Subject: Re: [PATCH v3 00/16] make rpc_pipefs be mountable multiple time
Posted by Rob Landley on Thu, 20 Jan 2011 13:57:52 GMT
View Forum Message <> Reply to Message

On 01/20/2011 05:35 AM, Kirill A. Shutemov wrote:
> On Mon, Jan 17, 2011 at 06:30:16AM -0600, Rob Landley wrote:
>> On 01/14/2011 07:48 AM, Kirill A. Shutemov wrote:
>>> Prepare nfs/sunrpc stack to use multiple instances of rpc_pipefs.
>>> Only for client for now.
>>
>> Ok, Google is being really unhelpful here.
>
> It's better if you read the code. :)

I did. That doesn't necessarily help if I don't know what it's _trying_ to do and this implementation is just one small piece of a much larger mechanism.

>> What is rpc_pipefs for? What uses it, and to do what exactly? Is it
>> used by nfs server code, or by the client code, or both? Is it a way
>> for userspace to talk to the kernel, or for the kernel to talk to
>> itself? Is it used at mount time, or during filesystem operation?
>
> Ok, It try to answer. Please correct me, if I'm wrong.
>
> rpc_pipefs is a userland/kernel interface (I don't see kernel-kernel
> usecases, but it's possible, I guess).

Then why is the host context treated specially with an init_rcp_pipefs instance that isn't necessarily mounted anywhere? (How does userspace access it if it isn't mounted anywhere? Why doesn't the host context just use the same does-not-exist code as container context? Presumably there's a reason for this?)

I've also mounted NFS shares in test environments I never mounted rpc_pipefs in. (It's possible that the nfs.mount command was doing so for me, and apparently unmounting it again afterwards.) My test environment is exporting with the userspace NFS daemon so it's not using the kernel cacheing infrastructure.

> There is client dir (nfs/clntX) in rpc_pipefs for every sunrpc client.
> Both client and server (see fs/nfsd/nfs4callback.c) can create sunrpc
> client. So we rpc_pipefs on both side.

Ok, both client and server can create instances. And use them to do what?

I understand that the kernel needs to handle RPC calls to service NFS requests, what I'm not figuring out is how userspace is involved after

the initial mount except on the other side of filesystem-agnostic system
calls.

> rpc_pipefs uses not only on mount time. See old idmapper, for example.

Um, does this sentence say that the old idmapper uses rpc_pipefs?
(Presumably not the kernel infrastructure controlled by
CONFIG_NFS_USE_NEW_IDMAPPER because you said it's not a kernel-kernel
communication mechanism... So you're saying that whatever userspace
infrastructure interacts with that will also depend on rpc_pipefs. But
the new infrastructure doesn't, which is why I need to look at the old?)

(And again, CONFIG_NFS_USE_NEW_IDMAPPER needs /sbin/request-key from the
keyutils package implying this is primarily for authentication, which
would explain why I haven't seen it in my simple test environment...)

Rob

_____

Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containe rs