
Subject: Re: [patch 0/1][NETNS49] Make af_unix autobind per namespace

Posted by [Cedric Le Goater](#) on Wed, 03 Oct 2007 15:34:10 GMT

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

Eric W. Biederman wrote:

> Cedric Le Goater <clg@fr.ibm.com> writes:

>> my 2 cnts,

>>

>> when 'restarting' a socket bound to an abstract name, we will have

>> a EADDRINUSE if we try to rebind it to an abstract name which is

>> already in use by a socket in a another namespace ?

>

> No.

ok. i just went over your AF_UNIX patch and saw that 'struct net' was being checked for abstract sockets.

C.

>> it seems to me that this is an identifier and like any identifier

>> it should be private to the namespace, which probably means having

>> unix_abstract_socket_table[] per net namespace.

>

> Yes it is. It is a hash table so we are filter the hash chain

> and not having two copies of the table. But effectively it's

> the same thing.

>

> All this patch was suggesting was having a per network namespace

> copy of the data structure for the random number generator for

> generating the name.

>

> The ``random number generator'' is just a 16bit counter that loops

> through all 64k values seeing if a name is available and if so

> using it. Sharing our place in the loop between different namespaces

> may be ineffeicient but it should work fine.

>

> Eric

Containers mailing list

Containers@lists.linux-foundation.org

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