

File Services Administration Guide

Hitachi Virtual Storage Platform Gx00 and Fx00 with NAS
Modules

VSP N series

Hitachi NAS Platform

Release 14.2

© 2011, 2019 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Edge, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Contents

Preface	9
Related Documentation.....	9
Accessing product documentation.....	12
Getting help.....	12
Comments.....	12
Chapter 1: File systems	13
File system access protocols.....	13
File system tiers.....	13
Data spillage between tiers.....	14
Confining new tiered file systems.....	15
Supported file system limits.....	15
Creating a new file system.....	16
Read caches.....	20
Dedupe File Systems.....	20
Deduplication characteristics.....	21
Deduplication interoperability with existing NAS Platform features.....	21
Calculating deduplication space savings.....	22
Viewing deduped file system usage.....	23
Chapter 2: Managing file systems.....	25
Viewing available file systems.....	25
Displaying file system details.....	28
Formatting a file system.....	33
Standard bitmap support.....	34
Mounting a file system.....	35
Cloning files and directory trees.....	35
Cloning symlinks in directory trees.....	36
File clone commands.....	37
Tree clone commands.....	37
Deleting a tree directory with tree-delete.....	38
Important considerations when using tree-delete.....	38
Unmounting a file system and tree-delete.....	38
Undeletable directories.....	39
Using tree-delete.....	39

Submitting a tree-delete job.....	39
Troubleshooting tree-delete.....	39
Controlling file system space usage.....	40
File system utilization recommendations.....	41
Archive file systems.....	41
High activity file systems.....	41
Dynamic Superblocks (DSB).....	42
Increasing the size of a file system.....	43
Thin provisioning file systems.....	43
Managing file system expansion.....	45
Enabling and disabling file system auto-expansion.....	47
Expanding a file system.....	47
Moving a file system.....	49
File system relocation.....	50
Using system lock on file systems.....	52
Enabling and disabling system lock for a file system.....	52
Recovering a file system.....	53
Restoring a file system from a checkpoint	54
File system recovery from a snapshot	55
Automatic file system recovery	56
Using deduplication file system.....	56
Determining sufficient space for dedupe conversion.....	56
Preparing for dedupe conversion.....	58
Viewing the deduplication file system page.....	58
Enabling dedupe for a file system.....	60
Converting a file system to enable dedupe.....	60
Dedupe support for object replication targets.....	61
Managing file system quotas	61
Managing usage quotas	61
Setting user and group file system quota defaults.....	64
Adding a quota.....	66
Modifying a file system quota.....	69
Deleting a file system quota.....	71
Managing quotas on virtual volumes.....	71
Advertising NFS exports for Virtual Volumes.....	72
Viewing and modifying virtual volume quotas.....	73
Setting user/group defaults.....	75
Exporting quotas for a specific virtual volume.....	76
Managing virtual volumes.....	77
Viewing virtual volumes.....	77
Adding a virtual volume.....	79

Modifying a virtual volume.....	81
Deleting a virtual volume	84
Enabling and disabling file system capacity and free space reporting based on virtual volume quotas.....	84
Using the per-file system throttle feature.....	85
Creating a read cache file system.....	86
Monitoring file system operations.....	88
Chapter 3: Managing file system security.....	91
Viewing file system security	91
NFS security and Kerberos.....	92
Kerberos principal formats.....	93
Setting secure NFS.....	94
Mixed security mode.....	94
AES support for SMB	94
SMB access to native SMB files.....	95
NFS access to native NFS files.....	96
Client access to non-native files.....	97
UNIX security mode.....	98
Changing security modes	98
Mixed mode operation.....	99
File name representation.....	99
Symbolic links.....	99
Mixed mode operation and LDAP servers.....	102
Mandatory and advisory byte-range file locks in mixed mode.....	103
Opportunistic locks (oplocks).....	104
Exclusive and batch oplocks.....	105
Level II oplocks.....	106
User and group names in NFSv4.....	106
Configuring user and group mappings.....	106
Managing NFS user and group mapping.....	107
About importing user or group mappings from a file or an NIS LDAP server.....	114
File system auditing.....	117
About file system audit logs.....	118
Controlling file system auditing.....	119
Creating a file system audit policy.....	119
Configuring auditing on the Windows client.....	121
Displaying file system audit logs.....	127
Chapter 4: Sharing resources with NFS clients.....	128
Enabling NFS Protocol Support.....	128

Supported clients and protocols.....	129
Supported NFS versions.....	129
NFS statistics.....	129
Unicode support.....	130
NFS and NIS unicode support.....	130
Changing the character set.....	131
Enabling and disabling file services.....	131
Configuring NFS exports.....	132
The NFSv4 pseudo file system.....	133
Kerberos configuration.....	133
Viewing NFS exports.....	135
Adding an NFS export.....	137
IP address export qualifiers.....	142
Specifying clients by name.....	143
Modifying NFS Export Details.....	144
Deleting an NFS export.....	146
Backing up or restoring NFS exports.....	146
About the rquotad service.....	147
Restrictive mode operation.....	148
Matching mode operation.....	149
Chapter 5: Using SMB for Windows access.....	150
SMB protocol support.....	150
Supported SMB versions.....	151
Supported SMB3 functionality for Hyper-V.....	154
SMB3 Multichannel support.....	154
SMB3 Encryption support.....	155
SMB3 Encryption client file access configurations.....	156
Configuring SMB security.....	157
Assigning SMB names.....	158
Viewing SMB setup.....	158
Joining an Active Directory.....	161
Removing SMB server names.....	164
Configuring local groups.....	164
Adding a local group or local group members.....	165
Deleting a local group or local group members.....	166
Local user authentication for SMB and FTP users.....	167
Using local user authentication.....	167
SID mappings.....	168
Configuration.....	169
Configuring SMB shares.....	169
Adding an SMB share.....	169

