
Subject: Re: [PATCH 9/9] signal: Ignore signals sent to the pid namespace init
Posted by [ebiederm](#) on Tue, 18 Dec 2007 08:37:40 GMT

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

ebiederm@xmission.com (Eric W. Biederman) writes:

> ebiederm@xmission.com (Eric W. Biederman) writes:

>>

>> Yes. If we are not in process context (in_interrupt) we do infer

>> the sender incorrectly. Duh. I saw something in earlier patches

>> people had posted didn't understand it, and didn't get an answer

>> when I asked about it.

>

> Grr. Rather I didn't see a reply when I asked about it. I just

> spotted Suka's old reply telling me it was to attempt to prevent

> init from getting sigio. Clearly I have a big fat blind spot

> here you could drive a truck through.

>

> Kernel generated signals are a pain because they are neither fish nor

> fowl. They are sent privileged, but we don't necessarily want to give

> them all curtsies and deliver them. Grr.

I solved this in my thinking earlier but I should really mention it.

There are two kinds of kernel generated signals.

- Services like timers, SIGIO, etc where the kernel is doing something on behalf of the process.
- Process management where we send SIGKILL to make a process go away.

For services we clearly want to ignore the signal.

For the kernel flexing's it's muscle we want to accept the signal, and let init die and accept the consequences. Unfortunately we don't do that today.

Since most signals are services and since we want to preserve today semantics we want to drop kernel generated signals.

For those cases where it matters we can have the kernel send a signal by a different path, that doesn't give init the option of ignoring the signal.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
