
Subject: Re: [patch -mm 09/17] nsproxy: add namespace flags
Posted by [ebiederm](#) on Fri, 08 Dec 2006 19:40:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

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

```
>>> /*
>>> + * 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?
>
> well, because namespaces are a new kind in the kernel
```

Gratuitous incompatibility.

```
>> is the clone()/unshare() set of flags not sufficient
>> for that?
>
> because we are reaching the limits of the CLONE_ flags.
```

Not really. There are at least 8 bits that clone cannot use but that unshare can.

```
>> if so, shouldn't we switch (or even better change?
>> the unshare() too) to a new set of syscalls?
>
> unshare_ns() is a new syscall and we don't really need a
> clone anyway. nop ?
```

Huh? Clone should be the primary. There are certain namespaces that it are very hard to unshare, without creating a new process.

```
> we could make the clone flags and namespace flags compatible
> with :
>
> #define NS_MNT CLONE_NEWNS
> #define NS_UTS CLONE_NEWUTS
> #define NS_IPC CLONE_NEWIPC
>
```

> that shouldn't be a big issue but we could also
> remove/deprecate :
>
> #define CLONE_NEWNS 0x00020000 /* New namespace group? */
> #define CLONE_NEWUTS 0x04000000 /* New utsname group? */
> #define CLONE_NEWIPC 0x08000000 /* New ipc? */
>
> to make sure namespaces can not be unshared using
> unshare()

Nope. We got them now they can not be removed. It's part of the ABI.
And since we are not removing the functionality it makes not sense
at all to just change the name of the flags.

>> 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)
>
> so you would keep the unshare as is but change the set
> of flags its using, making sure the old ones are still
> compatible with the new ones.
>
> something like this :
>
> int sys_unshare(int unshare_flags)
> {
> int unshare_ns_flags;
>
> unshare_ns_flags = convert_flags(unshare_flags);
>
> return sys_unshare_ns(unshare_ns_flags);
> }
>
> ?

If necessary.

Eric

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