
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>
