
Subject: Re: [PATCH 3/3] memcgou: allow memory.failcnt to be reset

Posted by [Balbir Singh](#) on Tue, 11 Mar 2008 10:14:49 GMT

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

KAMEZAWA Hiroyuki wrote:

> On Tue, 11 Mar 2008 19:09:02 +0900

> Li Zefan <lizf@cn.fujitsu.com> wrote:

>

>> Allow memory.failcnt to be reset to 0:

>>

>> echo 0 > memory.failcnt

>>

>> And '0' is the only valid value.

>>

> Can't this be generic resource counter function ?

>

I was about to suggest a generic cgroup option, since we do reset values even for the cpu accounting subsystem.

--

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: [PATCH 3/3] memcgou: allow memory.failcnt to be reset

Posted by [KAMEZAWA Hiroyuki](#) on Tue, 11 Mar 2008 10:15:17 GMT

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

On Tue, 11 Mar 2008 19:09:02 +0900

Li Zefan <lizf@cn.fujitsu.com> wrote:

> Allow memory.failcnt to be reset to 0:

>

> echo 0 > memory.failcnt

>

> And '0' is the only valid value.

>

Can't this be generic resource counter function ?

Thanks,

-Kame

> This is useful when testing or observing the memory resource
> controller. Without this function, one will have to remember
> the previous failcnt to decide whether memory reclaim has
> happened *again*.

>
> Signed-off-by: Li Zefan <lizf@cn.fujitsu.com>
> ---
> Documentation/controllers/memory.txt | 4 +---
> mm/memcontrol.c | 15 ++++++-----
> 2 files changed, 18 insertions(+), 1 deletions(-)
>
> diff --git a/Documentation/controllers/memory.txt b/Documentation/controllers/memory.txt
> index 866b9cd..28f80e3 100644
> --- a/Documentation/controllers/memory.txt
> +++ b/Documentation/controllers/memory.txt
> @@ -194,7 +194,9 @@ this file after a write to guarantee the value committed by the kernel.
> 4096
>
> The memory.failcnt field gives the number of times that the cgroup limit was
> -exceeded.
> +exceeded. It can be reset.
> +
> +# echo 0 > memory.failcnt
>
> The memory.stat file gives accounting information. Now, the number of
> caches, RSS and Active pages/Inactive pages are shown.
> diff --git a/mm/memcontrol.c b/mm/memcontrol.c
> index 6145031..fd26dc2 100644
> --- a/mm/memcontrol.c
> +++ b/mm/memcontrol.c
> @@ -883,6 +883,20 @@ static int mem_force_empty_write(struct cgroup *cont, struct cftype
*cft,
> return ret;
> }
>
> +static int mem_failcnt_write(struct cgroup *cont, struct cftype *cft,
> + u64 val)
> +{
> + struct res_counter *counter;
> +
> + if (val != 0)
> + return -EINVAL;
> +
> + counter = &mem_cgroup_from_cont(cont)->res;
> + res_counter_write_u64(counter, cft->private, 0);
> +

```
> + return 0;
> +}
> +
> static const struct mem_cgroup_stat_desc {
>     const char *msg;
>     u64 unit;
> @@ -934,6 +948,7 @@ static struct cftype mem_cgroup_files[] = {
>     {
>         .name = "failcnt",
>         .private = RES_FAILCNT,
> +        .write_u64 = mem_failcnt_write,
>         .read_u64 = mem_cgroup_read,
>     },
>     {
>     --
> 1.5.4.rc3
>
```

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

Subject: Re: [PATCH 3/3] memcgoup: allow memory.failcnt to be reset
Posted by [Li Zefan](#) on Thu, 13 Mar 2008 23:37:10 GMT

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

Balbir Singh wrote:
> KAMEZAWA Hiroyuki wrote:
>> On Tue, 11 Mar 2008 19:09:02 +0900
>> Li Zefan <lizf@cn.fujitsu.com> wrote:
>>
>>> Allow memory.failcnt to be reset to 0:
>>>
>>> echo 0 > memory.failcnt
>>>
>>> And '0' is the only valid value.
>>>
>> Can't this be generic resource counter function ?
>>
>
> I was about to suggest a generic cgroup option, since we do reset values even
> for the cpu accounting subsystem.
>

It won't help. You still have to write the write function, and you have to call
some res_counter routines to reset the value, and maybe also do some other

work.

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