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

Posted by [Dave Hansen](#) on Wed, 29 Aug 2007 22:36:40 GMT

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

On Thu, 2007-08-30 at 03:57 +0530, Balbir Singh wrote:

> True, mmap() is a good example of such an interface for developers, I
> am not sure about system admins though.
>
> To quote Andrew
> <quote>
> Reporting tools could run getpagesize() and do the arithmetic, but we
> generally try to avoid exposing PAGE_SIZE, HZ, etc to userspace in this
> manner.
> </quote>

Well, rounding to PAGE_SIZE exposes PAGE_SIZE as well, just in a non-intuitive fashion. :)

If we're going to modify what the user specifies, we should probably at least mandate that writes are only a "suggestion" and users must read back the value to ensure what actually got committed.

If we're going to round in any direction, shouldn't we round up? If a user specifies 4097 bytes and uses two pages, we don't want to complain when they hit that second page.

-- Dave

Containers mailing list

Containers@lists.linux-foundation.org

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

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

Posted by [Balbir Singh](#) on Wed, 29 Aug 2007 22:44:15 GMT

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

Dave Hansen wrote:

> On Thu, 2007-08-30 at 03:57 +0530, Balbir Singh wrote:
>> True, mmap() is a good example of such an interface for developers, I
>> am not sure about system admins though.
>>
>> To quote Andrew
>> <quote>
>> Reporting tools could run getpagesize() and do the arithmetic, but we
>> generally try to avoid exposing PAGE_SIZE, HZ, etc to userspace in this
>> manner.

>> </quote>

>

> Well, rounding to PAGE_SIZE exposes PAGE_SIZE as well, just in a
> non-intuitive fashion. :)

>

Agreed, but the user might choose to ignore it altogether.

> If we're going to modify what the user specifies, we should probably at
> least mandate that writes are only a "suggestion" and users must read
> back the value to ensure what actually got committed.

>

Agreed, excellent suggestion!

> If we're going to round in any direction, shouldn't we round up? If a
> user specifies 4097 bytes and uses two pages, we don't want to complain
> when they hit that second page.

>

Absolutely, I used rounding to mean round up, truncation for rounding down.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

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