Viewing and modifying SMB shares details.....	176
Controlling access to shares using qualifiers.....	181
Controlling access to shares using permissions.....	182
Adding or changing SMB share access permissions.....	183
About Home Directories.....	184
Using home directories with cluster EVS name spaces.....	184
Offline file access modes.....	185
Backing up and restoring SMB shares.....	185
Considerations when using Hyper-V.....	186
Configuring the Service Witness Protocol.....	187
Configuring a witness EVS.....	187
Using Windows server management.....	188
Using the computer management tool.....	188
Restoring a previous version of a file.....	189
Chapter 6: Transferring files with FTP.....	190
FTP protocol support.....	190
Configuring FTP preferences.....	190
Displaying FTP users.....	192
Adding an FTP user.....	193
Importing an FTP user.....	195
Modifying FTP users.....	196
FTP statistics.....	197
Configuring FTP audit logging.....	198
Chapter 7: Block-level access through iSCSI.....	199
iSCSI support.....	199
iSCSI MPIO.....	200
iSCSI access statistics.....	201
iSCSI prerequisites.....	201
Supported iSCSI initiators.....	201
Offload engines.....	201
Configuring iSCSI.....	201
Configuring iSNS.....	202
Viewing iSNS servers.....	202
Configuring iSCSI Logical Units.....	203
Logical unit management.....	203
Logical unit security.....	204
Concurrent access to logical units.....	204
Taking snapshots of logical units.....	204
Volume full conditions.....	205
Managing iSCSI logical units.....	205

Viewing the properties of iSCSI logical units.....	205
Adding iSCSI logical units.....	207
Modifying an iSCSI logical unit.....	210
Deleting an iSCSI logical unit.....	211
Backing up iSCSI logical units.....	212
Restoring iSCSI logical units.....	212
Setting up iSCSI targets.....	213
Viewing the properties of iSCSI targets.....	213
Adding iSCSI targets.....	214
Adding a logical unit to an iSCSI target.....	217
Modifying the properties of an iSCSI target.....	219
Deleting an iSCSI target.....	221
Configuring iSCSI security (mutual authentication).....	221
Configuring the storage server for mutual authentication.....	222
Changing the storage server's mutual authentication configuration	223
Configuring the Microsoft iSCSI initiator for mutual authentication.....	225
Accessing iSCSI storage	226
Using iSNS to find iSCSI targets	226
Using target portals to find iSCSI targets.....	227
Accessing available iSCSI targets	227
Verifying an active connection.....	228
Terminating an active connection.....	228
Using Computer Manager to configure iSCSI storage.....	229
Chapter 8: Hitachi Dynamic Provisioning.....	230
HDP high-level process.....	230
Understanding HDP thin provisioning.....	231
Understanding how HDP works with HNAS.....	233

Preface

This guide explains about file system formats, and provides information about creating and managing file systems, and enabling and configuring file services (file service protocols). Note that some features apply only to individual platforms and may not be applicable to your configuration.

Virtual Storage Platform G400, G600, G800 and Virtual Storage Platform F400, F600, F800 storage systems can be configured with NAS modules to deliver native NAS functionality in a unified storage platform. The term 'NAS module' in this document also applies to VSP F series, VSP G series, and VSP N series. The unified VSP Gx00 models, VSP Fx00 models, and VSP N series models automatically form a two-node cluster in a single chassis upon installation, with no external cabling required.

Related Documentation

Release Notes provide the most up-to-date information about the system, including new feature summaries, upgrade instructions, and fixed and known defects.

Command Line References

The Command Line Reference provides information on the commands used to manage your system, and includes relevant information on the operation of your hardware and software. Depending on the model of your server or cluster node, refer to the Command Line Reference that is appropriate for your system.

- *NAS Module Server Command Line Reference*
- *Command Line Reference for models 4060, 4080, and 4100*
- *Command Line Reference for models 3080, 3090 and 4040*
- *Command Line Reference for models 5200 and 5300*

