

---

Subject: Re: [RFC][PATCH] allow "unlimited" limit value.  
Posted by [David Rientjes](#) on Tue, 25 Sep 2007 20:58:29 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Tue, 25 Sep 2007, Paul Menage wrote:

> > 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.  
>  
> API != implementation.  
>

Having the limit expressed and configurable in bytes suggests that it can be defined on that granularity which is obviously wrong.

> > 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 is should have to be.  
>  
> Yes, that's true. With a 64-bit count in bytes, we can only limit  
> people to 16 exabytes of memory. We're going to bump up against that  
> limit in no time.  
>

So, by your logic, it would be fine to express it in bits too.

> > And this controller owns the memory.limit file so it can express its  
> > memory limits in whatever unit it wants.  
> >  
>  
> Right, but it would be nice to have different memory controllers be  
> API-compatible with one another. Bytes is the lowest common  
> denominator.  
>

Please cite examples of other memory controllers that you can imagine would actually support (not expose to userspace, but support) memory limits in terms of anything smaller than kilobytes and how you plan on charging for that memory as a fraction of a page size and that has any reasonable hope of ever being efficient.

---

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

---