
Subject: [RFC][mm] [2/2] Simple stats for memory resource controller

Posted by [Balaji Rao](#) on Sat, 05 Apr 2008 18:10:02 GMT

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

This patch implements trivial statistics for the memory resource controller.

Signed-off-by: Balaji Rao <balajirao@gmail.com>

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

CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>

diff --git a/mm/memcontrol.c b/mm/memcontrol.c

index a860765..ca98b21 100644

--- a/mm/memcontrol.c

+++ b/mm/memcontrol.c

@@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {
 */

MEM_CGROUP_STAT_CACHE, /* # of pages charged as cache */

MEM_CGROUP_STAT_RSS, /* # of pages charged as rss */

+ MEM_CGROUP_STAT_PGPIN_COUNT, /* # of pages paged in */

+ MEM_CGROUP_STAT_PGPGOUT_COUNT, /* # of pages paged out */

MEM_CGROUP_STAT_NSTATS,

};
@@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup *mem,
int flags,

__mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);

else

__mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);

+

+ if (charge)

+ __mem_cgroup_stat_add_safe(stat,

+ MEM_CGROUP_STAT_PGPIN_COUNT, 1);

+ else

+ __mem_cgroup_stat_add_safe(stat,

+ MEM_CGROUP_STAT_PGPGOUT_COUNT, 1);

}

static struct mem_cgroup_per_zone *

@@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {
} mem_cgroup_stat_desc[] = {

[MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },

[MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },

+ [MEM_CGROUP_STAT_PGPIN_COUNT] = { "pgpin", 1, },

+ [MEM_CGROUP_STAT_PGPGOUT_COUNT] = { "pgpgout", 1, },

};

static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,

--

regards,
Balaji Rao
Dept. of Mechanical Engineering,
National Institute of Technology Karnataka, India

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

Subject: Re: [RFC][mm] [2/2] Simple stats for memory resource controller
Posted by [Balbir Singh](#) on Sun, 06 Apr 2008 06:23:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

Balaji Rao wrote:

```
> This patch implements trivial statistics for the memory resource controller.
>
> Signed-off-by: Balaji Rao <balajirao@gmail.com>
> CC: Balbir Singh <balbir@linux.vnet.ibm.com>
> CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>
>
> diff --git a/mm/memcontrol.c b/mm/memcontrol.c
> index a860765..ca98b21 100644
> --- a/mm/memcontrol.c
> +++ b/mm/memcontrol.c
> @@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {
>  */
>  MEM_CGROUP_STAT_CACHE, /* # of pages charged as cache */
>  MEM_CGROUP_STAT_RSS, /* # of pages charged as rss */
> + MEM_CGROUP_STAT_PGPIN_COUNT, /* # of pages paged in */
> + MEM_CGROUP_STAT_PGPGOUT_COUNT, /* # of pages paged out */
>
>  MEM_CGROUP_STAT_NSTATS,
> };
> @@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup *mem,
int flags,
>  __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);
>  else
>  __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);
> +
> + if (charge)
> +  __mem_cgroup_stat_add_safe(stat,
> +  MEM_CGROUP_STAT_PGPIN_COUNT, 1);
> + else
> +  __mem_cgroup_stat_add_safe(stat,
> +  MEM_CGROUP_STAT_PGPGOUT_COUNT, 1);
> }
```

```

>
> static struct mem_cgroup_per_zone *
> @@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {
> } mem_cgroup_stat_desc[] = {
> [MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },
> [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
> + [MEM_CGROUP_STAT_PGPGIN_COUNT] = {"pgpgin", 1, },
> + [MEM_CGROUP_STAT_PGPGOUT_COUNT] = {"pgpgout", 1, },
> };
>
> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
>

```

Looks good to me. I wonder if pgpin/pgpgout per second will be a useful metric, to see how fast things are changing within the cgroup.

Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

Let me test this patch right away.

--

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: [RFC][-mm] [2/2] Simple stats for memory resource controller
Posted by [Balbir Singh](#) on Mon, 14 Apr 2008 14:39:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

Balaji Rao wrote:

```

> This patch implements trivial statistics for the memory resource controller.
>
> Signed-off-by: Balaji Rao <balajirao@gmail.com>
> CC: Balbir Singh <balbir@linux.vnet.ibm.com>
> CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>
>
> diff --git a/mm/memcontrol.c b/mm/memcontrol.c
> index a860765..ca98b21 100644
> --- a/mm/memcontrol.c
> +++ b/mm/memcontrol.c
> @@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {
> */

