



Go further, faster®

Data ONTAP Server-side-Copy

Manjunath Shankararao
(rudra@netapp.com)

-Presented By
Pranoop Erasani
(pranoop@netapp.com)





Outline

- We implemented a Server-side-Copy with additions from Space reservation draft in 2010
 - Prototype based on [draft-lentini-nfsv4-server-side-copy-06](#) & [draft-iyer-nfsv4-space-reservation-ops-01](#)
 - Extended NFSv4.1
 - Synthetic client implementation
 - Packet parser extensions
- Basic prototype
 - Limited args handling
 - Only full file copy
 - Basic RPC Auth – NULL & UNIX
 - Copy only within the same cluster



Server Requirements

- We wanted an architecture that can support NFS, SMB & iSCSI
 - SMB, iSCSI also has Server-side-Copy
 - Support space-reservation semantics
 - Built on top of Cluster architecture
 - Copy Engines to perform copy – same volume, across volumes, same node, across nodes, etc.
 - Support Intra-cluster copy only



Server Solution

- Implementation based on the latest draft versions
- Server frontend
 - Server detects the type of Copy
 - Client creates the destination file before triggering copy
 - Copy Engines supporting Protocols without namespace, cannot create files
 - Need Error value to return *“no support for file creation”*
 - Client reserves space
 - Only basic state handling – no preservation of states across data/interface migrations etc.



Server Solution (contd.)

- Copy Manager
 - Shim Copy Manager for managing copies
 - Support multiple Protocols – different copy models
 - Support multiple Copy Engines
 - Handle Statistics & Flow Control
- Copy Engines
 - Use existing Copy Engines to perform copy for intra-volume, inter-volume, etc.
 - Multiple threads to perform copy in parallel



Server Solution (contd.)

- Need an option to enforce serial copy so that recovery is easier – just resume Copy again from last-block
- Hole awareness
- Support both Synchronous & Asynchronous copies
- Error handling
 - Any errors during Copy simply implies “***Start Copy Afresh***”

