Subject: Re: [PATCH 0/6] SUNRPC: make RPC clients use network-namespace-aware PipeFS routines
Posted by Stanislav Kinsbursky on Wed, 23 Nov 2011 17:58:25 GMT
View Forum Message <> Reply to Message

> On Wed, Nov 23, 2011 at 02:51:10PM +0300, Stanislav Kinsbursky wrote:
>> This patch set was created in context of clone of git
>> branch: git://git.linux-nfs.org/projects/trondmy/nfs-2.6.git.
>> tag: v3.1
>>
>> This patch set depends on previous patch sets titled:
>> 1) "SUNRPC: initial part of making pipefs work in net ns"
>> 2) "SUNPRC: cleanup PipeFS for network-namespace-aware users"
>>
>> This patch set is a first part of reworking SUNPRC PipeFS users.
>> It makes SUNRPC clients using PipeFS nofitications for directory and GSS pipes
>> dentries creation. With this patch set RPC clients and GSS auth creations
>> routines doesn't force SUNRPC PipeFS mount point creation which actually means,
>> that they now can work without PipeFS dentries.
>
> I'm not following very well. (My fault, I haven't been paying
> attention.) Could you summarize the itended behavior of pipefs after
> all this is done?
>
> So there's a separate superblock (and separate dentries) for each
> namespace?
>

Yes, you right.
So, here is a brief summary of what will be at the end:

1) PipeFS superblock will be per net ns.

2) Superblock holds net ns, which is taken from current. Struct net will have link ot pipefs superblock (it can be NULL, if PipeFS wasn't mounted yet or already unmounted).

3) Notifications will be send on superblock creation and destruction.

4) All kernel mount point creation and destruction calls (rpc_get_mount() and rpc_put_mount()) will be removed. I.e. this superblock will be created only from user-space.

5) Kernel pipes and dentries will be created or destroyed:
1. During per-net operations (only for static NFS stuff: dns_resolve cache, pnfs blocklayout and idmap pipes).
2. On notification events (all directories, files and pipes in proper

callbacks). Notification subscribers:
a. rpc clients (responsible for client dentries and gss pipes creation),
b. nfs clients (responsible nfs idmap pipes),
c. nfs dns_resolve cache,
d. pnfs blocklayout pipes,

6) PipeFS dentries creation logic:
a) All directories and files creators - will try to create them as usual. But if fail - then no problem here. I.e. dentries will be created on PipeFS mount notification call.
b) Pipes creators - will create new structure rpc_pipe (all pipe stuff from rpc inode) nad then try to create pipe denties. If fail - then, again, no problem. Dentries will be be created on PipeFS mount notification call.

Almost all (exept 5.2.b - forgot about it during debasing and resending) is done and ready to send.

> What decides which clients are visible in which network namespaces?
>

Clients dentries will be created in proper superblock from the beginning. I.e. rpc clients transports have struct net reference. NFS clients will have such reference too. Struct net will have reference to pipefs superblock.
Currently, for dentry creation all we need is parent dentry. Some of creators (like GSS pipes) takes parent dentry from associated struct (like rpc_clnt). For others parent dentry can be found by simple d_lookup() starting from sb->root (reminder: sb can be taken from net).

--
Best regards,
Stanislav Kinsbursky