
Subject: Re: [RFC][PATCH] allow "unlimited" limit value.
Posted by [David Rientjes](#) on Tue, 25 Sep 2007 19:40:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, 25 Sep 2007, Paul Menage wrote:

> > nit pick, should be memory.limit_in_bytes
> >
>
> Can we reconsider this? I do think that plain "limit" would enable you
> to have a more consistent API across all resource counters users.
>

Why aren't limits expressed in kilobytes? All architectures have
PAGE_SIZE defined on that order.

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.

I think it would be best to express memory.limit in terms of KB, divide
that by PAGE_SIZE to store internally in res_counter.limit, deal with
charging for memory internally in terms of number of pages, and exposing
it back to userspace in terms of res_counter.limit * PAGE_SIZE (KB).

David

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