



Quota and Space Management Best Practices

Dell Compellent FS8600 Network-Attached Storage (NAS)

FluidFS System Engineering
April 2015

Revisions

| Revision | Date | Author | Description |
|----------|------------|-----------------|---------------------------------|
| A | April 2015 | Jonathan Shalev | Initial Release, for FluidFS v4 |

© 2015 Dell Inc. All Rights Reserved. Dell, the Dell logo, and other Dell names and marks are trademarks of Dell Inc. in the US and worldwide. All other trademarks mentioned herein are the property of their respective owners.



Table of contents

| | |
|--|----|
| Revisions..... | 2 |
| 1 Preface..... | 4 |
| 1.1 Audience..... | 4 |
| 1.2 Purpose..... | 4 |
| 1.3 Disclaimer..... | 4 |
| 1.4 Customer Support | 4 |
| 2 Introduction..... | 5 |
| 3 Space Management..... | 6 |
| 3.1 Physical vs. Logical space measurement..... | 6 |
| 3.2 The NAS Pool..... | 7 |
| 3.3 NAS Volumes | 7 |
| 3.4 Quotas..... | 8 |
| 4 Managing Quotas | 9 |
| 4.1 Default User/Group Quota Rules..... | 9 |
| 4.2 User/Group Quota Rules | 9 |
| 4.2.1 Quota Rules for "Entire Group" | 9 |
| 4.2.2 Quota Rules for "Any User in Group"..... | 10 |
| 4.3 Directory Quotas..... | 10 |
| 4.3.1 Guidelines for Creating Directory Quota Rules..... | 10 |
| 4.4 Hard and Soft Quotas..... | 10 |
| 4.5 Relationship between Quota Types..... | 11 |
| 4.6 Quotas in a Multiprotocol Environment | 12 |
| 4.6.1 NAS Volumes with NTFS Security Style..... | 12 |
| 4.6.2 NAS Volumes with UNIX Security Style..... | 12 |
| 4.6.3 NAS Volumes with Mixed Security Style | 13 |
| 4.7 Notes..... | 13 |
| 5 Reporting..... | 14 |
| 5.1 CLI..... | 14 |
| 5.2 Enterprise Manager..... | 15 |
| 6 Examples | 16 |
| 7 Summary | 17 |



1 Preface

1.1 Audience

This document is intended for systems, networking, or storage administrators who are responsible for the day-to-day management responsibilities of a Dell Compellent FS8600 FluidFS NAS solution.

1.2 Purpose

The purpose of this document is to help the storage administrator understand the FS8600 space management and quota mechanisms and choose the appropriate mechanisms for her environment. This document is not intended to be a primer or Dell Compellent FS8600 introductory resource for any of the subject matters involved, and it assumes at least basic knowledge of many of the subjects covered in this document.

This document should be used in conjunction with other Dell Compellent resources, such as the Dell Compellent Storage Center Connectivity Guide, FS8600 Admin Guide and Hardware Manual, Enterprise Manager User Guide, or any other available documentation resources.

1.3 Disclaimer

The information contained within this document is intended to provide best practices and general recommendations only. Actual configurations in customer environments may need to vary due to individual circumstances, budget constraints, service level agreements, applicable industry-specific regulations, or other factors. Configurations should be tested before implementing them in a production environment.

1.4 Customer Support

Dell Compellent provides live support at 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during standard business hours (Monday to Friday, US Central Time).



2 Introduction

FluidFS is an enterprise-class scale-out NAS system that allows customers to easily and efficiently manage file data. FluidFS removes the scaling limitations of traditional file systems. It also supports scale-out performance and scale-up capacity expansion, all within a single NAS pool for simple administration. Because FluidFS optimizes performance and scalability, it is an excellent choice for a wide range of use cases and deployment environments.

The FS8600 NAS appliance provides a number of ways to manage space usage. From top to bottom, these are the NAS pool, NAS volumes, and quotas.

