

Federated Filesystem (FedFS)

James Lentini
jlentini@netapp.com

IETF '77 NFSv4 WG Meeting
March 23, 2010

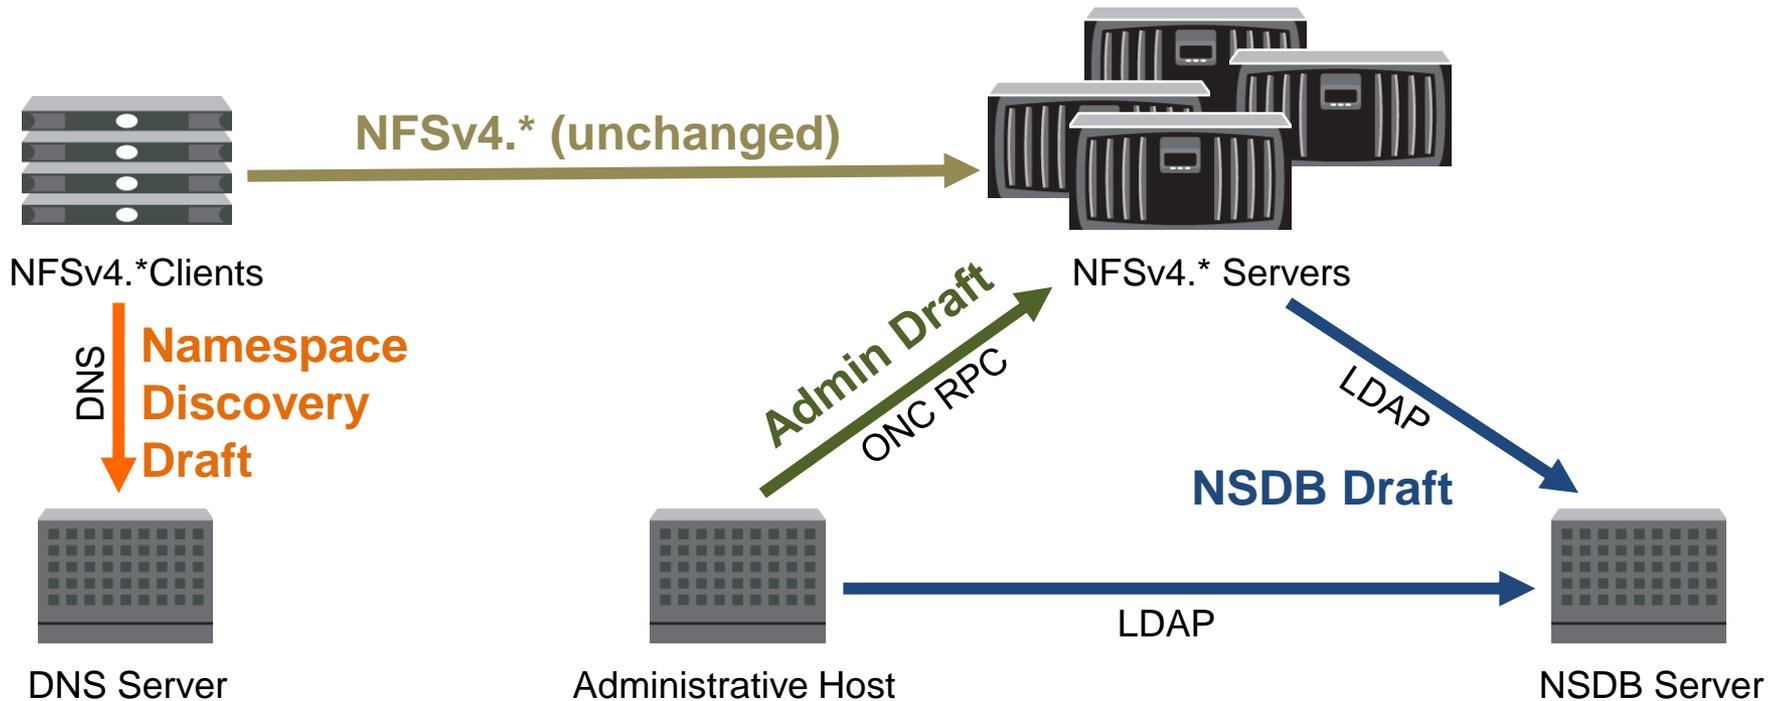
Summary

Requirements published as **RFC 5716** in January.

