
Subject: Re: [RFC][PATCH] another swap controller for cgroup
Posted by [Paul Menage](#) on Wed, 14 May 2008 03:27:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, May 13, 2008 at 8:21 PM, YAMAMOTO Takashi
<yamamoto@valinux.co.jp> wrote:

>
> note that a failure can affect other subsystems which belong to
> the same hierarchy as well, and, even worse, a back-out attempt can also fail.
> i'm afraid that we need to play some kind of transaction-commit game,
> which can make subsystems too complex to implement properly.
>

I was considering something like that - every call to `can_attach()` would be guaranteed to be followed by either a call to `attach()` or to a new method called `cancel_attach()`. Then the subsystem would just need to ensure that nothing could happen which would cause the attach to become invalid between the two calls.

Or possibly, since for some subsystems that might involve holding a spinlock, we should extend it to:

After a successful call to `can_attach()`, either `abort_attach()` or `commit_attach()` will be called; these calls are not allowed to sleep, and `cgroup.c` will not sleep between calls.. If `commit_attach()` is called, it will be followed shortly by `attach()`, which is allowed to sleep.

Paul

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