

Federated File System (FedFS) Status

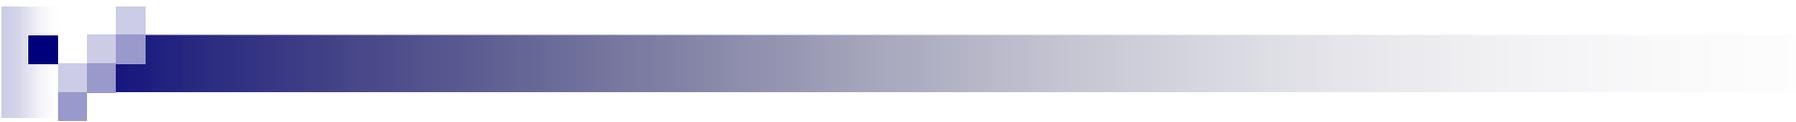
IETF'73 NFSv4 WG Meeting
November 21, 2008

James Lentini
jlentini@netapp.com



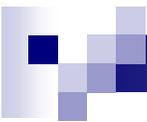
Overview

- FedFS is a set of open, portable protocols that permit the construction of a cross-platform, federated file system namespace accessible to unmodified NFSv4 clients.
- Key points:
 - Unmodified clients
 - Open: cross-platform, multi-vendor
 - Federated: participants retain control of their systems



Federated Namespace Benefits

- Simplified Management
 - Clients only need to mount to root (or some other part) of the namespace
- Provides a basis for future enhancements that are transparent to the client (outside the scope of this effort):
 - Replication: may be used for load balancing or high availability
 - Migration: may be used for moving data closer to compute or decommissioning systems



