



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFS Version 4

Spencer Shepler

Senior Software Engineer

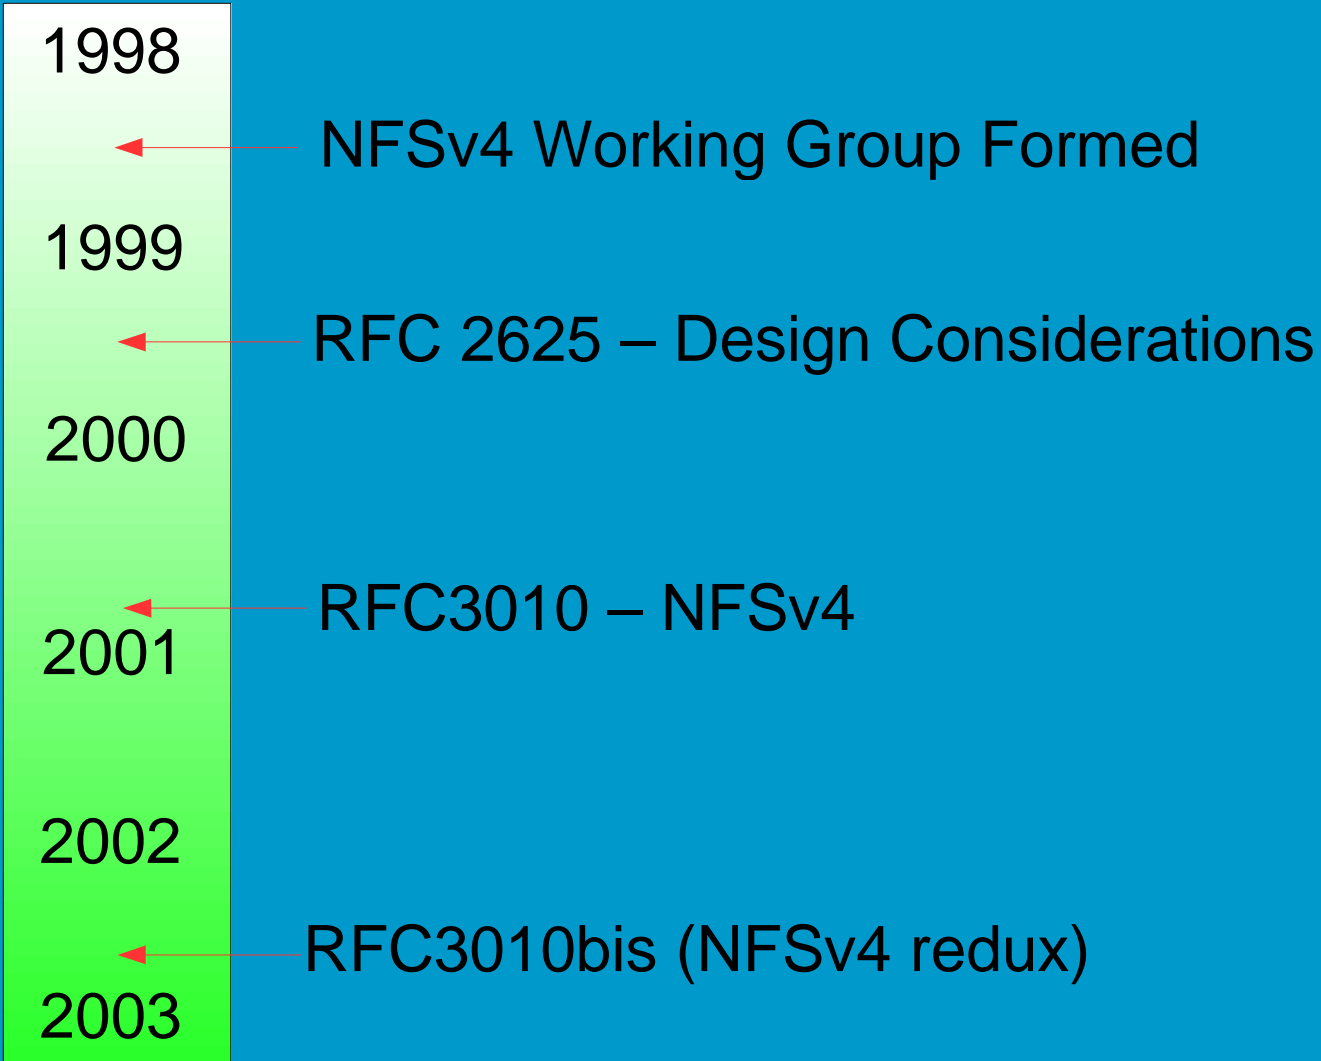
Sun Microsystems, Inc.

