



Education

# Global Namespaces For Summer

Matthew Geddes, Parascale

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - ◆ Any slide or slides used must be reproduced in their entirety without modification
  - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

**NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

## ➤ Global Namespaces for Summer

- ◆ Distributed filesystems and global namespaces are becoming a more common genus in the enterprise ecology. Using the Microsoft DFS species as an example, we'll explore the generic concepts behind these creatures and dissect the DFS species to help us understand some of the specifics.

# •Overview

- Dataset spread over multiple hosts
- Presented to user as a single filesystem
- Implementations generally aim to be somewhat efficient
- Clients are generally redirected to communicate directly with other hosts in DFS

# • Implementations

## ➤ NFSv3

- ◆ Has no built-in capability
- ◆ Generally a roll-your-own solution involving AutoFS and symbolic links

## ➤ NFSv4 and pNFS

- ◆ Support for *referrals* to refer clients to other hosts mid-filessystem-tree

## ➤ CIFS

- ◆ Called DFS, or *Distributed File System*
- ◆ *Supported by Microsoft CIFS clients from NT 4.0*
- ◆ *Uses referrals to redirect clients*
- ◆ *Can refer to multiple targets*

# • *NFSv3 Solution*


- *Client mounts primary top-level host*
- *Client traverses filesystem and hits symlink to AutoFS mount to another host*
- *Client mounts and accesses second host*
- *Top-level host does not act as proxy*

# • *MS-DFS as an Example*

- *Distributed File System (DFS to its friends)*
- *An implementation of a shared/global namespace for CIFS*
- *Works transparently to the user*
- *Core protocol and management protocol*