```

```

> MEM_CGROUP_STAT_CACHE, /* # of pages charged as cache */
> MEM_CGROUP_STAT_RSS, /* # of pages charged as rss */
> + MEM_CGROUP_STAT_PGPIN_COUNT, /* # of pages paged in */
> + MEM_CGROUP_STAT_PPGOUT_COUNT, /* # of pages paged out */
>
> MEM_CGROUP_STAT_NSTATS,
> };
> @@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup *mem,
int flags,
> __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);
> else
> __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);
> +
> + if (charge)
> + __mem_cgroup_stat_add_safe(stat,
> + MEM_CGROUP_STAT_PGPIN_COUNT, 1);
> + else
> + __mem_cgroup_stat_add_safe(stat,
> + MEM_CGROUP_STAT_PPGOUT_COUNT, 1);
> }
>
> static struct mem_cgroup_per_zone *
> @@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {
> } mem_cgroup_stat_desc[] = {
> [MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },
> [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
> + [MEM_CGROUP_STAT_PGPIN_COUNT] = {"pgpin", 1, },
> + [MEM_CGROUP_STAT_PPGOUT_COUNT] = {"pgpgout", 1, },
> };
>
> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
>

```

Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>

Hi, Andrew,

Could you please include these statistics in -mm.

Balbir

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

Subject: Re: [RFC][mm] [2/2] Simple stats for memory resource controller
Posted by [Balbir Singh](#) on Mon, 14 Apr 2008 14:40:52 GMT

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

Balbir Singh wrote:

> Balaji Rao wrote:

>> This patch implements trivial statistics for the memory resource controller.

>>

>> Signed-off-by: Balaji Rao <balajirao@gmail.com>

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

>> CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>

>>

>> diff --git a/mm/memcontrol.c b/mm/memcontrol.c

>> index a860765..ca98b21 100644

>> --- a/mm/memcontrol.c

>> +++ b/mm/memcontrol.c

>> @@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {

>> */

>> MEM_CGROUP_STAT_CACHE, /* # of pages charged as cache */

>> MEM_CGROUP_STAT_RSS, /* # of pages charged as rss */

>> + MEM_CGROUP_STAT_PGPGIN_COUNT, /* # of pages paged in */

>> + MEM_CGROUP_STAT_PGPGOUT_COUNT, /* # of pages paged out */

>>

>> MEM_CGROUP_STAT_NSTATS,

>> };

>> @@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup
*mem, int flags,

>> __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);

>> else

>> __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);

>> +

>> + if (charge)

>> + __mem_cgroup_stat_add_safe(stat,

>> + MEM_CGROUP_STAT_PGPGIN_COUNT, 1);

>> + else

>> + __mem_cgroup_stat_add_safe(stat,

>> + MEM_CGROUP_STAT_PGPGOUT_COUNT, 1);

>> }

>>

>> static struct mem_cgroup_per_zone *

>> @@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {

>> } mem_cgroup_stat_desc[] = {

>> [MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },

```
>> [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
>> + [MEM_CGROUP_STAT_PGPIN_COUNT] = {"pgpin", 1, },
>> + [MEM_CGROUP_STAT_PPGOUT_COUNT] = {"pgpgout", 1, },
>> };
>>
>> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
>>
>
> Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>
>
> Hi, Andrew,
>
> Could you please include these statistics in -mm.
>
> Balbir
>
>
```

Andrew,

This should have been addressed to you in the first place. Resending.

--

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: [RFC][-mm] [2/2] Simple stats for memory resource controller
Posted by [Balaji Rao](#) on Mon, 28 Apr 2008 16:00:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Monday 14 April 2008 08:09:48 pm Balbir Singh wrote:

```
> Balaji Rao wrote:
> > This patch implements trivial statistics for the memory resource controller.
> >
> > Signed-off-by: Balaji Rao <balajirao@gmail.com>
> > CC: Balbir Singh <balbir@linux.vnet.ibm.com>
> > CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>
> >
> > diff --git a/mm/memcontrol.c b/mm/memcontrol.c
> > index a860765..ca98b21 100644
> > --- a/mm/memcontrol.c
```

```

>> +++ b/mm/memcontrol.c
>> @@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {
>>  */
>>  MEM_CGROUP_STAT_CACHE, /* # of pages charged as cache */
>>  MEM_CGROUP_STAT_RSS, /* # of pages charged as rss */
>> + MEM_CGROUP_STAT_PGPGIN_COUNT, /* # of pages paged in */
>> + MEM_CGROUP_STAT_PGPGOUT_COUNT, /* # of pages paged out */
>>
>>  MEM_CGROUP_STAT_NSTATS,
>> };
>> @@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup
*mem, int flags,
>>  __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);
>>  else
>>  __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);
>> +
>> + if (charge)
>> + __mem_cgroup_stat_add_safe(stat,
>> + MEM_CGROUP_STAT_PGPGIN_COUNT, 1);
>> + else
>> + __mem_cgroup_stat_add_safe(stat,
>> + MEM_CGROUP_STAT_PGPGOUT_COUNT, 1);
>> }
>>
>> static struct mem_cgroup_per_zone *
>> @@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {
>> } mem_cgroup_stat_desc[] = {
>> [MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },
>> [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
>> + [MEM_CGROUP_STAT_PGPGIN_COUNT] = {"pgpgin", 1, },
>> + [MEM_CGROUP_STAT_PGPGOUT_COUNT] = {"pgpgout", 1, },
>> };
>>
>> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
>>
>
> Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>
>
> Hi, Andrew,
>
> Could you please include these statistics in -mm.
>
> Balbir
>
>
Hi Andrew,

```

Now that Balbir Singh has ACKed it, could you please include it in -mm ?

