
Subject: Re: [PATCH 11/15] Signal semantics

Posted by [Sukadev Bhattiprolu](#) on Fri, 27 Jul 2007 20:23:37 GMT

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

Serge E. Hallyn [serue@us.ibm.com] wrote:

| Quoting sukadev@us.ibm.com (sukadev@us.ibm.com):

| > Pavel Emelianov [xemul@openvz.org] wrote:

| > | Oleg Nesterov wrote:

| > | > (What about ptrace_attach() btw? If it is possible to send a signal to init

| > | > from the "parent" namespace, perhaps it makes sense to allow ptracing as

| > | > well).

| > |

| > | ptracing of tasks fro different namespaces is not possible at all, since

| > | strace utility determines the fork()-ed child pid from the parent's eax

| > | register, which would contain the pid value as this parent sees his child.

| > | But if the strace is in different namespace - it won't be able to find

| > | this child with the pid value from parent's eax.

| > |

| > | Maybe it's worth disabling cross-namespaces ptracing...

| > |

| > | I think so too. Its probably not a serious limitation ?

|

| Several people think we will implement 'namespace entering' through a

| ptrace hack, where maybe the admin ptraces the init in a child pidns,

| makes it fork, and makes the child execute what it wants (i.e. ps -ef).

|

| You're talking about killing that functionality?

No. I was only thinking in terms of debugging container init and missed the namespace entering part.

Pavel, I am not sure I understand your comment about being unable to ptrace() a child ns.

BTW, I am able to gdb a process (incl container-init) from parent ns now.

|

| -serge