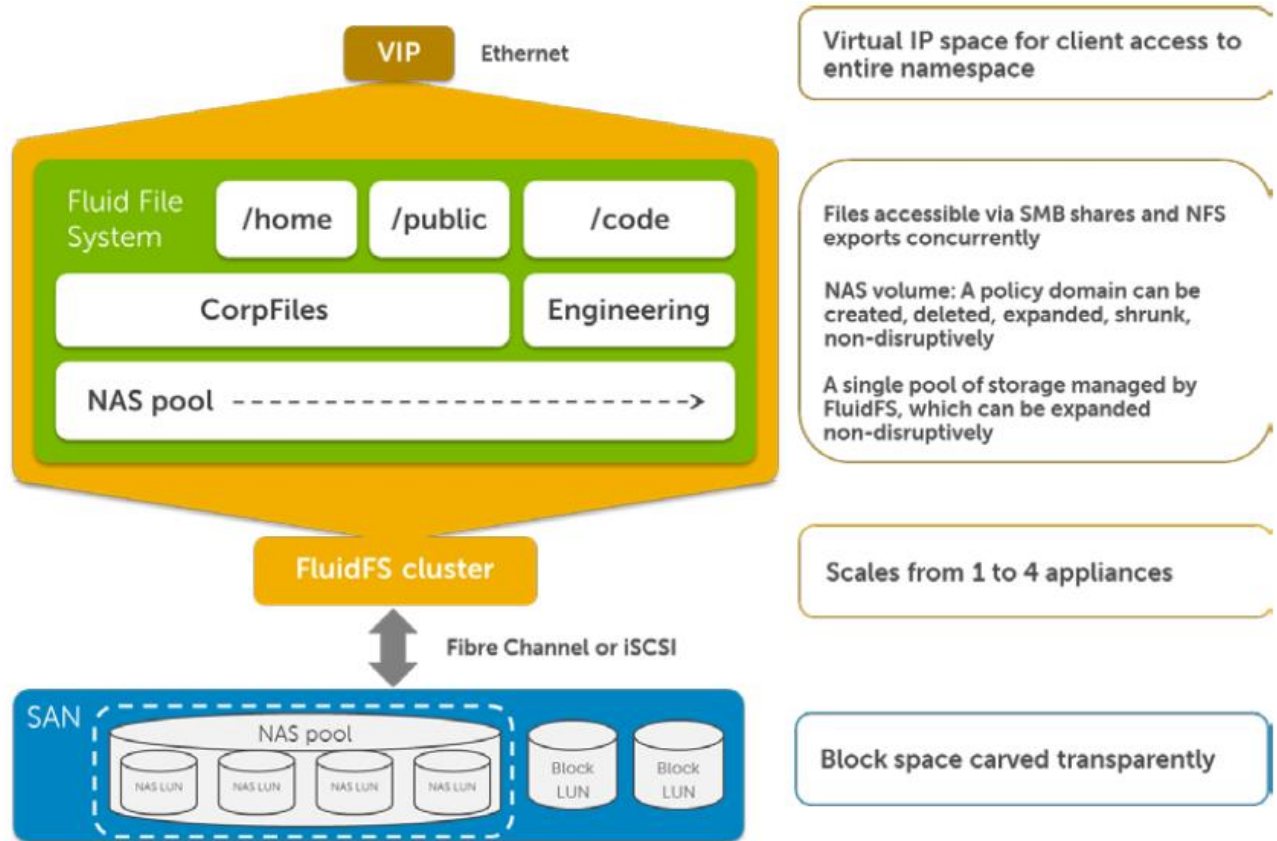
This best practices paper explains the different ways of managing space usage within a NAS volume, with emphasis on quota management.



3 Space Management

Space management on FluidFS is at three levels – the NAS pool, NAS volumes and quotas.

The following diagram shows a NAS pool containing two NAS volumes – CorpFiles and Engineering.



