
Subject: Re: [-mm PATCH] Memory controller improve user interface

Posted by [Paul Menage](#) on Wed, 29 Aug 2007 16:17:40 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 8/29/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

> >

> > This seems a bit inconsistent - if you write a value to a limit file,
> > then the value that you read back is reduced by a factor of 1024?

> > Having the "(kB)" suffix isn't really a big help to automated
> > middleware.

> >

>

> Why is that? Is it because you could write 4M and see it show up
> as 4096 kilobytes? We'll that can be fixed with another variant
> of the memparse() utility.

I was thinking the other way around - you can write 1048576 (i.e. 1MB)
to the file and read back 1024. It just seems to me that it's clearer
if you write X to the file to get X back.

>

> 64 bit might be an overkill for 32 bit machines. 32 bit machines with
> PAE cannot use 32 bit values, they need 64 bits.

How is using a 64-bit value for consistency overkill?

As someone pointed out, 4TB machines probably aren't that far around
the corner (if they're not here already) so even if you use KB rather
than bytes, userspace needs to be using an int64 for this value in
case it ends up running as a 32-bit-compiled app on a 64-bit kernel
with lots of memory.

Paul

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