Three remaining drafts being prepared for WG last call:

- Namespace Discovery
draft-ietf-nfsv4-federated-fs-dns-srv-namespace
- NSDB Protocol
draft-ietf-nfsv4-federated-fs-protocol
- Admin Protocol
draft-ietf-nfsv4-federated-fs-admin

Protocols



Namespace Discovery: how a client locates a namespace's root.

NSDB: (a) database management and (b) junction resolution.

Admin: junction management (create/delete/query).

Highlights since IETF'76

- NSDB and Admin drafts updated in January.
- Three week review period for all three drafts from January 21 to February 11.
 - All feedback was related to Admin protocol.
- Implementations of FedFS NSDB protocol were successfully tested at Connectathon in February.

Namespace Discovery

draft-ietf-nfsv4-federated-fs-dns-srv-namespace

Summary: A DNS record format for publishing the location of a namespace's root.

Proposed

Category: Standards Track

Status: Only cosmetic changes discussed since IETF'76 in November. idnits will be addressed in next update.

Next Steps:

- Respond to any future mailing list feedback.
- WG Last Call (April or May?).
- Final DNS Directorate Review.

NSDB Protocol

draft-ietf-nfsv4-federated-fs-protocol

Summary: Defines the LDAP schema and operations for a Namespace Database (NSDB).

Proposed

Category: Standards Track

Status: January -05 version included:

- Recommendations for generating a referral from an FSL (2.4.3).
- Expanded junction resolution example (3.2).
- Clarified format of fedfsNetAddr [DNS name, IPv4, or IPv6] (4.2.1.2) and fedfsNsdbName [only DNS name] (4.2.1.4).

Next Steps:

- WG Last Call (May or June?).
- Final LDAP Expert Review (passed previous reviews).

Admin Protocol

draft-ietf-nfsv4-federated-fs-admin

Summary: Describes an ONC RPC protocol to create/delete/query a junction on a fileserver.

Proposed

Category: Standards Track

Status: January -04 version included:

- Expanded junction query (see new error codes and functionality for FEDFS_LOOKUP_FSN).
- NSDB Trust Anchor Management (see FEDFS_*_NSDB_PARAMS).

Next Steps:

- Resolve questions about junction size, pathnames, and international strings (DNS names, paths, etc).
- WG last call (June or July?).

Meetings

Open meetings are held each week to resolve 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

Acknowledgements

Many people have contributed! Including:

George Amvrosiadis (Univ. of Ioannina)

Andy Adamson (NetApp)

Dan Ellard (BBN Technologies)

Craig Everhart (NetApp)

Sorin Faibish (EMC)

Paul Lemahieu (EMC)

James Lentini (NetApp)

Chuck Lever (Oracle)

Trond Myklebust (NetApp)

Pavan Mettu (Oracle)

Manoj Naik (IBM)

Chris Stacey (EMC)

Renu Tewari (IBM)

Robert Thurlow (Oracle)

Nicolas Williams (Oracle)

Mario Würzl (EMC)

BACKGROUND INFORMATION

What is FedFS?

- FedFS is a set of open protocols that permit the construction of a scalable, cross-platform federated file system namespace accessible to unmodified NFSv4[.1] clients.
- Key points:
 - Unmodified clients
 - Open: cross-platform, multi-vendor
 - Federated: participants retain control of their systems
 - Scalable: supports large namespaces with many clients and servers in different geographies

