



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

NFSv4 Replication and Migration Update

Robert Thurlow

Staff Engineer

Sun Microsystems, Inc.

robert.thurlow@sun.com



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

A reminder of the problem

What was that about again?

- ◆ Replication: ability to create and maintain a copy of an entire filesystem on one or more servers
- ◆ Migration : ability to move an entire filesystem from one server to another
- ◆ ... both with client able to track and adapt to new sources of data



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

Motivations

Why are we doing this?

- ◆ Better availability
- ◆ Better scalability
- ◆ Better resource utilization
- ◆ Disaster recovery
- ◆ Better flexibility
- ◆ Load balancing



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

Motivations

Why are we doing this?

- ◆ Multivendor environments
- ◆ Less vendor lock-in
- ◆ Try before you buy?
- ◆ New: AFS (and DCE/DFS?) EOL



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

Since last year ...

What has been going on?

◆ Repl-mig in NFSv4 WG charter

- Oct 02: strawman protocol
- Nov 02: requirements document / AD review
- Mar 03: prototypes at Connectathon?

◆ Design team formed

- Engineers from Sun, NetApp, Umich/CITI, IBM, Storigen



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

Since last year ...

What has been going on?

- ◆ Skeleton protocol discussed
- ◆ Internet Drafts:
 - June: design considerations
draft-thurlow-nfsv4-repl-mig-design-00.txt
 - October: protocol strawman
draft-thurlow-nfsv4-repl-mig-proto-00.txt
 - November: requirements document
draft-pawlowski-nfsv4-repl-mig-requirements-00.txt
- ◆ Interest picking up!



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

Protocol directions

- ◆ Simple bulk protocol – XDR+TCP
- ◆ Moves V4 named attrs, ACLs
- ◆ Complete or incremental
- ◆ Assumes source knows changes



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

Protocol messages

- OPEN_SESSION_NEW / RESUME,
OPEN_SESSION_CONFIRM / DENY
- SEND_OBJECT_METADATA, SEND_FILE_DATA,
SEND_LOCK_STATE, SEND_SHARE_STATE,
SEND_RENAME, SEND_REMOVE,
SEND_DIR_CONTENTS, SEND_CHECKPOINT,
CLOSE_OBJECT
- CONFIRM_MESSAGE
- ABORT_SESSION, CLOSE_SESSION



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

Open issues

- ◆ Per-message security or TLS?
- ◆ Compression
- ◆ "rsync"-like operation
- ◆ What NFSv4 base protocol changes do we need?
- ◆ How to manage: SNMP or CIM/WBEM?



Related work

- ◆ Namespace work – time to revisit?
 - Will discuss this within IETF WG
- ◆ Possible work with NDMP
 - Could define standard payload

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



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

Get involved!



- ◆ Please help us define requirements and the protocol!
- ◆ NFSv4 WG alias is:
nfsv4-wg@sunroof.eng.sun.com
- ◆ To join: send "subscribe" to:
majordomo@sunroof.eng.sun.com
- ◆ Archives are located at:
<http://playground.sun.com/pub/nfsv4/nfsv4-wg-archive>