Administration Guides

- *System Access Guide* (MK-92HNAS014)—Explains how to log in to the system, provides information about accessing the NAS server/cluster CLI and the SMU CLI, and provides information about the documentation, help, and search capabilities available in the system.
- *Server and Cluster Administration Guide* (MK-92HNAS010)—Provides information about administering servers, clusters, and server farms. Includes information about licensing, name spaces, upgrading software, monitoring servers and clusters, and backing up and restoring configurations.
- *Storage System User Administration Guide* (MK-92HNAS013)—Explains user management, including the different types of system administrator, their roles, and how to create and manage these users.
- *Network Administration Guide* (MK-92HNAS008)—Provides information about the server's network usage, and explains how to configure network interfaces, IP addressing, name and directory services.
- *File Services Administration Guide* (MK-92HNAS006)—Explains about file system formats, and provides information about creating and managing file systems, and enabling and configuring file services (file service protocols).
- *Data Migrator Administration Guide* (MK-92HNAS005) —Provides information about the Data Migrator feature, including how to set up migration policies and schedules.
- *Storage Subsystem Administration Guide* (MK-92HNAS012)—Provides information about managing the supported storage subsystems (RAID arrays) attached to the server/cluster. Includes information about tiered storage, storage pools, system drives (SDs), SD groups, and other storage device related configuration and management features and functions.
- *Snapshot Administration Guide* (MK-92HNAS011)—Provides information about configuring the server to take and manage snapshots.
- *Replication and Disaster Recovery Administration Guide* (MK-92HNAS009)—Provides information about replicating data using file-based replication and object-based replication, provides information on setting up replication policies and schedules, and using replication features for disaster recovery purposes.
- *Antivirus Administration Guide* (MK-92HNAS004)—Describes the supported antivirus engines, provides information about how to enable them, and how to configure the system to use them.
- *Backup Administration Guide* (MK-92HNAS007)—Provides information about configuring the server to work with NDMP, and making and managing NDMP backups.



Note: For a complete list of Hitachi NAS open source software copyrights and licenses, see the *System Access Guide*.

Hardware References

- *Hitachi NAS Platform 3080 and 3090 G2 Hardware Reference* (MK-92HNAS017) — Provides an overview of the second-generation server hardware, describes how to resolve any problems, and replace potentially faulty parts.
- *Hitachi NAS Platform and Hitachi Unified Storage Series 4000 Hardware Reference* (MK-92HNAS030)—Provides an overview of the Hitachi NAS Platform Series 4000 server hardware, describes how to resolve any problems, and how to replace potentially faulty components
- *Hitachi NAS 5000 Series Hardware Reference* MK-92HNAS089—Provides an overview of the Hitachi NAS Platform Series 5000 server hardware, describes how to resolve any problems, and how to replace potentially faulty components.
- *Hitachi NAS Platform System Manager Unit (SMU) Hardware Reference* (MK-92HNAS065)—This document describes the usage and replacement instructions for the SMU 300/400.

Best Practices

- *Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere* (MK-92HNAS028)—This document covers best practices specific to using VMware vSphere with the Hitachi NAS platform.
- *Hitachi NAS Platform Deduplication Best Practice* (MK-92HNAS031)—This document provides best practices and guidelines for using deduplication.
- *Hitachi NAS Platform Best Practices for Tiered File Systems* (MK-92HNAS038)—This document describes the Hitachi NAS Platform feature that automatically and intelligently separates data and metadata onto different Tiers of storage called Tiered File Systems (TFS).
- *Hitachi NAS Platform Data Migrator to Cloud Best Practices Guide* (MK-92HNAS045)—Data Migrator to Cloud allows files hosted on the HNAS server to be transparently migrated to cloud storage, providing the benefits associated with both local and cloud storage.
- *Best Practices for Hitachi NAS Universal Migrator* (MK-92HNAS047)—The Hitachi NAS Universal Migrator (UM) feature provides customers with a convenient and minimally disruptive method to migrate from their existing NAS system to the Hitachi NAS Platform. The practices and recommendations outlined in this document describe how to best use this feature.
- *Hitachi Data Systems SU 12.x Network File System (NFS) Version 4 Feature Description* (MK-92HNAS056)—This document describes the features of Network File System (NFS) Version 4.
- *Hitachi NAS HDP Best Practices* (MK-92HNAS057)—This document lists frequently asked questions regarding the use of Hitachi Dynamic Provisioning.
- *Hitachi Multi-tenancy Implementation and Best Practice Guide* (MK-92HNAS059)—This document details the best practices for configuring and using Multi-Tenancy and related features, and EVS security.
- *Hitachi NAS Platform HDP Best Practices* (MK-92HNAS063)—This document details the best practices for configuring and using storage pools, related features, and Hitachi Dynamic Provisioning (HDP).

- *Hitachi NAS Platform System Manager Unit (SMU) Hardware Reference* (MK-92HNAS065)—This document describes the usage and replacement instructions for the SMU 300/400.
- *NAS Platform ICC with Cisco Nexus Reference Guide* (MK-92HNAS085)—This document describes how to configure Cisco Nexus series switches for use as an ICC (intra-cluster communication) switch.
- *File System Snapshots Operational Best Practice* (MK-92HNAS068)—This document provides operational guidance on file system snapshots.
- *Virtual Infrastructure Integrator for Hitachi Storage Platforms Operational Best Practice* (MK-92HNAS069)—This document provides operational guidance on Hitachi Virtual Infrastructure Integrator for the HNAS platform.
- *Hitachi NAS Platform Replication Best Practices Guide* (MK-92HNAS070)—This document details the best practices for configuring and using HNAS Replication and related features.
- *Hitachi NAS Platform to Hitachi Virtual Storage Platform Unified Gx00 Models Migration Guide* (MK-92HNAS075)—This best practice guide describes how to perform a data-in-place migration of the Hitachi NAS Platform and Virtual Storage Platform (VSP) Gx00 File solution to the VSP Gx00 platform.