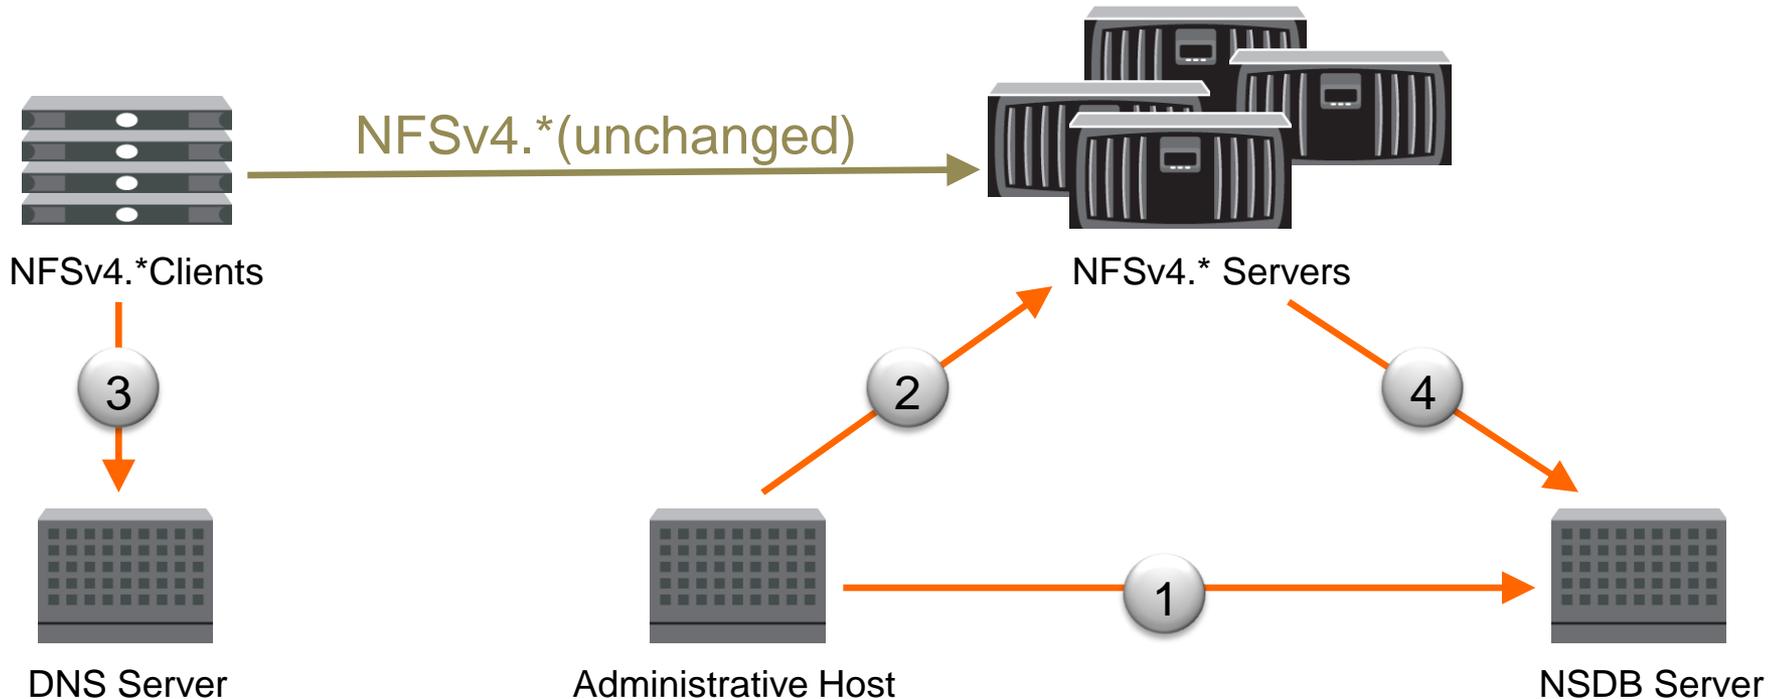
FedFS Protocols

Namespace Management

- 1 NSDB Management (LDAP)
- 2 Junction Management (ONC RPC)

Namespace Navigation

- 3 Namespace discovery (DNS)
- 4 Junction resolution (LDAP)

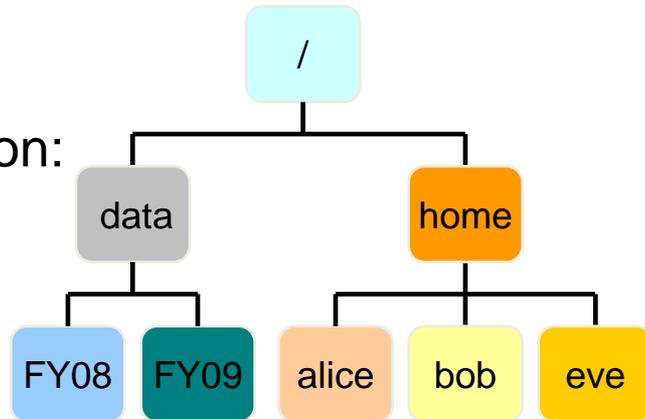


What are the benefits?

- Simplified management: eliminates complicated software such as the automounter.
- Separates logical and physical data location: allows data movement for cost/performance tiering, worker mobility, and application mobility.
- Enhances:
 - Data Replication: for load balancing or high availability
 - Data Migration: for moving data closer to compute or decommissioning systems
 - Cloud Storage: for the dynamic data center, enterprise clouds, or private internet clouds.

Federated Namespace Example

The illusion:



- The user and application software see a simple, hierarchical namespace.