[spencer.shepler@sun.com](mailto:spencer.shepler@sun.com)



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFSv4 History



**Bakeoffs/Connectathons**



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# What is NFSv4?



# NFSv4 is...

**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

## NFSv4 (RFC3010bis)

KerberosV5 (RFC1510)  
Kerberos GSSAPI (RFC1964)  
SPKM-3/LIPKEY (RFC2847)

RPC (RFC1831)  
XDR (RFC1832)

RPCSEC\_GSS (RFC2203)  
GSSAPI (RFC2743)

TCP\*



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFSv4 is...

- Secure
  - RPCSEC\_GSS allows for per request authentication/integrity/privacy
  - Private Key (KerberosV5)
  - Public Key (LIPKEY/SPKM3)
  - Negotiation within protocol



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFSv4 is...

## One Protocol

PORTMAP	Port 111	→
MOUNT	Dynamic	→
NFSv2/v3	Port 2049	→
LOCK/NLM	Dynamic	→
STATUS	Dynamic	→
ACL*	Dynamic	→

---

NFSv4	Port 2049 TCP	→
-------	------------------	---



# NFSv4 is...

- Internet Accessible
  - Use of TCP
  - Well defined port (2049) - traverse firewalls
  - Delegation
  - COMPOUND procedure – roundtrips reduced
  - Use of UTF-8
  - User/Group Names (not numbers)

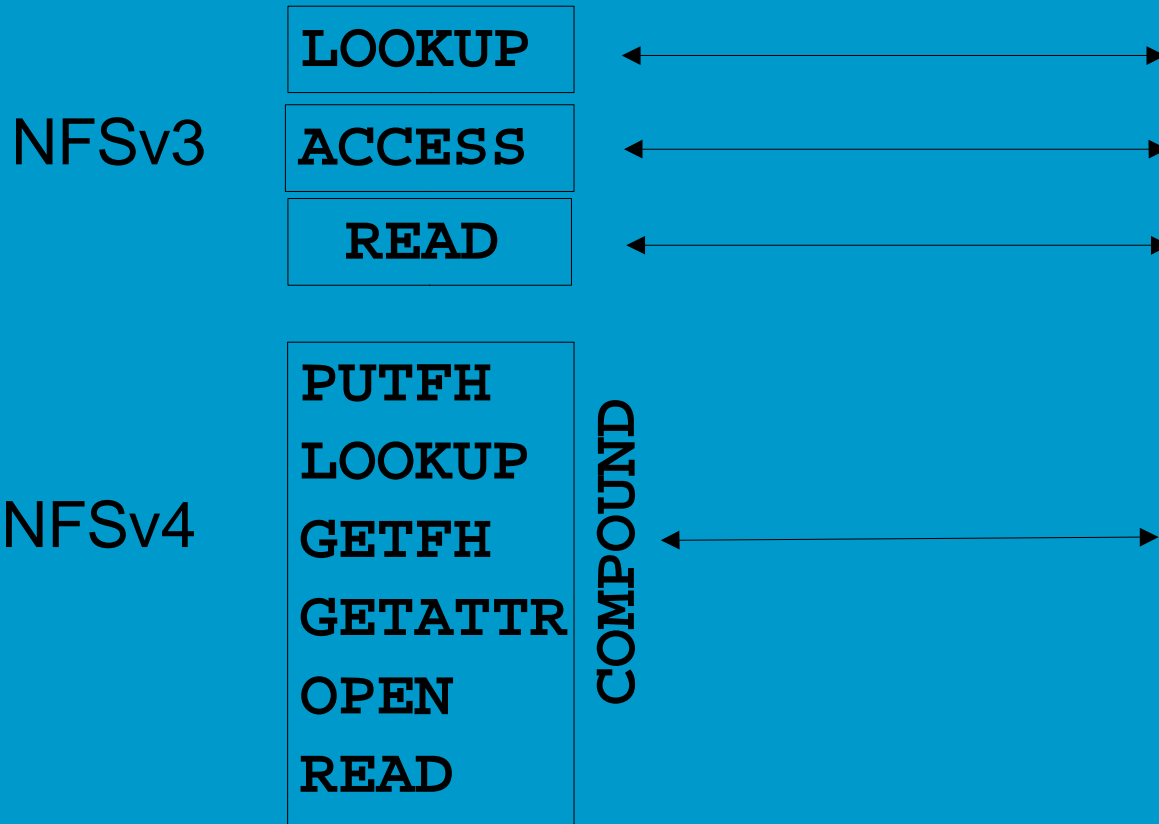
**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**



