

Federated file system status IETF'72 NFSv4 WG meeting

Daniel Ellard, Theresa Raj, Amy Weaver
NetApp



Acknowledgments

Many people have contributed!

- Craig Everhart: NetApp ATG
- Renu Tewari, Manoj Naik: IBM/Almaden
- Paul Lemahieu, EMC/Rainfinity
- Mario Würzl: EMC
- Rob Thurlow: SUN

Outline

- Overview
- Current status

Goals

- An open, portable protocol that permits the construction of cross-platform, federated file systems accessible to unmodified NFSv4 clients.
 - Open and portable: anyone can use it
 - Cross-platform: cross-vendor, cross-system, cross-version
 - Federated: federation members don't have to give up control of their systems
- Not meant to replace existing cluster protocols!
 - Cluster-of-clusters protocol

Project history

- 2/2007: Public “federated-fs” community formed at FAST’07
 - Mailing list and weekly conference calls
- 6/2007: First draft of requirements doc
 - Includes common terms and definitions
 - Presented at IETF’69
- 12/2007: First draft of protocol specs
 - Presented at IETF’70 and IETF’71
- 6/2008: Proof-of-concept implementation

Terms and definitions

- Fileset: a directory tree (volume)
- Junction: a fileset object that provides a way for one fileset to reference the root of another
- FSL (fileset location): the location of a fileset instance
- FSN (fileset name): an opaque fileset identifier
- NSDB (namespace database): a service that tracks the mapping between FSNs and FSLs
 - Typically one NSDB per administrative domain
- Each FSN contains an FsnUuid (a UUID) and an NSDB location

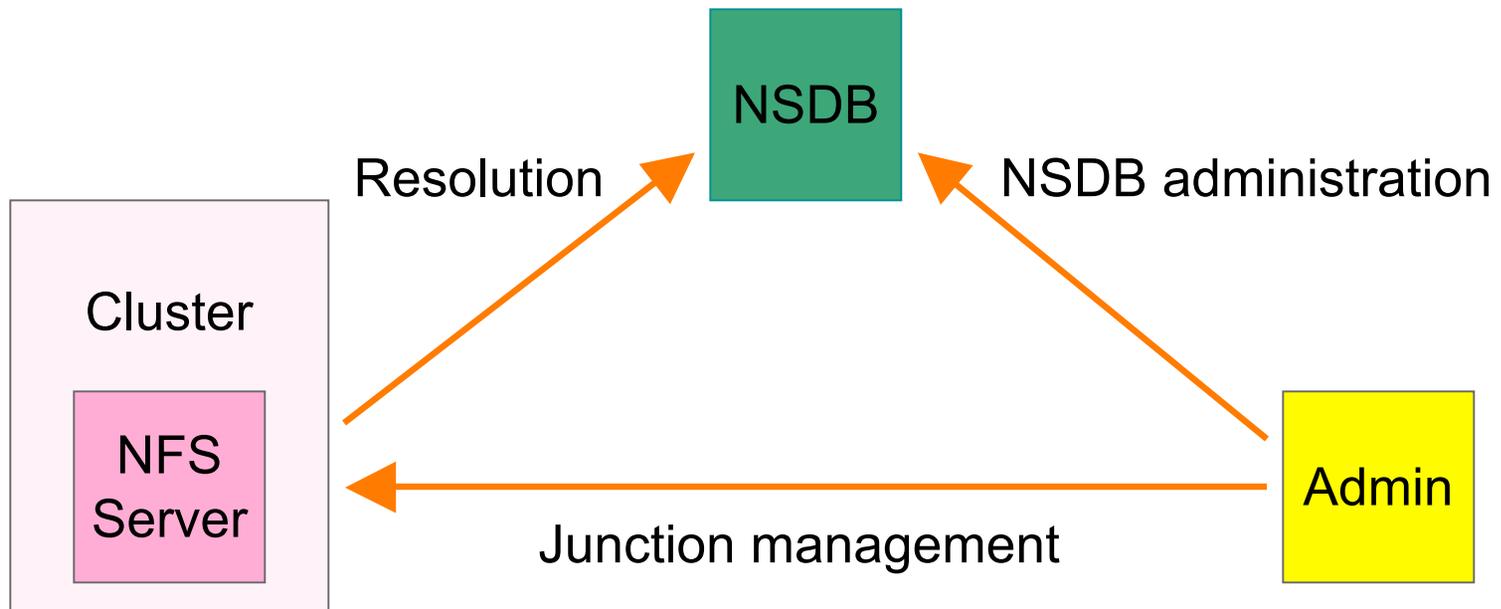
Protocol overview

- NFSv4 servers know where the junctions are
 - Can detect when a client accesses a junction
 - Can figure out the FSN of the target fileset
 - How both of these are accomplished are implementation details – not part of the protocol
- NSDBs know where each FSN is implemented
- NFSv4 servers contact the NSDBs to find where to redirect clients when the client accesses a junction
- Administrators manage the junctions and the FSN->FSL mappings

Three sub-protocols are being proposed

1. Resolution (server to NSDB): where are the filesets?
2. NSDB administration (admin to NSDB): FSN->FSL map
3. Junction management (admin to server): create, delete, manage junctions

Note: no changes to client protocols



Status

- Requirements draft fairly stable
 - Common terms and definitions agreed upon
- Three sub-protocol drafts in progress
 - Resolution protocol: farthest along
 - NSDB admin protocol: personal draft
 - Server admin protocol: straw man proposal

Open issues

- Much discussion over a possible fourth sub-protocol for root fileset discovery
 - May require changes to the client (bad)
 - May simplify configuration of the client (good)
- Current NSDB based on LDAP
 - LDAP is a convenient platform for a directory
 - LDAP might not be adequately extensible
- Cross-enterprise authentication/identity
 - How effectively can we federate identity without mutual trust?
- We could really use a fileset migration protocol

Conclusions

- We believe that this work is ready to be picked up by the NFSv4 WG and added to charter
 - Proposed on the nfsv4 mailing list
 - All responses were positive
- For more information, join the federated-fs mailing list:
federated-fs@sdsc.edu