
Subject: Re: netns : close all sockets at unshare ?

Posted by [Daniel Lezcano](#) on Wed, 03 Oct 2007 08:40:49 GMT

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

Eric W. Biederman wrote:

> Daniel Lezcano <dlezcano@fr.ibm.com> writes:

>

>> Hi,

>>

>> I was looking at some cornercases and trying to figure out what happens if
>> someone does:

>>

>> 1 - fd = socket(...)

>> 2 - unshare(CLONE_NEWNET)

>> 3 - bind(fd, ...) / listen(fd, ...)

>>

>> There is here an interaction between two namespaces.

>> Trying to catch all these little tricky paths everywhere with the network

>> namespace is painful, perhaps we should consider a more radical solution.

>

> Huh?

>

> socket() puts the namespace on struct sock.

> bind/listen etc just look at that namespace.

>

> Unless I'm blind it is simple and it works now.

Yes, it will work.

Do we want to be inside a network namespace and to use a socket
belonging to another network namespace ? If yes, then my remark is
irrelevant.

>> Shall we close all fd sockets when doing an unshare ? like a close-on-exec
>> behavior ?

>

> I think adopting that policy would dramatically reduce the usefulness
> of network namespaces.

>

> Making the mix and match cases gives the implementation much more flexibility
> and it doesn't appear that hard right now.

I am curious, why such functionality is useful ?

Containers mailing list

Containers@lists.linux-foundation.org

