



# Using Dell FluidFS Global Namespace

Dell Storage Engineering  
February 2016

## Revisions

Date	Description
February 2016	Initial draft

## Acknowledgements

Author: Mordekhay Shushan

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2016 Dell Inc. All rights reserved. Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.



# Table of contents

- Revisions ..... 2
- Acknowledgements..... 2
- Executive summary ..... 4
- Objective ..... 4
- Audience..... 4
- 1 Dell FluidFS and FS8600 overview ..... 5
- 2 FluidFS namespace ..... 6
  - 2.1 FluidFS filesystem namespace ..... 6
  - 2.2 FluidFS Global Namespace..... 7
    - 2.2.1 FluidFS Global Namespace operation ..... 8
    - 2.2.2 Configuring redirection folders ..... 9
    - 2.2.3 Global Namespace and FluidFS features ..... 9
    - 2.2.4 Global Namespace and FluidFS usage scenarios ..... 10
    - 2.2.5 Global Namespace considerations ..... 18
- 3 Clients and Global Namespace ..... 19
  - 3.1 Connecting to a Global Namespace ..... 19
  - 3.2 Global Namespace and permissions ..... 20
- A Additional resources..... 21
  - A.1 Technical support and resources..... 21
  - A.2 Related documentation ..... 21



## Executive summary

This document describes the Global Namespace feature of Dell FluidFS v5. The topics included provide fundamental knowledge and use cases for FluidFS Global Namespace.

## Objective

This document does not provide step-by-step procedures for recommended configurations, which are described in other referenced documents. It is recommended to use this document in conjunction with the *Dell FluidFS NAS Solutions Administrator's Guide*, available on [Dell.com/support](https://dell.com/support), which includes detailed descriptions of FluidFS features and their configuration procedures.

## Audience

This document is intended for system, network, and storage administrators and integrators who plan to use the FluidFS Global Namespace feature.

It is assumed throughout the document that the reader is familiar with the Dell FluidFS network attached storage (NAS) platform functionality, features, installation, user interface, and operation.



# 1 Dell FluidFS and FS8600 overview

FluidFS is an enterprise-class, fully distributed file system that provides the tools necessary to manage file data in an efficient and simple manner. The underlying software architecture leverages a symmetric clustering model with distributed metadata, native load balancing, advanced caching capabilities, and a rich set of enterprise-class features. FluidFS removes the scalability limitations such as limited volume size associated with traditional file systems, and supports high capacity, performance-intensive workloads by scaling up (adding capacity to the system) and by scaling out (adding nodes, or performance, to the system).

Powered by FluidFS, the Dell FS8600 scale-out NAS solution consists of one to four FS8600 appliances configured as a FluidFS cluster. Each NAS appliance is a rack-mounted 2U chassis that contains two hot-swappable NAS controllers in an active-active configuration. In a NAS appliance, the second NAS controller that has one paired NAS controller is called the peer controller. The solution supports expansion through the addition of NAS appliances to the FluidFS cluster as needed to increase performance.

The FS8600 shares a back-end infrastructure with the Dell SC Series SAN. The SAN network connects the FS8600 to the SC array and carries the block-level traffic. The FS8600 communicates with SC storage using either the internet small computer system interface (iSCSI) or Fibre Channel protocol, depending on the NAS appliance configuration.



## 2 FluidFS namespace

### 2.1 FluidFS filesystem namespace

Understanding FluidFS Global Namespace first requires examining the FluidFS single namespace architecture, as depicted in Figure 1.

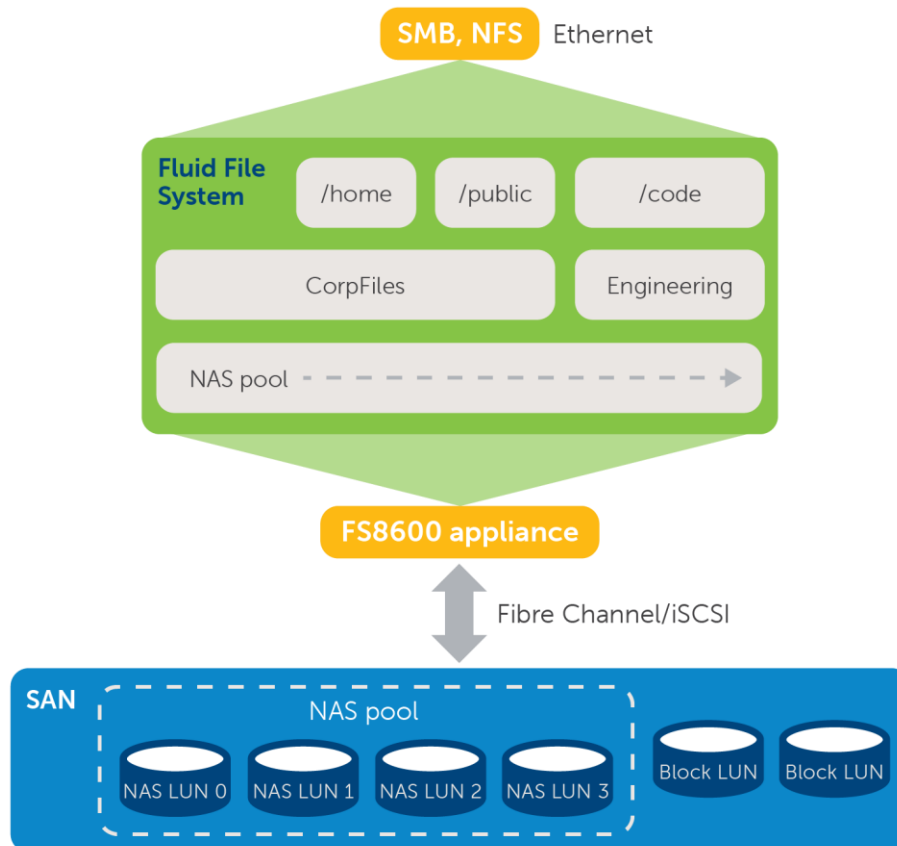


Figure 1 FluidFS single namespace architecture

At its core, FluidFS creates a single distributed file system which spans across all SAN block device LUNs that are presented to the system. This layer, represented as the NAS pool, provides file system services to the NAS volumes defined by the administrator in an optimized manner.

