
Subject: Add pid_namespace to nsproxy
Posted by [Sukadev Bhattiprolu](#) on Fri, 08 Sep 2006 05:41:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Add a notion of pid namespace to nsproxy (and by extension, to task_struct).
Currently there is only one pid namespace, init_pid_ns and all tasks belong
to this pid namespace. When a new task is created, it inherits its parent's
pid namespace (in copy_process()).

This is based on Eric Biederman's patch: <http://lkml.org/lkml/2006/2/6/285>

Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
Cc: Eric Biederman <ebiederm@xmission.com>
Cc: Cedric Le Goater <clg@fr.ibm.com>
Cc: Dave Hansen <haveblue@us.ibm.com>
Cc: Serge Hallyn <serue@us.ibm.com>
Cc: containers@lists.osdl.org

```
include/linux/init_task.h | 2 ++
include/linux/nsproxy.h | 1 +
2 files changed, 3 insertions(+)
```

Index: lx26-18-rc5-mm1/include/linux/init_task.h

```
=====
--- lx26-18-rc5-mm1.orig/include/linux/init_task.h 2006-09-07 22:21:06.000000000 -0700
+++ lx26-18-rc5-mm1/include/linux/init_task.h 2006-09-07 22:29:58.000000000 -0700
@@ -7,6 +7,7 @@
#include <linux/utsname.h>
#include <linux/lockdep.h>
#include <linux/ipc.h>
+#include <linux/pid_namespace.h>
```

```
#define INIT_FDTABLE \
{ \
@@ -72,6 +73,7 @@
```

```
extern struct nsproxy init_nsproxy;
#define INIT_NS_PROXY(nsproxy) { \
+ .pid_ns = &init_pid_ns, \
. count = ATOMIC_INIT(1), \
. nslock = SPIN_LOCK_UNLOCKED, \
. uts_ns = &init_uts_ns, \
```

Index: lx26-18-rc5-mm1/include/linux/nsproxy.h

```
=====
--- lx26-18-rc5-mm1.orig/include/linux/nsproxy.h 2006-09-07 22:21:06.000000000 -0700
+++ lx26-18-rc5-mm1/include/linux/nsproxy.h 2006-09-07 22:30:22.000000000 -0700
@@ -23,6 +23,7 @@ struct ipc_namespace;
struct nsproxy {
```

```
atomic_t count;
spinlock_t nslock;
+ struct pid_namespace *pid_ns;
  struct uts_namespace *uts_ns;
  struct ipc_namespace *ipc_ns;
  struct namespace *namespace;
```

Containers mailing list

Containers@lists.osdl.org

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

Subject: Re: Add pid_namespace to nsproxy

Posted by [Cedric Le Goater](#) on Fri, 08 Sep 2006 10:30:05 GMT

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

Sukadev Bhattiprolu wrote:

<snip>

```
> +++ lx26-18-rc5-mm1/include/linux/nsproxy.h 2006-09-07 22:30:22.000000000 -0700
> @@ -23,6 +23,7 @@ struct ipc_namespace;
> struct nsproxy {
>   atomic_t count;
>   spinlock_t nslock;
> + struct pid_namespace *pid_ns;
>   struct uts_namespace *uts_ns;
>   struct ipc_namespace *ipc_ns;
>   struct namespace *namespace;
```

argh that patches conflicts with the user namespace patchset and the patch from kirill changing namespace to mnt_namespace.

I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.

- * kthread conversion patches not yet accepted by maintainers
- * ro bind mount, from dave
- * user beancounter, from kirill and pavel
- * network namespace, from andrey
- * some extra stuff from daniel (in progress)
- * mnt namespace, from kirill
- * pid namespace, from suka.
- * user namespace, from cedric

the last 3 patches conflicts because of nsproxy. I'm going to merge them all, and resend on the list for inclusion in -mm.

thanks,

C.

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

Subject: Re: Add pid_namespace to nsproxy
Posted by [dev](#) on Fri, 08 Sep 2006 17:44:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

Cedric Le Goater wrote:

> Sukadev Bhattiprolu wrote:

>
> <snip>
>
>>+++ lx26-18-rc5-mm1/include/linux/nsproxy.h 2006-09-07 22:30:22.000000000 -0700
>>@@ -23,6 +23,7 @@ struct ipc_namespace;
>> struct nsproxy {
>> atomic_t count;
>> spinlock_t nslock;
>>+ struct pid_namespace *pid_ns;
>> struct uts_namespace *uts_ns;
>> struct ipc_namespace *ipc_ns;
>> struct namespace *namespace;

>
>
> argh that patches conflicts with the user namespace patchset and the patch
> from kirill changing namespace to mnt_namespace.

>
> I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.

>
> * kthread conversion patches not yet accepted by maintainers
> * ro bind mount, from dave
> * user beancounter, from kirill and pavel
> * network namespace, from andrey
> * some extra stuff from daniel (in progress)
> * mnt namespace, from kirill
> * pid namespace, from suka.
> * user namespace, from cedric

Oh, good to know.

Is it downloadable from somewhere?

(just in case we can provide FTP/WWW whatever you need)

> the last 3 patches conflicts because of nsproxy. I'm going to merge them
> all, and resend on the list for inclusion in -mm.

