

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/next_id.

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/next_id
next object to be created will have an id set to XX
- . echo "LONG<n> X0 ... X<n-1>" > /proc/self/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.

The objects covered here are ipc objects and processes.

Today, the ids are specified as long, but having a type string specified in the next_id file makes it possible to cover more types in the future, if needed.

The patches are against 2.6.25-rc3-mm1, in the following order:

[PATCH 1/4] adds the procfs facility for next object to be created, this object being associated to a single id.

[PATCH 2/4] enhances the procfs facility for objects associated to multiple ids (like processes).

[PATCH 3/4] makes use of the specified id (if any) to allocate the new IPC

object (changes the ipc_addid() path).
[PATCH 4/4] uses the specified id(s) (if any) to set the upid nr(s) for a newly
allocated process (changes the alloc_pid()/alloc_pidmap() paths).

Any comment and/or suggestions are welcome.

Regards,
Nadia

--
--

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