
Subject: Re: [RFC][PATCH 6/6]: Enable unsharing pid namespace.

Posted by [ebiederm](#) on Tue, 13 Mar 2007 09:01:26 GMT

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

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

```
>> | > Index: lx26-20-mm2b/kernel/nsproxy.c
>> | > =====
>> | > --- lx26-20-mm2b.orig/kernel/nsproxy.c 2007-03-09 14:56:12.000000000 -0800
>> | > +++ lx26-20-mm2b/kernel/nsproxy.c 2007-03-09 15:03:05.000000000 -0800
>> | > @@ -83,13 +83,16 @@ int copy_namespaces(int flags, struct ta
>> | > struct nsproxy *old_ns = tsk->nsproxy;
>> | > struct nsproxy *new_ns;
>> | > int err = 0;
>> | > + int ns_all;
>> | >
>> | > if (!old_ns)
>> | > return 0;
>> | >
>> | > get_nsproxy(old_ns);
>> | >
>> | > - if (!(flags & (CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC)))
>> | > + ns_all = CLONE_NEWNS | CLONE_NEWUTS | CLONE_NEWIPC | CLONE_NEWPID;
>> | > +
>> |
>> | This doesn't quite seem to make sense why the extra intermediate variable?
>>
>> Will drop ns_all variable.
>
> well, in the patch reserving the pid namespace clone flags,
> we could define :
>
> #define NS_ALL (CLONE_NEWNS|CLONE_NEWUTS|CLONE_NEWIPC|CLONE_NEWPID)
>
> which is useful in many ways.
```

Yes. Especially if we structured it something like:

```
#ifdef CONFIG_PID_NS
#define CLONE_NEWPID_NSALL CLONE_NEWPID
#else
#define CLONE_NEWPID_NSALL
#endif

#define
NS_ALL (CLONE_NEWNS|CLONE_NEWUTS|CLONE_NEWIPC|CLONE_NEWPID_NSALL)
```

So we have a natural way of disabling the pid namespace, until we are certain

it is complete and remove it from CONFIG_EXPERIMENTAL

Eric

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