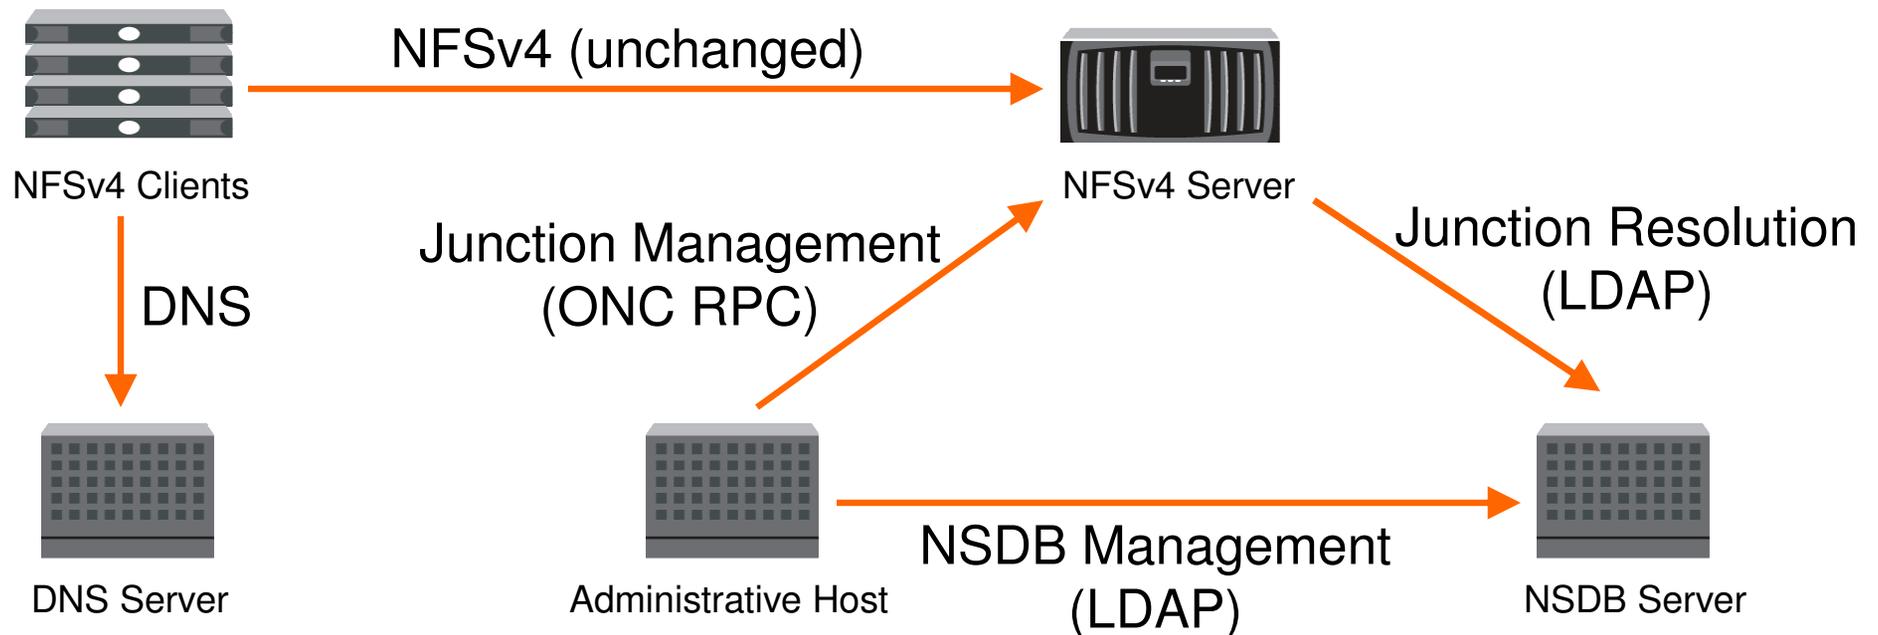
Terms

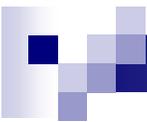
- **Fileset**: a directory tree (volume)
- **FSN** (fileset name): a fileset identifier that is independent of the representation of the fileset
 - Each FSN contains an FsnUuid (a UUID) and an NSDB location
- **FSL** (fileset location): network location of a fileset instance
- **Junction**: an object that provides a way for one fileset to reference another
- **NSDB** (namespace database): a service that tracks the mapping between FSNs and FSLs; implemented with LDAP

Design

FedFS defines:

- DNS SRV record format for root namespace information
- NSDB's LDAP schema, used for
 - Junction Resolution: mapping an FSN to an FSL
 - NSDB Management: create FSN to FSL mapping
- ONC RPC administration protocol used for creating a junction at a given point in the namespace to a specified FSN





Client Support

NFSv4 clients supporting referrals are available on many platforms. For example:

- **AIX:** referrals and replication (including failover) supported since 5.3 (released August, 2004)
- **HPUX:** referrals supported in HP-UX 11iv3 with ONCplus B.11.31.03 (released May, 2008)
- **Linux:** referrals supported since 2.6.18 (released September, 2006)
 - Migration/replication support under development



Status

- Prototype demonstrated at the summer WG meeting in Dublin
- Four drafts published as NFSv4 WG documents
- Addition to NFSv4 WG charter proposed, waiting for update from WG



Drafts

Four drafts published as working group documents:

- Requirements for Federated File Systems
- NSDB Protocol for Federated Filesystems
- Admin Protocol for Federated Filesystems
- Using DNS SRV to Specify a Global File Name Space with NFS version 4



Meetings

- Open meetings are held each week to discuss issues and review proposals.
- Thursdays, 1:30 – 2:30 PM Eastern
(10:30 - 11:30 AM Pacific)
 - Conference Number: 1-888-765-3653
 - Conference ID: 2354843



Open Issues

- Support for other file access protocols (e.g. CIFS, NFSv3)
 - Assumption is that namespaces will be independent, no referrals from one protocol to another.
- NSDB implemented with LDAP: Is LDAP flexible enough? Are the semantics correct?
 - Define NSDB replication semantics
 - Participation and review from LDAP experts would be appreciated.
- Validate assumptions about group membership in a federated system.
- Define a mechanism for distributing the root fileset (absent a fileset replication protocol)



Acknowledgements

Many people have contributed!

- Dan Ellard (BBN Technologies)
- Craig Everhart, James Lentini (NetApp)
- Renu Tewari, Manoj Naik (IBM)
- Paul Lemahieu, Mario Würzl (EMC)
- Rob Thurlow (Sun)

Special thanks to Mike Eisler for presenting these slides.