--

Warm Regards,

Balaji Rao
Dept. of Mechanical Engineering
NITK

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

Subject: Re: [RFC][mm] [2/2] Simple stats for memory resource controller
Posted by [akpm](#) on Mon, 28 Apr 2008 16:40:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 28 Apr 2008 21:30:29 +0530 Balaji Rao <balajirao@gmail.com> wrote:

```
> On Monday 14 April 2008 08:09:48 pm Balbir Singh wrote:
> > Balaji Rao wrote:
> > > This patch implements trivial statistics for the memory resource controller.
> > >
> > > Signed-off-by: Balaji Rao <balajirao@gmail.com>
> > > CC: Balbir Singh <balbir@linux.vnet.ibm.com>
> > > CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>
> > >
> > > diff --git a/mm/memcontrol.c b/mm/memcontrol.c
> > > index a860765..ca98b21 100644
> > > --- a/mm/memcontrol.c
> > > +++ b/mm/memcontrol.c
> > > @@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {
> > >     */
> > >     MEM_CGROUP_STAT_CACHE,    /* # of pages charged as cache */
> > >     MEM_CGROUP_STAT_RSS,      /* # of pages charged as rss */
> > > + MEM_CGROUP_STAT_PGPGIN_COUNT, /* # of pages paged in */
> > > + MEM_CGROUP_STAT_PGPGOUT_COUNT, /* # of pages paged out */
> > >
> > >     MEM_CGROUP_STAT_NSTATS,
> > > };
> > > @@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup
> > > *mem, int flags,
> > >     __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);
> > >     else
> > >     __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);
> > > +
> > > + if (charge)
> > > +     __mem_cgroup_stat_add_safe(stat,
```

```

>>> + MEM_CGROUP_STAT_PGPIN_COUNT, 1);
>>> + else
>>> + __mem_cgroup_stat_add_safe(stat,
>>> + MEM_CGROUP_STAT_PGPGOUT_COUNT, 1);
>>> }
>>>
>>> static struct mem_cgroup_per_zone *
>>> @@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {
>>> } mem_cgroup_stat_desc[] = {
>>> [MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },
>>> [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
>>> + [MEM_CGROUP_STAT_PGPIN_COUNT] = {"pgpgin", 1, },
>>> + [MEM_CGROUP_STAT_PGPGOUT_COUNT] = {"pgpgout", 1, },
>>> };
>>>
>>> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
>>>
>>
>> Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>
>>
>> Hi, Andrew,
>>
>> Could you please include these statistics in -mm.
>>
>> Balbir
>>
>>
> Hi Andrew,
>
> Now that Balbir Singh has ACKed it, could you please include it in -mm ?

```

<looks>

I guess we can add this one, sure. But [patch 1/2] needs work.

- The local_irq_save()-around-for_each_possible_cpu() locking doesn't make sense.
- indenting is busted in account_user_time() and account_system_time()
- The use of for_each_possible_cpu() can be grossly inefficient. It would be preferred to use for_each_possible_cpu() and add a cpu-hotplug notifier.
- The proposed newly-added userspace interfaces are undocumented
- The changelogs don't explain why we might want this feature in Linux.

- Generally: there are a heck of a lot of different ways of accounting for things in core kernel and it's really sad to see yet another one being added.

Actually, [patch 2/2] adds new kernel->user interfaces and doesn't document them. But afaict the existing memcgroup stats are secret too.

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

Subject: Re: [RFC][mm] [2/2] Simple stats for memory resource controller
Posted by [Balbir Singh](#) on Mon, 28 Apr 2008 18:03:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

