

Storage Quota Management and Provisioning with Dell PowerScale SmartQuotas

December 2024

H10575.21

White Paper

Abstract

Most file systems are a thin layer of organization on top of a block device and cannot efficiently address data at a large scale. This paper focuses on OneFS, a modern file system that meets the unique needs of big data. OneFS includes SmartQuotas, a native data management capability that enables enterprises to reduce storage costs with a simple-to-use, highly scalable, and flexible storage quota and provisioning application for scale-out storage environments.

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2013-2024 Dell Inc. or its subsidiaries. Published in the USA December 2024 H10575.21.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Overview	5
Architecture	6
SmartQuotas licensing	15
Quota notifications	15
Quota accounting.....	19
Quota reports	24
Quota performance tuning	29
SmartQuotas best practices and considerations.....	30
SmartQuotas and OneFS feature integration	32
SmartQuotas use cases	34
Conclusion.....	35

Executive summary

Overview

Unstructured data continues to grow at an astonishing rate, making the need for optimized file-based data storage and its simplified and automated management more crucial than ever.

To help enterprises maximize the long-term value of their critical business data and drive down storage management cost and complexity, Dell PowerScale offers SmartQuotas—a simple, scalable, and flexible quota management and provisioning software application that integrates seamlessly with the OneFS operating system.

Audience and scope

This paper presents best practices for deploying and managing storage quotas in a PowerScale cluster. It also offers configuration and tuning recommendations to help achieve optimal performance for different workloads. This paper does not intend to provide a comprehensive background to the OneFS architecture.

For more information about the OneFS architecture, see the [OneFS Technical Overview](#) white paper.

The target audience for this white paper is anyone configuring and managing storage quotas in a OneFS powered clustered storage environment. It is assumed that the reader has an understanding and working knowledge of the OneFS components, architecture, commands, and features.

For more information about OneFS commands and feature configuration, see the [OneFS Administration Guide](#).

Revisions

Date	Part Number	Description
November 2013	H10575	Initial release for OneFS 7.1
June 2014	H10575.1	Updated for OneFS 7.1.1
November 2014	H10575.2	Updated for OneFS 7.2
June 2015	H10575.3	Updated for OneFS 7.2.1
November 2015	H10575.4	Updated for OneFS 8.0
September 2016	H10575.5	Updated for OneFS 8.0.1
April 2017	H10575.6	Updated for OneFS 8.1
November 2017	H10575.7	Updated for OneFS 8.1.1
February 2019	H10575.8	Updated for OneFS 8.1.3
April 2019	H10575.9	Updated for OneFS 8.2
August 2019	H10575.10	Updated for OneFS 8.2.1
December 2019	H10575.11	Updated for OneFS 8.2.2
June 2020	H10575.12	Updated for OneFS 9.0
September 2020	H10575.13	Updated for OneFS 9.1

Date	Part Number	Description
April 2021	H10575.14	Updated for OneFS 9.2
September 2021	H10575.15	Updated for OneFS 9.3
April 2022	H10575.16	Updated for OneFS 9.4
January 2023	H10575.17	Updated for OneFS 9.5
January 2024	H10575.18	Updated for OneFS 9.7
April 2024	H10575.19	Updated for OneFS 9.8
August 2024	H10575.20	Updated for OneFS 9.9
December 2024	H10575.21	Updated for OneFS 9.10

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Nick Trimbee

Note: For links to other documentation for this topic, see the [PowerScale Info Hub](#).

Overview

PowerScale SmartQuotas enables administrators to understand, predict, control, and limit storage usage across their organization and provision a cluster to best meet their storage needs.




Figure 1. SmartQuotas user quotas

SmartQuotas also facilitates thin provisioning, or the ability to present more storage capacity to applications and users than is physically present (overprovisioning). Thin provisioning allows customers to buy and provision storage as they grow rather than having to make large, speculative purchasing decisions ahead of time.

This paper explores how SmartQuotas can be configured by organizations to establish storage usage quotas and how these quotas can be enforced. The document also illustrates the reporting and notification options available to simplify the management of scale-out NAS environments and data lakes.

Architecture

Architectural overview

The following table shows the primary elements of SmartQuotas, from the file system point of view:

Table 1. SmartQuotas primary elements (file system point of view)

Element	Description
Domain	The files and directories that belong to a quota.
Resource	The quantity that is limited.
Enforcement	Quota limits and what actions are taken when those thresholds are exceeded.

Each quota domain includes usage levels, limits, and configuration options. Most of this information is organized and managed by the file system and stored in the quota database. This database is represented in a B-tree structure, which is known as the quota tree, and provides both scalability and fast random access. Because of its importance, the quota database is protected at the highest level for metadata in OneFS. The quota accounting blocks (QABs) within individual records are protected at the same level as the associated directory.

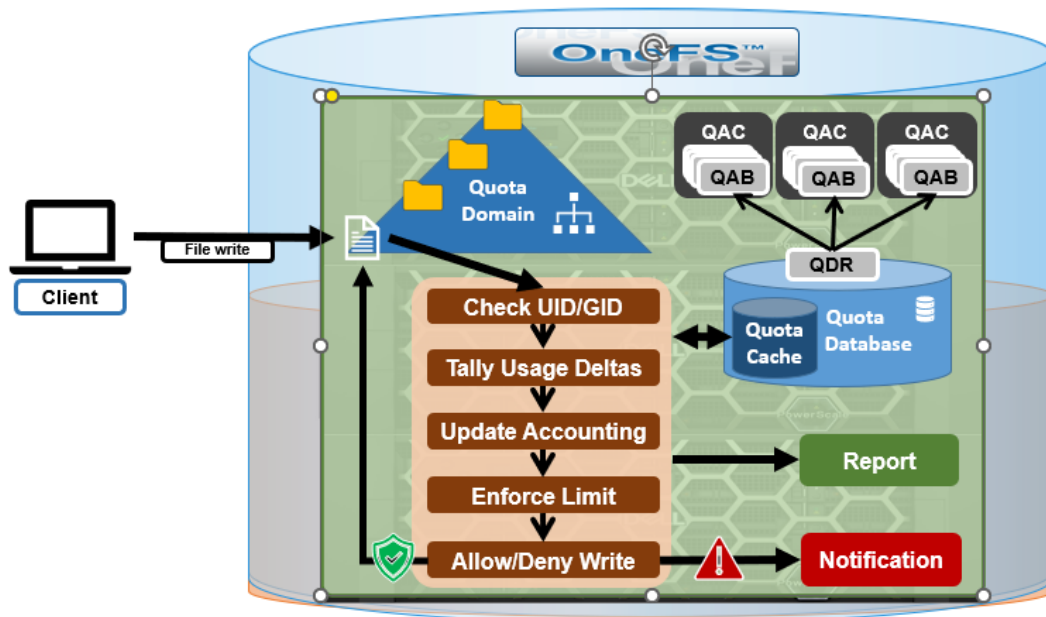


Figure 2. SmartQuotas architectural overview

The following table outlines the principal parts of a quota domain:

Table 2. Quota domain components

Component	Description
Quota domain key	Where the unique identifier for the domain is stored.
Quota domain header (QDH)	Contains various state and configuration information that affects the domain as a whole.
Quota domain enforcements	Manages quota limits, including whether they have been reached or exceeded, notification information, and the quota grace period.
Quota domain account (QDA)	Handles tracking of usage levels for the domain. The QDA tracks physical, logical, and file resource types for each domain.

Quota database

The quota database (QDB) is a data structure that stores the quota domain record (QDR). Resource allocation and governance changes are recorded in the quota operation that is associated with a transaction, totaled and applied persistently to the QDRs.

Quota domain record

The QDR stores all configurations and state information that is associated with a domain. The record includes the following components:

Table 3. Quota domain record components

Component	Description
Configuration	Fields within quota configuration, such as whether the domain is a container. Despite the name, including some state fields like the Ready flag.
Enforcements	A list of quota enforcements that includes the limit, grace period, and notification state. Although the structure is flexible, only three enforcements are allowed and only for a single resource.
Account	The quota account for the domain.

The following figure shows the on-disk format of the QDR. The structure is dynamic, based on the configured enforcements and state of the account, so the on-disk structures look different than the in-memory structures. Subsequent sections of this paper further describe the main components.

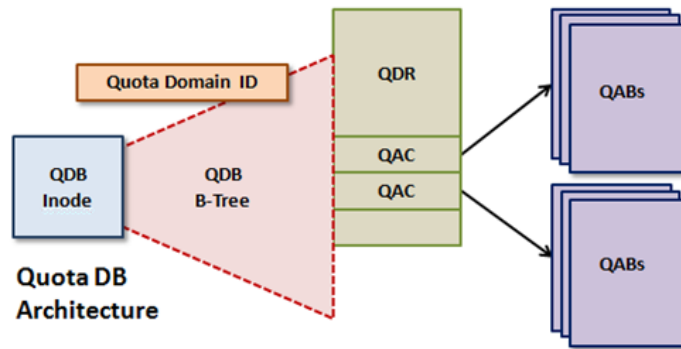


Figure 3. Quota domain architecture

Quota domain locks

Quota domain locks synchronize access to quota domain records in the QDB.

The main challenge for quota domain locks is that the need to lock quota domains exclusively is not known until the accounting is fully determined. It might not be until responses from transaction deltas are received before this requirement is reported to the initiator. To address this challenge, quota domain locks use optimistic restarts.

