

---

Subject: Re: [PATCH 2/2] Make res\_counter hierarchical  
Posted by [KAMEZAWA Hiroyuki](#) on Tue, 11 Mar 2008 09:39:03 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Tue, 11 Mar 2008 14:46:58 +0530  
Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

> Paul Menage wrote:  
> > On Tue, Mar 11, 2008 at 2:13 AM, KAMEZAWA Hiroyuki  
> > <kamezawa.hiroyu@jp.fujitsu.com> wrote:  
> >> or remove all relationship among counters of \*different\* type of resources.  
> >> user-land-daemon will do enough jobs.  
> >>  
> >  
> > Yes, that would be my preferred choice, if people agree that  
> > hierarchically limiting overall virtual memory isn't useful. (I don't  
> > think I have a use for it myself).  
> >  
>  
> Virtual limits are very useful. I have a patch ready to send out.  
> They limit the amount of paging a cgroup can do (virtual limit - RSS limit).  
> Some times end users want to set virtual limit == RSS limit, so that the cgroup  
> OOMs on cross the RSS limit.  
>  
I have no objection to adding virtual limit itself.  
(It can be considered as extended ulimit.)

But if you'd like to add relationship between virtual-limit/memory-usage-limit,  
please take care to make it clear that relationship is reasonable.

- memory-usage includes page-cache.
- memory-usage doesn't include hugepages.
- How to treat MAP\_NORESERVE is depends on over-commit-memory type.  
how cgroup does ?
- shared memory will be counted per mmap.

Thanks,  
-Kame

---

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

---