

---

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>

---