Constituents and quota account blocks

Quota account blocks (QABs) enable high-performance accounting using transaction deltas. Because the quota usage info is stale when viewed anyway, locking is simplified by using an exclusive domain lock for coherent reads of usage.

Quota accounts also have a mechanism to avoid hot spots on the nodes storing QABs. It is easy to imagine all nodes in a cluster producing accounting changes in the same domain, for example when an ALL domain is configured on a top-level directory. This problem is solved using quota account constituents, which parallelize the accounting to include additional QABs. [Quota performance tuning](#) provides more information.

Quota types

SmartQuotas consists of two types of capacity quotas:

- Accounting quotas

- Enforcement quotas

Accounting quotas monitor and report on the amount of storage that is consumed but do not take any limiting action or intervention. Instead, they are primarily used for auditing, planning, or billing purposes. For example, SmartQuotas accounting quotas can be used to:

- Generate reports to analyze and identify storage usage patterns and trends. The information can then be used to define storage policies for the business.
- Track the amount of disk space used by various users, groups, or departments to bill each entity for only the storage capacity they consume (charge-back).
- Intelligently plan for capacity expansions and future storage needs.

Enforcement quotas, on the other hand, include all the functionality of the accounting option plus the ability to limit disk storage and send notifications. Using enforcement limits, you can logically partition a cluster to control or restrict how much storage that a user, group, or directory can use. For example, you can set capacity limits to ensure that adequate space is always available for key projects and critical applications and to ensure that users of the cluster do not exceed their allotted storage capacity.

Optionally, real-time email quota notifications can be sent to users, group managers, or administrators when they are approaching or have exceeded a quota limit.

SmartQuotas can have one of four enforcement types:

Table 4. SmartQuotas enforcement settings

Enforcement setting	Description
Hard	A limit that cannot be exceeded. If an operation such as a file write causes a quota target to exceed a hard quota, the operation fails, an alert is logged to the cluster and a notification is sent to any specified recipients. Writes resume when the usage falls below the threshold.
Soft	A limit that can be exceeded until a grace period has expired. When a soft quota is exceeded, an alert is logged to the cluster, and a notification is issued to any specified recipients. However, data writes are permitted during the grace period. If the soft threshold is still exceeded when the period expires, writes are blocked and a hard-limit notification is issued to any specified recipients.
Advisory	An informal limit that can be exceeded. When an advisory quota threshold is exceeded, an alert is logged to the cluster, and a notification is issued to any specified recipients. Reaching an advisory quota threshold does not prevent data writes.
None	No enforcement. The quota is accounting only.

All quota enforcement types have both a limit, or threshold, and a grace period. In OneFS 8.2 and later, a soft quota and advisory quota threshold can be specified as a percentage as well as a specific capacity. For example:

```
# isi quota quotas create /ifs/quota directory --percent-advisory-
threshold=80 --percent-soft-threshold=90 --soft-grace=1d --hard-
threshold=100G
```

A hard quota has a zero-time grace period, an advisory quota has an infinite grace period, and a soft quota has a configurable grace period. When a quota limit and grace period have been exceeded, client write operations to anywhere within that quota domain will fail with EDQUOT. Although enforcements are implemented generically in the quota databases, only one resource may be limited per domain, either logical or physical space.

Even when a hard quota limit is reached, operations are not blocked in certain instances. These instances include administrative control through root (UID 0), system maintenance activities, and the ability of a blocked user to free up space.

The following table describes the three SmartQuotas enforcement states:

Table 5. SmartQuotas enforcement states

Enforcement state	Description
Under (U)	If the usage is less than the enforcement threshold, the enforcement is in the state U.
Over (O)	If the usage is greater than the enforcement threshold, the enforcement is in state O.
Expired (E)	If the usage is greater than the soft threshold, and the usage has remained over the enforcement threshold past the grace period expiration, the soft threshold is in state E. If an administrator modifies the soft threshold but not the grace period, and the usage still exceeds the threshold, the enforcement is in state E.

Exceptions to enforcement of quotas include the following scenarios:

- If a domain has only an accounting quota, enforcements for the domain are not applied.
- Any administrator action may push a domain over quota. Examples include changing protection, taking a snapshot, and removing a snapshot. The administrator may write into any domain without obeying enforcements.
- Any system action might push a domain over quota, including repair. OneFS maintenance processes are as powerful as the administrator.

Governance is the mechanism by which SmartQuotas determines which domains apply to a given file or directory. After a sequence of domain configuration changes, a persistent record is required to know where a file had been accounted. As such, quotas use tagging, and the governing domains are recorded in a dynamic attribute of the inode.

A quota domain account tracks the usage and limits of a particular domain. For scalability reasons, the QDA system dynamically breaks up the quota domain's account of the quota into some number of quota account constituents (QACs), each of which tracks a part of

the account. Modifications to the account are distributed at random among these constituents. Each QAC is stored in a set of mirrored quota accounting blocks (QABs). QABs track quota usage and consist of several level counters for different tracked resource types and level limits for advisory, soft, and hard quotas.

The quota domain record stores all configuration and state information that is associated with a domain. The record can be subdivided into three components:

Table 6. Quota domain record components

Component	Description
Configuration	Quota configuration
Enforcement	The grace period, limit, and notification state
Account	A mechanism for space utilization accounting

With SmartQuotas, the primary ways of tracking, enforcing, and reporting resource usage are as follows:

Table 7. SmartQuotas tracking methods

Tracking method	Description
Physical size	Simple to track because it includes all the data and metadata resources used, including the data-protection overhead. The quota system can also track the difference before and after the operation.
File system logical size	Slightly more complex to calculate and track but provides users with a more comprehensible means of understanding their usage.
File accounting	The most straightforward tracking method, because whenever a file is added to a domain, the file count is incremented.
Application logical size	The method that reports total logical data stored across different tiers, including CloudPools, to account for the exact file sizes. This method allows users to view quotas and free space as an application would view them, in terms of how much capacity is available to store logical data, regardless of data reduction or tiering technology.

Before OneFS 8.2, SmartQuotas size accounting metrics typically used a count of the number of 8 KB blocks that are required to store file data on a cluster. Accounting based on block count can result in challenges, such as small file over-reporting. For example, a 4 KB file would be logically accounted for as 8 KB. Similarly, block-based quota accounting only extends to on-premises capacity consumption. This means that a 100 MB file that is stored within a CloudPools tier is only accounted for as an 8 KB SmartLink stub file, rather than its actual size.

To directly address this issue, in OneFS 8.2 and later, application logical quotas provide a new quota accounting metric. Application logical size accounts for, reports, and enforces the actual space that is consumed and available for storage, independent of whether files

are cloud-tiered, sparse, deduplicated, or compressed. Application logical quotas can be configured from the CLI with the following syntax:

```
# isi quota quotas create <dir> directory --thresholds-  
on=applogicalsize
```

Legacy quotas created on earlier OneFS versions can easily be converted to use application logical size after upgrading to OneFS 8.2.

For logical space accounting, some inode attributes, such as ACLs and symbolic links, are included in the resource count. This uses the same data that is displayed in the **logical size** field by the `isi get -DD <file>` CLI command.

QuotaScan job

The QuotaScan job updates quota accounting for domains created on an existing directory path. Although the job is typically run without any intervention, the administrator has the option of manually controlling it if necessary or preferred. By default, QuotaScan runs with a “low” impact policy and a low-priority value of 6.

If quotas are created on empty directories, governance will instantaneously propagate from parent to child incrementally. If the directory is not empty, the QuotaScan job is used to update the governance.

A domain that is created on a nonempty directory will not be marked as ready, which triggers a QuotaScan job to be started. QuotaScan is run by the OneFS job engine and is subject to the general scheduling and prioritization of jobs. The QuotaScan performs a tree walk to traverse the directory tree under the domain root.

The QuotaScan job is the cluster maintenance process responsible for scanning the cluster, and performing accounting activities to bring the determined governance to each inode. In essence, the job is a distributed tree walk that is performed based on the state of the domain.

For more information about the job engine, see the [OneFS Job Engine white paper](#).

Quota domains

SmartQuotas is based on the concept of domains—the linchpins of quota accounting. Since OneFS is a single file system, it relies on domains for defining the scope of a quota in place of the typical volume boundaries found in most storage systems. As such, a domain defines which files belong to a quota, accounts for each resource type in that set, and defines the top-level directory configuration point.

The following table lists the primary SmartQuotas resource types:

Table 8. SmartQuotas resource types

Resource type	Description
Directory	A specific directory and all its subdirectories
User	A specific user
Group	All members of a specific group

A domain defined as `<name>@<folder>` is the set of files under “folder,” owned by “name,” which could be either a user or a group. The files accounted include all files reachable from the given path, without traversing any soft links. The owner name can be ALL; `/ifs`, the OneFS root directory, is also an effective ALL for “folder.”

With SmartQuotas, you can quickly and easily create traditional domain types by using ALL. Here are a few examples of domain types:

- All files belonging to user Jane: `user:Jane@ifs`
- All files under `/ifs/home` belonging to any user: `ALL@ifs/home`
- All files under `/ifs/home` belonging to user Jane: `user:Jane@ifs/home`

Domains cannot be created on anything but directories. More specifically, domains are associated with the actual directories themselves, not directory paths. For example, if the domain is `ALL@ifs/home/data`, but `/ifs/home/data` gets renamed to `/ifs/home/files`, the domain stays with the directory.