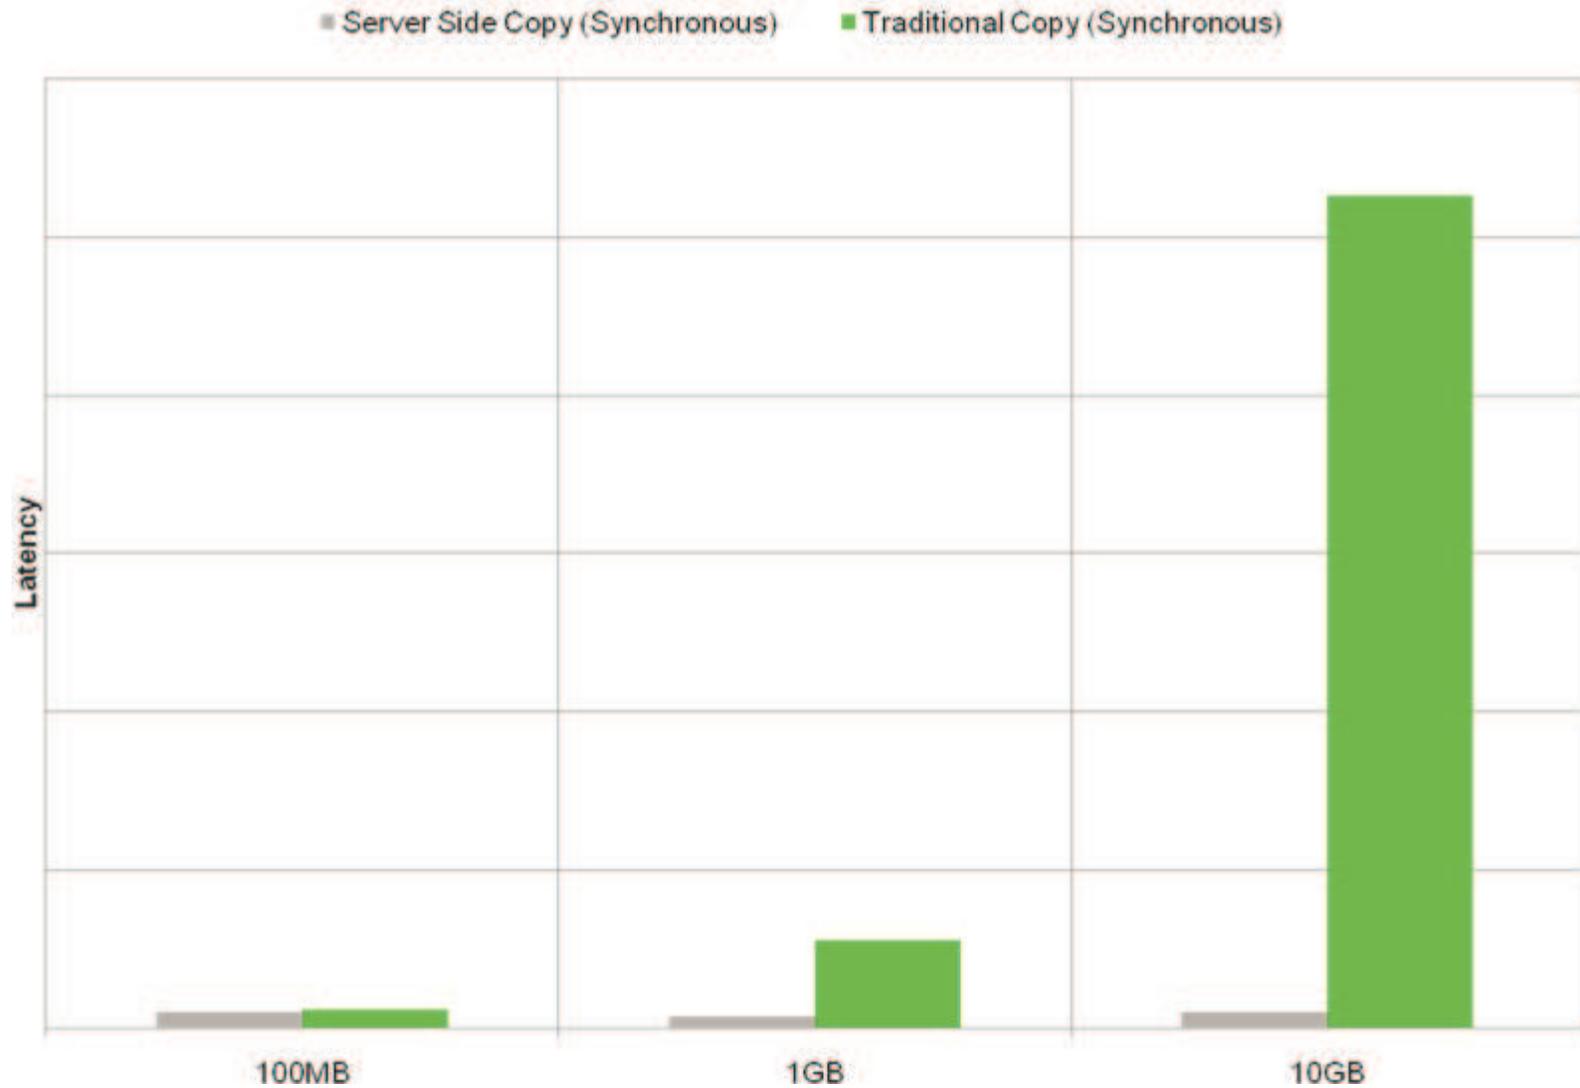


Client Solution

- Synthetic client to prototype Copy
 - Basic functionality only
 - Creates File before copy
 - Doesn't implement locking/delegation semantics, no delegations taken on source file or destination file
- Wireshark Packet Parser
 - Copy & virtualization XDR incorporated into NFSv4.1



Performance (Data ONTAP & Synthetic Client)





Prototype Future

- Implement Sub-file Copy
- Implement Locking/Delegation semantics
- Better Error Handling/Recovery



Draft Extensions

- Better Argument error handling instead of just NFS4ERR_INVALID
 - Error codes for handling **netloc4** types
 - Need Error value to return “*no support for file creation*”
- Option to enforce serial copy for easier recovery
- Consider sending IP address along with hostname
- Copying large/changing files in a heterogeneous server environment needs to be considered
 - Consider sub-file delegations i.e server-side copy is a good use case for sub-file delegations



Thank you!