

---

Subject: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper  
Posted by [Sukadev Bhattiprolu](#) on Sat, 10 Mar 2007 03:59:24 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: Sukadev Bhattiprolu <sukadev@us.ibm.com>  
Subject: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper.

Statically initialize a struct pid\_nr for the swapper process.

Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.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 | 7 +++++++  
kernel/pid.c | 1 +  
2 files changed, 8 insertions(+)

Index: lx26-20-mm2b/include/linux/init\_task.h

```
=====
--- lx26-20-mm2b.orig/include/linux/init_task.h 2007-03-08 17:56:05.000000000 -0800
+++ lx26-20-mm2b/include/linux/init_task.h 2007-03-09 14:56:11.000000000 -0800
@@ -96,6 +96,12 @@ extern struct group_info init_groups;
#define INIT_PREEMPT_RCU
#endif

+#define INIT_STRUCT_PID_NR { \
+ .node = { .next = NULL, .pprev = NULL }, \
+ .nr = 0, \
+ .pid_ns = &init_pid_ns, \
+}
+
#define INIT_STRUCT_PID { \
. count = ATOMIC_INIT(1), \
. nr = 0, \
@@ -106,6 +112,7 @@ extern struct group_info init_groups;
{ .first = &init_task.pids[PIDTYPE_PGID].node }, \
{ .first = &init_task.pids[PIDTYPE_SID].node }, \
}, \
+ .pid_nrs = { .first = &init_struct_pid_nr.node }, \
. rcu = RCU_HEAD_INIT, \
}
```

Index: lx26-20-mm2b/kernel/pid.c

```
=====
--- lx26-20-mm2b.orig/kernel/pid.c 2007-03-08 17:56:57.000000000 -0800
+++ lx26-20-mm2b/kernel/pid.c 2007-03-09 14:56:33.000000000 -0800
```

```
@@ -34,6 +34,7 @@ static struct hlist_head *pid_hash;
static int pidhash_shift;
static struct kmem_cache *pid_cachep;
static struct kmem_cache *pid_nr_cachep;
+struct pid_nr init_struct_pid_nr = INIT_STRUCT_PID_NR;
struct pid init_struct_pid = INIT_STRUCT_PID;

int pid_max = PID_MAX_DEFAULT;
```

---

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

---

---

Subject: Re: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper  
Posted by [Herbert Poetzl](#) on Sat, 10 Mar 2007 22:09:08 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Fri, Mar 09, 2007 at 07:59:24PM -0800, sukadev@us.ibm.com wrote:

```
>
> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
> Subject: [RFC][PATCH 4/6] Initialize struct pid_nr for swapper.
>
> Statically initialize a struct pid_nr for the swapper process.
```

does that actually satisfy procfs needs too, or  
just the abstract internal references?

because if it is enough to make procfs happy, this  
would be a viable solution for the lightweight guest  
case (as fake init process) too, which doesn't require  
a blend through functionality anymore, and would allow  
to make the pid isolation complete without wasting  
any resources ...