3.1 Physical vs. Logical space measurement

To be able to correctly manage space, it is important to understand the difference between logical and physical space measurement. In this document, **physical size** refers to space used in disk blocks. This includes overhead for internal allocations, snapshots, and metadata (e.g. file inodes and mapping data). Each file and directory has a **logical size** (a number of bytes, which is zero for an empty file), and this size is used for quota calculations. This is the size displayed by "`ls -l`" on UNIX, and in Windows Explorer or "`dir`" in Windows. The logical size of a file is stored in the inode of the file (part of the metadata).

Some examples:

- An empty file has a logical size of zero bytes, even though its physical size is larger because of the inode space and the space for the filename in its parent directory.

- The logical size of a sparse file¹ is usually larger than its physical size, as the “holes” in the file do not need to be stored on the disk.
- A file that was deduplicated and/or compressed uses less space on the disk (physical), while its logical size remains unchanged.
- Snapshot space is counted in the volume used size (this is also physical – snapshots are composed of disk blocks).

3.2 The NAS Pool

The NAS Pool is at the top of the tree. It contains all of the space provided to the NAS cluster by the backend SAN. The size of the NAS pool can be increased – if the SAN has free space – but it cannot be decreased. The size of the NAS pool is **physical**, i.e. it measures the capacity in disk space. This includes overhead for internal allocations, snapshots, and metadata. The total size of data in the NAS pool can be reduced by using Fluid Data Reduction (deduplication and compression). The size of the NAS pool can be defined between 2TB and 4PB (these numbers refer to FluidFS v4 with the appropriate SCOS version on the Compellent SAN Storage Centers). Because all SAN space is thin provisioned, the disk space is not used until actually written.

3.3 NAS Volumes

NAS volumes are flexible file systems defined in the NAS pool, which allow convenient management of various policies. The following definitions are at the NAS volume level (those relating to space management are in **boldface**):

- **Space provisioning type (thin or thick), space reservation, and size**
- **Snapshots and snapshot schedules**
- Replication and replication schedules
- **User and group quotas**
- **Data reduction policy**
- Security style (NTFS, UNIX or mixed)
- SACL policies for generating events
- Restricting access to certain client subnets or VLANs.

NAS volume used space is **physical**, measured in the same way as the NAS pool. The NAS volume size can be increased or decreased simply and flexibly, with no impact on users accessing files on the NAS volume. There is almost no internal reservation (just a small portion that contains backup configuration data which can be used for bare metal recovery in case of disaster recovery). The *volume used space* includes

¹ A sparse file is a file with “holes”. These holes do not contain data and do not need to occupy physical disk space.



snapshot data and metadata (e.g. inodes), and it can often be reduced by enabling Fluid Data Reduction (deduplication and compression).

NAS volumes can be *thin provisioned*, (this is the default for new NAS volumes) allowing for overcommitting space that has not yet been purchased. Space can be *reserved* from the NAS pool for a NAS volume if desired. If a NAS volume or the NAS pool is full, writes to the NAS volume will fail (just like on a full local disk).

3.4 Quotas

Quotas are used to manage space within a NAS volume. Quotas are either user/group quotas for the entire NAS volume (sections **Error! Reference source not found.** and **Error! Reference source not found.**), or directory quotas for a directory subtree (section **Error! Reference source not found.**). Quotas operate at the **logical** level; space usage is calculated based on the file logical size. This means that on the one hand, data in snapshots and metadata is not counted against the quota. On the other hand, sparse files and data reduction do not contribute to reducing the used space for quota enforcement.

4 Managing Quotas

There are a number of types of quotas that can be defined for a NAS volume, described below, and it is important to understand the relationship between them to be able to plan, manage, and troubleshoot quota usage.

User and group quotas are defined for each NAS volume and apply to all data written to the NAS volume.

Directory quotas are defined for a directory within a NAS volume and apply to all data written to the directory tree starting at the specified directory.

User quotas apply to a user (defined locally, in Active Directory, or LDAP/NIS). Group quotas apply to a group (also local, AD, or LDAP/NIS). Directory quotas apply to a directory tree. If more than one of these apply for a certain write operation, the most restrictive of them is used.