NAS volumes contained within a NAS pool are virtual entities that provide policy-based management of snapshots, replication, quotas, deduplication, backup, and security style as appropriate for different workloads. NAS volumes can be created on demand and can shrink or grow in capacity, non-disruptively, up to the physical limits of the back-end storage.

Each NAS volume can have one or more SMB shares or NFS exports. Access to NFS exports and SMB shares is done using a single or multiple virtual IP addresses.

A single FluidFS filesystem and its namespace can grow up to 4 PB in capacity using multiple SC arrays.

## 2.2 FluidFS Global Namespace

FluidFS Global Namespace provides the ability to connect multiple FluidFS systems to present a single global file namespace under single virtual IPs (VIPs). In addition, FluidFS Global Namespace can connect any other NAS device which supports SMB 2 or NFSv4 protocols.

Using FluidFS Global Namespace, storage administrators can greatly simplify access to large amounts of data spread over one or more NAS devices. Users only need to know the top-level SMB share or NFS mount point, which they can seamlessly traverse to any storage which is part of the global namespace.

Key benefits of FluidFS Global Namespace include the following:

- Simplify user access to tens of petabytes of data by joining multiple devices under a single addressable namespace, removing the need for users to remember the location of different datasets.
- Simplify administrator maintenance of multiple client machines to create consistent views of multiple datasets
- Allow simultaneous access from SMB and NFS to a single namespace spanning multiple discrete NAS devices
- Allow the namespace to span across NAS devices spread across multiple geographic locations
- Simplify data migration to a FluidFS cluster



Figure 2 shows a FluidFS Global Namespace configuration. FluidFS cluster1 is configured with a single NAS volume (named Common), and this volume has local folders and a remote folder that redirect to different FS8600 appliances (FluidFS cluster2 and cluster3) and additional servers (Windows and Linux).

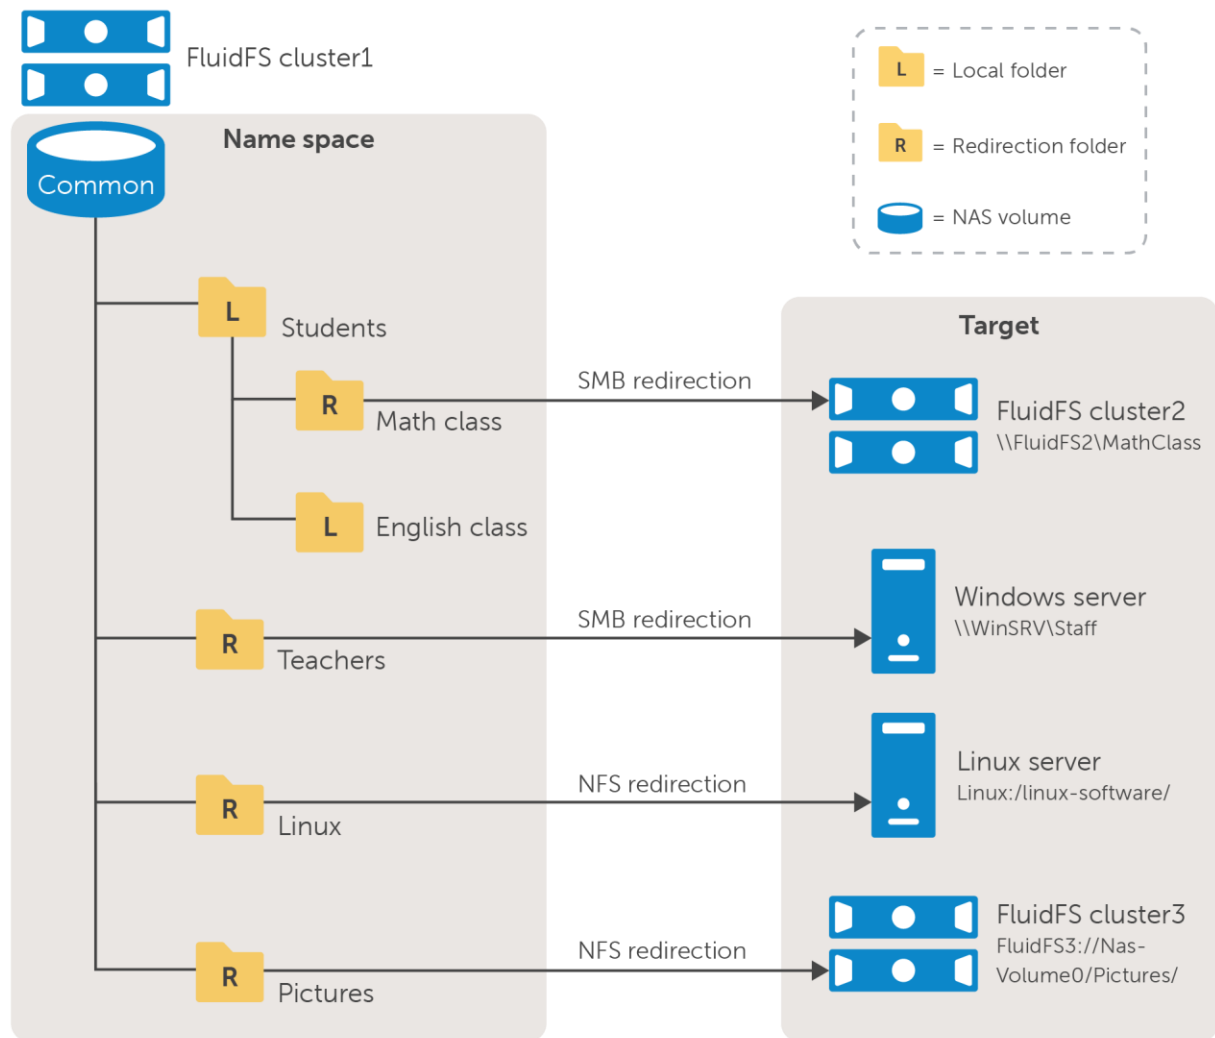


Figure 2 FluidFS Global Namespace design

## 2.2.1 FluidFS Global Namespace operation

A FluidFS Global Namespace uses redirection folders that are configured on a FluidFS NAS volume and redirect to an external SMB share or NFS export.