# NFSv4 is...

## Fewer Round Trips

**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**







# NFSv4 is...

- Filesystem Flexibility
  - Non-Unix specific
  - File attribute model
    - Extensible numbered attributes
    - Named attributes
  - Single server namespace

**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

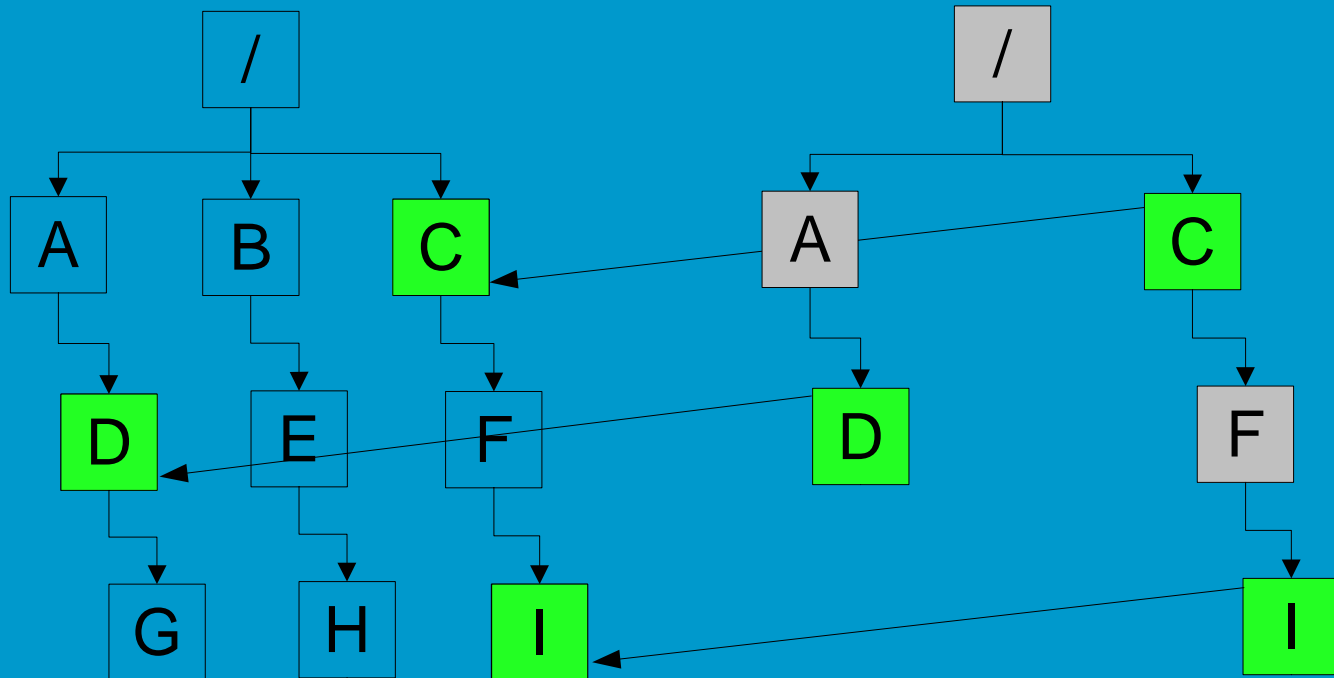


# NFSv4 is...

## Single Namespace

Local FS

Pseudofs

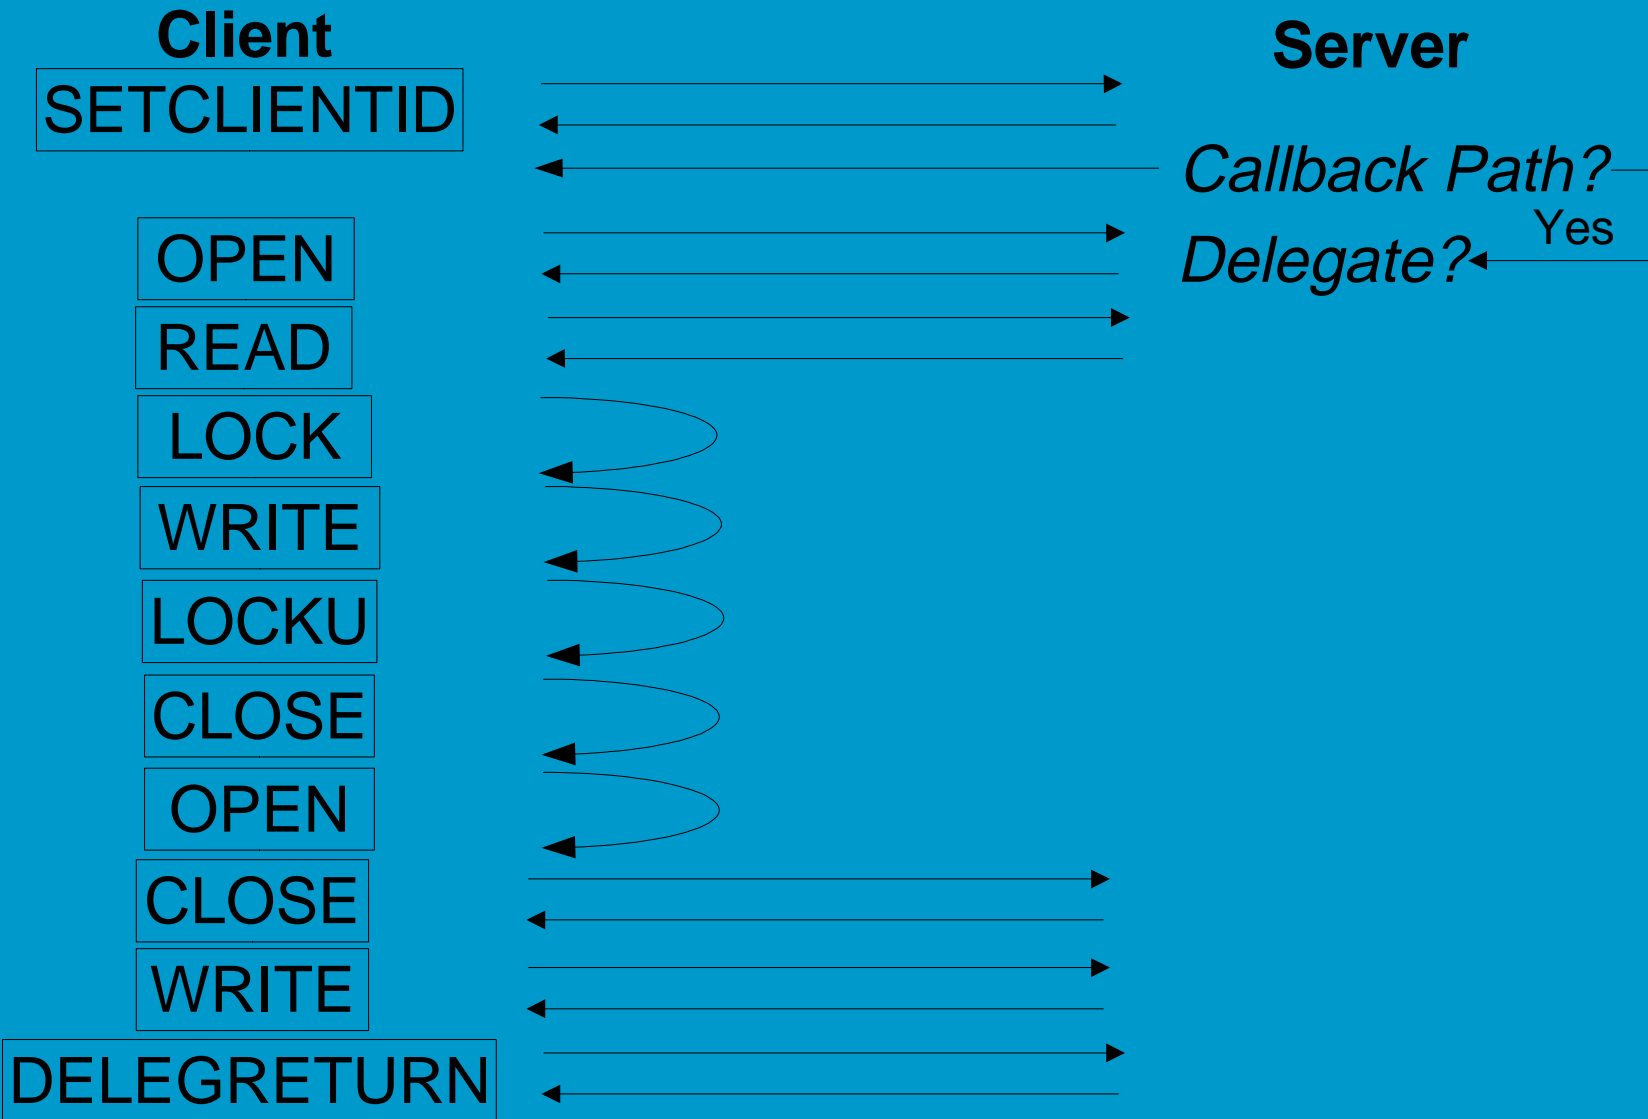