We start in Sections 4.1 to 4.3 by describing quota rule types (user, group, and directory). In Section 4.4 we present the two types of quota thresholds (hard and soft). In Section 4.5 we provide more details about the relationship and precedence of the different types of quota rules. Section 4.6 gives some guidelines to follow when using quotas in a multiprotocol environment (access using NFS and SMB). Section 4.7 lists other items to consider when using quotas.

4.1 Default User/Group Quota Rules

There are two types of default quota rules – default user quota rules and default group quota rules.

A default user quota rule defined for a volume applies to any user that does not have a specific user quota rule and her primary group does not have an “any user in group” quota rule.

A default group quota rule defined for a volume applies to any user whose primary group does not have a group quota rule.

4.2 User/Group Quota Rules

User quota rules at the NAS volume level apply to specific users. These can be local users, UNIX users (defined in LDAP or NIS) or Active Directory (AD) users.

Group quota rules at the NAS volume level refer to directory groups. These could be local groups, AD groups (for NTFS style volumes) or NIS/LDAP groups (for UNIX style volumes).

4.2.1 Quota Rules for “Entire Group”

A quota rule for the “Entire Group” limits the space that can be used on a NAS volume by the group.

For the **UNIX security style**, each file or directory has a GID property that indicates which group owns the object.



For the **NTFS security style** what would be limited is the size of files owned by users whose primary group is specified in the rule. In most cases, the default primary group is "Domain Users" for all users in the AD domain. Therefore, it is recommended to set the users primary groups in AD appropriately when using group quotas.

4.2.2 Quota Rules for "Any User in Group"

A quota rule on a NAS volume for "Any User in Group" for a given group is actually a user-level quota rule. It applies to users whose primary group is the given group, that do not have a specific user quota rule for the NAS volume.

4.3 Directory Quotas

Directory quota rules limit logical file usage in a directory tree. This means that you can limit the total space usage of a directory (and its subdirectories) within a NAS volume. This is useful for scenarios where you want to limit space consumption within a NAS volume based on directory structure and not based on user/group membership. Example use cases for directory quota include:

- Chargeback to departments that pay according to their logical space allocation.
- Directories used by a number of users (e.g. drop-boxes, temporary areas)
- Project directories that should be limited when a number of projects share a volume
- Multiple shares on the same volume (which simplifies management of snapshots, replication, etc.) that should each be limited in size

4.3.1 Guidelines for Creating Directory Quota Rules

Directory quota structures should be planned in advance. Directory quotas can only be defined for empty directories. Afterwards, the thresholds can be modified at any time.

If a directory tree is moved (e.g. with the UNIX `mv` command) into a quota directory the file sizes are summed to add to the quota usage. This can take longer than a move between two directories, neither of which is a quota directory.

4.4 Hard and Soft Quotas

Any type of quota that can be configured with FluidFS has Hard and Soft limits.

Hard quotas are limits to the logical size of files. A write that would cause this limit to be exceeded will fail with a quota-exceeded error.

Soft quotas are limits that cause a FluidFS event to be generated if they are exceeded. Users are not notified of soft quota violations.

Any of the quota rule types (user, group, and directory) can include hard and/or soft quotas.

4.5 Relationship between Quota Types

When a write to a file is requested and a quota threshold is exceeded, the request is denied (for a hard quota) or the administrator receives an alert (for a soft quota). There are six flavors of quota rules that potentially apply – default user, specific user, default group, specific group, any user in group, and directory. It is important to understand the relationship between these types of quota rules, as some types override other types.

There are three levels of quota rules.

The first level is the **directory** level – if the write is to a file in or under a **quota directory**, this quota rule is checked.

The second level is the **user** level. If the user has a specific quota rule for the NAS volume, this is checked. If not, and the user's primary group has an "any user in group" quota rule for this NAS volume, this is used. If neither of these is applicable, and there is a default user quota rule for the NAS volume, this is used.

The third level is the **group** level. If the user's primary group has a specific quota rule for the NAS volume, this is used. If it does not have a specific quota rule, and there is a default group quota rule for the NAS volume, this is used.