For the SMB protocol to support FluidFS Global Namespace, the client needs to use SMB 2.x protocol or higher. SMB 2.x uses symbolic links which point to a remote server location. For the NFS protocol to support FluidFS Global Namespace, the client needs to use NFSv4.x or higher.

**Note:** NFSv4 is disabled by default on FluidFS clusters and must be enabled for Global Namespace to be used.



Once a client is redirected by FluidFS to another NAS device, the client OS connects to the remote NAS device directly. All further communication for the contents under the redirection folder happens between the client and the remote NAS device, without interaction with the original FluidFS system.

General	SMB Shares	NFS Exports	Snapshots & Clones	Replications	Quotas	Redirection Folders	Historical Storage Usage
Path		Remote Host		Enable SMB		SMB Share	
/account-folder		dell-remote		True		acct-share	
/acctredir1/acctredir2		172.16.30.51		False			
/market-redirect		dell-remote		True		marketing	
/mkt1/mkt2		Dell-Remote		True		Marketing	
/NFSTest/NFStest2		172.16.30.51		True		acct-share	
/oneUp/twoUp		dell-remote		True		Marketing	
/redirect2/subdir1/subdir2		Dell-Remote		True		Communications_share	

Figure 3 FluidFS Global Namespace redirection folder

## 2.2.2 Configuring redirection folders

Configuration of a redirection folder can be done using Dell Enterprise Manager or the Dell FS8600 CLI. The NAS administrator defines redirection folders by providing the remote SMB shares and remote NFS exports. Once a folder becomes a redirection folder, it can only be deleted by the administrator only.

Permissions to access the redirected data are defined by the destination share, export, destination volume subnet limitations, and remote system security.

## 2.2.3 Global Namespace and FluidFS features

Using FluidFS Global Namespace requires considering the implications on other NAS-volume-based features, such as snapshots, volume cloning, replication, and NDMP backup.

**FluidFS volume cloning:** The cloned NAS volume will contain the redirection folders, and will represent the same global namespace as the original volume. Note that remote data will not be cloned.

**FluidFS replication:** FluidFS replication will replicate the redirection folder information to the destination NAS volume but data from remote folders will not be replicated.

**FluidFS NDMP:** Using FluidFS NDMP backup and restore, the redirection folder information is supported, but not the remote content.

**Note:** NAS Volume Backup, Restore, Replication, or Snapshot policies will not apply on the remote target data but only to the redirection folders (including the redirection data information) which reside inside the local volume data.

## 2.2.4 Global Namespace and FluidFS usage scenarios

This section discusses how FluidFS Global Namespace can be beneficial in several scenarios, including capacity expansion greater than 4 PB of single FluidFS filesystem and data migration from a third-party NAS to FluidFS storage.

The following example (depicted in Figure 4) shows a typical configuration in which several NAS volumes stored in multiple Dell FS8600 FluidFS clusters can be presented to a client as a single namespace through a single share. The following example shows a typical school class that can utilize FluidFS global namespace.

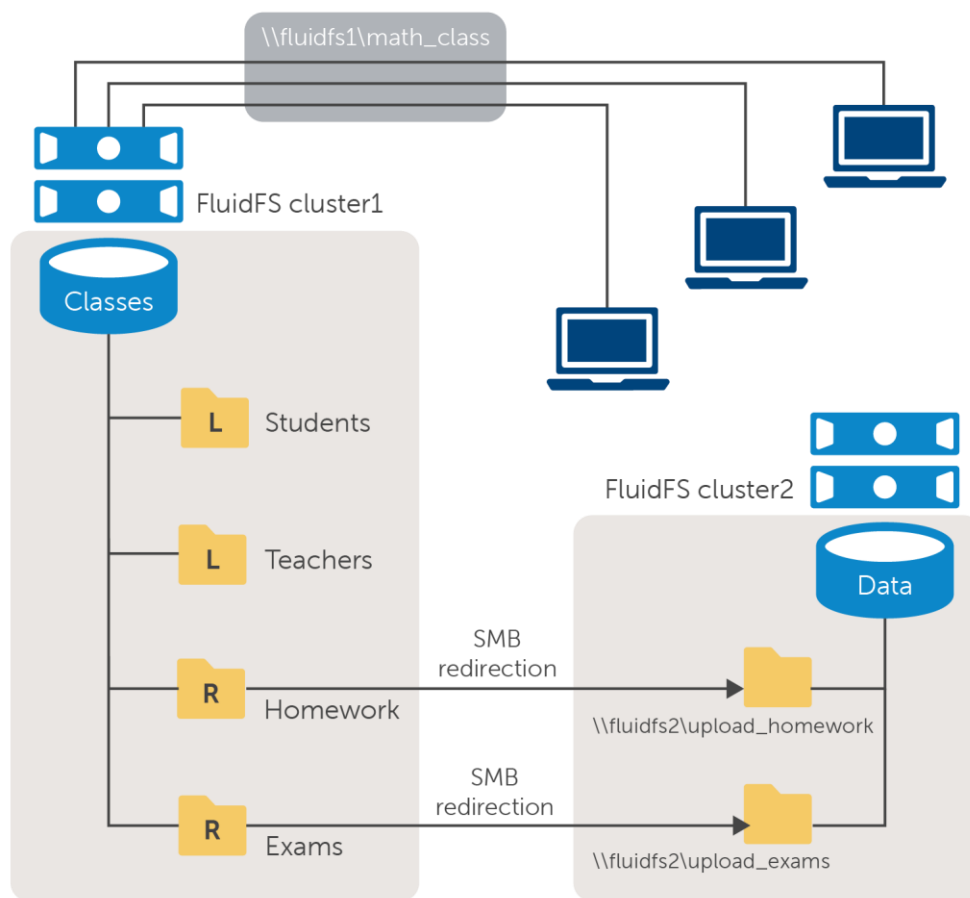


Figure 4 Global Namespace usage example

FluidFS cluster1 is configured with the main Windows share, `math_class`, under the `Classes` volume. The UNC for this share is `\\fluidfs1\math_class`. The NAS volume contains two local directories and two redirected folders pointing to FluidFS cluster2.

