
Subject: Re: [PATCH] Masquerade sender information
Posted by [Cedric Le Goater](#) on Fri, 02 Nov 2007 13:40:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

Serge E. Hallyn wrote:

> Quoting Eric W. Biederman (ebiederm@xmission.com):
>> sukadev@us.ibm.com writes:
>>
>>> From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
>>> Subject: [PATCH] Masquerade sender information
>>>
>>> With multiple pid namespaces, sender of a signal could be in an ancestor
>>> namespace of the receiver and so the sender will not have a valid 'pid_t'
>>> in the receiver's namespace.
>>>
>>> In this case, masquerade the 'siginfo' for the signal to pretend that the
>>> signal originated from the kernel.
>> At first glance this looks ok. I think the only case where we can
>> be sending a signal from inside a pid namespace to something not
>> in a child pid namespace is if we are the kernel. In which case
>
> Are we now blocking F_SETOWN|F_SETSIG signals to outside our pid
> namespace? mq_notify? (I didn't think we were)

My understanding is that we're not blocking and that a process killing
another process in a sibling pid namespace will have a si_pid = 0.

C.

>
>> we also want si_pid = 0.
>>
>> If that holds this problem is easier then I was thinking it would
>> be.
>>
>> Eric
>>

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