Each of these three levels is checked separately. At each level, only the first defined quota rule at that level is used. A violation at any of the three levels causes a write failure (hard quota) or an event (soft quota).

The following table shows how the effective quota is calculated:

| Quota calculation for user U whose primary group is G , writing to directory D in NAS volume V | | | | |
|--|---|--|---|---|
| Quota Type | Priority 1 (highest) | Priority 2 | Priority 3 | Quota for type |
| User | <i>Specific user quota for user U in NAS volume V</i> | <i>"any user in group" quota for group G in NAS volume V</i> | <i>Default user quota in NAS volume V</i> | Highest priority entry in the row, or unlimited if none |
| Group | <i>Specific group quota for group G in NAS volume V</i> | <i>Default group quota in NAS volume V</i> | | Highest priority entry in the row, or unlimited if none |
| Directory | <i>Directory quota for directory D</i> | | | Highest priority in the row, or unlimited if none |
| Effective Quota | | | | Most restrictive (minimum) of the quotas for the three types |

4.6 Quotas in a Multiprotocol Environment

Even if a user has a Windows identity and a UNIX identity that are connected by a mapping (manual or automatic), the quota management for each identity is separate.

4.6.1 NAS Volumes with NTFS Security Style

If a NAS volume has NTFS security style, then the quota rules for the volume should be defined for the AD users. If the NAS volume is accessed through an NFS export, a mapping (manual or automatic) is performed to an AD user, and this mapped user is used for user quotas. The primary group of the mapped user is used for group quotas.

4.6.2 NAS Volumes with UNIX Security Style

If the NAS volume has UNIX security style then the quota rules for the volume should be defined for the LDAP or NIS users. If the NAS volume is accessed through an SMB share, a mapping (manual or automatic)



is performed to a UNIX user, and this mapped user is used for user quotas. The primary group of the mapped user is used for group quotas.

4.6.3 NAS Volumes with Mixed Security Style

If the NAS volume has the Mixed security style, there may be cases where user quota rules need to be defined for both types of users. The security style of each file or directory determines which quota rules are used for that file or directory.

4.7 Notes

- The root UNIX user is not limited by directory quotas on directories in NFS exports that trust "everybody" (i.e. no root squash).
- Since the quota checks are performed every 10 seconds, there could be a slight overrun of quota rules until writes are blocked.

5 Reporting

The CLI and the GUI can report the defined quota rules. The current usage can be accessed from the CLI.

5.1 CLI

The following examples show the reporting of quota rules and usage from the CLI. Note that for directory quotas, the rule and usage are displayed in the same query (there is only one usage number per rule).

```
CLI> NAS-volumes quota rules users list
```

| Volume Name | User Domain | User Name | Is Required Alert | SoftLimit | Is Quota Limited | Hard Limit |
|-------------|-------------|-----------|-------------------|-----------|------------------|------------|
| jvol1 | QA | j-user | Yes | 10.00 GB | Yes | 12.00 GB |
| ohadvol | DELL-IDC | ohadf | Yes | 4.00 GB | Yes | 10.00 GB |

```
CLI> NAS-volumes quota rules users-in-groups list
```

| Volume Name | Group Domain | Group Name | Is Required Alert | Soft Limit | Is Quota Limited | Hard Limit |
|-------------|--------------|------------|-------------------|------------|------------------|------------|
| jvol1 | QA | j-group | No | 0.00 MB | Yes | 200.00 GB |

```
CLI> NAS-volumes quota rules groups list
```

| Volume Name | Group Domain | Group Name | Is Required Alert | Soft Limit | Is Quota Limited | Hard Limit |
|-------------|--------------|------------|-------------------|------------|------------------|------------|
| jvol1 | QA | j-group | No | 0.00 MB | Yes | 500.00 GB |

```
CLI> NAS-volumes quota directory list jvol1
```