Accessing product documentation

Product user documentation is available on the Hitachi Support Website: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The [Hitachi Support Website](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi. To contact technical support, log on to the Hitachi Support Website for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Community](https://community.hitachivantara.com) is a global online community for Hitachi customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi.

Thank you!

Chapter 1: File systems

This section describes NAS server file system characteristics.

File system access protocols

The server supports the SMB (CIFS), NFS, and FTP protocols for client file access, as well as iSCSI for block-level access to storage. All supported protocols can be enabled or disabled.

The server allows NFS, SMB, and FTP users to access the same file space; however, although iSCSI Logical Units (LUs) reside on file systems, it is not possible to access folders and files located on an iSCSI target through the server's file services (for example, SMB or NFS).

These protocols, with the exception of FTP, require a license key for activation.

File system tiers

A file system typically consists of files and directories. Data about the files and directories (as well as many other attributes) is kept; this information is the metadata. The data within the file system (both user data and metadata) is stored on the storage media of a storage subsystem.

Different storage subsystems have varying characteristics, in terms of both performance and cost. Based on their performance, storage subsystems are classified into "tiers," which are then used by administrators to manage the storage resources in the organization.

Storage pools can be created using storage from different tiers (up to two tiers are currently supported). These storage pools are called tiered storage pools. File system data (metadata and user data) may be stored in a single tier, or in multiple tiers. When file system data is stored on storage subsystems on multiple tiers, the file system is called a tiered file system.

For more information on storage tiers and tiered storage pools, refer to the *Storage Subsystem Administration Guide*.

In a tiered file system, metadata is stored on the highest performance tier of storage, and user data is stored on a lower-performance tier. A tiered file system may provide the following:

- **Performance benefits:** Storing metadata on the higher-performance tier provides system performance benefits over storing both the metadata and user data on the same, lower, tier of storage. The performance gain is seen because metadata is accessed more often than the user data, so storing it on higher-performance storage increases overall performance.
- **Reduced expenses for storage:** Storing metadata on the higher-performance storage (which is usually more expensive than the lower performance storage) and user data on lower performance (and less expensive) storage may provide cost benefits. This is because metadata typically consumes a relatively small amount of storage, while the user data consumes the bulk of the storage. Because the higher-performance storage is used only to hold the metadata, less of the expensive storage is used than if both the metadata and the user data were on the higher-performance storage. Also, because user data can be kept on lower performance storage while achieving better performance than keeping both metadata and user data on the lower performance storage, you may not have to upgrade storage as often (or you may be able to repurpose aging storage when you do upgrade).

A tiered file system has the following characteristics:

- Maintains a single file system view while providing data separation. This separation allows the file system to store file system metadata (which is critical to system performance) on very high-performance devices, while storing user data on cheaper, lower-performance storage subsystems.
- The use of multiple tiers of storage is completely transparent to applications or clients. No environmental tweaking or effort is required. All file system functionality (such as snapshots, replication, quotas, cluster name space, and virtual volumes) are preserved.
- File system management activities, such as mounting, unmounting, sharing, and exporting, for tiered file systems are the same as for untiered file systems.
- The file system block size (4 KiB or 32 KiB) is maintained across all tiers of storage.
- Cross volume links are treated as metadata.

Data spillage between tiers

It is possible for data to spill between the user data and metadata tiers in both directions.

Tier 1 data (user data) can spill over into Tier 0 (metadata). This only occurs if the Tier 1 file system is completely full, and additional data is written to the file system. Users are alerted if this type of spillage occurs, enabling them to better allocate data.

Tier 0 data (metadata) can also spill over into Tier 1 if necessary. Once there is space again on Tier 0, the NAS server returns the metadata to its original tier. If Tier 0 data spills over to Tier 1, performance can be degraded, including reduced write performance.

Use the `fs-analyze-data-usage` command to confirm data spillage.

Confining new tiered file systems

If the host span is tiered, all new file systems created on it are automatically tiered as well.

In this case, no file system can be created unless the span has enough space on both tiers. If you specify an initial size for a new tiered file system, the NAS server allocates the minimum possible amount of space from Tier 0 (metadata) and the rest from Tier 1 (user-data). Otherwise, it allocates the smallest possible amount from both tiers.

If you specify a confining capacity for a new tiered file system, it applies only to Tier 1 (user-data). If you need Tier 0 (metadata) to also be confined, use the `filesystem-confine` command to confine it manually. As user data is normally bigger than metadata, the two tiers are typically confined to different capacities.

Supported file system limits

The number of supported file systems on a NAS server depends on the model.

By default, a server can support up to 128 filesystems at a time.

The maximum number of supported file systems for each model is:

Model	Maximum number of supported file systems
4040	128
4060 / 4080	500
4100 / NAS module	500
5200 / 5300	500

To increase the number of supported file systems from the default to the maximum, contact customer support.

For further information, see the CLI command `filesystem-enable-max-count` man page.

