
Subject: Re: - merge-sys_clone-sys_unshare-nsproxy-and-namespace.patch
removed from -mm tree

Posted by [Herbert Poetzl](#) on Sat, 16 Jun 2007 19:17:42 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, May 08, 2007 at 07:45:35PM -0700, akpm@linux-foundation.org wrote:

>
> The patch titled
> Merge sys_clone()/sys_unshare() nsproxy and namespace handling
> has been removed from the -mm tree. Its filename was
> merge-sys_clone-sys_unshare-nsproxy-and-namespace.patch
>
> This patch was dropped because it was merged into mainline or a subsystem tree
>
> -----
> Subject: Merge sys_clone()/sys_unshare() nsproxy and namespace handling
> From: Badari Pulavarty <pbadari@us.ibm.com>
>
> sys_clone() and sys_unshare() both makes copies of nsproxy and its associated
> namespaces. But they have different code paths.
>
> This patch merges all the nsproxy and its associated namespace copy/clone
> handling (as much as possible). Posted on container list earlier for
> feedback.
>
>
> - Create a new nsproxy and its associated namespaces and pass it back to
> caller to attach it to right process.
>
> - Changed all copy_*_ns() routines to return a new copy of namespace
> instead of attaching it to task->nsproxy.
>
> - Moved the CAP_SYS_ADMIN checks out of copy_*_ns() routines.
>
> - Removed unnecessary !ns checks from copy_*_ns() and added BUG_ON()
> just incase.
>
> - Get rid of all individual unshare_*_ns() routines and make use of
> copy_*_ns() instead.
>

.. [zapped] ...

> + * Called from unshare. Unshare all the namespaces part of nsproxy.
> + * On success, returns the new nsproxy and a reference to old nsproxy
> + * to make sure it stays around.
> + */
> +int unshare_nsproxy_namespaces(unsigned long unshare_flags,

```
> + struct nsproxy **new_nsp, struct fs_struct *new_fs)
> +{
```

this makes sys_unshare leak and nsproxy (reference)

can be tested with the following command sequence:

```
vcmd -nu ^17 -- vcmd -nu ^17 -- sleep 10
```

(and some nsproxy accounting/debugging as used in Linux-VServer)

we probably want to drop the reference to the old nsproxy in sys_unshare() but I do not see a good reason to take the reference in the first place (at least not with the code in mainline 2.6.22-rc4)

HTH,
Herbert

PS: vcmd can be found here:

<http://vserver.13thfloor.at/Experimental/TOOLS/vcmd-0.09.tar.bz2>

... [more zapped] ...

Containers mailing list

Containers@lists.linux-foundation.org

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