

---

Subject: Re: controlling mmap()'d vs read/write() pages  
Posted by [Nick Piggin](#) on Fri, 23 Mar 2007 05:57:40 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Eric W. Biederman wrote:

> Dave Hansen <[hansendc@us.ibm.com](mailto:hansendc@us.ibm.com)> writes:

>

>

>>So, I think we have a difference of opinion. I think it's \_all\_ about  
>>memory pressure, and you think it is \_not\_ about accounting for memory  
>>pressure. :) Perhaps we mean different things, but we appear to  
>>disagree greatly on the surface.

>

>

> I think it is about preventing a badly behaved container from having a  
> significant effect on the rest of the system, and in particular other  
> containers on the system.

That's Dave's point, I believe. Limiting mapped memory may be mostly OK for well behaved applications, but it doesn't do anything to stop bad ones from effectively DoSing the system or ruining any guarantees you might proclaim (not that hard guarantees are always possible without using virtualisation anyway).

This is why I'm surprised at efforts that go to such great lengths to get accounting "just right" (but only for mmaped memory). You may as well not even bother, IMO.

Give me an RSS limit big enough to run a couple of system calls and a loop...

--

SUSE Labs, Novell Inc.

Send instant messages to your online friends <http://au.messenger.yahoo.com>

---

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

---