Domains can also be nested and may overlap. For example, a hard quota is set on `/ifs/data/marketing` for 5 TB, and 1 TB soft quotas are then placed on individual users in the marketing department. These quotas ensure that the marketing directory as a whole never exceeds 5 TB, while the marketing department users are limited to 1 TB each.

Default quota domains

A default quota domain does not account for any specific set of files but instead specifies a policy for new domains that match a specific trigger. In other words, default domains are configuration templates for actual domains. SmartQuotas uses the identity notations `default-user`, `default-group`, and `default-directory` to describe domains with default policies. For example, the domain `default-user@ifs/home` becomes `specific-user@ifs/home` for each specific user that is not otherwise defined. All enforcement on `default-user` are copied to `specific-user` when the `specific-user` allocates within the domain and the new inherited domain quota is a “linked quota”. There may be overlapping defaults (`default-user@ifs` and `default-user@ifs/home` may both be defined).

Default quota domains help drastically simplify quota management for large environments by providing a mechanism to define top-level template configurations from which many actual quotas are cloned, or linked. When a default quota domain is configured on a directory, any subdirectories created directly underneath the directory automatically inherit the quota limits specified in the parent domain. This mechanism streamlines the provisioning and management quotas for large enterprise environments. Furthermore, default directory quotas can co-exist with the user and/or group quotas and legacy default quotas.

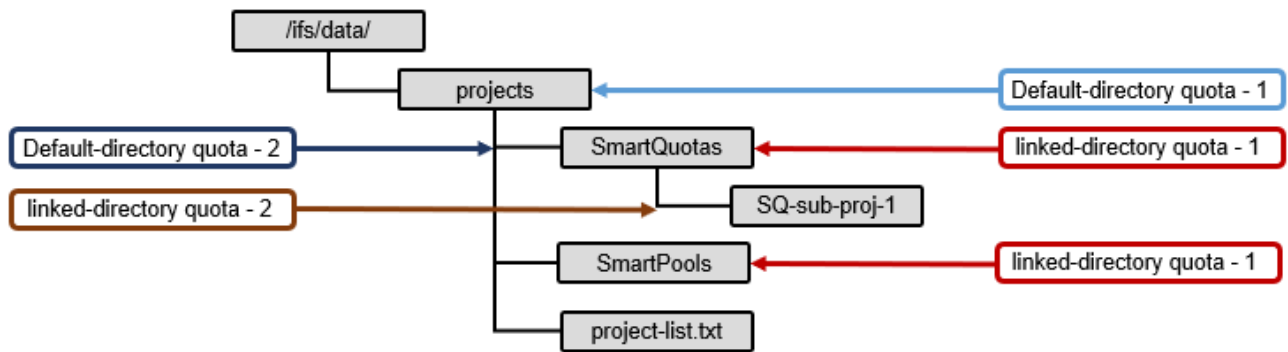


Figure 4. Default quota domains

Default directory quotas are available in OneFS 8.2 and later, in addition to the default user and group quotas available in earlier releases.

If the enforcements on a default domain change, SmartQuotas automatically propagates the changes to the linked quota domains. If a default quota domain is deleted, SmartQuotas deletes all children marked as inherited. An administrator may also choose to delete the default quota domain without deleting the children, but this will break inheritance on all inherited children.

A domain can be in one of three accounting states, as described in the following table:

Table 9. Domain accounting states

Domain accounting state	Description
Ready	A domain in the ready state is fully accounted. SmartQuotas displays “ready” domains in all interfaces and all enforcements apply to such domains.
Accounting	A domain is placed in the accounting state when it is waiting on accounting updates.
Deleting	After a request to delete a domain, SmartQuotas places the domain in the deleting state until the tear-down is complete. Domain removal might be a lengthy process.

SmartQuotas displays accounting domains in all interfaces including usage data but indicates they are in the process of being “accounted.” SmartQuotas applies all enforcements to accounting domains, even when it might reject an allocation that would have proceeded if it had completed the QuotaScan.

Domains in the deleting state are hidden from all interfaces, and the top-level directory of a domain may be deleted while the domain is still in the deleting state (assuming that no domains are in a ready or accounting state, defined on the directory). No enforcements are applied for domains in the deleting state.

A quota scan is performed when the domain is in an accounting state. This can occur during quota creation to account for the new domain, if a quota has been set for the domain, and quota deletion to un-account the domain. A QuotaScan is required when creating a quota on a nonempty directory. If quotas are created upfront on an empty directory, no QuotaScan is necessary.

In addition, a QuotaScan job may be started from the WebUI or the command-line interface using the `isi job` command. Any path specified on the command line is treated as the root of a tree to be processed. This functionality is provided primarily as a means to rescan a directory or for maintenance reasons.

Quota daemons The main processes or daemons associated with SmartQuotas are:

- `isi_quota_notify_d`
- `isi_quota_sweeper_d`
- `isi_quota_report_d`

The job of the notification daemon, `isi_quota_notify_d`, is to listen for “limit exceeded” and “link denied” events and generate notifications for each. The notification daemon also responds to configuration change events and instructs the QDB to generate “expired” and “violated” over-threshold notifications.

A quota sweeper daemon, `isi_quota_sweeper_d`, is responsible for several quota housekeeping tasks, such as propagating default changes, collecting domain and notification rule garbage, and initiating QuotaScan jobs when necessary.

Finally, the reporting daemon, `isi_quota_report_d`, is responsible for generating quota reports. Since the QDB produces only real-time resource usage, reports are necessary for providing point-in-time views of a quota domain’s usage. These historical reports are useful for trend analysis of quota resource usage.

OneFS 8.2 and later versions also include the `rpc.quotad` service to facilitate client-side quota reporting on UNIX and Linux clients using native quota tools. The service, which runs on tcp/udp port 762, is enabled by default, and control is under NFS global settings.

Also, in OneFS 8.2 and later, users can view their available user capacity set by the soft or hard user and group quotas rather than the entire cluster capacity or parent directory-quotas. The ability to view capacity based on quotas allows users to avoid the illusion of seeing available space that might not be associated with their quotas.

SmartQuotas licensing

SmartQuotas is included as a core component of OneFS but requires a valid product license key for activation. This license key can be purchased through your Dell Technologies account team. An unlicensed cluster shows a SmartQuotas warning until a valid product license has been purchased and applied to the cluster.

License keys can be easily added in the OneFS WebUI at **Cluster Management > Licensing > Activate License**.

Quota notifications

Overview

A crucial part of the quota system is to provide user notifications regarding enforcement violations, both when a violation event occurs and while the violation state persists on a scheduled basis. An enforcement quota may have several notification rules that are

associated with it. Each notification rule specifies a condition and an action to be performed when the condition is met. Notification rules are considered a part of the enforcements. Clearing an enforcement also clears any notification rules associated with it.

Enforcement quotas support the notification settings, with the following types being available:

Table 10. Quota notification settings

Quota notification setting	Description
Global default	Uses the global default notification for the specified type of quota.
Custom – basic	Enables the creation of basic custom notifications that apply to a specific quota. Can be configured for any or all the threshold types (hard, soft, or advisory) for the specified quota.
Custom – advanced	Enables the creation of advanced, custom notifications that apply to a specific quota. Can be configured for any or all threshold types (hard, soft, or advisory) for the specified quota.
None	Disables all notifications for the quota.

