Subject: Re: [RFC PATCH] SUNRPC: connect local transports synchronously
Posted by Stanislav Kinsbursky on Fri, 17 Feb 2012 08:25:45 GMT
View Forum Message <> Reply to Message

> On Thu, 2012-02-16 at 19:06 +0400, Stanislav Kinsbursky wrote:
>> Local tranports uses UNIX sockets and connecting of these sockets is done in
>> context of file system namespace (i.e. task file system root).
>> Currenly, all sockets connect operations are performed by rpciod work queue,
>> which actually means, that any service will be registered in the same rpcbind
>> instance regardless to process file system root.
>> This is not containers, which usually have it's own nested root. There are 2
>> approaches, how to solve the problem. First one is to store proper root in
>> tranport and switch to it in rpciod workqueue function for connect operations.
>> But this looks ugly. The second one is to connect to unix sockets
>> synchronously. This aptch implements the last one.
>
> That approach can fall afoul of the selinux restrictions on the process
> context. Processes that are allowed to write data, may not be allowed to
> create sockets or call connect(). That is the main reason for doing it
> in the rpciod context, which is a clean kernel process context.
>

Thanks for explanation, Trond.
So, this connect have to be done in kernel process context.
Now I can see 2 ways how to meet this requirement and reach the goal:

1) Change the fs root for rpciod while connecting.
2) Do not touch rpciod and launch special "connect" kernel thread to perform
connect operations for unix sockets.

What do you think about this 2 ways above? Which one is less worse from your POW?
Maybe you have even a better solution for the problem?

--
Best regards,
Stanislav Kinsbursky