# •Core DFS Protocol

➤ *Client traverses remote filesystem*



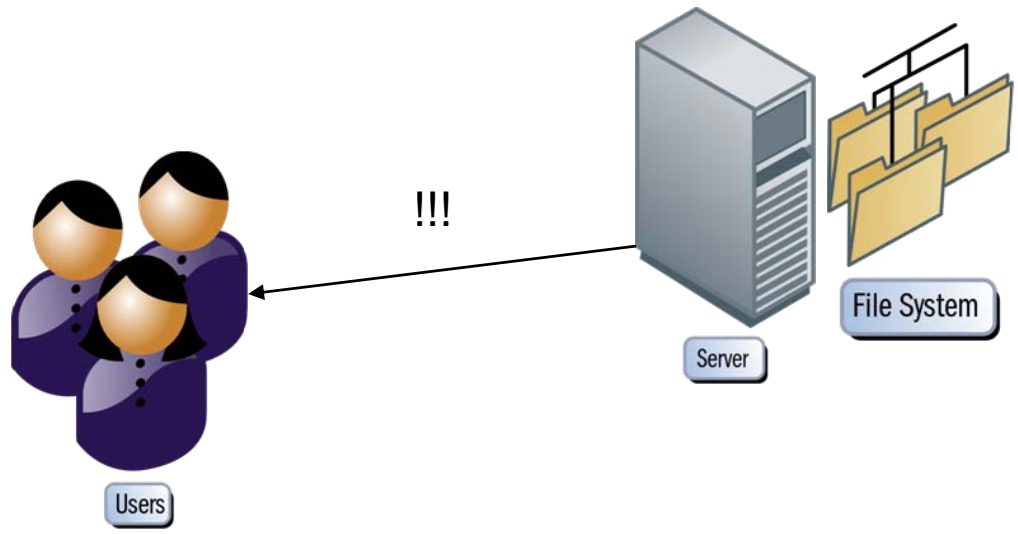
# •Core DFS Protocol

➤ *Client requests an object that is a DFS referral*




# •Core DFS Protocol

➤ Server responds with *NT\_STATUS\_PATH\_NOT\_COVERED*



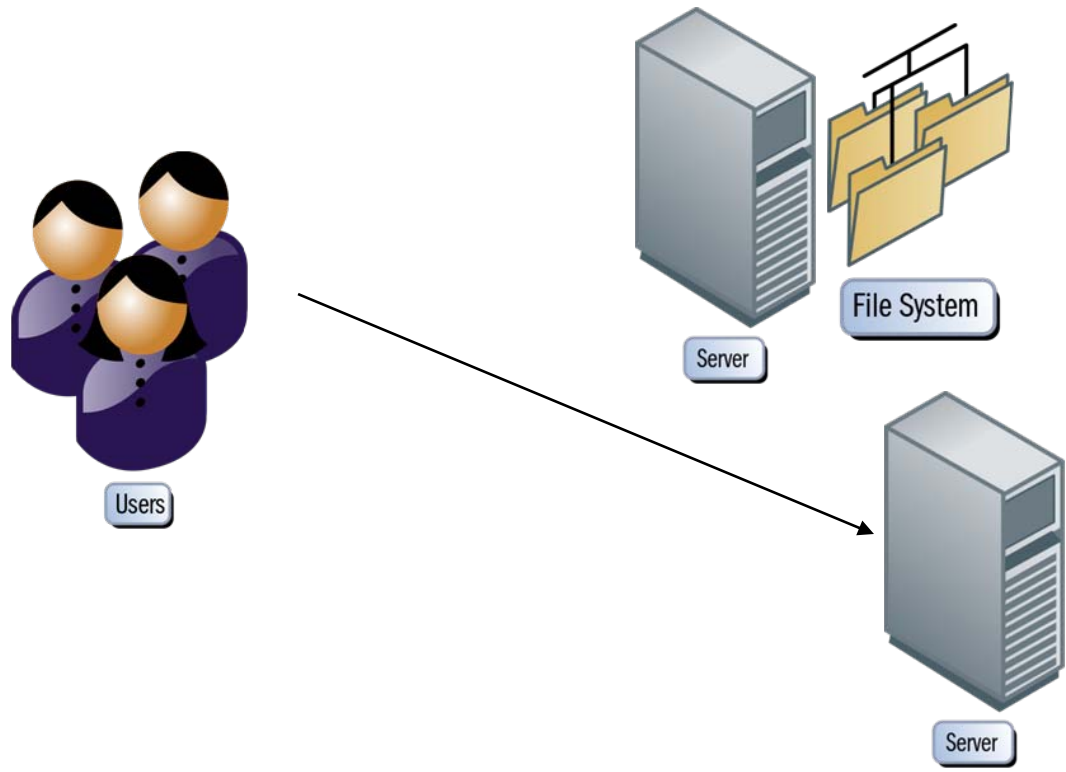
# •Core DFS Protocol

➤ Client requests DFS referral for object



# •Core DFS Protocol

➤ *Client establishes new session with new server*



# •Core DFS Protocol

- Referral points to UNC path (\\SERVER\SHARE\PATH)
- Path can point to anything:
  - ◆ CIFS (SMB and SMB2.0)
  - ◆ NFS
  - ◆ Itself
- Client contacts new server, never the referring server
- Multiple targets are allowed for fault tolerance and crude load balancing
- Clients cache referrals