The configuration of the directories includes:

- NAS volume:
  - Name: classes
  - Root share: math\_class
- Students: Local directory on NAS volume Classes on FluidFS1
- Teachers: Local directory on FluidFS cluster1
- Homework: Redirected directory from FluidFS cluster2
- Exams: Redirected directory from FluidFS cluster2

The workflow:

1. Students have access to student directory and upload\_homework.
2. Only teachers have access to the entire Classes NAS volume.
3. Students have study materials on FluidFS cluster1 under the Students folder.
4. Students upload homework to \\fluidfs1\math\_class\homework, which redirects to the actual data that reside on FluidFS cluster2 under the redirected folder name, upload\_homework.

In case of a maintenance scenario requiring FluidFS cluster2 to be shut down, the following will happen:

- Access to \\fluidfs1\math\_class remains available.
- Access to the redirected folders from FluidFS cluster2 stops (no access to upload\_homework and upload\_exams).
- Under FluidFS cluster1, Students and Teachers folders are accessible.

This configuration provides some ideas on how FluidFS Global Namespace can be beneficial to any organization that will like to use multiple FS8600 clusters under a single global namespace.

#### 2.2.4.1 Scenario 1: Expanding FluidFS addressable capacity beyond a single cluster limit

The FluidFS NAS pool size can support up to 4 PB, however, FluidFS Global Namespace allows the total namespace size to increase to tens of petabytes by using multiple FS8600 clusters or third-party NAS devices.

In the following scenario, ABC Company is using the Dell FS8600 and the current used space is 3.2 PB. The company expects to reach 4 PB in the coming month and will like to keep using the FS8600. In addition, the company wants to maintain the same namespace, which includes share names and directory structures that been used in the current FS8600.

##### **The solution**

By adding an additional FS8600 and utilizing the global namespace feature, the company can keep the same namespace and expand the namespace capacity beyond 4 PB. This can be done without interruption to users or applications connecting to the system.



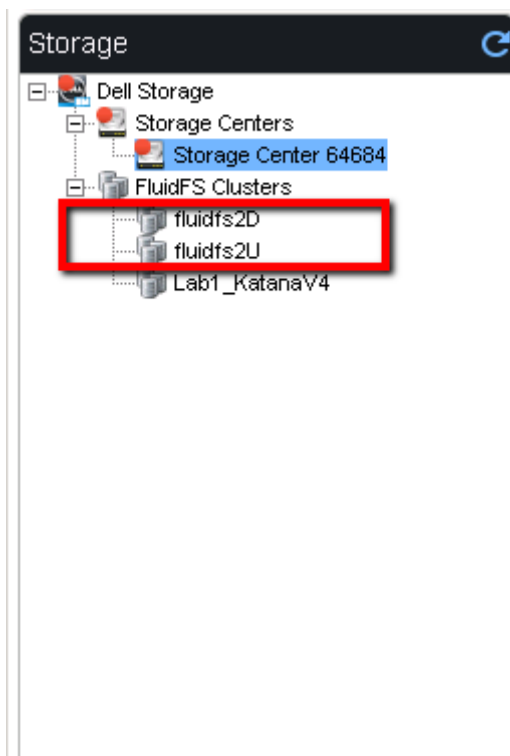
Users will be able to access the system through the same shares used before without knowing that they are working on multiple systems.

### Benefits of the solution

The namespace can grow beyond 4 PB in a way that is seamless for clients. Also, portions of the namespace remain available even if all FluidFS clusters are not up (for example, during software upgrades).

The following example demonstrates this solution:

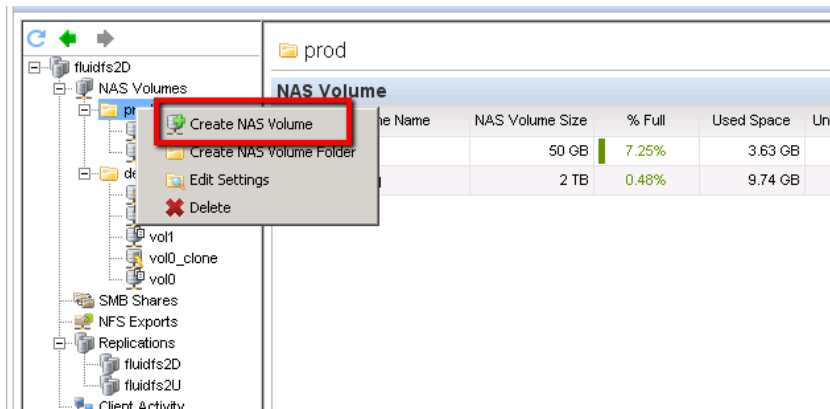
1. Begin by using two FluidFS clusters:
  - Cluster 1:
    - > Local cluster: FluidFS2D
    - > DNS name: \\fluidfs2d.fluidfs.lab.com
  - Cluster 2:
    - > Remote cluster: FluidFS2U
    - > DNS name: \\fluidfs2u.fluidfs.lab.com



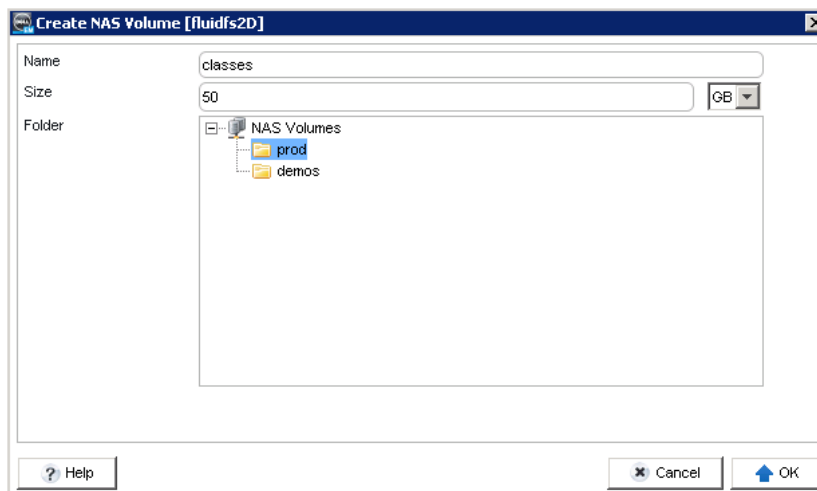
2. Create a volume on each of the clusters:
  - Cluster1 volume name: classes
  - Cluster2 volume name: vol0

To create a NAS volume:

- a. In the **Filesystem** tab, right-click **NAS Volume** or the NAS volume folder and click **Create NAS volume**.



