Subject: Re: [PATCH 0/4] - v2 - Object creation with a specified id
Posted by Pavel Machek on Wed, 23 Apr 2008 14:23:51 GMT
View Forum Message <> Reply to Message

Hi!

> When restarting a process that has been previously checkpointed, that process
> should keep on using some of its ids (such as its process id, or sysV ipc ids).
>
> This patch provides a feature that can help ensuring this saved state reuse:
> it makes it possible to create an object with a pre-defined id.
>
> A first implementation had been proposed 2 months ago. It consisted in
> changing an object's id after it had been created.
>
> Here is a second implementation based on Oren Ladaan's idea: Oren's suggestion
> was to force an object's id during its creation, rather than 1. create it,
> 2. change its id.
>
> A new file is created in procfs: /proc/self/task/<my_tid>/next_id.
> This makes it possible to avoid races between several threads belonging to
> the same process.

Ugly...

> When this file is filled with an id value, a structure pointed to by the
> calling task_struct is filled with that id.
>
> Then, when an object supporting this feature is created, the id present in
> that new structure is used, instead of the default one.
>
> The syntax is one of:
>   . echo "LONG XX" > /proc/self/task/<my_tid>/next_id
>   next object to be created will have an id set to XX
>   . echo "LONG<n> X0 ... X<n-1>" > /proc/self/task/<my_tid>/next_id
>   next object to be created will have its ids set to XX0, ... X<n-1>
>   This is particularly useful for processes that may have several ids if
>   they belong to nested namespaces.

So ugly it is not even funny.

Create fork_with_pid(int pid) and corresponding ipc primitives, if
you have to...
--
(english) http://www.livejournal.com/~pavelmachek
(cesky, pictures) http://atrey.karlin.mff.cuni.cz/~pavel/picture/horses/blog.html

Containers mailing list

Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers