

---

Subject: Re: [-mm PATCH] Memory controller improve user interface  
Posted by [Dave Hansen](#) on Wed, 29 Aug 2007 22:18:14 GMT

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

---

On Thu, 2007-08-30 at 03:34 +0530, Balbir Singh wrote:

> I've thought about this before. The problem is that a user could  
> set his limit to 10000 bytes, but would then see the usage and  
> limit round to the closest page boundary. This can be confusing  
> to a user.

True, but we're lying if we allow a user to set their limit there,  
because we can't actually enforce a limit at 8,192 bytes vs 10,000.  
They're the same limit as far as the kernel is concerned.

Why not just -EINVAL if the value isn't page-aligned? There are plenty  
of interfaces in the kernel that require userspace to know the page  
size, so this shouldn't be too difficult.

-- Dave

---

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

---

---

Subject: Re: [-mm PATCH] Memory controller improve user interface  
Posted by [Paul Menage](#) on Wed, 29 Aug 2007 22:20:09 GMT

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

---

On 8/29/07, Dave Hansen <[haveblue@us.ibm.com](mailto:haveblue@us.ibm.com)> wrote:

> On Thu, 2007-08-30 at 03:34 +0530, Balbir Singh wrote:  
> > I've thought about this before. The problem is that a user could  
> > set his limit to 10000 bytes, but would then see the usage and  
> > limit round to the closest page boundary. This can be confusing  
> > to a user.

>  
> True, but we're lying if we allow a user to set their limit there,  
> because we can't actually enforce a limit at 8,192 bytes vs 10,000.  
> They're the same limit as far as the kernel is concerned.

>  
> Why not just -EINVAL if the value isn't page-aligned? There are plenty  
> of interfaces in the kernel that require userspace to know the page  
> size, so this shouldn't be too difficult.

I'd argue that having the user's specified limit be truncated to the  
page size is less confusing than giving an EINVAL if it's not page  
aligned.

Paul

---

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 [Dave Hansen](#) on Wed, 29 Aug 2007 22:25:52 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2007-08-29 at 15:20 -0700, Paul Menage wrote:

>  
> I'd argue that having the user's specified limit be truncated to the  
> page size is less confusing than giving an EINVAL if it's not page  
> aligned.

Do we truncate mmap() values to the nearest page so to not confuse the user? ;)

Imagine a careful application setting and accounting for limits on a long-running system. Might its internal accounting get sufficiently misaligned from the kernel's after a while to cause a problem? Truncating values like that would appear reserve significantly less memory than desired over a long period of time.

-- 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 22:27:07 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Dave Hansen wrote:

> On Thu, 2007-08-30 at 03:34 +0530, Balbir Singh wrote:  
>> I've thought about this before. The problem is that a user could  
>> set his limit to 10000 bytes, but would then see the usage and  
>> limit round to the closest page boundary. This can be confusing  
>> to a user.  
>  
> True, but we're lying if we allow a user to set their limit there,

> because we can't actually enforce a limit at 8,192 bytes vs 10,000.  
> They're the same limit as far as the kernel is concerned.  
>  
> Why not just -EINVAL if the value isn't page-aligned? There are plenty  
> of interfaces in the kernel that require userspace to know the page  
> size, so this shouldn't be too difficult.

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>

--

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>

---

---

Subject: Re: [-mm PATCH] Memory controller improve user interface  
Posted by [Balbir Singh](#) on Wed, 29 Aug 2007 22:37:11 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Dave Hansen wrote:

> On Wed, 2007-08-29 at 15:20 -0700, Paul Menage wrote:  
>> I'd argue that having the user's specified limit be truncated to the  
>> page size is less confusing than giving an EINVAL if it's not page  
>> aligned.  
>  
> Do we truncate mmap() values to the nearest page so to not confuse the  
> user? ;)  
>

I think rounding to the closest page size is a better option, but again it can be a bit confusing. I am all for using memparse() to parse the user input as a specification of the memory limit.

The second question of how to store it internally without truncation/rounding is something we need to agree upon. We also need to see

how to display the data back to the user.

I chose kilobytes for two reasons

1. Several people recommended it
2. Herbert mentioned that they've moved to that interface and it was working fine for them.

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

PS: I am going off to the web to search for some CUI/CLI guidelines.

---

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 [KAMEZAWA Hiroyuki](#) on Thu, 30 Aug 2007 05:38:59 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Thu, 30 Aug 2007 04:07:11 +0530  
Balbir Singh <balbir@linux.vnet.ibm.com> wrote:  
> 1. Several people recommended it  
> 2. Herbert mentioned that they've moved to that interface and it  
> was working fine for them.  
>

I have no strong opinion. But how about Mega bytes ? (too big ?)  
There will be no rounding up/down problem.

-Kame.

---

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 Thu, 30 Aug 2007 09:13:15 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

KAMEZAWA Hiroyuki wrote:

> On Thu, 30 Aug 2007 04:07:11 +0530

> Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>> 1. Several people recommended it

>> 2. Herbert mentioned that they've moved to that interface and it

>> was working fine for them.

>>

>

> I have no strong opinion. But how about Mega bytes ? (too big ?)

> There will be no rounding up/down problem.

>

Here is what I am thinking, allow the user to input bytes/kilobytes/  
megabytes or gigabytes. Store the data internally in kilobytes or  
PFN. I prefer kilobytes (no rounding issues), but while implementing  
limits we round up to the closest PFN.

--

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>

---