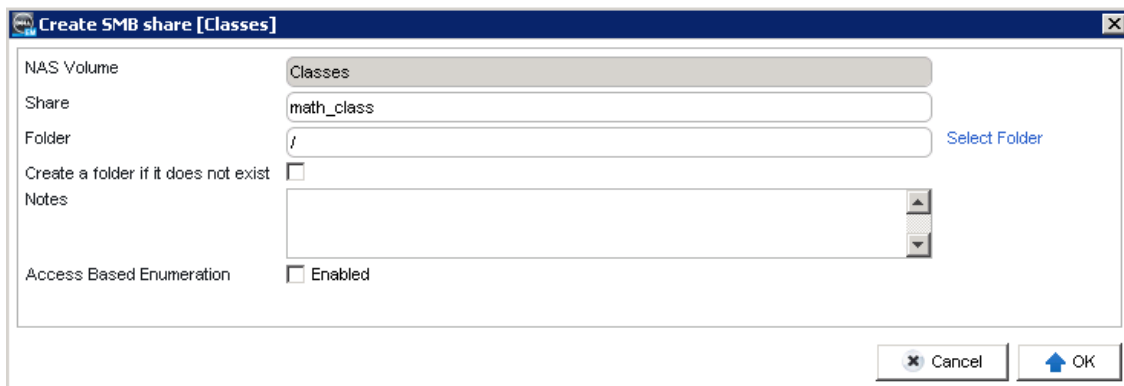
- b. Provide the NAS volume name and size.
- c. Repeat step a and step b on cluster2.



Once the volume is ready, SMB shares are created on both clusters:

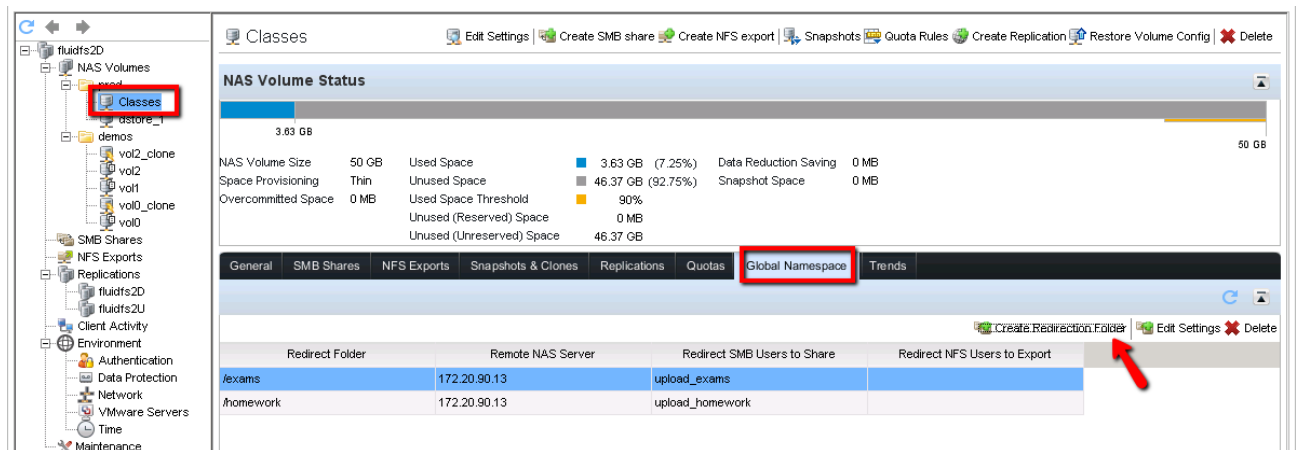
- Cluster1 SMB share: math\_class
- Cluster2 SMB shares: upload\_exams and upload\_homework.

- To create an SMB share:
  - Right-click the volume created in the previous step and click **Create SMB share**.
  - Type the share name and choose a folder (you can also select existing folder).
  - Repeat step a and step b on cluster2 for the two shares.

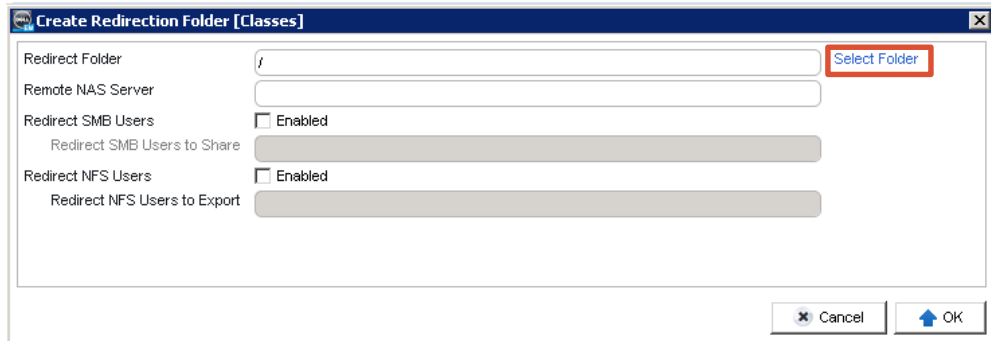


Now that the shares are ready, redirection folders are created on cluster1.

