

---

Subject: Re: [RFC][PATCH] allow "unlimited" limit value.  
Posted by [David Rientjes](#) on Tue, 25 Sep 2007 20:32:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Tue, 25 Sep 2007, Paul Menage 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.  
>

If you're fine with rounding up to the nearest page, then what's the point of exposing it as a number of bytes?? You'll never get a granularity finer than a kilobyte.

So by expressing it in terms of bytes instead of kilobytes, you're just making the largest amount of memory allowed via this interface smaller that it should have to be. That is absolutely horrid in terms of scalability and you're never going to be able to get rid of it because everything that interfaces with it by then will have been written in terms of bytes.

> > That fundamental unit being charged are pages,  
>  
> No, that just happens to be the implementation mechanism in this controller.  
>

And this controller owns the memory.limit file so it can express its memory limits in whatever unit it wants.

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---