Subject: Re: [patch 0/1][NETNS49] Make af_unix autobind per namespace
Posted by Daniel Lezcano on Wed, 03 Oct 2007 08:35:31 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> Daniel Lezcano <dlezcano@fr.ibm.com> writes:
>
>> Eric W. Biederman wrote:
>>> Daniel Lezcano <dlezcano@fr.ibm.com> writes:
>>>
>>>> The following patch change autobind fonction to use the ordernum
>>>> from the network namespace instead of using the local static variable.
>>> Why do we care?
>>> Information leak?
>>> Some application is expecting a predictable autobind value?
>>>
>>> Just skimming the code it looks like it will work correctly without
>>> this.
>> I think my summary is ... too short :)
>>
>> I don't see any applications taking care of this. If they ask for an abstract
>> socket, then they don't care about the bind result. So probably, the patchset is
>> totally useless.
>>
>> But from the POV of the checkpoint/restart, we should check if this value is
>> somewhere visible from userspace and so storable by an application.
>
> Right. And we already can already specifically select this result.
> My point is that the semi random sequence generator logic does not
> need to be per namespace, because people don't care what the sequence.
> That sequence is not exported to user space.
>
>> It appears this is the case with /proc/net/unix, where an abstract socket is
>> symbolized by the path pattern "@". Example:
>>
>> cat /proc/net/unix
>>
>> Num   RefCount Protocol Flags   Type St Inode Path
>> c6a27710: 00000002 00000000 00000000 0002 01 4357 @00003
>
> Right, and that part we should definitely preserve for checkpoint/restart
> purposes.
>
>> I agree by the fact that can be considered as a detail and the probability to
>> have an application storing this informaton is very small (eg. checkpointing
>> while doing netstat in the container). But IMHO, the paradigm "never seen from
>> userspace" fails and that justifies to have the ordernum variable relative to a
>> namespace.

>
> My point was that ordernum itself is not seen. It is just an arbitrary number
> and we are allowed to change the algorithm for selecting a new abstract
> namespace name at will.

Hmm, right. That makes sense.

> If there is something in userspace that depends on the algorithm for selecting
> the abstract name then making ordernum per namespace make sense.

Ok, fair enough. Let's forget this patch. It is small enough to rewrite
it if unexpectedly something bad happens with ordernum.

  -- Daniel