
Subject: Re: Re: [RFC/PATCH] cgroup swap subsystem
Posted by [Pavel Emelianov](#) on Thu, 06 Mar 2008 08:50:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On Thu, Mar 6, 2008 at 12:38 AM, Pavel Emelyanov <xemul@openvz.org> wrote:

>> > Hierarchical res_counter makes sense.

>> > Making it in simple/reasonable style will be our challenge.

>>

>> I have this in my TODO list. Since this is not so urgent, then if you

>> don't mind I can prepare the patches next week - after I set the git

>> tree up. This change doesn't seem that big.

>>

>

> The change that you're referring to is allowing a cgroup to have a

> total memory limit for itself and all its children, and then giving

> that cgroup's children separate memory limits within that overall

> limit?

Yup. Isn't this reasonable?

Without this, if I'm a task in a 1GB limited cgroup, I can create a new one, set 2GB limit and spawn a kid into it (or move there myself) and be happy with 2GB of memory... With the proposed change, even if I set a 2GB for a subgroup it will not pass _my_ (1GB) limit.

> Paul

>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