The reality:



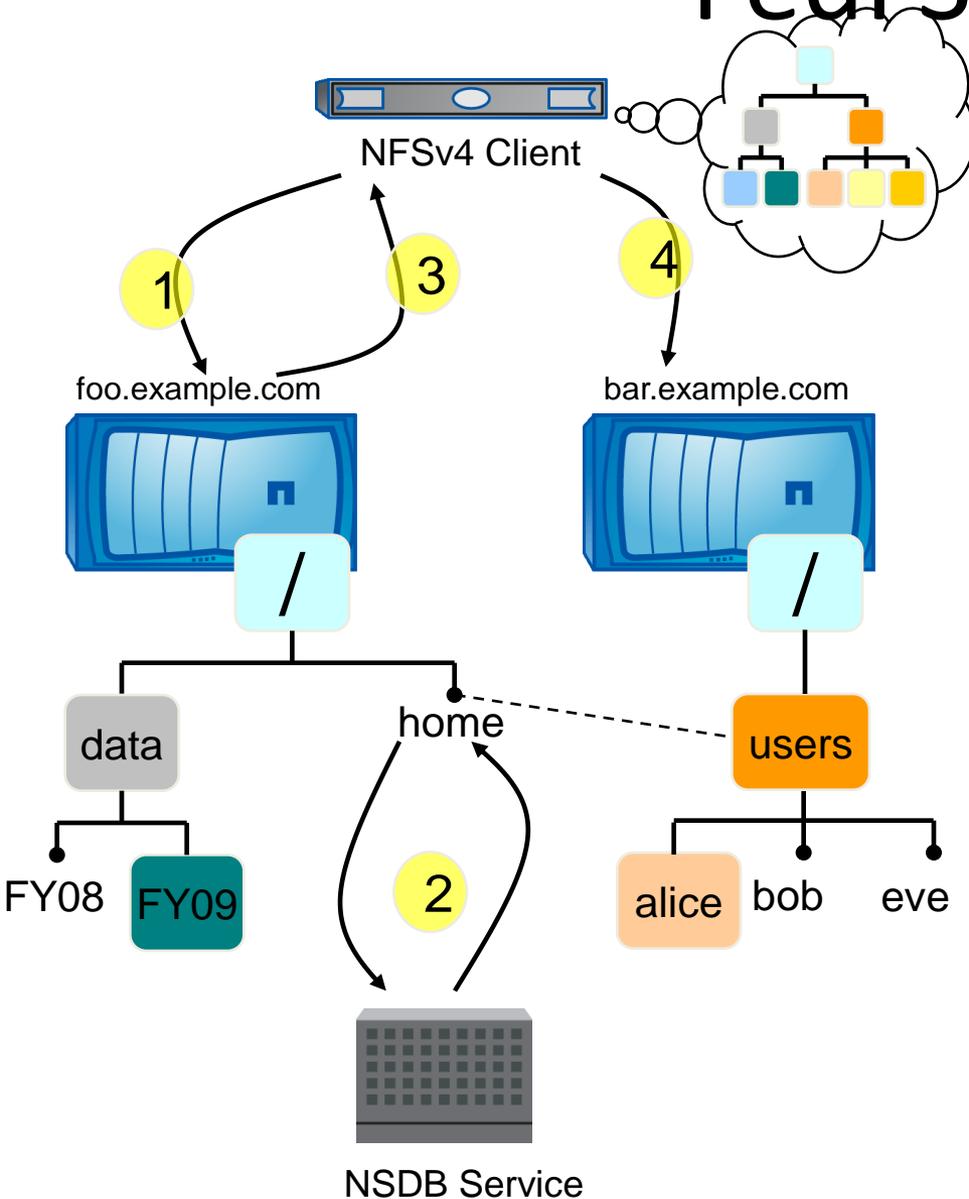
L.A.

New York

Boston

- Behind the scenes, simple management operations allow data mobility for high performance, high reliability, and high availability.

FedFS in Action



The user requests */home/alice*:

1. The client attempts to access */home/alice* on server foo.
2. Server foo discovers that *home* is a namespace junction and determines its location using the FedFS NSDB service.
3. Server foo returns an NFSv4 referral to the client directing it to server bar's */users*.
4. The client accesses */users/alice* on server bar.

Client Support for Referrals

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 since 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
- **OpenSolaris:** referrals supported since build 131 (released January, 2010).

Past Milestones

- Prototype of NSDB protocols demonstrated at the summer WG meeting in Dublin (Summer, 2008).
- Four drafts published as NFSv4 WG documents (Fall, 2008).
- Federated namespace added to the NFSv4 WG charter (Spring, 2009).
- Requirements
 - passed WG last call (May, 2009).
 - approved for publication (October, 2009).
 - published as RFC 5716 (January, 2010).