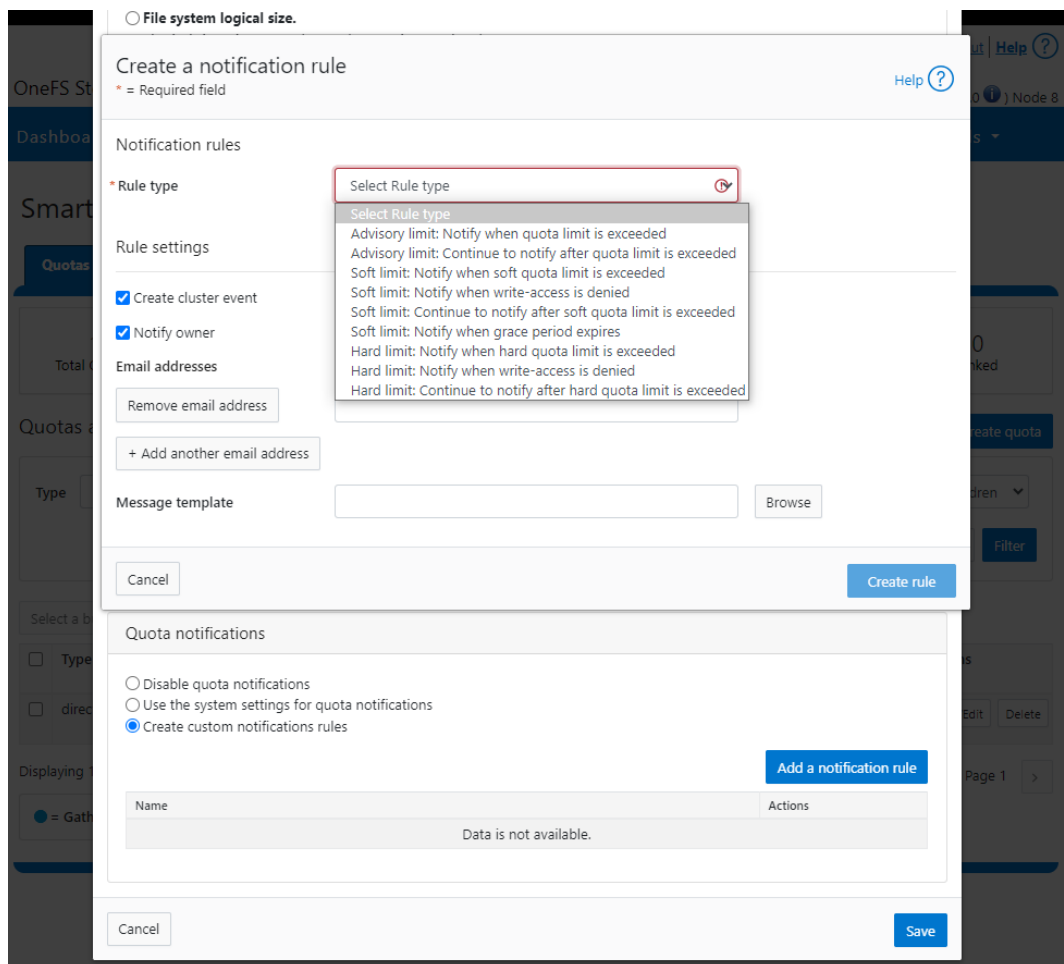


Figure 5. Quota notification setting

Note: With OneFS versions earlier than OneFS 8.2, a quota can have only a single notification setting.

A quota notification condition is an event that can trigger an action defined by a notification rule. These notification rules can specify a schedule for performing an action (for example, "every day at 5:00 AM") or immediate notification of a certain condition.

Examples of notification conditions include:

- When a threshold is exceeded; at most, once every 5 minutes
- When allocation is denied; at most, once an hour
- While over threshold, daily at 2 AM
- While grace period has been expired, weekly, on Sundays at 2 AM

Notifications are grouped within these categories:

- **Instant notifications**, which include:
 - Write-denied notifications that are triggered when a hard threshold denies a write
 - Threshold-exceeded notifications that are triggered at the moment a hard, soft, or advisory threshold is exceeded.

These instant notifications are one-time notifications because they represent a discrete event in time.

- **Ongoing notifications:** Generated on a scheduled basis to indicate a persisting condition, such as a hard, soft, or advisory threshold being over a limit or a soft threshold's grace period being expired for a prolonged period.

Notification actions

Each notification rule can perform either one or none of the following notification actions.

Table 11. Quota notification actions

Quota notification action	Description
Alert	Sends an alert for one of the quota actions
Email Manual Address	Sends email to a specific address, or multiple addresses (OneFS 8.2 and later).
Email Owner	Emails an owner mapping based on its identity source.

The email owner mapping is as follows:

- **Active Directory:** Lookup is performed against the domain controller (DC). If the user does not have an email setting, a configurable transformation from username and DC fully qualified domain name is performed to generate an email address.
- **LDAP:** LDAP user email resolution is similar to AD users. In this case, only the email attribute looked up in the LDAP server is configurable by an administrator based on the LDAP schema for the user account information.
- **NIS:** Only the configured email transformation for the NIS fully qualified domain name is used.

- **Local users:** Only the configured email transformation is used.

The actual quota notification is handled by a daemon, `isi_quota_notify_d`, which performs the following functions:

- Processes kernel notification events that get sent out. They are matched to notification rules to generate instant notifications (or other actions as specified in the notification rule).
- Processes notification schedules. The daemon checks notification rules on a scheduled basis. These rules specify what violation condition triggers a notification on a regular scheduled basis.
- Performs notifications based on rule configuration to generate email messages or alert notifications.
- Manages persistent notification states so that pending events are processed in the event of a restart.
- Handles rescan requests when quotas are created or modified.

Notification email templates

SmartQuotas provides email templates for advisory, grace, and regular notification configuration, which can be found under `/etc/ifs`. The advisory limit email template (`/etc/ifs/quota_email_advisory_template.txt`) for example, displays:

```
Subject: Disk quota exceeded
```

```
The <ISI_QUOTA_DOMAIN_TYPE> quota on path <ISI_QUOTA_PATH> owned by
<ISI_QUOTA_OWNER> has exceeded the <ISI_QUOTA_TYPE> limit.
```

```
The quota limit is <ISI_QUOTA_THRESHOLD>, and <ISI_QUOTA_USAGE> is
currently in use. <ISI_QUOTA_HARD_LIMIT> Contact your system
administrator for details.
```

An email template contains text and, optionally, variables that represent quota values. The following table lists the SmartQuotas variables that may be included in an email template.

Table 12. SmartQuotas variables for email templates

Variable	Description	Example
ISI_QUOTA_DOMAIN_TYPE	Quota type. Valid values are: directory, user, group, default-directory, default-user, default-group	default-directory
ISI_QUOTA_EXPIRATION	The expiration date of grace period	Fri Jan 8 12:34:56 PST 2021
ISI_QUOTA_GRACE	The grace period, in days	5 days
ISI_QUOTA_HARD_LIMIT	Includes the hard limit information of the quota to make advisory/soft email notifications more informational	You have 30 MB left until you reach the hard quota limit of 50 MB.
ISI_QUOTA_NODE	The hostname of the node on which the quota event occurred	us-wa-1
ISI_QUOTA_OWNER	The name of quota domain owner	jsmith

Variable	Description	Example
ISI_QUOTA_PATH	The path of quota domain	/ifs/home/jsmith
ISI_QUOTA_THRESHOLD	The threshold value	20 GB
ISI_QUOTA_TYPE	The threshold type	Advisory
ISI_QUOTA_USAGE	The disk space in use	10.5 GB

System alerts

Various system alerts are sent out to the standard cluster alerting system when specific events occur. These include:

Table 13. System alert types

Alert type	Level	Event description
NotifyFailed	Warning	An attempt to process a notification rule failed externally, such as an undelivered email.
NotifyConfig	Warning	A notification rule failed due to a configuration issue, such as a nonexistent user or missing email address.
NotifyExceed	Warning	A child quota's advisory/soft/hard limit is greater than any of the parent quota's hard limit.
ThresholdViolation	Info	A quota threshold was exceeded. The conditions under which this alert is triggered are defined by notification rules.
DomainError	Error	An invariant was violated, resulting in a forced domain rescan.

Quota accounting

Overview

SmartQuotas has four main resources that are used in quota accounting:

Table 14. SmartQuotas accounting resources

Accounting resource	Description
Physical size	Physical size includes all the on-disk storage associated with files and directories, except for some metadata objects including the LIN tree and snapshot tracking files (STFs). For deduplicated data and file clones, each file's 8 KB reference to a shadow store is included in the physical space calculation.
File system logical size	The file system logical size calculation approximates disk usage on typical storage arrays by ignoring the erasure code, or FEC, protection overhead that OneFS employs. For regular files, the logical data space is the amount of storage required to house a particular file if it was 1x mirrored. Logical space also incorporates a file's metadata resources.
Application logical size	Application logical size reports total logical data store size across different tiers, including CloudPools. This allows users to view quotas and free space as an application would view it, in terms of how much capacity is available to store logical data regardless of data reduction or tiering technology.
Inodes	SmartQuotas counts the number of the logical inodes, allowing accounting for files without any ambiguity from hard links or protection.

During quota configuration, these accounting options are available as enforcement limits. For example, from the OneFS WebUI:

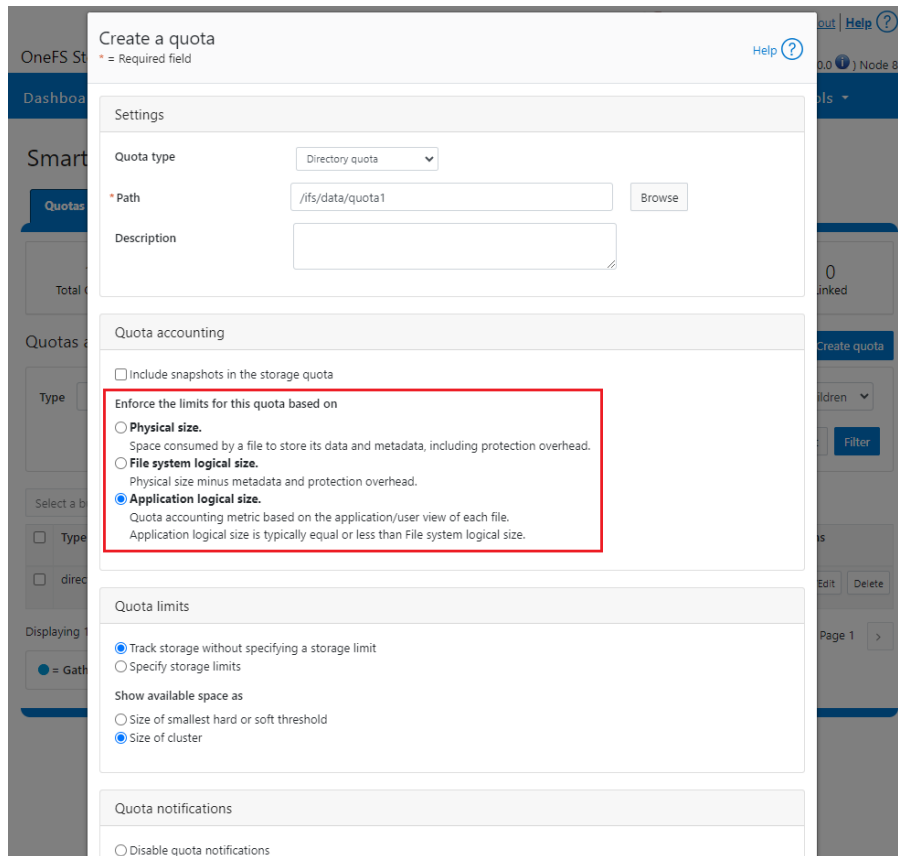


Figure 6. Quota configuration accounting options

Application logical size quotas are available in OneFS 8.2 and later. Existing quotas can easily be configured to use application logical size upon upgrading from an earlier OneFS version. The benefits of application logical size quotas include:


