Subject: Re: [RFC][PATCH] allow "unlimited" limit value.
Posted by Balbir Singh on Tue, 25 Sep 2007 19:21:59 GMT
View Forum Message <> Reply to Message

David Rientjes wrote:
> On Wed, 26 Sep 2007, KAMEZAWA Hiroyuki wrote:
>
>>>> #define RES_COUNTER_INIFINITY (~0ULL)
>>>> or some nice name
>>> Why do we need this at all? We can simply push -1 there and be happy.
>>>
>> Hm, can this work now ?
>> ==
>> echo -1 > /cgroup/memory.limit_in_bytes
>> ==
>> Or users have to do following for unlimit resource ?
>> ==
>> echo some-very-very-big-number > /cgroup/memory.limit_in_bytes
>>
>>
>> I just think when some special value "-1" has a nice nick name, users will
>> be happy. If I'm a novice user, I don't imagine I can write -1 to limit value.
>> (but ok, tools can hide it for them.)
>>
>
> Please simply use 0 to denote unconstrained memory, it's quite obvious
> that nobody will sanely attach tasks to a cgroup that has no bytes of
> memory allowed.
>

Yes, I prefer 0 as well and had that in a series in the Lost World
of my earlier memory/RSS controller patches. I feel now that 0 is
a bit confusing, we don't use 0 to mean unlimited, unless we
treat the memory.limit_in_bytes value as boolean. 0 is false,
meaning there is no limit, > 0 is true, which means the limit
is set and the value is specified to the value read out.

> diff --git a/kernel/res_counter.c b/kernel/res_counter.c
> --- a/kernel/res_counter.c
> +++ b/kernel/res_counter.c
> @@ -16,12 +16,15 @@
> void res_counter_init(struct res_counter *counter)
> {
>   spin_lock_init(&counter->lock);
> - counter->limit = (unsigned long)LONG_MAX;

So, we create all containers with infinite limit?

> }
>
> int res_counter_charge_locked(struct res_counter *counter, unsigned long val)
> {
> - if (counter->usage + val > counter->limit) {
> + /*
> + * If 'memory.limit' is set to 0, there is no charge to this

nit pick, should be memory.limit_in_bytes

> + * res_counter.
> + */
> + if (counter->limit && counter->usage + val > counter->limit) {
> counter->failcnt++;
> return -ENOMEM;
> }

--
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