Per-span limits

By default, there is a limit of 32 filesystems per span. If you require a greater number of filesystems per span, it is possible to increase this number using the `filesystem-create` CLI command with the `--exceed-safe-count` option. This option must not be used when creating up to 32 filesystems. It must only be used when creating filesystems beyond the 32nd one.



Note: This option is only available on the CLI. The NAS Manager does not permit you to create more than 32 filesystems.

Creating too many filesystems fills up the filesystem catalogue. The filesystem catalogue usage is dependent on the filesystem name lengths; the longer the name of the filesystem, the more space it consumes in the catalogue. Use the `span-dump` command to check the remaining space in the FS catalogue.

For example:

Cod areas:	Used	Maximum	Full
Span Cod	1049	14336	7%
Fs catalogue	3096	4096	75%
Chunk table	479	40960	1%

If the filesystem catalogue is full, this will cause the recycle bin to become unavailable. In this case, in order to delete a filesystem, it is necessary to use the `filesystem-delete --no-undeletion-information` option to bypass the recycle bin.



Note: There is a limit of 128 filesystems that can be assigned to a single EVS.

Creating a new file system

This procedure creates a new file system. A storage pool is required before a file system can be created.



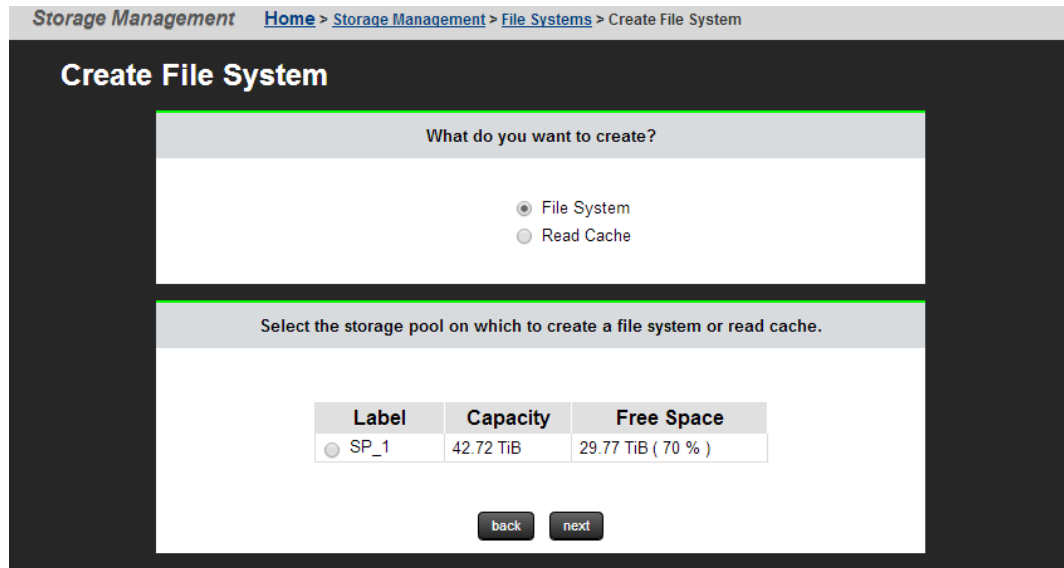
Note: If Dynamic Write Balancing is not enabled, or if your system does not support Dynamic Write Balancing, when expanding a storage pool, use as many disk drives as possible and keep SDs as large as possible to attain optimal performance. For more information on Dynamic Write Balancing, refer to the *Storage Subsystem Administration Guide*.




Note: The maximum size of a file system is 1 PiB but a 1 PiB file system is only supported on an HDP storage pool.


Procedure

1. From the Home page, navigate to **Storage Management > File Systems** to display the **File System** page.
2. Click **create** to display the **Create File System** page.



3. Click **File System**.
4. Select the required storage pool, and then click next to start the Storage Pool Wizard. For more information about the Storage Pool Wizard, refer to the *Storage Subsystem Administration Guide*.
5. Select a storage pool to contain the file system, then click **next** to display the **Create File System** page.



Field/Item	Description
Storage Pool	Displays the name of the storage pool in which the file system or read cache is being created.
Free Capacity Amount	Displays the available space in the storage pool that can be used by the file systems.
Tier 0 Meta-data and Tier 1 User-data	<p>Displays the size of the storage pool's metadata tier (Tier 0) and the user data tier (Tier 1).</p> <div data-bbox="711 533 1393 653" style="border: 1px solid #add8e6; padding: 5px;">  Note: This information applies only to tiered file systems, and is not displayed for untiered file systems. </div>
Guideline Chunk Size	Displays the approximate size of the chunks used in the selected storage pool.
Size Limit	<p>If Auto-Expansion is enabled:</p> <ul style="list-style-type: none"> ▪ For an untiered file system, this is the maximum size to which a file system will be allowed to expand. ▪ For a tiered file system this is the maximum size to which the user data tier (Tier 1) of the file system will be allowed to expand. <p>If Auto-Expansion is disabled:</p> <ul style="list-style-type: none"> ▪ Specifies the capacity with which the new file system should be created. The file system will be created and this amount of space is immediately allocated to the file system. This means that the currently unused space is reserved for this file system, and is not available for any other use.
Rounded to nearest chunk	Click to read about how the file system created and expansion is affected by rounding to the nearest chunk size.
Allocate on Demand or Allocate Now	<p>Available if the storage pool allows auto expansion.</p> <ul style="list-style-type: none"> ▪ Fill the Allocate on Demand check box to create a small file system and allow it to automatically expand to the size limit. ▪ Fill the Allocate Now check box to create the file system to the full size specified by the size limit.
Label	The label (name) by which the file system or read cache should be referenced.
Initial Capacity	The initial capacity of the file system, or the user data tier of a tiered file system. As auto expansion occurs, the file