TIA,  
Herbert

```
> Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.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 | 7 +++++++
> kernel/pid.c              | 1 +
> 2 files changed, 8 insertions(+)
>
```

```

> Index: lx26-20-mm2b/include/linux/init_task.h
> =====
> --- lx26-20-mm2b.orig/include/linux/init_task.h 2007-03-08 17:56:05.000000000 -0800
> +++ lx26-20-mm2b/include/linux/init_task.h 2007-03-09 14:56:11.000000000 -0800
> @@ -96,6 +96,12 @@ extern struct group_info init_groups;
> #define INIT_PREEMPT_RCU
> #endif
>
> +#define INIT_STRUCT_PID_NR { \
> + .node = { .next = NULL, .pprev = NULL }, \
> + .nr = 0, \
> + .pid_ns = &init_pid_ns, \
> +}
> +
> #define INIT_STRUCT_PID { \
> .count = ATOMIC_INIT(1), \
> .nr = 0, \
> @@ -106,6 +112,7 @@ extern struct group_info init_groups;
> { .first = &init_task.pids[PIDTYPE_PGID].node }, \
> { .first = &init_task.pids[PIDTYPE_SID].node }, \
> }, \
> + .pid_nrs = { .first = &init_struct_pid_nr.node }, \
> .rcu = RCU_HEAD_INIT, \
> }
>
> Index: lx26-20-mm2b/kernel/pid.c
> =====
> --- lx26-20-mm2b.orig/kernel/pid.c 2007-03-08 17:56:57.000000000 -0800
> +++ lx26-20-mm2b/kernel/pid.c 2007-03-09 14:56:33.000000000 -0800
> @@ -34,6 +34,7 @@ static struct hlist_head *pid_hash;
> static int pidhash_shift;
> static struct kmem_cache *pid_cachep;
> static struct kmem_cache *pid_nr_cachep;
> +struct pid_nr init_struct_pid_nr = INIT_STRUCT_PID_NR;
> struct pid init_struct_pid = INIT_STRUCT_PID;
>
> int pid_max = PID_MAX_DEFAULT;
>
> 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: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper  
Posted by [ebiederm](#) on Sun, 11 Mar 2007 11:27:59 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Herbert Poetzl <herbert@13thfloor.at> writes:

> On Fri, Mar 09, 2007 at 07:59:24PM -0800, sukadev@us.ibm.com wrote:  
>>  
>> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>  
>> Subject: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper.  
>>  
>> Statically initialize a struct pid\_nr for the swapper process.  
>  
> does that actually satisfy procfs needs too, or  
> just the abstract internal references?  
>  
> because if it is enough to make procfs happy, this  
> would be a viable solution for the lightweight guest  
> case (as fake init process) too, which doesn't require  
> a blend through functionality anymore, and would allow  
> to make the pid isolation complete without wasting  
> any resources ...

Herbert I'm not quite certain what you are asking but  
largely I think the answer is yes. Making procfs work  
on top of something like this patchset is pretty straight forward.

Eric

---

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

---

---

Subject: Re: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper  
Posted by [Herbert Poetzl](#) on Sun, 11 Mar 2007 14:36:09 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Sun, Mar 11, 2007 at 05:27:59AM -0600, Eric W. Biederman wrote:  
> Herbert Poetzl <herbert@13thfloor.at> writes:  
>  
> > On Fri, Mar 09, 2007 at 07:59:24PM -0800, sukadev@us.ibm.com wrote:  
> >>  
> >> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>  
> >> Subject: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper.  
> >>  
> >> Statically initialize a struct pid\_nr for the swapper process.  
> >

> > does that actually satisfy procfs needs too, or  
> > just the abstract internal references?  
> >  
> > because if it is enough to make procfs happy, this  
> > would be a viable solution for the lightweight guest  
> > case (as fake init process) too, which doesn't require  
> > a blend through functionality anymore, and would allow  
> > to make the pid isolation complete without wasting  
> > any resources ...  
>  
> Herbert I'm not quite certain what you are asking but  
> largely I think the answer is yes. Making procfs work  
> on top of something like this patchset is pretty straight  
> forward.

okay, then please lets make sure that this actually works, because I think it might solve most of the lightweight guest issues the suggested pid spaces would introduce ...

I tried that some time back, but the procfs really provides \_a lot\_ of deep linked details for each process, and I postponed that approach back then when I realized that I would have to fill in quite a lot of static data to make procfs happy (with a static initialized fake init)

TIA,  
Herbert

> Eric

---

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

---

---

Subject: Re: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper  
Posted by [ebiederm](#) on Sun, 11 Mar 2007 18:17:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Herbert Poetzl <herbert@13thfloor.at> writes:

> okay, then please lets make sure that this actually  
> works, because I think it might solve most of the  
> lightweight guest issues the suggested pid spaces  
> would introduce ...  
>

> I tried that some time back, but the procfs really  
> provides \_a lot\_ of deep linked details for each  
> process, and I postponed that approach back then  
> when I realized that I would have to fill in quite  
> a lot of static data to make procfs happy (with a  
> static initialized fake init)

Ok. I now see what the real question is. The idle thread never shows up in /proc so I don't know if this is quite complete. The bits we fill in are designed to be enough for fork/clone and other pieces so we don't need special cases in the code to deal with the idle thread.

I think the idle thread quite possibly has enough information to show up in /proc but that is a separate case.

Definitely something to discuss when we come back to unshare of the pid namespace.

Eric

---

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

---

---

Subject: Re: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper  
Posted by [ebiederm](#) on Sun, 11 Mar 2007 18:22:23 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

sukadev@us.ibm.com writes:

> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>  
> Subject: [RFC][PATCH 4/6] Initialize struct pid\_nr for swapper.  
>  
> Statically initialize a struct pid\_nr for the swapper process.

I see the sense in this bit but where does this come up?  
That is can we get away without having a struct pid\_nr for the swapper process.

pid\_nr would continue to operate correctly.

I think only task\_pid\_ns would have problems and I'm not at all certain we ever call that on the swapper process.

Eric

---

Containers mailing list

