



NetApp®

NFSv4.2 Inter SSC Prototype

Andy Adamson
andros@netapp.com
IETF 90, Toronto Canada

Presenter Title

Date

Motivation

- NFSv4.2 inter SSC is complicated.
- The destination server needs to become an NFS client
- There is no OPEN of the “file to be copied”
- Add in GSS3 and there is a lot going on
- A working prototype gives us confidence that the wire protocol is complete.

What is Working

- Anna Schumaker from NetApp wrote an *Intra* SSC prototype for Linux client and servers
 - Reviewed by Trond
 - Included new VFS interfaces for the do-splice call
- My *Inter* SSC prototype starts with Anna's code
- NFSv4.2 Inter SSC copy is working
 - Caveat one change in the protocol
- GSS3 is started
 - I can negotiate and run with GSS3 or GSS1

NFSv4.2 Inter SSC Tasks

- Destination server needs to launch an NFS client
 - Wrote a module that is called by the destination server when an Inter SSC COPY is received
- Module mounts the source server
 - NL4_NETADDR : not enough info to mount the source server as an export path is required.
 - Replaced NL4_NETADDR with NL4_EXPORT

```
struct fs_export4 {  
    utf8str_cis  server;  
    pathname4    export;  
}
```

NFSv4.2 Inter SSC Tasks

- Module then needs to cobble together (e.g. hack) an “OPEN” filling in both VFS and NFSv4 data structures
 - Do a GETATTR on the COPY_SAVE_FH file handle
 - Create an inode, a dentry, a struct file, an nfs_opencontext (with ca_src_stateid) all bypassing most of the normal interfaces
 - Place the “open file” in the directory client namespace directly under the mount point with a fake name

NFSv4.2 Inter SSC Tasks

- The module ends up with a struct file with sufficient data to pass to the do-splice function
 - Same server interface used by Anna's Intra SSC which does the “read from write to”
- The do-splice function calls NFSv4.1 READ
- Then the module cleans up and umounts
 - Took a while to get the reference counting right as it is done by hand, not by the VFS.
 - Caveat: The NFS clean up function assumes that an OPEN was called, and calls CLOSE which gets an NFS4ERR_BAD_STATEID

NFSv4.2 Inter SSC

- There are many loose ends in the code, but the prototype served its purpose of verifying most of the protocol and identifying a protocol change
- I will clean up what I have, and forward port it to Anna's Intra SSC git tree, as well as post the code to the linux-nfs@vger.kernel.org list.

GSS3 Prototype Tasks (done)

- Add the ability to negotiate GSS3, and fall back to GSS1 if GSS3 is not supported
 - Required changes to libtirpc which does the negotiation and rpc.gssd which passes the negotiated version to the kernel.
 - Client kernel changes to parse the GSS version and store it in the gss_context
 - Server changes to allow GSS3 and GSS1 then grab the GSS version and store it with the server context

GSS3 Prototype Tasks (not done)

- Code the new reply verifier for GSS3
- Add the RPCSEC_GSS_CREATE call
- Add the Inter SSC copy_to_auth, copy_from_auth, and copy_confirm_auth payloads
- Verify the Inter SSC RPCSEC_GSS_CREATE payloads
- Use GSS3 on the Inter SSC READ

GSS3 Prototype Tasks (not done)

- Code the responses to errors – e.g. when one of the three GSS3 handles is destroyed or invalidated.
- Check that the Inter SSC copy stops etc.



NetApp®

Thank you