
Subject: [PATCH] Define CLONE_NEWPID flag
Posted by [Sukadev Bhattiprolu](#) on Wed, 21 Mar 2007 19:41:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

This was discussed on containers and we thought it would be useful
to reserve this flag.

From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
Subject: [PATCH] Define CLONE_NEWPID flag

Define CLONE_NEWPID flag that will be used to clone pid namespaces.

Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>

include/linux/sched.h | 1 +
1 file changed, 1 insertion(+)

Index: lx26-21-rc3-mm2/include/linux/sched.h

```
=====
--- lx26-21-rc3-mm2.orig/include/linux/sched.h 2007-03-20 20:13:19.000000000 -0700
+++ lx26-21-rc3-mm2/include/linux/sched.h 2007-03-21 11:10:33.000000000 -0700
@@ -26,6 +26,7 @@
#define CLONE_STOPPED 0x02000000 /* Start in stopped state */
#define CLONE_NEWUTS 0x04000000 /* New utsname group? */
#define CLONE_NEWIPC 0x08000000 /* New ipcns */
+#define CLONE_NEWPID 0x10000000 /* New pid namespace */

/*
 * Scheduling policies
```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [PATCH] Define CLONE_NEWPID flag
Posted by [akpm](#) on Wed, 21 Mar 2007 20:39:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 21 Mar 2007 12:41:03 -0700
sukadev@us.ibm.com wrote:

>
> This was discussed on containers and we thought it would be useful
> to reserve this flag.
> ---
>

> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> Subject: [PATCH] Define CLONE_NEWPID flag
>
> Define CLONE_NEWPID flag that will be used to clone pid namespaces.
>
> Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> ---
> include/linux/sched.h | 1 +
> 1 file changed, 1 insertion(+)
>
> Index: lx26-21-rc3-mm2/include/linux/sched.h
> =====
> --- lx26-21-rc3-mm2.orig/include/linux/sched.h 2007-03-20 20:13:19.000000000 -0700
> +++ lx26-21-rc3-mm2/include/linux/sched.h 2007-03-21 11:10:33.000000000 -0700
> @@ -26,6 +26,7 @@
> #define CLONE_STOPPED 0x02000000 /* Start in stopped state */
> #define CLONE_NEWUTS 0x04000000 /* New utsname group? */
> #define CLONE_NEWIPC 0x08000000 /* New ipcns */
> +#define CLONE_NEWPID 0x10000000 /* New pid namespace */
>

Do we actually have any need to reserve it at this time? I'd have thought that we could defer adding this until we have some code in-kernel which uses it.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [PATCH] Define CLONE_NEWPID flag
Posted by [ebiederm](#) on Wed, 21 Mar 2007 20:57:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

Andrew Morton <akpm@linux-foundation.org> writes:

>> Index: lx26-21-rc3-mm2/include/linux/sched.h
>> =====
>> --- lx26-21-rc3-mm2.orig/include/linux/sched.h 2007-03-20 20:13:19.000000000
>> -0700
>> +++ lx26-21-rc3-mm2/include/linux/sched.h 2007-03-21 11:10:33.000000000 -0700
>> @@ -26,6 +26,7 @@
>> #define CLONE_STOPPED 0x02000000 /* Start in stopped state */
>> #define CLONE_NEWUTS 0x04000000 /* New utsname group? */
>> #define CLONE_NEWIPC 0x08000000 /* New ipcns */
>> +#define CLONE_NEWPID 0x10000000 /* New pid namespace */
>>

>
> Do we actually have any need to reserve it at this time? I'd have thought
> that we could defer adding this until we have some code in-kernel which
> uses it.

In practice this is pretty much reserved already but I understand the sentiment.

Currently the plan is to work on the core pid namespace and get those functions merged at least to -mm with a big fat CONFIG_EXPERIMENTAL so people can really understand what we are talking about when we say a pid_namespace. Then go through and finish up all converting the rest of the pid uses that we haven't tackled yet. There aren't that many left and most of the remaining conversions only make sense in the context of a pid namespace.

Currently we are one or two review cycles away from being able to push out the core pid namespace code.

Further it is actually critical that we have a clone flag for the pid namespace because unshare is very much harder if it is possible at all.

Personally if you want to delay it a week or two, until the rest of the code is ready that is fine. Mostly getting it out now is about release early and release often.

Eric

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [PATCH] Define CLONE_NEWPID flag
Posted by [Herbert Poetzl](#) on Fri, 23 Mar 2007 01:36:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Mar 21, 2007 at 01:39:38PM -0700, Andrew Morton wrote:

> On Wed, 21 Mar 2007 12:41:03 -0700
> sukadev@us.ibm.com wrote:
>
> > This was discussed on containers and we thought it would be useful
> > to reserve this flag.
> > ---
> >
> > From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> > Subject: [PATCH] Define CLONE_NEWPID flag
> >

```

> > Define CLONE_NEWPID flag that will be used to clone pid namespaces.
> >
> > Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> > ---
> > include/linux/sched.h | 1 +
> > 1 file changed, 1 insertion(+)
> >
> > Index: lx26-21-rc3-mm2/include/linux/sched.h
> > =====
> > --- lx26-21-rc3-mm2.orig/include/linux/sched.h 2007-03-20 20:13:19.000000000 -0700
> > +++ lx26-21-rc3-mm2/include/linux/sched.h 2007-03-21 11:10:33.000000000 -0700
> > @@ -26,6 +26,7 @@
> > #define CLONE_STOPPED 0x02000000 /* Start in stopped state */
> > #define CLONE_NEWUTS 0x04000000 /* New utsname group? */
> > #define CLONE_NEWIPC 0x08000000 /* New ipc */
> > + #define CLONE_NEWPID 0x10000000 /* New pid namespace */
>
> Do we actually have any need to reserve it at this time?
> I'd have thought that we could defer adding this until we
> have some code in-kernel which uses it.

```

FWIW, I'm fine with the reservation, we won't get around it for the pid space, so we can as well register it now
YMMV

best,
Herbert

```

> _____
> Containers mailing list
> Containers@lists.linux-foundation.org
> https://lists.linux-foundation.org/mailman/listinfo/containers

```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
