
Subject: Re: [RFC,PATCH] fix /sbin/init signal handling
Posted by [Oleg Nesterov](#) on Tue, 21 Aug 2007 10:40:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 08/21, Pavel Emelyanov wrote:

```
>
> >+static int sig_init_ignore(struct task_struct *tsk)
> >+{
> >+ // Currently this check is a bit racy with exec(),
> >+ // we can _simplify_ de_thread and close the race.
> >+ if (likely(!is_init(tsk->group_leader)))
> >+ return 0;
> >+
> >+ // ----- Multiple pid namespaces -----
> >+ // if (current is from tsk's parent pid_ns && !in_interrupt())
> >+ // return 0;
> >+
> >+ return 1;
> >+}
> >+
> >+static int sig_task_ignore(struct task_struct *tsk, int sig)
> >+{
> >+ void __user * handler = tsk->sigband->action[sig-1].sa.sa_handler;
> >+
> >+ if (handler == SIG_IGN)
> >+ return 1;
> >+
> >+ if (handler != SIG_DFL)
> >+ return 0;
> >+
> >+ return sig_kernel_ignore(sig) || sig_init_ignore(tsk);
> >+}
>
> These two look like the init ignores "less" than a usual task,
> i.e. the decision of whether a task has to ignore a signal depends
> on whether the init has and some more. This is... strange :)
```

Strange, indeed... Unless you misread the code or I misunderstood your message ;)

Could you clarify? The intended behaviour is: the SIG_DFL signal is ignored if sig_kernel_ignore(sig) or we are /sbin/init. This means init ignores "more", not "less". Unless I am terribly confused...

```
> >@@ -569,6 +590,9 @@ static void handle_stop_signal(int sig,
> > /*
> > return;
> >
```

```
> >+ if (sig_init_ignore(p))
> >+ return;
> >+
>
> Why do we need for explicit stop handling for init? Shouldn't
> it be automatically checked in get_signal_to_deliver()?
```

Again, I don't quite understand what you mean.

The current behaviour is not good, we shouldn't do things like
rm_from_queue(SIGCONT) or ->signal->flags = 0 for /sbin/init.

This becomes worse with multiple namespaces if /sbin/init is ptraced
from the parent namespace (yes, such a ptracing is questionable).

Oleg.

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