



Linux Support of NFS

v4.1 and v4.2

Steve Dickson
steved@redhat.com
Mar Thu 23, 2017

Agenda

- NFS v4.0 issues
- NFS v4.1 Supported Features
- NFS v4.2 Supported Features
- V4.2 Not Supported Features
- V4.x Other Features



NFS v4.0 – Issues

- Performance issues
- Callback issues
- Delegation issues
- Locking Issues
- Mount negotiations start with v4.2



NFSv4.1 - Features

- Sessions
 - Exactly Once Semantics.
 - Callback problem fixed.
- pNFS
 - File layout
 - Netapp
 - SCSI Layout
 - Linux
 - Flex File layouts
 - Primary Data



NFS V4.2 – Feature

- **Sparse File Support** – Sparse files are ones which have allocated or uninitialized blocks in the file.
(1.4.3)
 - SEEK - Find the Next Data or Hole
 - NFS: Implement SEEK (Sep 2014)
 - READ_PLUS - READ Data or Holes from a File
 - No published upstream patches yet
 - Prototype patches improve performance with reading holes (See Anna's keynote graph).



NFS V4.2 – Feature

- **Space Reservation** - Creating holes in files by transferring just the metadata over the network
 - ALLOCATE - Reserve Space in A Region of a File
 - nfs: Add ALLOCATE support (Jul 2015)
 - DELLOCATE - Unreserve Space in a Region of a File
 - nfs: Add DEALLOCATE support (Jul 2015)



NFS V4.2 - Feature

- **Server-side copy** – Provides a mechanism for the NFS client to perform a file copy on the server without the data being transmitted back and forth over the network.
 - COPY - Initiate a server-side copy
 - NFS: Add COPY nfs operation (Feb 2017)
 - Clone - nfs42: add CLONE proc functions (Sep 26 2015)
 - Huge performance gain!



NFS v4.2 - Feature

- **Security labels** - A file object attribute allows the server to store labels on files, which the client retrieves and uses to enforce data access. (1.4.6)
 - SECURITY_LABEL
 - NFS: Client implementation of Labeled-NFS (May 2013)



NFS v4.2 – Not Supported

- **IO Advise** - Applications and clients want to advise the server as to expected I/O behavior. (1.4.2)
 - IO_ADVICE - Application I/O access pattern hints
 - not implemented
- **Application Data Block (ADB) Support** - Some applications treat a file as if it were a disk and as such want to initialize (or format) the file image. (1.4.5)
 - WRITE_SAME - WRITE an ADB Multiple Times to a File
 - not implemented



NFSv4.x – Other features

- LAYOUTSTATS
 - Performance statistics from DS to MDS
 - FlexFile use to load balance between DS(s)
- CB_NOTIFY_LOCK - Notify Client of Possible Lock Availability
 - Improve locking performance
- Session Trunking
 - The use of multiple connections between a client and server in order to increase the speed of data transfer.
 - pNFS File and Flexfile layout
- NFSv4.1 RDMA
 - NFS4.1 support
 - Kerberos Support



Questions

References

- RFC 5661 - <https://tools.ietf.org/html/rfc5661>
- RFC 7862 - <https://tools.ietf.org/html/rfc7862>
- Vault 2015 Talk - <http://events.linuxfoundation.org/sites/events/files/slides/vault2015.pdf>

