Subject: Re: [PATCH 2/3] make clone_children a flag Posted by Glauber Costa on Wed, 14 Dec 2011 07:09:14 GMT

View Forum Message <> Reply to Message

On 12/14/2011 06:29 AM, Li Zefan wrote:

- > Tejun Heo wrote:
- >> On Sun, Dec 11, 2011 at 03:45:37PM +0100, Glauber Costa wrote:
- >>> There is no reason to have a flags field, and then a separate
- >>> bool field just to indicate if the clone children flag is set.
- >>> Make it a flag
- >>>
- >>> Signed-off-by: Glauber Costa<glommer@parallels.com>
- >>
- >> Doesn't this change how remount conditions are checked?

>>

Well, I was thinking it wouldn't, because I patched all callers. But I forget life is not always that simple: After you mentioned, I checked and we do test for changes in the flag field explicitly on remount. So I missed that, indeed.

- > Right. Currently we can do this:
- >
- > # mount -t cgroup xxx /mnt
- > # mount -o remount, clone_children /mnt

>

> with this patch, the above remount will fail.

>

- > But..the current bevaiour of remount is a bit confusing in that remount
- > with/without "clone_children" has no effect on anything:

>

- > # mount -t cgroup -o clone_children xxx /mnt
- > # cat /mnt/cgroup.clone_children
- > 1
- > # mount -o remount xxx /mnt
- > # mount | grep cgroup
- > xxx on /mnt type cgroup (rw,clone_children)
- > # cat /mnt/cgroup.clone_children
- > 1

That's indeed confusing, and it comes from the fact that we always inherit clone_children from the parent - which is sane, IMHO. So this flag only has any value in establishing the initial behaviour of the top root cgroup. I wonder then if it wouldn't better to just be explicit and fail in this case?