- Prevents quota consumption being affected by snapshots, protection overhead, deduplication, compression, and location of files
- Removes previous limitation where SmartQuotas reported only on-cluster storage, ignoring cloud consumption
- Presents view that aligns with Windows storage accounting
- Enables accounting and enforcing quota on actual file sizes
- Precisely accounts for small files
- Enables enforcing quotas on a path irrespective of the physical location of file

The following table describes how SmartQuotas accounts for a 1 KB file with the various datatypes:

Table 15. Quota accounting datatypes

Datatype	Accounting
File: physical size	Every non-sparse 8 KB disk block that a file consumes, including protection
File: file system logical size	Every non-sparse 8 KB disk block that a file consumes, excluding protection
File: application logical size	The actual size of file (rather than total of 8 KB disk blocks consumed)
CloudPools file: file system logical size	Size of CloudPools SmartLink stub file (8 KB)
CloudPools file: application logical size	The actual size of file on cloud storage (rather than local stub file)
Directories	Sum of all directory entries
Symlinks	Data size
ACL and similar	Data size
Alternate data stream	Each ADS is charged as a file and a container as a directory

The following figure shows each method of accounting for a 1 KB file.

**Figure 7. Quota accounting datatypes**

Logical size reports 8 KB, or one block, physical size reports 24 KB (file with 3x mirroring protection), and application logical shows the actual size of 1 KB.

Other resources that are encountered during quota accounting include:

- **Hard links**—Each logical inode is accounted exactly once in every domain to which it belongs. If an inode is present in multiple domains, it is accounted in multiple domains. Alternatives such as shared accounting were considered. However, if inodes are not accounted once in every domain, it is possible for the deletion of a hard link in one domain to put another domain over quota.
- **Alternate data streams (ADS)**—A file with an alternate data stream or resource fork is accounted as the sum of the resource usage of the individual file, the usage for the container directory, and the usage for each ADS. SmartQuotas handles the

renaming of a file with ADS synchronously, even though the ADS container is just a directory. SmartQuotas stores an accounting summary on the ADS container to handle renames.

- **Directory rename**—A directory rename presents a unique challenge to a per-directory quota system. Renames of directories within a domain are trivial—if both the source and target directories have the same domain membership, no accounting changes. However, moving nonempty directories is not permitted when the SmartQuotas configuration is different on the source and the target parent directories. If a user trusts the client operating systems to copy files and preserve all the necessary attributes, and then the user may set `dir_rename_errno` to `EXDEV`, which causes most UNIX and Windows clients to do a copy and delete of the directory tree to affect the move.
- **Snapshot accounting**—Optionally, a quota domain can also include snapshot usage in its accounting. SmartQuotas snapshot accounting supports only snapshots that are created after the quota domain was created because determining quota governance (including QuotaScan job) for existing snapshots consumes significant time and resources. Because most administrators cycle their snapshots through timed expirations, SmartQuotas will eventually accrue enough accounting information to include the entire set of relevant snapshots on the system.

Storage quotas and usage info

SmartQuotas software supports flexible reporting options that enable administrators to more effectively manage cluster resources and analyze usage statistics. The goal of quota reporting is to provide a summarized view of the past or present state of the quota domains. SmartQuotas supports the following methods of data collection and reporting:

- **Scheduled reports**—Generated and saved on a regular interval
- **Ad hoc reports**—Generated and saved per request of the user
- **Live reports**—Generated for immediate and temporary viewing

A summary of general quota usage info can be viewed from the WebUI at **File System > SmartQuotas > Quotas & Usage**.

Smart Quotas

Quotas and usage
Generated reports archive
Settings

3
Total Quotas

3
Directory

0
User

0
Default-user

0
Group

0
Default-group

0
Linked

Quotas and usage
Create quota

Type
All ▼
User/Group name
Exceeded
All ▼
Path
Path
Recursive path
Exclude children ▼

Reset
Filter

Select a bulk action ▼

<input type="checkbox"/>	Type	Path	Usage	FS Physical	Reduction ?	Efficiency ?	Advisory Limit	Soft Limit	Hard Limit	Actions
<input type="checkbox"/>	directory	/ifs/data	● 1.13 TiB	111.98 MiB	--	--	500 GiB / 231.59 % used	--	--	View/Edit Delete
<input type="checkbox"/>	directory	/ifs/data/quot a2	● ≡ 111.92 MiB	111.92 MiB	--	--	--	100 MiB / 111.92 % used	--	View/Edit Delete
<input type="checkbox"/>	directory	/ifs/data/quot a1	● ≡ 32 KiB	32 KiB	--	--	--	100 MiB / < 1% used	--	View/Edit Delete

Displaying 1 to 3 entries

«
<
Page 1
>
»

● = Gathering information(values may be inaccurate) ≡ = Usage includes data-protection overhead

Figure 8. General quota usage summary

For each quota profile, click **View/Edit** to see additional information and context:

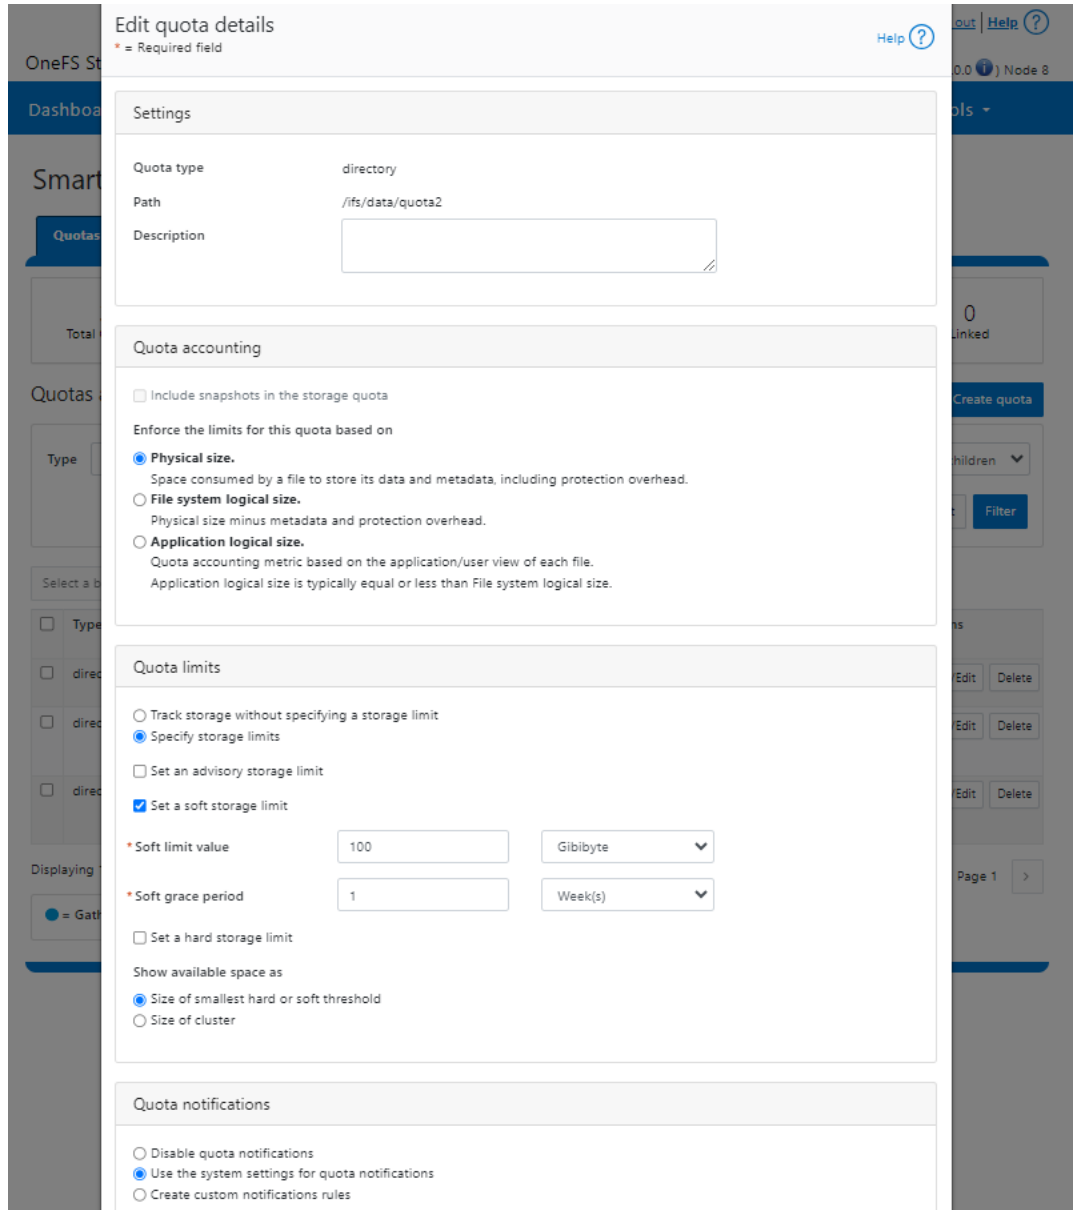


Figure 9. Quota information

Client-side quota views

Client-side quota reporting in OneFS 8.2 and later includes:

- Support for `rpc.quotad`, which allows NFS clients to view quota consumption for both hard and soft quotas using the native Linux and UNIX quota CLI utilities.
- The ability to view available user capacity set by soft and/or hard user or group quotas, rather than the entire cluster capacity or parent directory-quotas.

