

---

Subject: Re: [patch -mm 09/17] nsproxy: add namespace flags

Posted by [Herbert Poetzl](#) on Tue, 05 Dec 2006 23:41:58 GMT

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

---

On Tue, Dec 05, 2006 at 11:28:01AM +0100, clg@fr.ibm.com wrote:

> From: Cedric Le Goater <clg@fr.ibm.com>

>

> Signed-off-by: Cedric Le Goater <clg@fr.ibm.com>

> ---

> include/linux/nsproxy.h | 11 ++++++++

> 1 file changed, 11 insertions(+)

>

> Index: 2.6.19-rc6-mm2/include/linux/nsproxy.h

> =====

> --- 2.6.19-rc6-mm2.orig/include/linux/nsproxy.h

> +++ 2.6.19-rc6-mm2/include/linux/nsproxy.h

> @@ -14,6 +14,17 @@ struct net\_namespace;

> struct user\_namespace;

>

> /\*

> + \* namespaces flags

> + \*/

> +#define NS\_MNT 0x00000001

> +#define NS\_UTS 0x00000002

> +#define NS\_IPC 0x00000004

> +#define NS\_PID 0x00000008

> +#define NS\_NET 0x00000010

> +#define NS\_USER 0x00000020

> +#define NS\_ALL (NS\_MNT|NS\_UTS|NS\_IPC|NS\_PID|NS\_NET|NS\_USER)

hmm, why \_another\_ set of flags to refer to the namespaces? is the clone()/unshare() set of flags not sufficient for that? if so, shouldn't we switch (or even better change? the unshare() too) to a new set of syscalls?

note: even if they are just for internal purpose and will never get exposed to userspace, we should think twice before we create just another set of flags, and if we do so, please let us change them all, including certain clone flags (and add a single compatibility wrapper for the 'old' syscalls)

best,  
Herbert

> +

> +/\*

> \* A structure to contain pointers to all per-process  
> \* namespaces - fs (mount), uts, network, sysvipc, etc.  
> \*  
>  
> --

> \_\_\_\_\_  
> Containers mailing list  
> Containers@lists.osdl.org  
> <https://lists.osdl.org/mailman/listinfo/containers>

\_\_\_\_\_  
Containers mailing list  
Containers@lists.osdl.org  
<https://lists.osdl.org/mailman/listinfo/containers>

---