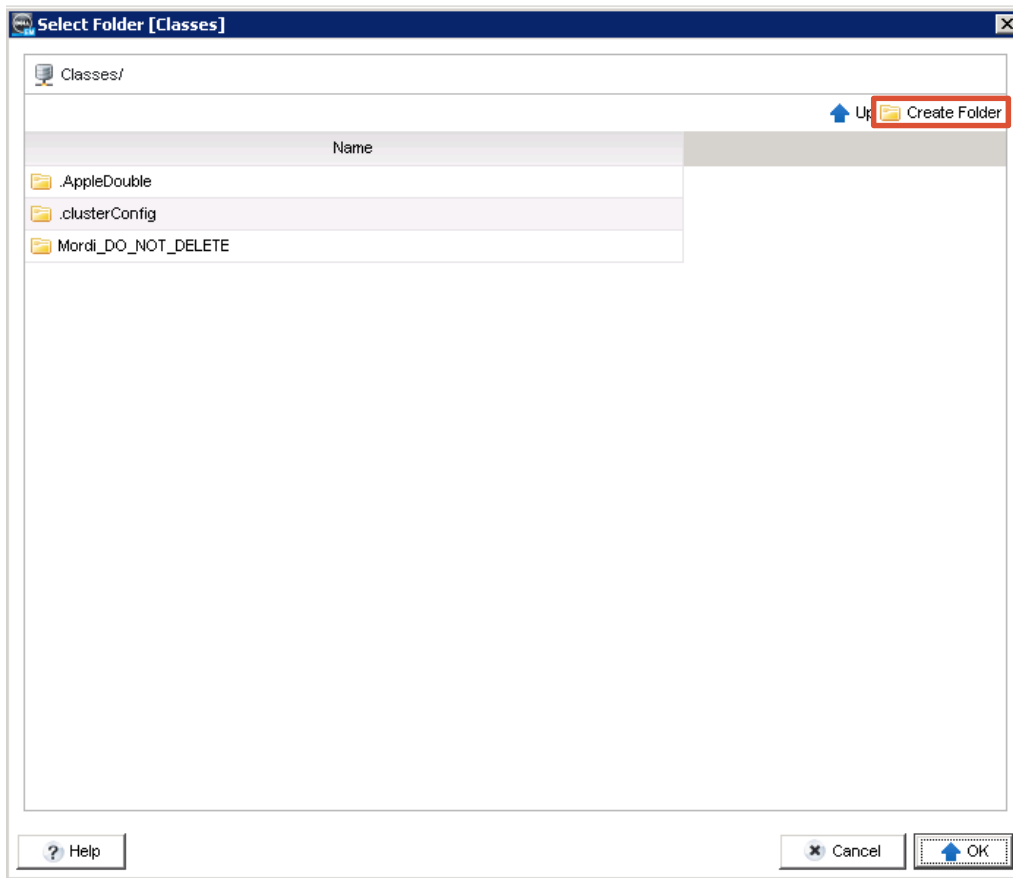
- Create the redirection folder:
  - Click the volume and click the sub tab, **Global Namespace**.
  - Click **Create Redirection Folder**.



- c. Click **Select Folder**.



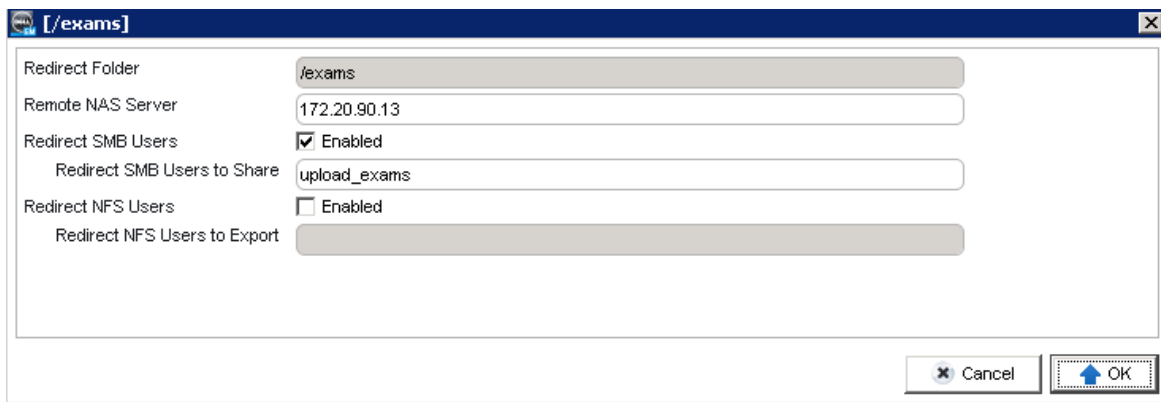
- d. Click **Create Folder** and provide a folder name.





The 'Create Folder [Classes]' dialog box has a title bar with a close button. It contains three input fields: 'NAS Volume' with the value 'Classes', 'Parent Folder' with the value '/', and 'Folder Name' which is empty. At the bottom right, there are 'Cancel' and 'OK' buttons.

5. Assign a share to a specific folder:
  - a. For **Redirect Folder**, choose the redirected folder created in step 4.
  - b. For **Remote NAS Server**, enter the cluster2 VIP.
  - c. Click the check box for **Redirect SMB Users** and provide the SMB share from cluster2.



The '[/exams]' configuration dialog box has a title bar with a close button. It contains several fields: 'Redirect Folder' with the value '/exams', 'Remote NAS Server' with the value '172.20.90.13', 'Redirect SMB Users' with a checked 'Enabled' checkbox and a sub-field 'Redirect SMB Users to Share' containing 'upload\_exams', 'Redirect NFS Users' with an unchecked 'Enabled' checkbox, and 'Redirect NFS Users to Export' which is empty. At the bottom right, there are 'Cancel' and 'OK' buttons.

- d. Repeat step a – step c for the second SMB share.




Figure 5 Capacity expansion solution



### 2.2.4.2 Scenario 2: Simplifying data migration

One of the benefits of FluidFS Global Namespace is to simplify data migration from third-party NAS devices to a FluidFS solution. In the following scenario, Company ABC migrates 9 TB of data from three old Linux servers to a Dell FS8600 while minimizing downtime of data accessibility to users.

#### The solution

The storage administrator creates redirection folders on the FS8600 that point to all three Linux exports. The storage administrator can use software such as Linux rsync to copy all the data to the newly created export on Dell FS8600. Once all data is copied from each Linux server to Dell FluidFS file system, the administrator can redirect the users to the new VIP and run the last sync delta.