Field/Item	Description
	system will grow to the Size Limit specified in the Size Limit field.
Assign to EVS	The EVS to which the file system or read cache is assigned.
Object Replication Target	When this check box is filled, the file system or read cache will be formatted to allow shares and exports.
Support and Enable Dedupe	When this check box is filled, the file system can support and enable deduplication on the file system.
Block Size	Sets the optimal block size for the file system or read cache.

6. Enter a **Size Limit** for the file system.
This defines the maximum size to which the file system or user data tier can grow through Auto-Expansion. Once the file system has been created, this value can be changed on the **File System Details** page. This limit is not enforced for manual file system expansions performed through the CLI.
7. The value in the **Rounded Size Limit** field is calculated automatically, but can be changed. For more information, click **Rounded to nearest chunk**. If the specified size is not a multiple of the chunk size, the server rounds down to the nearest chunk boundary.
8. Use these radio buttons to **enable** or **disable** Auto-Expansion, which allows or constrains growth of this file system or user data tier (for tiered file systems).
Be aware that storage pools can be configured to prevent the growth of file systems. A file system can never shrink; once space is allocated to a file system, the space cannot be recovered, and the file system cannot be reduced in size. When expanding, the file system will use the storage pool's chunk size as its growth increment. File systems configured to automatically expand will do so when they are about 80 percent full. File systems can be expanded manually through the CLI. File system expansion does not interrupt file services or require the file system to be unmounted.
9. In the **Initial Capacity** field, set the initial size for the file system or user data tier.
10. In the **Label** field, enter the name of the file system.
File system labels are not case sensitive, but they do preserve case (labels will be kept as entered, in any combination of upper and lowercase characters). Also, file system labels cannot contain spaces or any of the following special characters: "&*/;:<>?|. File system labels must be unique within a server or cluster. Also, a file system cannot have the same label as a storage pool.
11. From the EVS list, select the **EVS** to which the file system should be assigned.
12. Fill in the **Object Replication Target** check box if the file system is intended to be the target of an object replication. When this check box is filled, the file system will be formatted to allow shares and exports.
13. Fill in the **Support and Enable Dedupe** check box to support and enable deduplication on the file system.
14. In the **Block Size** field, enter the desired file system block size.
15. Click **OK**.

Read caches

A read cache is a special read-only file system that stores copies of individual files outside of their local file systems, enabling a server or a node to have a cached copy of the file. When NFS v2 or NFS v3 clients submit a read request for a file in the read cache, the server or node can serve the read request from the copy in the read cache. Note that a read cache does not benefit CIFS clients, and that read caches have special characteristics and limitations.

Dedupe File Systems

Deduplication is a file system feature that incorporates enhancements to the file system and the application layer. Deduplication features the ability to reduce redundancy in stored data blocks. All data in the specified file system is scanned at intervals and duplicate blocks are removed, resulting in reclaimed disk space. All dedupe activity and the elimination of redundant blocks is transparent to the user.

Base deduplication is enabled by default and does not require a license key. This is a dedupe feature with a single SHA-256 engine, capable of indexing data at a rate of up to 120 MB per second.

Premium deduplication is a licensed feature and must be installed before deduplication can be performed. This is a dedupe feature with four SHA-256 engines, capable of indexing data at a much faster rate. Contact your Hitachi representative for more information.

For license key information, see the *Server and Cluster Administration Guide*.



Note: Do not use NAS deduplication and storage-based deduplication (capacity saving) on the same LUs, as the additional processing reduces I/O performance.



Note: Deduplication support for object replication target filesystems is supported from release 13.6.

Deduplication characteristics

The deduplication feature has the following characteristics:

- Only user data blocks are deduplicated.
- Dedupe is a post-process that is performed as a fixed block-based deduplication. It is not an inline dedupe process.
- Data is deduped within a given file system and not across multiple file systems.
- Dedupe has been designed with quality of service (QoS) as a key component. File system activity takes precedence over dedupe activity when file serving load goes beyond 50 percent of the available IOPS or throughput capacity. The deduplication process throttles back to allow the additional file serving traffic load to continue without impacting performance.
- You can configure a new file system to support and enable dedupe.
- An existing WFS-2 file system can be converted to be dedupe-enabled.
- File systems with support for dedupe can be dedupe-enabled or dedupe-disabled.

Deduplication interoperability with existing NAS Platform features

The following table lists the applications that are compatible with Dedupe.