| Volume Name | Path | Is Required Alert | Soft Limit | Is Quota Limited | Hard Limit | Usage |
|-------------|----------------|-------------------|------------|------------------|------------|-----------|
| jvol1 | /1G | No | 0.00 MB | Yes | 1.00 GB | 0.00 MB |
| jvol1 | /Data/4G | Yes | 3.30 GB | Yes | 40.00 GB | 4.00 GB |
| jvol1 | /Data/qdir500G | No | 0.00 MB | Yes | 500.00 GB | 227.00 MB |

```
CLI> NAS-volumes quota usage users list
```

| Volume Name | User Domain | User Name | Usage |
|-------------|-------------|---------------|-----------|
| GvolClone | cp1sup3 | Administrator | 999.00 MB |
| cfs3 | cp1sup3 | nobody | 1.66 GB |
| jvol1 | cp1sup3 | Administrator | 64.04 GB |



```
CLI> NAS-volumes quota usage groups list
```

| Volume Name | Group Domain | Group Name | Usage |
|-------------|--------------|-------------|-----------|
| GvolClone | cp1sup3 | Local Users | 999.00 MB |
| cfs3 | cp1sup3 | Local Users | 65.10 GB |
| jvol2 | cp1sup3 | Local Users | 227.00 MB |
| ohadvol | QA | ohads | 0.00 MB |

5.2 Enterprise Manager

The following screenshot shows the user rules, group rules, and directory rules for a chosen volume (in this case, volume jvol1). For the Directory quotas, the usage is also displayed. For the user and group quotas, this information is available via the CLI.

The screenshot displays the Dell Enterprise Manager interface for volume jvol1. The left sidebar shows the storage hierarchy, including Storage Centers, FluidFS Clusters, and various NAS Volumes. The main panel shows the 'Quotas' tab for volume jvol1, which has a size of 12.77 GB and 5 TB of total space. The 'User Rules' section shows a default hard quota of 3 GB for the user domain QA. The 'Group Rules' section shows a default hard quota of 500 GB for the group domain QA. The 'Directory Rules' section shows quotas for various paths, including /1G, /Data/4G, and /Data/qdr500G.

| Path | Soft Quota | Hard Quota | Quota Limited | Required Alert | Used Space |
|---------------|------------|------------|---------------|----------------|------------|
| /1G | 0 MB | 1 GB | Yes | No | 0 MB |
| /Data/4G | 3.3 GB | 40 GB | Yes | Yes | 4 GB |
| /Data/qdr500G | 0 MB | 500 GB | Yes | No | 227 MB |



6 Examples

- NAS Volume containing home folders– Creating a NAS volume for home folders and setting the default user quota to a limit makes it easy to control the maximum size of home folders. If exceptions are needed, they can be handled by specific user or group quotas for the NAS volume.
- Home folders on a NAS volume with different quotas per department – In an environment where different departments have different home folder quotas, for example in a scenario where the finance department employees have a 1GB home folder quota and the project management department have a 2GB of home folder quota. For this case, creating a group quota (any user in group) for the finance group and the project management group will limit the home folder's maximum size based on department group membership.
- Limiting space consumption in “scratch” NAS volumes – In a development environment it is common to provide scratch space for programmers. To avoid a scenario where individuals use all available space you can create a default user quota on the scratch volume that will apply to all users writing data in the volume.
- Multiple projects directories in the same volume – Occasionally storage admins are required to allocate storage space for various projects the are being worked on by the same users. Due to other considerations (such as common snapshot and replication policies), the storage admin allocates space for the projects in the same NAS volume. In these scenarios user or group quotas are not applicable as the same user may work on multiple projects in parallel. For such cases the admin can limit projects storage space consumption by creating quota directories per project in the NAS volume.



7 Summary

This best practices paper explains the different ways of managing space usage, with emphasis on quota management.

Space management comes at three levels. NAS Pool (physical), NAS Volume (physical), and quotas (logical).

There are 3 general types of quotas – user, group, and directory. User and group quotas can have defaults for a volume. Group quotas can be for the entire group, or for each user in the group.