Quota reports

Overview

Quota reports and summaries are typically stored in the `/ifs/.isilon/smartquotas/reports` directory. This location is configurable.

Each generated report includes the quota domain definition, state, usage, and global configuration settings. By default, ten reports and ten summaries are kept at one time, and older versions are purged. This setting can be configured from the WebUI, by going to **File System > SmartQuotas > Settings**:

Smart Quotas

Quota reports and summaries are typically stored in the `/ifs/.isilon/smartquotas/reports` directory. This location is configurable.

Each generated report includes the quota domain definition, state, usage, and global configuration settings. By default, ten reports and ten summaries are kept at one time, and older versions are purged. This setting can be configured from the WebUI, by going to **File System > SmartQuotas > Settings**:

Smart Quotas

Quotas and usage | Generated reports archive | **Settings**

Edit quota settings

Scheduled reporting

Archive directory:

Number of scheduled reports retained:

Manual
 Scheduled

Scheduled:

Run schedule every: week(s)

Run schedulers on:

- Sunday
- Monday
- Tuesday
- Wednesday
- Thursday
- Friday
- Saturday

Run one schedule per specified day

Run schedule at:

Run multiple schedules per specified day

Manual reporting

Archive directory:

Number of live reports retained:

Email mapping

Provider type	Current domain	Map-to domain	Actions
---------------	----------------	---------------	---------

Figure 10. Quota report configuration

On-demand reports can also be created at any time to view the current state of the storage quotas system. These live reports can be saved manually.

Reports and summaries are prefixed by either `ad hoc` or `scheduled` to aid with identification.

The OneFS CLI export functionality uses the same data generation and storage format as quota reporting but should not require any extra requirements beyond the three types of reports. After the collection of the raw reporting data, data summaries can be produced based on a set of filtering parameters and sorting type.

Reports can be viewed from historically sampled data or a live system. In either case, the reports are views of usage data at a given time. SmartQuotas does not provide reports on aggregated data over time (trending reports). However, the raw data can be used by a quota administrator to answer trending questions.

Quota report format

A quota report is a timestamped XML file that starts with global configuration settings and global notification rules:

```
# cat scheduled_quota_report_1465786800.xml
  <global-config>
    <quota-global-config>
      <reports>
        <schedule-pattern>1100000000|every sunday at 11pm</schedule-
pattern>
        <schedule-dir>/ifs/.isilon/smartquotas/reports</schedule-dir>
        <schedule-copies>10</schedule-copies>
        <adhoc-dir>/ifs/.isilon/smartquotas/reports</adhoc-dir>
        <adhoc-copies>10</adhoc-copies>
      </reports>
    </quota-global-config>
  </global-config>
  <global-notify>
</global-notify>
  <domains>
    <domain type="default-group" snaps="0" lin="0x0000000100020006">
      <path>/ifs/home</path>
      <inactive/>
      <enforcements default-resource="logical">
</enforcements>
      <notifications use="global"/>
    </domain>
    <domain type="group" snaps="0" lin="0x0000000100020006" id="0">
      <inherited/>
      <id-name>wheel</id-name>
      <usage resource="physical">109568</usage>
      <usage resource="logical">32929</usage>
      <usage resource="inodes">6</usage>
      <path>/ifs/home</path>
      <inactive/>
      <enforcements default-resource="logical">
</enforcements>
      <notifications use="default"/>
    </domain>
    <domain type="group" snaps="0" lin="0x0000000100020006" id="10">
      <inherited/>
      <id-name>admin</id-name>
      <usage resource="physical">28160</usage>
      <usage resource="logical">8208</usage>
      <usage resource="inodes">2</usage>
      <path>/ifs/home</path>
      <inactive/>
      <enforcements default-resource="logical">
</enforcements>
      <notifications use="default"/>
    </domain>
```

```

<domain type="group" snaps="0" lin="0x0000000100020006" id="1800">
  <inherited/>
  <id-name>Isilon Users</id-name>
  <usage resource="physical">1811456</usage>
  <usage resource="logical">705620</usage>
  <usage resource="inodes">42</usage>
  <path>/ifs/home</path>
  <inactive/>
  <enforcements default-resource="logical">
</enforcements>
  <notifications use="default"/>
</domain>
<domain type="user" snaps="0" lin="0x0000000100020596" id="2002">
  <id-name>nick</id-name>
  <usage resource="physical">1001984</usage>
  <usage resource="logical">483743</usage>
  <usage resource="inodes">12</usage>
  <path>/ifs/home/nick</path>
  <enforcements default-resource="logical">
    <enforcement type="soft" resource="logical">
      <limit>10485760</limit>
      <grace>7776000</grace>
    </enforcement>
    <enforcement type="advisory" resource="logical">
      <limit>5242880</limit>
    </enforcement>
  </enforcements>
  <notifications>
    <quota-notify-map tag="1"></quota-notify-map>
  </notifications>
</domain>
</domains>
</quota-report>

```

Figure 11. XML quota report

Both the inode and path, as well as name and ID, are stored with each domain listing. Quota notification rules are read and inserted into a domain entry only if the domain is not inherited to avoid any performance impact of reading the quota notification rules with each domain.

Quota report management

SmartQuotas can be configured to produce scheduled reports to help monitor, track, and analyze storage usage on a OneFS powered cluster.

Quota reports are managed by configuring settings that provide control over when reports are scheduled, how they are generated, where and how many are stored, and how they are viewed. The maximum number of scheduled reports that are available for viewing in the web-administration interface can be configured for each report type. When the maximum number of reports is stored, the system automatically deletes the oldest reports to make space for new reports as they are generated.

Creating a quota report schedule

SmartQuotas can be easily configured to generate a quota report on a specified schedule. These settings determine whether and when scheduled reports are generated, and where and how the reports are stored. Even if scheduled reports are disabled, you can still run unscheduled reports at any time.

To run unscheduled reports:

1. From the OneFS WebUI, go to **File System Management > SmartQuotas > Settings**.
2. (Optional) On the **Quota settings** page, for **Scheduled Reporting**, click **On**. The **Report Frequency** option appears.
3. Click **Change schedule** and select the report frequency that you want to set from the list.
4. Select the reporting schedule options that you want.
5. Click **Save**.

Reports are generated according to your criteria and can be viewed in the **Generated Reports Archive**.

Generating a quota report

In addition to scheduled quota reports, you can generate a report to capture usage statistics at a point in time. Before you can generate a quota report, quotas must exist and no QuotaScan jobs can be running.

To generate a quota report:

1. Click **File System Management > SmartQuotas > Generated Reports Archive**.
2. Under **Generated Quota Reports Archive**, click **Generate a quota report**.
3. Click **Generate Report**.

The new report appears in the **Quota Reports** list.

Locating a quota report

You can locate quota reports, which are stored as XML files, and then use your own tools and transforms to view them. This task can only be performed from the OneFS command-line interface, as follows:

1. Open a secure shell (SSH) connection to any node in the cluster and log in.
2. Go to the directory where quota reports are stored.

The default quota report location is:

```
/ifs/.isilon/smartquotas/reports
```

If quota reports are not in the default directory, you can run the `isi quota settings` command to find the directory where they are stored.

3. At the command prompt, run the `ls` command.

To view a list of all quota reports in the directory, run the following command:

```
# ls -a *.xml
```

To view a specific quota report in the directory, run the following command:

```
# ls <filename>.xml
```

Quota performance tuning

Within OneFS, quota data is maintained in quota accounting blocks (QABs). Each QAB contains many quota accounting records, which need to be updated whenever a particular user adds or removes data from an area of the file system on which quotas are enabled (quota domain). If many clients are simultaneously accessing the quota domain, these blocks can become highly contended and a potential bottleneck. Similarly, if a single client (or small number of clients) consistently makes a large number of small writes to files within a single quota, write performance could again be impacted.

To address this issue, quota accounts have a mechanism to help avoid hot spots on the nodes storing QABs. Quota account constituents (QACs) help parallelize the quota accounting by including additional QAB mirrors distributed across other nodes in the cluster.

Configuration is managed through a `sysctl`, `efs.quota.reorganize.qac_ratio`, which increases the number of quota accounting constituents. This provides better scalability and reduces latencies on heavy create/delete activities when quotas are used.