Thanks,
Kirill

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

Subject: Re: Add pid_namespace to nsproxy
Posted by [Cedric Le Goater](#) on Fri, 08 Sep 2006 17:51:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev wrote:

> Cedric Le Goater wrote:

>> Sukadev Bhattiprolu wrote:

>>

>> <snip>

>>

>>> +++ lx26-18-rc5-mm1/include/linux/nsproxy.h 2006-09-07 22:30:22.000000000 -0700

>>> @@ -23,6 +23,7 @@ struct ipc_namespace;

>>> struct nsproxy {

>>> atomic_t count;

>>> spinlock_t nslock;

>>> + struct pid_namespace *pid_ns;

>>> struct uts_namespace *uts_ns;

>>> struct ipc_namespace *ipc_ns;

>>> struct namespace *namespace;

>>

>> argh that patches conflicts with the user namespace patchset and the patch

>> from kirill changing namespace to mnt_namespace.

>>

>> I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.

>>

>> * kthread conversion patches not yet accepted by maintainers

>> * ro bind mount, from dave

>> * user beancounter, from kirill and pavel

>> * network namespace, from andrey

>> * some extra stuff from daniel (in progress)

>> * mnt namespace, from kirill

>> * pid namespace, from suka.

>> * user namespace, from cedric

> Oh, good to know.

> Is it downloadable from somewhere?

> (just in case we can provide FTP/WWW whatever you need)

thanks,

It will.

I just finished the port on rc6-mm1 and a fast compile/boot on i386, x86_64 and s390. Let us give it a first round of test.

C.

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

Subject: Re: Add pid_namespace to nsproxy
Posted by [Cedric Le Goater](#) on Tue, 12 Sep 2006 16:31:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

>> I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.
>>
>> * kthread conversion patches not yet accepted by maintainers
>> * ro bind mount, from dave
>> * user beancounter, from kirill and pavel
>> * network namespace, from andrey
>> * some extra stuff from daniel (in progress)
>> * mnt namespace, from kirill
>> * pid namespace, from suka.
>> * user namespace, from cedric
> Oh, good to know.
> Is it downloadable from somewhere?

<http://www.sr71.net/patches/2.6.18/2.6.18-rc6-mm2/2.6.18-rc6-mm2-lxc1/>

It's mostly stable. you can give it a try.

thanks,

C.

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

Subject: Re: Add pid_namespace to nsproxy
Posted by [Herbert Poetzl](#) on Tue, 12 Sep 2006 17:00:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, Sep 12, 2006 at 06:31:54PM +0200, Cedric Le Goater wrote:

>
> >> I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.
> >>
> >> * kthread conversion patches not yet accepted by maintainers
> >> * ro bind mount, from dave
> >> * user beancounter, from kirill and pavel
> >> * network namespace, from andrey
> >> * some extra stuff from daniel (in progress)
> >> * mnt namespace, from kirill
> >> * pid namespace, from suka.
> >> * user namespace, from cedric
> > Oh, good to know.
> > Is it downloadable from somewhere?
>
> <http://www.sr71.net/patches/2.6.18/2.6.18-rc6-mm2/2.6.18-rc6-mm2-lxc1/>
>
> It's mostly stable. you can give it a try.

great! thanks!

> thanks,

>

>

> C.

>

> 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>

Subject: Re: Add pid_namespace to nsproxy
Posted by [dev](#) on Wed, 13 Sep 2006 14:13:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

Cedric Le Goater wrote:

>>> I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.
>>>
>>> * kthread conversion patches not yet accepted by maintainers
>>> * ro bind mount, from dave
>>> * user beancounter, from kirill and pavel
>>> * network namespace, from andrey
>>> * some extra stuff from daniel (in progress)
>>> * mnt namespace, from kirill
>>> * pid namespace, from suka.

>>>* user namespace, from cedric
>>
>>Oh, good to know.
>>Is it downloadable from somewhere?
>
>
> <http://www.sr71.net/patches/2.6.18/2.6.18-rc6-mm2/2.6.18-rc6-mm2-lxc1/>
>
> It's mostly stable. you can give it a try.
thanks!
added this info at:
<http://wiki.openvz.org/Containers>

Kirill

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

Subject: Re: Add pid_namespace to nsproxy
Posted by [Cedric Le Goater](#) on Wed, 13 Sep 2006 14:16:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev wrote:
> Cedric Le Goater wrote:
>>>> I'm maintaining a patchset on 2.6.18-rc5-mm1 with the following input.
>>>>
>>>> * kthread conversion patches not yet accepted by maintainers
>>>> * ro bind mount, from dave
>>>> * user beancounter, from kirill and pavel
>>>> * network namespace, from andrey
>>>> * some extra stuff from daniel (in progress)
>>>> * mnt namespace, from kirill
>>>> * pid namespace, from suka.
>>>> * user namespace, from cedric
>>> Oh, good to know.
>>> Is it downloadable from somewhere?
>>
>> <http://www.sr71.net/patches/2.6.18/2.6.18-rc6-mm2/2.6.18-rc6-mm2-lxc1/>
>>
>> It's mostly stable. you can give it a try.
> thanks!
> added this info at:
> <http://wiki.openvz.org/Containers>

good.

I'm adding a few signed-off-by: and from: in the patches i fixed. Next version should be cleaner.

C.

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