
Subject: Re: [PATCH 5/5] net: Make AF_UNIX per network namespace safe.
Posted by [ebiederm](#) on Sat, 29 Sep 2007 17:03:05 GMT

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

Patrick McHardy <kaber@trash.net> writes:

> Eric W. Biederman wrote:

>> Because of the global nature of garbage collection, and because of the
>> cost of per namespace hash tables `unix_socket_table` has been kept
>> global. With a filter added on lookups so we don't see sockets from
>> the wrong namespace.

>>

>> Currently I don't fold the namespace into the hash so multiple
>> namespaces using the same socket name will be guaranteed a hash
>> collision.

>

>

> That doesn't sound like a good thing :) Is there a reason for
> not avoiding the collisions?

Two reasons. Minimizing the size of the changes to make review easier, and I don't know if hash collisions are likely in practice or if they matter. I don't believe we can't physically collide and have the same inode because we make a node in the filesystem. The abstract domain is local to linux and so people don't use it as much.

All of which boils down to. I don't see it matter a heck of a lot especially initially. So I did the traditional unix thing and started with a simple and stupid implementation. But it didn't quite feel right to me either so I documented it.

Whipping up a patch to take the namespace into account in `mkname` doesn't look too hard though.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

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