Subject: Re: [RFC PATCH] SUNRPC: connect local transports synchronously
Posted by Myklebust, Trond on Thu, 16 Feb 2012 15:13:53 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.

--
Trond Myklebust
Linux NFS client maintainer

NetApp
Trond.Myklebust@netapp.com
www.netapp.com