Using this parameter, the internally calculated QAC count for each quota is multiplied by the specified value. If a workflow experiences write performance issues, and it has many writes to files or directories that are governed by a single quota, and then increasing the QAC ratio can significantly improve write performance.

The `sysctl efs.quota.reorganize.qac_ratio` can be reconfigured to its maximum value of 8 from its default value of 1 using the following CLI command:

```
# isi_sysctl_cluster efs.quota.reorganize.qac_ratio=8
```

To verify the persistent change, run:

```
# cat /etc/mcp/override/sysctl.conf | grep qac_ratio
efs.quota.reorganize.qac_ratio=8 #added by script
```

Although increasing the QAC count through this `sysctl` can improve performance on write-heavy quota domains, some amount of experimentation might be required until you find the ideal QAC ratio value. Adjusting the parameter can adversely affect write performance, if you apply a value that is too high or if you apply the parameter in an environment that does not have diminished write performance due to quota contention.

Also, OneFS provides a CLI command that can restripe the QABs to improve their performance:

```
# isi_restripe_qabs retune
```

This utility can be run either on demand or periodically to randomly redistribute QABs for all existing quotas. It does this by ignoring the default `rebalance` layout and running a `retune` layout strategy instead, thereby alleviating the performance impact from an imbalanced QAB layout.

For more information about OneFS commands and feature configuration, see the [OneFS Administration Guide](#).

SmartQuotas best practices and considerations

SmartQuotas best practices

SmartQuotas best practices include:

- Avoid creating quotas on the root directory of the default OneFS share (`/ifs`). A root-level quota might result in performance degradation.
- Avoid governing a single directory with overlapping quotas, which can degrade performance.
- Consider using directory quotas to alert you to and constrain runaway jobs, preventing them from consuming massive amounts of storage space.
- Avoid using enforcement quotas for snapshot-tracking quota domains.
- Before using quota data for analysis or other purposes, verify that no QuotaScan jobs are running.
- Rather than editing the quota email notification templates directly, copy them from `/etc/ifs` to another directory to edit and deploy them.
- If quota reports are not in the default directory, you can run the `isi quota settings reports view` command to find the directory where they are stored.
- Use the `isi quota quotas notifications disable` command to disable all notifications for a quota.
- Use the following CLI syntax to configure default directory quotas (in addition to default user and group quotas), available in OneFS 8.2 and later:

```
# isi quota quotas create -path=</ifs/parent-dir> --  
type=default-directory -hard-threshold=<limit>
```

- Use the TreeDelete job, as follows, to delete directories with quotas:

```
# isi job start TreeDelete --paths=</ifs/quota_dir>
```

SmartQuotas considerations

- Do not exceed 500,000 quotas per cluster in OneFS 8.2 and 20,000 quotas per cluster in earlier releases.
- SmartQuotas schedule notification rules are limited to 20,000.
- In OneFS versions earlier than OneFS 8.2, SmartQuotas only reports the true quota container size for directory quotas (user and group quotas report the entire file system size). Beginning with OneFS 8.2, this functionality is extended to user and group quotas.
- NFS quota support allows customers to view quota soft and hard limits using UNIX and Linux-based quota tools (rpc.quotad protocol).
- With CloudPools data, the quota is calculated based on the size of the data local to the cluster. For example, for a 100 MB file tiered to a cloud provider, SmartQuotas would calculate just the size of the local stub file (8 K).

- SmartQuotas reports the logical capacity of the files, whether they are deduplicated or not.
- The QuotaScan job runs after the creation of a quota, but not after a change. However, it does run on a schedule and incorporates any changes then.
- If two quotas are created on the same directory—for example, an accounting quota without snapshots and a hard quota with snapshots—the quota without snapshot data overrules the limit from the quota with snapshot data.
- SmartQuotas also provide a low impact way to provide directory file count reports.
- Configuration changes for linked quotas must be made on the parent quota that the linked quota is inheriting from. Changes to the parent quota are propagated to all children. To override configuration from the parent quota, you must first unlink the quota.
- If a quota type uses the accounting-only option, enforcement limits cannot be used for that quota.
- Cloned and deduplicated files are treated as ordinary files by quotas. If the quota includes data-protection overhead, the data-protection overhead for shared data is not included in the usage calculation.
- Moving quota directories across quota domains is not supported.
- You can edit or delete a quota report only when the quota is not linked to a default quota.
- A quota can only be unlinked when it is linked to a default quota. Configuration changes for linked quotas must be made on the parent (default) quota that the linked quota is inheriting from. Changes to the parent quota are propagated to all children. If you want to override configuration from the parent quota, you must first unlink the quota.
- Disabling all quota notifications also disables all system notification behavior. Use the `--clear` options to remove specific quota notification rules and fall back to the system default.
- With support for multiple email recipients for notifications and alerts, the maximum supported size of the email address list (comma-separated) is 1,024 characters. Multiemail notifications can be configured with the following CLI syntax:


```
# isi quota notifications create </ifs/quota> directory advisory
exceeded --action-email-
address=user1@mail.com,user2@mail.com,user3@mail.com --holdoff=0
```
- Quota containers compartmentalize `/ifs`, so that a directory with a container appears as its own separate file system slice. To configure a directory quota with a 4 TB container on `/ifs/data/container1`, you could use the following CLI command:


```
# isi quota quotas create /ifs/data/container1 directory --
hard-threshold 4T --container true
```
- Only soft or advisory limits can be defined as a percentage of a hard limit. Defining the same type of limit with both percent-based and absolute is not permitted in a single request.

- While OneFS 8.2 and later versions support an overall quota limit of up to 500,000, support for schedule notification rules is still limited to 20,000.

SmartQuotas and OneFS feature integration

Snapshots and SmartQuotas

In addition to including data-protection overhead, you can opt to include snapshot data when calculating a quota's usage limits.

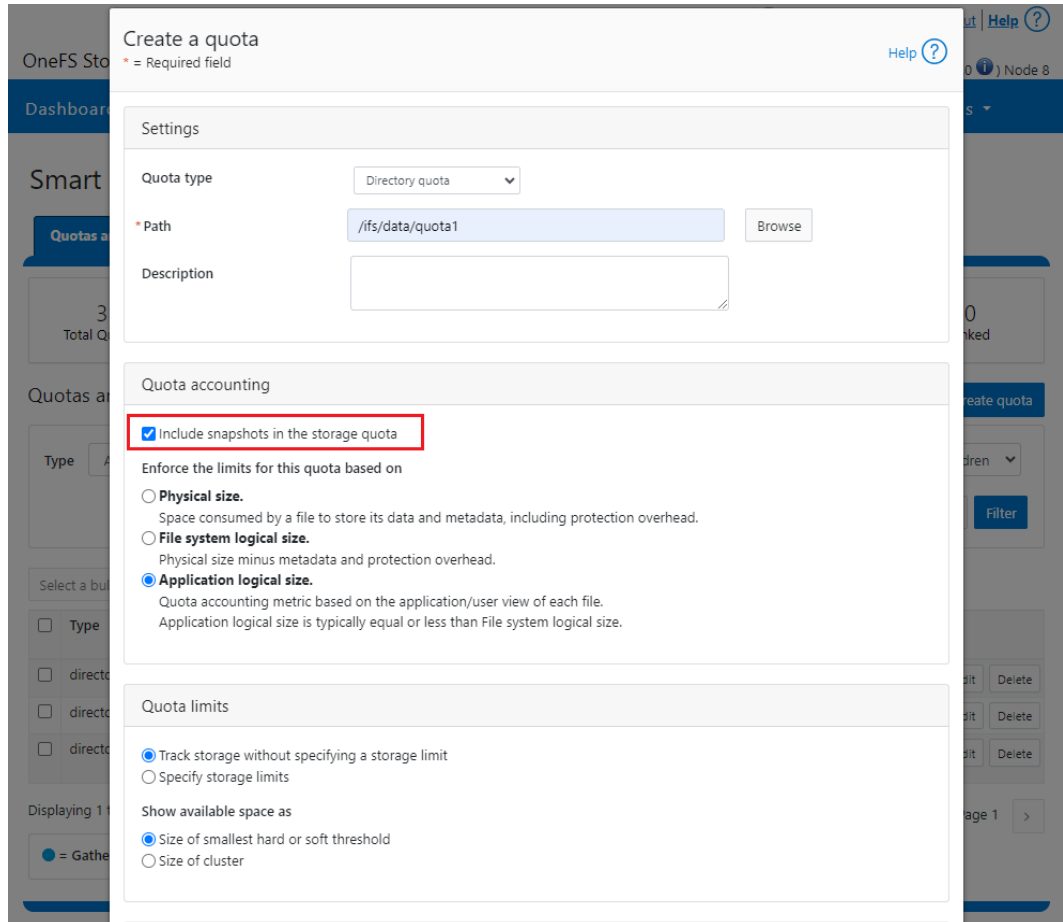


Figure 12. Quota snapshot reporting

Note: SmartQuotas reports only on snapshots that are created after the quota domain was created because determining quota governance (including QuotaScan job) for existing snapshots consumes significant time and resources. However, as snapshots age out, SmartQuotas will gradually accrue accounting information for the entire set of relevant snapshots.

