Sep 2007 20:07:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 25 Sep 2007, Paul Menage wrote:

> > If I echo -n 8191 > memory.limit_in_bytes, I'm still only going to be able
> > to charge one page on my x86_64. And then my program's malloc(5000) is
> > going to fail, which leads to the inevitable head scratching.
>
> This is a very unrealistic argument. Page-size rounding really has no
> effect on any reasonable-sized memory cgroup.
>

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.

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.

> Expressing it in bytes seems reasonable to me, since they are after
> all the fundamental unit that's being counted ("kilobytes" are
> explicitly an aggregation of "bytes").
>

That fundamental unit being charged are pages, so any memory limit that has a finer granularity than kilobytes is just plain wrong.

David

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