Andrew Morton wrote:

```
> On Mon, 28 Apr 2008 21:30:29 +0530 Balaji Rao <balajirao@gmail.com> wrote:
>
>> On Monday 14 April 2008 08:09:48 pm Balbir Singh wrote:
>>> Balaji Rao wrote:
>>>> This patch implements trivial statistics for the memory resource controller.
>>>>
>>>> Signed-off-by: Balaji Rao <balajirao@gmail.com>
>>>> CC: Balbir Singh <balbir@linux.vnet.ibm.com>
>>>> CC: Dhaval Giani <dhaval@linux.vnet.ibm.com>
>>>>
>>>> diff --git a/mm/memcontrol.c b/mm/memcontrol.c
>>>> index a860765..ca98b21 100644
>>>> --- a/mm/memcontrol.c
>>>> +++ b/mm/memcontrol.c
>>>> @@ -47,6 +47,8 @@ enum mem_cgroup_stat_index {
>>>>  */
>>>>  MEM_CGROUP_STAT_CACHE,    /* # of pages charged as cache */
>>>>  MEM_CGROUP_STAT_RSS,     /* # of pages charged as rss */
>>>> + MEM_CGROUP_STAT_PGPIN_COUNT, /* # of pages paged in */
>>>> + MEM_CGROUP_STAT_PGPGOUT_COUNT, /* # of pages paged out */
>>>>
>>>>  MEM_CGROUP_STAT_NSTATS,
>>>> };
>>>> @@ -198,6 +200,13 @@ static void mem_cgroup_charge_statistics(struct mem_cgroup
>>>> *mem, int flags,
>>>>  __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_CACHE, val);
>>>>  else
>>>>  __mem_cgroup_stat_add_safe(stat, MEM_CGROUP_STAT_RSS, val);
>>>> +
```

```

>>>> + if (charge)
>>>> + __mem_cgroup_stat_add_safe(stat,
>>>> + MEM_CGROUP_STAT_PGPIN_COUNT, 1);
>>>> + else
>>>> + __mem_cgroup_stat_add_safe(stat,
>>>> + MEM_CGROUP_STAT_PPGOUT_COUNT, 1);
>>>> }
>>>>
>>>> static struct mem_cgroup_per_zone *
>>>> @@ -897,6 +906,8 @@ static const struct mem_cgroup_stat_desc {
>>>> } mem_cgroup_stat_desc[] = {
>>>> [MEM_CGROUP_STAT_CACHE] = { "cache", PAGE_SIZE, },
>>>> [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },
>>>> + [MEM_CGROUP_STAT_PGPIN_COUNT] = {"pgpgin", 1, },
>>>> + [MEM_CGROUP_STAT_PPGOUT_COUNT] = {"pgpgout", 1, },
>>>> };
>>>>
>>>> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
>>>>
>>> Acked-by: Balbir Singh <balbir@linux.vnet.ibm.com>
>>>
>>> Hi, Andrew,
>>>
>>> Could you please include these statistics in -mm.
>>>
>>> Balbir
>>>
>>>
>> Hi Andrew,
>>
>> Now that Balbir Singh has ACKed it, could you please include it in -mm ?
>
> <looks>
>
> I guess we can add this one, sure. But [patch 1/2] needs work.
>
> - The local_irq_save()-around-for_each_possible_cpu() locking doesn't
> make sense.
>

```

Yes, that needs re-work. Peter Zijlstra had detailed review comments for the patch

```

> - indenting is busted in account_user_time() and account_system_time()
>
> - The use of for_each_possible_cpu() can be grossly inefficient. It
> would be preferred to use for_each_possible_cpu() and add a cpu-hotplug
> notifier.
>

```

> - The proposed newly-added userspace interfaces are undocumented
>

Yes, we need more documentation

> - The changelogs don't explain why we might want this feature in Linux.
>

We need more accurate utime/stime per cgroup. Summing them in user space is insufficient, since tasks can move across groups and what we have is accumulated time per task.

> - Generally: there are a heck of a lot of different ways of accounting
> for things in core kernel and it's really sad to see yet another one
> being added.
>

We thought of summing up stuff in user space, we've look harder. The plan is to finally send all the data using cgroupstats.

>
> Actually, [patch 2/2] adds new kernel->user interfaces and doesn't document
> them. But afaict the existing memcgroup stats are secret too.
>

The statistics was added as a part of git commit
d52aa412d43827033a8e2ce4415ef6e8f8d53635. I'll go ahead and try to document
them. These patches piggy back on the statistics patches and add pagein/pageout
counts, which is a useful statistic for the memory controller.

--

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: [RFC][mm] [2/2] Simple stats for memory resource controller
Posted by [KAMEZAWA Hiroyuki](#) on Tue, 29 Apr 2008 05:39:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 28 Apr 2008 23:33:05 +0530

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

> The statistics was added as a part of git commit
> d52aa412d43827033a8e2ce4415ef6e8f8d53635. I'll go ahead and try to document
> them. These patches piggy back on the statistics patches and add pagein/pageout
> counts, which is a useful statistic for the memory controller.
>

I'm sorry for lack of documentation.

BTW, when you adds documentation on this pagein/pageout, please make it clear

1. pagein means a new page is newly accouted to this cgrop, doesn't means a page is read from disk.
2. pageout means a page is dropped from this cgroup, doesn't means a page is written into disk.

So, pagein/pageout is a bit ambiguous. I'm happy if you find better names.

Thanks,
-Kame

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