## • *Core DFS Protocol*

- *The DFS GET\_REFERRAL call has been replaced with an IOCTL in SMB2.0*

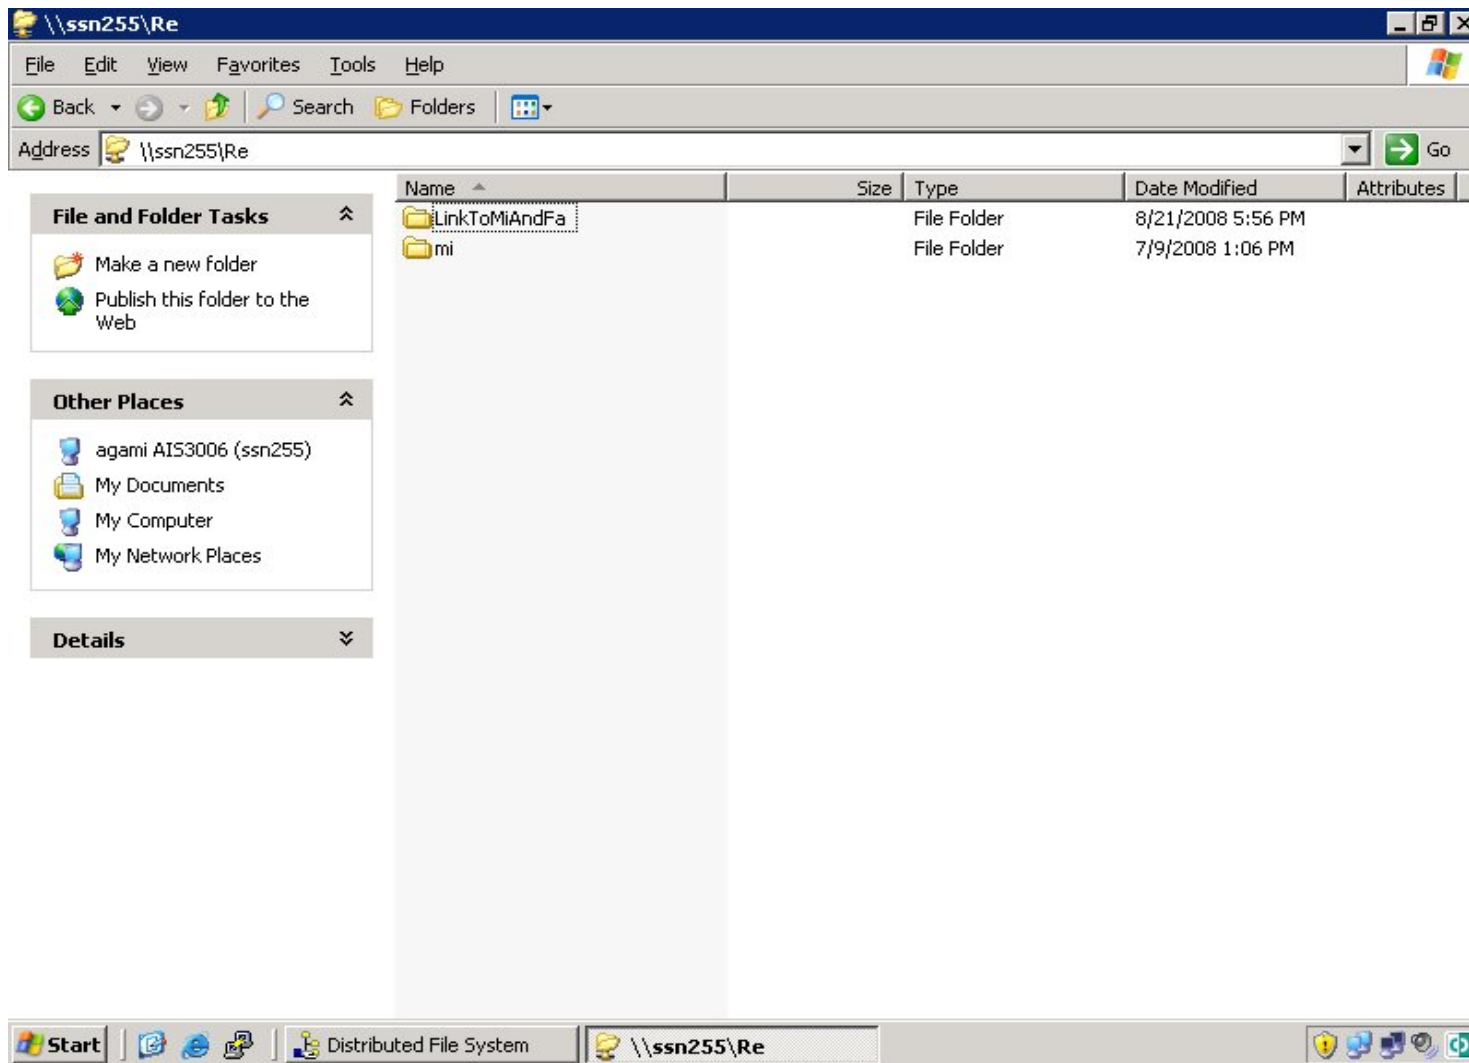
# • *DFS Management Protocol*

- *Allows remote DFS administration*
- *Used by MMC snapins and DFSCMD.EXE*
- *Allows creation and modification of:*
  - ◆ *DFS roots – shares containing DFS redirects*
  - ◆ *DFS links/redirects and their targets*

# •DFS Management Protocol



# •DFS Management Protocol



## • *Active Directory DFS Roots*

- ***It is also possible to have DFS links stored in Active Directory, instead of on the storage devices***
- **Needs client Active Directory support**
- **Allows redundancy at the DFS root level more easily than standalone DFS**

## •Replication

- Despite claims to the contrary, DFS does not handle replication
- It is often deployed with other replication mechanisms, such as FRS or the newer DFS-R

## •DFS Self-Referrals

- It is also possible to create DFS redirects at the share level
- Users connecting to such a DFS root share will immediately be redirected to the target
- This can aid in the migration of data from one system to another before users' share mappings are updated
- This may also aid fault tolerance if multiple targets are specified

## • Practical Uses of DFS

- Cost-effective hierarchical storage across multiple storage devices
- Assisting in transparent data migration


# • Hierarchical Storage With DFS

- Users see a single CIFS share
- Data may in fact be spread over multiple CIFS hosts
- DFS used to tie all data sets together into the single CIFS namespace seen by the users

# • Hierarchical Storage With DFS

- Data might be distributed across a range of hosts like this:
  - ◆ *Infrequently accessed data stored on lower-cost hardware. This hardware may even keep disks spun down when not being accessed*
  - ◆ *More frequently accessed data stored on reasonably-performing, but somewhat more expensive storage*
  - ◆ *Performance critical data might be stored on a host utilising expensive, but high-performance solid state storage, such as flash*

# • Hierarchical Storage With DFS



## •DFS To Aid Migration

- **Data migration takes time and often requires significant downtime**
- **Data may be migrated a little at a time by:**
  - ◆ **Blocking access to a subset of the data by users**
  - ◆ **Migrating the inaccessible data**
  - ◆ **Replacing the source copy with a DFS redirect to the migrated copy**
  - ◆ **Allow access by users**
  - ◆ **Lather, Rinse, Repeat...**
- **Downtime is minimal**

# •DFS To Aid Migration

- Migration may be done in stages using DFS

## •Other DFS Notes

- DFS has been around long enough to be stable in most implementations
- DFS is a useful tool and, combined with some lateral thinking, can solve a number of common problems

- Please send any questions or comments on this presentation to SNIA: [trackfilesystems@snia.org](mailto:trackfilesystems@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

**- SNIA Education Committee**

**Matt Geddes**