For more information, see the [OneFS SnapshotIQ](#) white paper.

Data reduction and SmartQuotas

Compressed and deduplicated files appear no differently than regular files to standard quota policies. However, for deduplicated files, if the quota is configured to include data-protection overhead, the additional space that is used by the shadow store will not be accounted for by the quota.

SmartQuotas reports both storage efficiency and data reduction as a ratio across the dataset as specified in the quota path field. These efficiency and data reduction ratios are for the full quota directory and its contents, including any overhead, and reflects the net efficiency of both compression and deduplication.

Smart Quotas

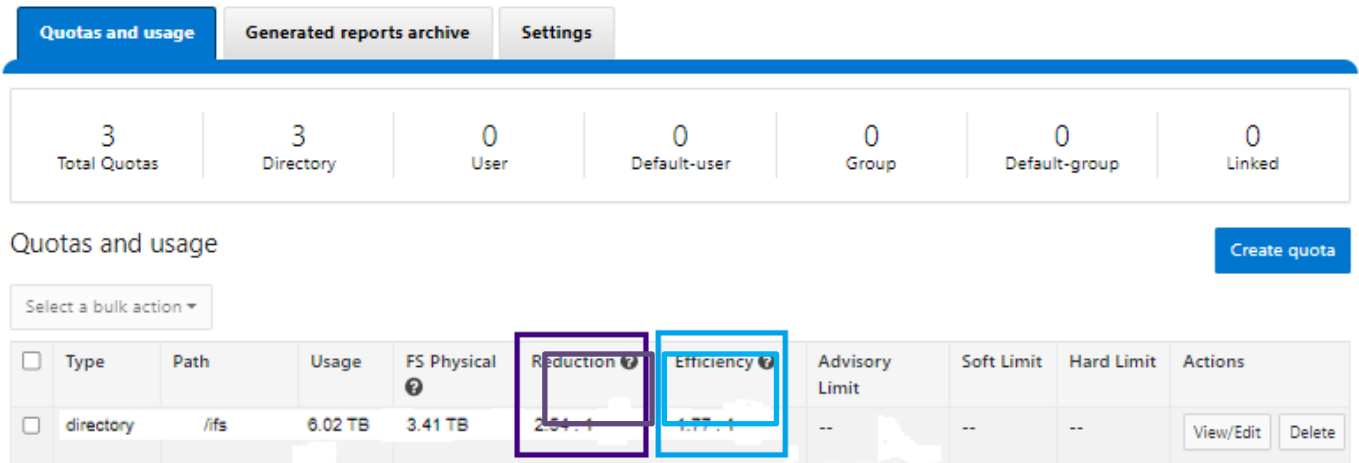


Figure 13. Quotas and usage status detailing efficiency ratios

The 'isi quota quotas view' CLI command provides considerably more detailed storage capacity and efficiency metrics. These include the following:

Metric	Description
AppLogical	The application data that can be written to the cluster, irrespective of where it is stored from.
FSLogical	Removing sparse data (data that was always sparse, was zero block eliminated, or data that has been moved to the cloud and so on.) results in filesystem logical, which is the non-sparse data stored on the filesystem.
AppPhysical	Data reduction techniques, such as compression and deduplication, reduce filesystem logical to application physical, or preprotected physical. This is the physical size application data on the filesystem disks, and does not include metadata, protection overhead, or data moved to the cloud.
FSPhysical	Application physical with data protection overhead added – including inode, mirroring and FEC blocks and so on. Filesystem physical is also seen as protected physical.
Reduction	The data reduction ratio is the amount that has been reduced from the filesystem logical down to the application physical.
Efficiency	The storage efficiency ratio is the filesystem logical divided by the filesystem physical.

The graphic below illustrates the relationship between the capacity, data reduction and storage efficiency elements in the OneFS storage taxonomy:

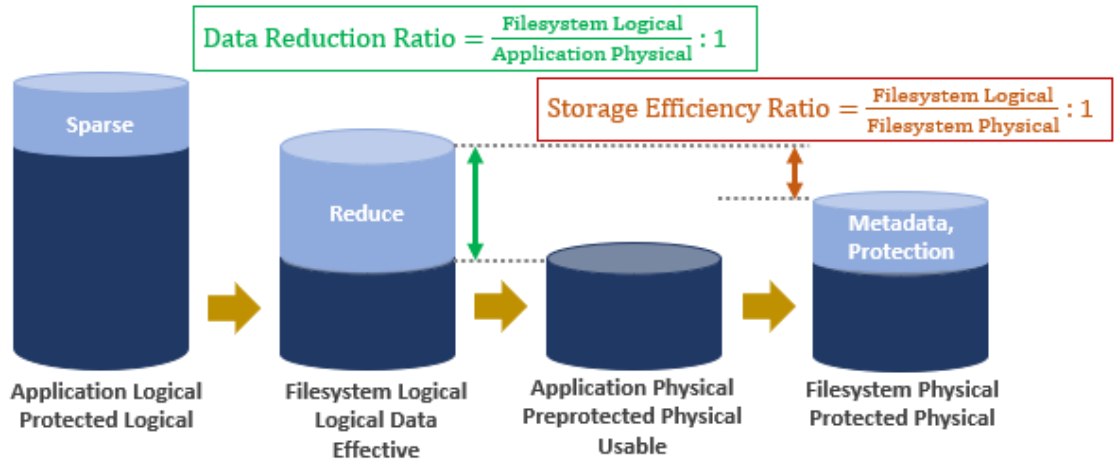


Figure 14. SmartQuotas capacity, data reduction, and storage efficiency relationships.

For more information about compression and inline deduplication, see the [OneFS In-line Data Reduction](#) white paper.

For more information about post-process deduplication, see the [OneFS SmartDedupe](#) white paper.

SyncIQ and SmartQuotas

Quotas are matched one-to-one across the replication set. Multiple quotas are supported within a source directory or domain structure, and the target directory is now included in a quota domain.

During replication SyncIQ ignores quota limits. However, if a quota is over limit, quotas still prevent users from adding data. SyncIQ will never automatically delete an existing target quota; instead, the SyncIQ operation fails. This might occur during an initial synchronization where the target directory has an existing quota under it, or if a source directory is deleted that has a quota on it on the target. The quota remains and requires administrative removal if needed.

For more information, see the [OneFS SyncIQ](#) white paper.

CloudPools and SmartQuotas

Application logical quotas provide a quota accounting metric, which accounts for, reports on, and enforces the actual space consumed and available for storage, independent of whether files are on-premises or cloud-tiered.

SmartQuotas use cases

Scenario 1: Quota management

A university wants to give its students and groups a fixed amount of storage to control and keep storage growth in check. The storage administrator wants to know how much each student is consuming and limit storage use. To accomplish this, the storage administrator:

- Sets default user hard or soft quotas
- Configures email alerts to students to encourage self-cleanup of files

Scenario 2: Thin provisioning and scale-out

A company in the media and entertainment industry wants to overprovision storage and only add capacity when needed. The company buys two 200 TB of capacity for its cluster and allocates 500 GB to each of 1,000 users—effectively thin provisioning 500 TB. To enable this, the storage administrator:

- Sets directory quotas for each user with a hard quota of 500 GB each.
- Sets up a system alert to notify the storage administrator to add capacity (nodes) when the 200 TB is 75 percent full.
- Scales out effectively by adding capacity only when needed.

Scenario 3: HPC compute farm constraining

A semiconductor company uses a large HPC compute cluster for parts of their EDA workflow, and it wants to guard against runaway jobs for consuming massive amounts of storage. The company runs heavy computation jobs from a large compute farm against a “scratch space” directory, housed on an S-node tier on their cluster, and garbage collection is run at midnight.

Throughout the workday, storage administrators have difficulty tracking storage utilization. Occasionally, jobs from the compute farm run amok, tying up large swathes of fast, expensive storage resources and capacity. To help prevent this, the storage administrator:

- Sets an advisory directory quota on the scratch space at 80 percent utilization for advanced warning of an issue.
- Configures a hard directory quota to prevent writes at 90 percent utilization.

Conclusion

Traditional systems with volume quotas are limited to a single storage device—and to a single volume. After setup, volume-level solutions are limited in management flexibility. For example, typical implementations require the management of quotas across different volumes, multiple storage systems, and multiple file systems. This approach creates challenges when making changes to resources, tracking specific users and groups, moving directory structures, or moving data between file systems and volumes. As more storage silos are added, the complexity only increases.

OneFS SmartQuotas helps to eliminate these management challenges. By the virtue of being fully cluster-aware and spanning all nodes, regardless of the cluster size, IT managers can administer SmartQuotas from a single point of view, provisioning storage exactly how they want, and changing quota policy options on demand.

Conclusion

Storage administrators can define named quotas for specific individual users or groups, or they can create default quotas that control disk usage for anyone accessing the cluster.

Administrators can manage storage across their enterprise so that specific users and groups are allowed to see only the storage that they have been provisioned. Hard, soft, and advisory limits can be set across the organization for specific users and groups, and across the various directory structures. Storage administrators can also configure alerts and send email notifications to users to advise them that quota limits are approaching, enforcing hard stops on writes or providing a grace period of several days before enforcing thresholds. When new users need to be added or directory structures need to be changed, SmartQuotas allows IT managers to immediately modify their quota policies, thus easily meeting the changing storage demands of the enterprise.