
Subject: Re: [RFC][PATCH 2/3] Signal semantics for /sbin/init
Posted by [Sukadev Bhattiprolu](#) on Thu, 30 Aug 2007 07:11:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

Oleg Nesterov [oleg@tv-sign.ru] wrote:

| On 08/29, sukadev@us.ibm.com wrote:

```
|>
|> --- 2.6.23-rc3-mm1.orig/kernel/signal.c 2007-08-29 22:53:20.000000000 -0700
|> +++ 2.6.23-rc3-mm1/kernel/signal.c 2007-08-29 23:10:16.000000000 -0700
|> @@ -26,6 +26,7 @@
|> #include <linux/freezer.h>
|> #include <linux/pid_namespace.h>
|> #include <linux/nsproxy.h>
|> +#include <linux/hardirq.h>
|>
|> #include <asm/param.h>
|> #include <asm/uaccess.h>
|> @@ -39,11 +40,42 @@
|>
|> static struct kmem_cache *sigqueue_cachep;
|>
|> +static int sig_init_ignore(struct task_struct *tsk)
|> +{
|>
|> -static int sig_ignored(struct task_struct *t, int sig)
|> + // Currently this check is a bit racy with exec(),
|> + // we can _simplify_de_thread and close the race.
|> + if (likely(!is_container_init(tsk->group_leader)))
|> + return 0;
|> +
|> +/*
|> + * If signal is from an ancestor pid namespace, do not
|> + * ignore the signal.
|> + */
|> + if (task_ancestor_pid_ns(current, tsk))
|> + return 0;
|> +
```

This patch was intended as a fix for the current behaviour, and a preparation for the pid_ns requirements. If possible, it should go ahead of all other pid_ns changes, imho.

In any case, we shouldn't mix all this in one patch.

If there are no objections to this approach, I'll re-send the patch on weekend.

Yes I think your patch looks good. I will apply the other two patches on top of this.

```
| > @@ -2320,6 +2346,13 @@ int do_sigaction(int sig, struct k_sigac
| >   k = &current->sighand->action[sig-1];
| >
| >   spin_lock_irq(&current->sighand->siglock);
| > +
| > + if (current->signal->flags & SIGNAL_GROUP_EXIT) {
| > +   spin_unlock_irq(&current->sighand->siglock);
| > + /* The return value doesn't matter, SIGKILL is pending */
| > + return -EINTR;
| > + }
```

This is not needed. Please look at

<http://marc.info/?l=linux-kernel&m=118763989718143>

I specially removed this chunk in a separate patch, no need to re-introduce.

Sorry I missed that. I will remove it.

Thanks,

Suka

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