
Subject: Re: [PATCH 4/16] Change data structures for pid namespaces

Posted by [Pavel Emelianov](#) on Tue, 10 Jul 2007 07:04:09 GMT

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

sukadev@us.ibm.com wrote:

> Cedric Le Goater [clg@fr.ibm.com] wrote:

> | Pavel Emelianov wrote:

> | > struct pid_namespace will have the kmem_cache to allocate

> | > the pids from, the parent, as they are hierarchical, and

> | > the level of nesting value.

> | >

> | > struct pid will have a variable length array of pid_number-s

> | > one for each namespace this pid lives in. The level value

> | > shows the level of the namespace this pid lives in and thus -

> | > the number of elements in the numbers array.

> | >

> | > Signed-off-by: Pavel Emelianov <xemul@openvz.org>

> | >

> | > ---

> | >

> | > include/linux/init_task.h | 6 ++++++

> | > include/linux/pid.h | 9 ++++++++

> | > include/linux/pid_namespace.h | 3 +++

> | > kernel/pid.c | 3 ++-

> | > 4 files changed, 20 insertions(+), 1 deletion(-)

> | >

> | > diff -upr linux-2.6.22-rc4-mm2.orig/include/linux/pid.h

linux-2.6.22-rc4-mm2-2/include/linux/pid.h

> | > --- linux-2.6.22-rc4-mm2.orig/include/linux/pid.h 2007-06-14 12:14:29.000000000 +0400

> | > +++ linux-2.6.22-rc4-mm2-2/include/linux/pid.h 2007-07-04 19:00:38.000000000 +0400

> | > @@ -40,6 +40,13 @@ enum pid_type

> | > * processes.

> | > */

> | >

> | > +struct pid_number {

> | > + /* Try to keep pid_chain in the same cacheline as nr for find_pid */

> | > + int nr;

> | > + struct pid_namespace *ns;

> | > + struct hlist_node pid_chain;

> | > +};

>

> We meant to go back and look at removing the extra 'struct pid *' we had

> here. Looks like you did that. Cool.

>

> | > +

> | > struct pid

> | > {

> | > atomic_t count;

```

> | > @@ -40,6 +40,8 @@ enum pid_type
> | > /* lists of tasks that use this pid */
> | > struct hlist_head tasks[PIDTYPE_MAX];
> | > struct rcu_head rcu;
> | > + int level;
> | > + struct pid_number numbers[1];
> | > };
> | >
> | > extern struct pid init_struct_pid;
> | > diff -upr linux-2.6.22-rc4-mm2.orig/include/linux/pid_namespace.h
linux-2.6.22-rc4-mm2-2/include/linux/pid_namespace.h
> | > --- linux-2.6.22-rc4-mm2.orig/include/linux/pid_namespace.h 2007-06-14
12:14:29.000000000 +0400
> | > +++ linux-2.6.22-rc4-mm2-2/include/linux/pid_namespace.h 2007-07-04 19:00:39.000000000
+0400
> | > @@ -16,7 +15,10 @@ struct pidmap {
> | > struct kref kref;
> | > struct pidmap pidmap[PIDMAP_ENTRIES];
> | > int last_pid;
> | > + int level;
> | > struct task_struct *child_reaper;
> | > + struct kmem_cache *pid_cachep;
> |
> | so, that looks like a good idea to have the cache in the pidmap. could you
> | push that independently to see how it all fits together ?
>
> Yes. I like this idea too.

```

OK. I will collect the patches you Aacked and send them to Andrew independently.

```

> |
> | thanks,
> |
> | C.
>

```

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