#### Benefits of the solution

With this method, the administrator can start the migration in the background and gradually move the data. Once data is migrated, the redirection folder to the other Linux server can be removed. This process allows data to be migrated piece-by-piece in a way that is transparent to the users.

For more information about migration and FluidFS Global Namespace, refer to the [Dell FluidFS Migration Guide](#).

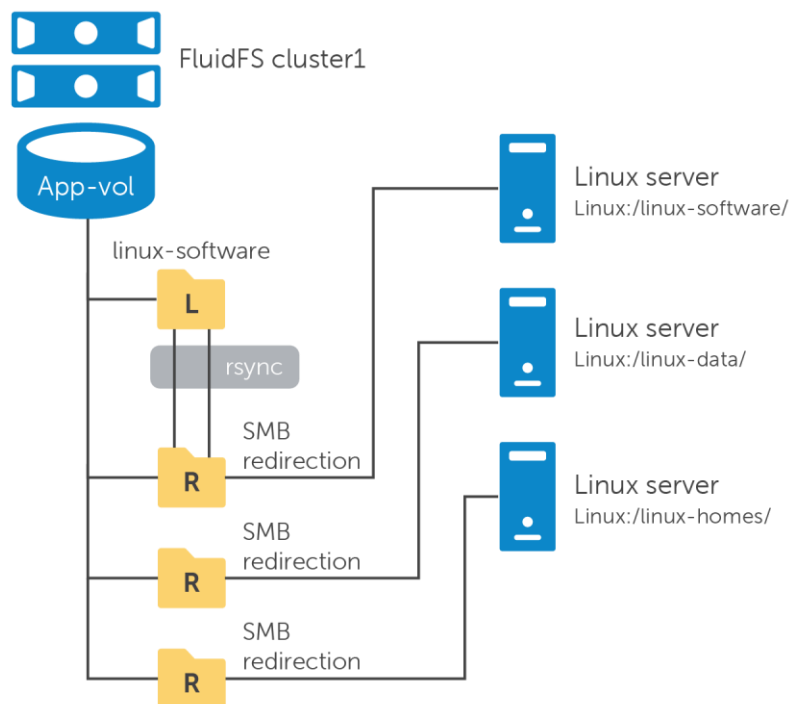


Figure 6 Migration solution example

## 2.2.5 Global Namespace considerations

FluidFS Global Namespace is designed to greatly simplify how users access unstructured data spread across an organization or multiple sites.

Consider the following points before using the FluidFS Global Namespace feature:

- FluidFS Global Namespace requires clients to use SMB 2.x (Windows), and NFSv4 or higher (UNIX or Linux); older clients or servers will not be able to use the feature.
- Administrators must monitor and manage capacity and load balancing between the individual NAS devices that are part of the global namespace.
- FluidFS Global Namespace does not increase the capacity limit of an existing share; while using a single share name under a global namespace configuration, this requires closely monitoring the capacity of other devices that include in the global namespace.
- Performance to a redirected share or export in the global namespace configuration depends on the server or storage hardware that hosts the target. If target shares or exports are hosted on different types of NAS devices, users may experience varying levels of performance depending on which redirected share they are accessing.
- NAS volume backup, restore, replication, or snapshot policies will not apply on the remote target data, but only on the redirection folder objects (including the redirection data information) which reside inside the local volume data. Hence, administrators must ensure the data in the target shares or exports is backed up by existing data-protection mechanisms.