Application	Interoperability
Object Replication	File data is not retained in the deduplicated state when replicating a file system using the Object Replication option. Deduplication is supported on object replication targets from release 13.6.
File Replication	File data is not retained in the deduplicated state during file replication. The target file system can have dedupe-enabled, in which case the files on the target will be eventually deduped.
Snapshots	Deduped file systems lack exact physical snapshot block usage reporting. There is no way to determine how much physical space is taken up by a given snapshot and, thus, how much will be freed up when they delete a snapshot. Additionally, running <code>kill-snapshots</code> on a deduped file system will leak all snapshot blocks and you will need to take further measures to reclaim the leaked space. <code>snapshot-delete-all</code> is the preferred tool for deleting all snapshots as it does not require the file system to be unmounted, no space is leaked, and it does not affect checkpoint selection.
Data Migration	With both local and external migrations, migrated files will be rehydrated.
NDMP file-based backup	Files to be deep copied during a NDMP file based backup are restored to single files from tape; a NDMP recovery however cannot be deduped supported or enabled.
Sync Image backup	Sync Image backups do not retain their deduplicated state.
Quotas	Quota calculations are based on the length of the file (size-based quotas), or the rehydrated usage of a file (usage-based quotas).
Tiered File Systems	User data can be deduplicated; however, metadata is not.

Calculating deduplication space savings

The following example describes how the deduplication space savings are calculated.

If the difference between physical and logical space of 100 TB of data before deduplication is as follows:

- Group A: 30 TB of distinct data
- Group B: 70 TB of duplicated data that contains only 10 TB of unique data blocks.
 - Given an arbitrary data block in Group B, there may be one or more identical data blocks in Group B, and not in Group A, but an arbitrary data block in Group A has no identical data block in either groups.

If both Group A and Group B have gone through the dedupe process:

- Group A had no duplicates removed and consumed the same 30 TB.
- Group B had duplicates removed and consumed only 10 TB to hold the unique data blocks.
- Group B (70 TB) = {Group C (10 TB raw remaining)} + {Group D (60 TB deduped and now sharing or pointing to physical blocks of group C)}
- The original 100 TB of data now requires only 40 TB (30 plus 10) of physical blocks because all duplicates were removed. However, the logical data size is 100 TB (30 plus 70), which is the amount of space needed if the data were not deduped. The results are outlined in the following table:

Used Space	The amount of physical disk space used by the file system, in this example, group A and group C = 30 + 10 = 40 TB
Deduped space	The amount of duplicate data that does not occupy its own physical disk space, but has been deduped to share existing physical blocks = group D = 60 TB
Logical space	The amount of physical disk space that would be required if the data were not deduped = {used space} + {deduped space} = 40 + 60 = 100 TB

Based on the example presented, the dedupe percentage gives the amount of physical disk space saved by removing the duplicates. The percentage measures against the amount of space that would be required if the data were not deduped.

$$\text{Dedupe Percentage} = \frac{\text{Dedupe}}{\text{Logical}} = \frac{\text{Dedupe}}{\text{Used} + \text{Dedupe}} = \frac{60}{40 + 60} = 60\%$$

Viewing deduped file system usage

The `df` command reports deduplication information for a file system. The Deduped column reports the amount of data that has been deduped in the file system, which is the amount of disk space that was saved by using this feature.

For example:

```
Example df output:
(MMB):$ df -f m9fs1
```

ID	Label	Size	Used	Snapshots	Deduped	Avail	Thin	ThinSize	ThinAvail	FS Type
1040	m9fs1	1.27 TB	561 GB (43%)	NA	750 GB (57%)	737 GB (57%)	No			4 KB,WFS-2,128 DSBs,dedupe enabled

- All columns except Snapshots and Deduped have the same meaning as a normal file system:
 - Size column: The formatted capacity of the file system.
 - Used column: The amount (and percentage) of formatted capacity used by live and snapshot data.
 - Avail column: The amount of formatted capacity available for further data.
- **Deduped column**
 - This column reports the amount (and percentage) of deduped data in the file system.
- **Snapshots column**
 - This column normally reports the amount (and percentage) of logical space used by snapshots.
 - On file systems that do not support dedupe, the logical size of data is equal to the physical space used by that data. However, on file systems that support dedupe, the logical space used by snapshots can exceed the physical used size due to block sharing through Dedupe.
 - In some cases, snapshot space usage can even exceed the total formatted capacity of file system size. To avoid confusion, the Snapshots column displays NA for file systems that support dedupe.

Chapter 2: Managing file systems

This section describes how to manage NAS server file systems.

Viewing available file systems

You can view the available file systems using the NAS Manager.


