
Subject: Re: [PATCH] Allow signalling container-init
Posted by [dev](#) on Thu, 09 Aug 2007 08:16:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Pittman wrote:

> sukadev@us.ibm.com writes:

>

>

>>Should we include this in the patchset ?

>

>

> [...]

>

>

>>Only the global-init process must be special - any other

>>container-init process must be killable to prevent run-away processes

>>in the system.

>

>

> One problem I hit while using OpenVZ is that some init processes --

> notable upstart used by Ubuntu -- depend on the special signal processing

> extended to init by the kernel.

>

> The problem here was that a signal the kernel would otherwise have

> blocked was sent to upstart, the default handler was invoked and init

> was terminated.

>

>

>>TODO: Ideally we should allow killing the container-init only from

>> ancestor containers and prevent it being killed from that or

>> descendant containers. But that is a more complex change and

>> will be addressed by a follow-on patch. For now allow the

>> container-init to be terminated by any process with sufficient

>> privileges.

>

>

> This will break, as far as I can see, by allowing the container root to

> send signals to init that it doesn't expect.

This was fixed in OpenVZ in recent kernels and Pavel tried to address this in pid namespace patches as well, but since no beautiful solution was found it was decided to postpone this issue.

NOTE: parent can still send signals to child container' init. This is convinient since you can terminate the whole container quickly from the host node.

Thanks,

Kirill

Containers mailing list

Containers@lists.linux-foundation.org

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