
Subject: [PATCH 2/2] signals: replace p->pid == 1 check with a check for task_child_reaper

Posted by [Daniel Hokka Zakrisso](#) on Thu, 17 Jul 2008 14:56:33 GMT

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

p->pid == 1 is insufficient when there are multiple pid namespaces.
Instead, check whether the task is in the current task's
child reaper's thread group.

Signed-off-by: Daniel Hokka Zakrisson <daniel@hozac.com>

```
diff --git a/kernel/signal.c b/kernel/signal.c
index 93713a5..be932b9 100644
--- a/kernel/signal.c
+++ b/kernel/signal.c
@@ -1142,10 +1142,20 @@ static int kill_something_info(int sig, struct
siginfo *info, int pid)
    pid ? find_vpid(-pid) : task_pgrp(current));
    } else {
        int retval = 0, count = 0;
- struct task_struct * p;
+ struct task_struct *p, *reaper = task_child_reaper(current);
+
+ /*
+  * The reaper has died, so there's probably a
+  * SIGKILL pending. Return.
+  */
+ if (unlikely(!reaper)) {
+     ret = -ESRCH;
+     goto out;
+ }

    for_each_process(p) {
- if (p->pid > 1 && !same_thread_group(p, current) &&
+ if (!same_thread_group(p, reaper) &&
+     !same_thread_group(p, current) &&
        task_in_pid_ns(p, current->nsproxy->pid_ns)) {
            int err = group_send_sig_info(sig, info, p);
            ++count;
@@ -1155,6 +1165,7 @@ static int kill_something_info(int sig, struct
siginfo *info, int pid)
    }
    ret = count ? retval : -ESRCH;
}
+out:
    read_unlock(&tasklist_lock);

    return ret;
```

--

1.5.5.1

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

Subject: Re: [PATCH 2/2] signals: replace p->pid == 1 check with a check for task_child_reaper

Posted by [ebiederm](#) on Thu, 17 Jul 2008 17:55:15 GMT

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

Daniel Hokka Zakrisson <daniel@hozac.com> writes:

> p->pid == 1 is insufficient when there are multiple pid namespaces.
> Instead, check whether the task is in the current task's
> child reaper's thread group.

We should just drop the check for init as it is redundant.

Eric

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

Subject: Re: [PATCH 2/2] signals: replace p->pid == 1 check with a check for task_child_reaper

Posted by [Daniel Hokka Zakrisso](#) on Thu, 17 Jul 2008 18:21:37 GMT

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

Eric W. Biederman wrote:

> Daniel Hokka Zakrisson <daniel@hozac.com> writes:
>
>> p->pid == 1 is insufficient when there are multiple pid namespaces.
>> Instead, check whether the task is in the current task's
>> child reaper's thread group.
>
> We should just drop the check for init as it is redundant.

I'm not sure what you mean? Without protecting init here, kill -s 9 -- -1 will kill it (i.e. the init in the pid namespace). E.g.:

```
# vspace --new --pid --mount -- bash
# bash -c 'kill -s 9 -- -1'
```

will kill off all those processes, and dispose of the pid namespace.

> Eric

--

Daniel Hokka Zakrisson

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

Subject: Re: [PATCH 2/2] signals: replace p->pid == 1 check with a check for task_child_reaper

Posted by [ebiederm](#) on Thu, 17 Jul 2008 18:51:38 GMT

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

"Daniel Hokka Zakrisson" <daniel@hozac.com> writes:

> Eric W. Biederman wrote:

>> Daniel Hokka Zakrisson <daniel@hozac.com> writes:

>>

>>> p->pid == 1 is insufficient when there are multiple pid namespaces.

>>> Instead, check whether the task is in the current task's

>>> child reaper's thread group.

>>

>> We should just drop the check for init as it is redundant.

Sorry that was a half truth. Outside of the context of pid namespaces it is true.

In the context of pid namespaces it is false because we haven't merged the patches to drop signals from inside the pid namespace on the way to init.

So it is a check that `_should_` be redundant.

Eric

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