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

Posted by [Balaji Rao](#) on Wed, 26 Mar 2008 18:18:59 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 eb681a6..84f3fe5 100644
--- a/mm/memcontrol.c
+++ b/mm/memcontrol.c
@@ -45,6 +45,8 @@ 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_PGIN_COUNT, /* # of pages paged in */
+ MEM_CGROUP_STAT_PGOOUT_COUNT, /* # of pages paged out */

 MEM_CGROUP_STAT_NSTATS,
};
@@ -196,6 +198,13 @@ 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_PGIN_COUNT, 1);
+ else
+     __mem_cgroup_stat_add_safe(stat,
+         MEM_CGROUP_STAT_PGOOUT_COUNT, 1);
}

static struct mem_cgroup_per_zone *
@@ -886,6 +895,8 @@ 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_PGIN_COUNT] = {"page_in_count", 1, },
+ [MEM_CGROUP_STAT_PGOOUT_COUNT] = {"page_out_count", 1, },
};

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

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 [Paul Menage](#) on Wed, 26 Mar 2008 18:52:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Mar 26, 2008 at 11:18 AM, Balaji Rao <balajirao@gmail.com> wrote:

```
>     [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },  
> +     [MEM_CGROUP_STAT_PGIN_COUNT] = {"page_in_count", 1, },  
> +     [MEM_CGROUP_STAT_PGOUT_COUNT] = {"page_out_count", 1, },  
> };  
>
```

Should these be called "pgpgin" and "pgpgout" for consistency with /proc/vmstat?

Paul

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 Wed, 26 Mar 2008 18:53:36 GMT

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

Paul Menage wrote:

> On Wed, Mar 26, 2008 at 11:18 AM, Balaji Rao <balajirao@gmail.com> wrote:

```
>>     [MEM_CGROUP_STAT_RSS] = { "rss", PAGE_SIZE, },  
>> +     [MEM_CGROUP_STAT_PGIN_COUNT] = {"page_in_count", 1, },  
>> +     [MEM_CGROUP_STAT_PGOUT_COUNT] = {"page_out_count", 1, },  
>> };  
>>  
>
```

> Should these be called "pgpgin" and "pgpgout" for consistency with /proc/vmstat?

>

> Paul

Yes, that is a good name. We also want to instantaneous metrics like page_in_count per second and page_out_count per second to help the administrator monitor what is going on within the tasks in the cgroup w.r.t. memory.

--
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 Wed, 26 Mar 2008 18:54:42 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 eb681a6..84f3fe5 100644
> --- a/mm/memcontrol.c
> +++ b/mm/memcontrol.c
> @@ -45,6 +45,8 @@ 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_PGIN_COUNT, /* # of pages paged in */
> + MEM_CGROUP_STAT_PGOUT_COUNT, /* # of pages paged out */
>
> MEM_CGROUP_STAT_NSTATS,
> };
> @@ -196,6 +198,13 @@ 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_PGIN_COUNT, 1);
> + else
> + __mem_cgroup_stat_add_safe(stat,
> + MEM_CGROUP_STAT_PGOUT_COUNT, 1);

```
> }
>
> static struct mem_cgroup_per_zone *
> @@ -886,6 +895,8 @@ 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_PGIN_COUNT] = {"page_in_count", 1, },
>     [MEM_CGROUP_STAT_PGOUT_COUNT] = {"page_out_count", 1, },
> };
>
> static int mem_control_stat_show(struct cgroup *cont, struct cftype *cft,
```

Looks simple and nice. Please see Paul's mail on the suggested name change as well

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

--
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>