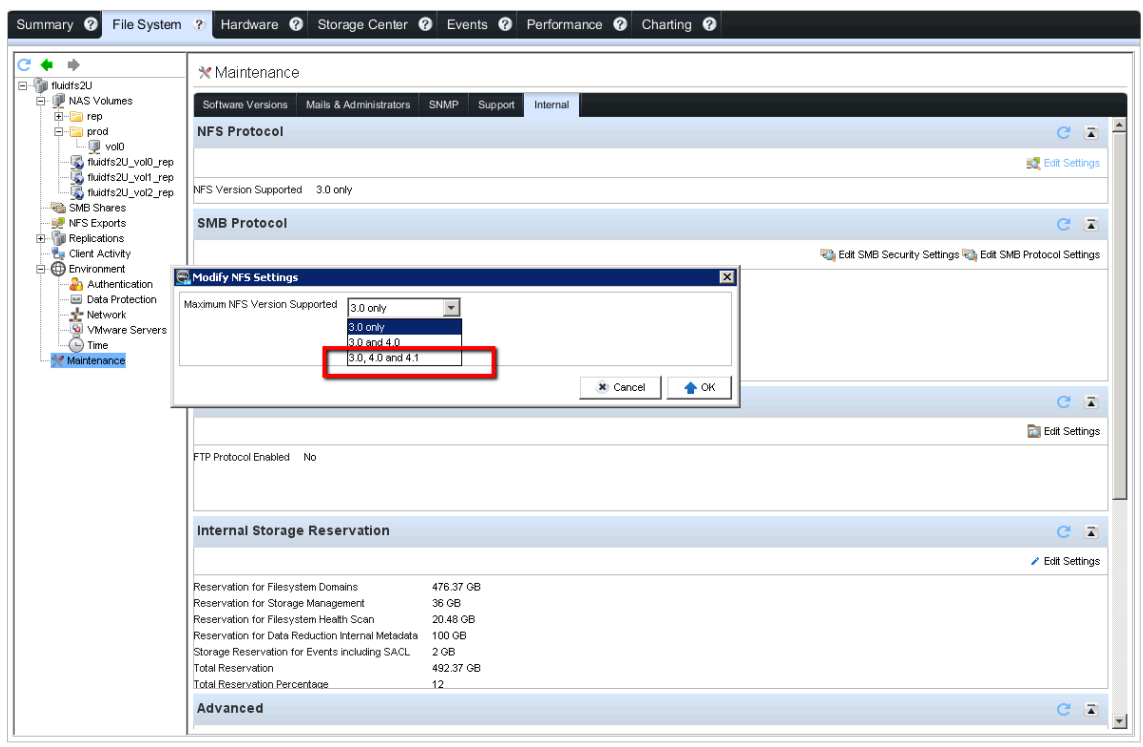
## 3 Clients and Global Namespace

### 3.1 Connecting to a Global Namespace

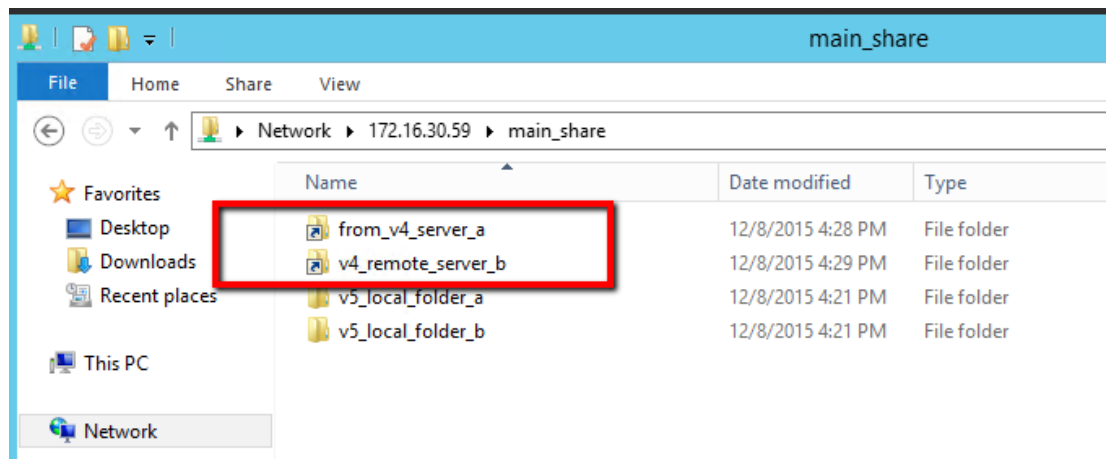
As mentioned in the previous section, client workstations and servers must support SMB 2.x or greater (Windows clients), or NFSv4.x or greater (UNIX or Linux clients).

To enable NFSv4 on a FluidFS cluster:

1. Under **Maintenance**, click the **File System** tab.
2. Click the **Internal** sub tab.
3. In the NFS protocol section, click **Edit Settings**.

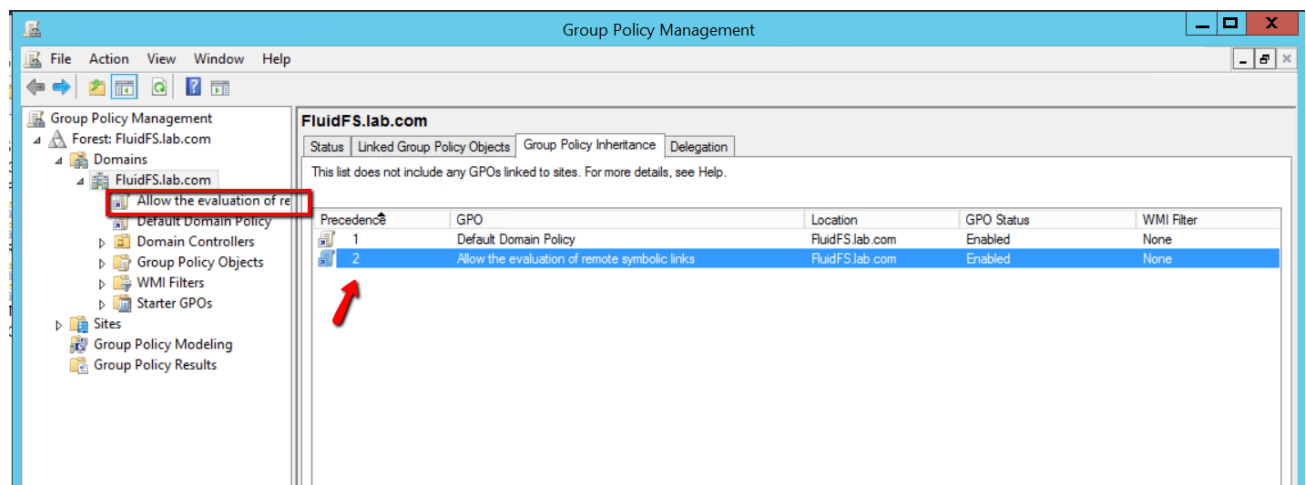


When a Windows client connects to Global Namespace, the client detects a shortcut-type directory which indicates a redirected folder. When a user clicks the redirected directory, the system follows the SMB2.x symbolic link to the target destination that resides on the remote system.



The symbolic link evaluation settings can be controlled using Group Policy.

To enable symbolic link evaluation: navigate to **Computer Configuration > Administrative Templates > System > Filesystem** and click **Allow the evaluation of remote symbolic links**.



## 3.2 Global Namespace and permissions

Since Global Namespace uses multiple servers and permissions can be different from server to server, the servers must be connected to the same Microsoft Active Directory®, and permissions and ACLs must be the same for the user to access multiple servers.

This is especially important when removing permissions for users. For specific files or directories, all permissions on all servers included in the global namespace must be removed.

## A Additional resources

### A.1 Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

For additional support information on specific array models, see the following table.

Dell Storage	Online support	Email	Phone support (US only)
SC Series and Compellent	<a href="https://customer.compellent.com">https://customer.compellent.com</a>	<a href="mailto:support@compellent.com">support@compellent.com</a>	866-EZ-STORE (866-397-8673)
SCv Series	<a href="http://www.dell.com/support">http://www.dell.com/support</a>	<a href="#">Specific to service tag</a>	800-945-3355
PS Series (EqualLogic)	<a href="http://eqlsupport.dell.com">http://eqlsupport.dell.com</a>	<a href="mailto:eglx-customer-service@dell.com">eglx-customer-service@dell.com</a>	800-945-3355

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.

### A.2 Related documentation

See the following referenced or recommended Dell publications:

- [FluidFS V4 \(FS8600\) Networking Best Practices](#)
- [Dell FluidFS Migration Guide](#)
- *Dell FluidFS NAS Solutions Administrator's Guide*, available on [Dell.com/support](https://dell.com/support)