**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**



# NFSv4 is... Delegation

**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

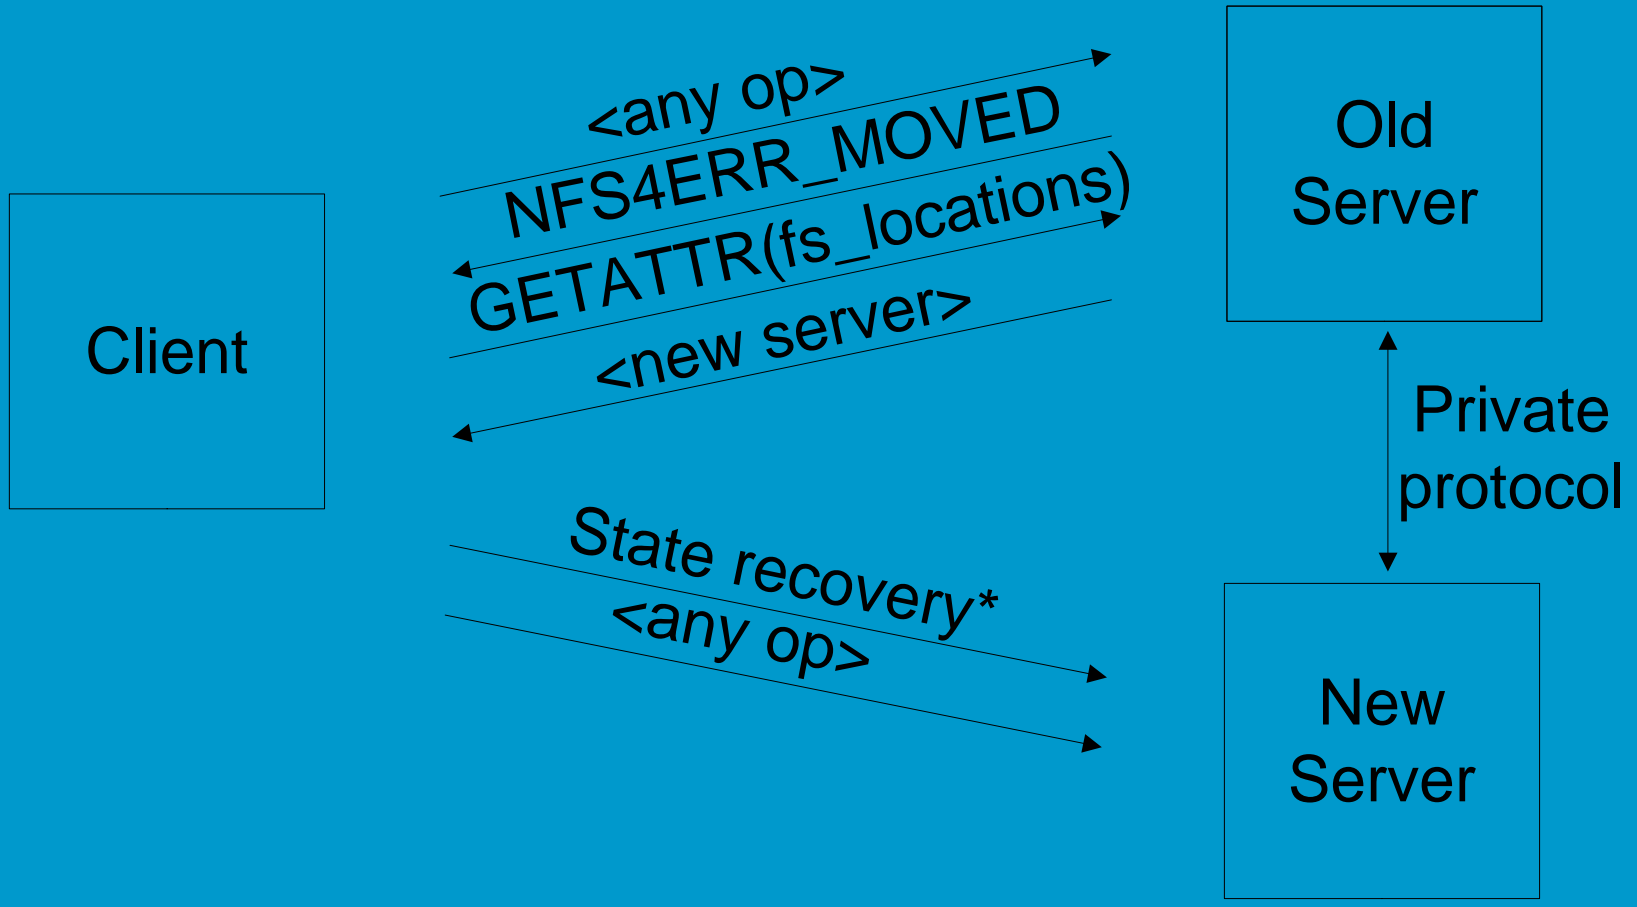




**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFSv4 is...

## Replication/Migration





# NFSv4 is...

- Extensible
  - Minor versions
  - Attribute addition
  - Operation addition

**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# What is rfc3010bis?



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# RFC3010bis

- NFSv4 again
- Reflects implementation work
- About 250 issues resolved
  - Typos -> Changes in state model
- IETF Last Call
  - `draft-ietf-nfsv4-rfc3010bis-04.txt`



# RFC3010bis Updates

- State model changed to allow for Unix
- ACL interpretation clarifications
- Limits set for some opaque data types
- SETCLIENTID clarifications
- ...and many others

**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**





**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

What next?



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFSv4 WG Next Steps

- Replication/Migration
  - In WG charter! Rob Thurlow to discuss
- Investigate Minor Version Items
  - Performance issues
  - Delegations of directories
  - Additional namespace mechanisms
  - Others



**N I C  
F N O  
S D N  
U F  
S E  
T R  
R E  
Y N  
C  
E**

# NFS Version 4

Spencer Shepler

Senior Software Engineer

Sun Microsystems, Inc.

[spencer.shepler@sun.com](mailto:spencer.shepler@sun.com)