Procedure

1. Navigate to **Home > Storage Management > File Systems** to display the **File Systems** page.

Label	Total	Used (%)	Used	Free	Storage Pool	Status	EVS
<input checked="" type="checkbox"/> data2	3.97 GiB	53%	2.12 GiB	1.85 GiB	sp1	Mounted	g5-evs2
<input type="checkbox"/> FileSystem2	3.97 GiB	54%	2.14 GiB	1.83 GiB	sp1	Mounted	g5-evs1
<input type="checkbox"/> FS	3.98 GiB	54%	2.15 GiB	1.83 GiB	sp1	Mounted	g5-evs1
<input type="checkbox"/> FS2	39.82 GiB	6%	2.48 GiB	37.34 GiB	sp1	Mounted	g5-evs1
<input type="checkbox"/> g5-fs1	200 GiB	2%	3.50 GiB	196.50 GiB	sp1	Mounted	g5-evs1
<input type="checkbox"/> TestFileSystem	3.75 GiB	64%	2.41 GiB	1.34 GiB	sp1	Mounted as Object Replication target	g5-evs1

Field/Item	Description
Filter	Click to open the Filter dialog box and enter one or more of the following filtering criteria: <ul style="list-style-type: none">▪ File System▪ Storage Pool▪ Status▪ EVS

Field/Item	Description
	Click OK to save the filter settings. File Systems that meet the specified criteria are displayed on the page. Click reset to remove all filter settings.
Label	Name of the file system, assigned upon creation and used to identify the file system when performing particular operations; for example, creating an export or taking a snapshot.
Total	Size of the file system.
Used (%)	<p>Percentage amount of space used by the file system.</p> <p>The following Usage Alerts criteria apply:</p> <ul style="list-style-type: none"> ▪ If Usage Alerts are enabled on the entire file system, the sliding bar turns yellow when the warning limit is exceeded and orange when the severe limit is exceeded. ▪ If Usage Alerts are not enabled, the sliding bar turns yellow when 85% capacity is reached and orange when the file system is full.
Used	Amount of space used by the file system.
Free	Amount of free space available on the file system.
Storage Pool	Name of the storage pool on which the file system resides.
Status	<p>The status can be one of the following:</p> <ul style="list-style-type: none"> ▪ Checking: The file system is being checked; during this check, the approximate percentage of completion for each phase of the check (note that some phases take longer than others, and that the percentage complete is only for the current phase, not for the overall procedure). ▪ Failing: The file system has failed, but is being checked, fixed, or recovered. ▪ Fixing: The file system is being repaired. When fixing, an approximate percentage complete is displayed for each phase of the repair (note that some repair phases take longer than others, and that the percentage complete is only for the current phase, not for the overall procedure). ▪ Formatting: The file system is being formatted. ▪ Initializing System Drive: The system drive is initializing. ▪ Mounted: The file system is mounted and is available for service. ▪ Mounted as Readonly: The file system is mounted, but is in read-only mode.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Mounted as Object Replication target: The file system is mounted and has been formatted as an object replication target. ▪ Mounting: The file system is being mounted and available for service. ▪ Not Assigned to EVS: The file system is not currently assigned to an EVS. ▪ Not Available for Mounting: The file system is not available. Make sure to enable the EVS to which the file system is assigned and make sure the file system is not marked "hidden". ▪ Not Mounted: The file system is not mounted. ▪ Not Mounted (System Drive initialization status unknown): The file system is not mounted, and the SD initialization status is not known. ▪ Not Mounted (System Drive is not initialized): The file system is not mounted. The SD initialization status is known, but SD initialization is not complete. ▪ Syslocked: The file system is syslocked.
EVS	EVS to which the file system is assigned.
details	Displays the File System Details page for the selected file system.
mount	<p>Select one or more unmounted file systems and click mount to mount the file system.</p> <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note: The server remembers which file systems were mounted when it shuts down, and mounts them automatically during system startup.</p> </div>
unmount	Select one ore more mounted file system and click unmount to unmount the file system.
create	Click to advance to the Create file system page.
Download File Systems	Downloads a spreadsheet containing information about all of the listed file systems.
System Drives	Displays the System Drives page.
Quotas by File System	Displays the Quotas by File Systems page.
Storage Pools	Displays the list of storage pools on the server.

Field/Item	Description
Active Tasks	Displays more information about active tasks on the Active Tasks page.
Dedupe File Systems	Displays the Deduplication page.

Displaying file system details



You can view file system details using the NAS Manager.


Procedure

1. Navigate to **Home > Storage Management > File Systems**.
2. Select a file system and click **details** to display the **File System Details** page.

The following table describes the fields in this page:


Field/Item	Description
Settings/Status	
Label	Name of the file system, assigned upon creation and used to identify the file system when performing particular operations; for example, creating an export or taking a snapshot. To change the name of the file system, enter a new name and click rename .
Capacity	<p>Displays information about the space allocation and usage for the file system. For a tiered file system, information about both the metadata and user data tiers is displayed. For an untiered file system, information about the total size of the file system is displayed.</p> <ul style="list-style-type: none"> ▪ % Total Used Space: Percentage of the file system's total allocated space that has been used. This total reflects data and snapshots, if any. ▪ Capacity: Total amount of formatted space (free + used space). ▪ Free: Total amount of file system space unused (free), in GiB or TiB and as a percentage of the total. ▪ Total Used: Total amount of file system space in use, in GiB or TiB and as a percentage of the total.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Expansion Limit: Defines the size limit up to which a file system can expand if auto-expansion is allowed, provided it accommodates the chunk size. ▪ Live File System: Total space used by the file system data, in GiB or TiB and as a percentage of the total. ▪ Snapshots: Total space used by snapshots of the file system, in GiB or TiB and as a percentage of the total. Where no snapshots exist, this is reported as '0 Bytes'. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;">  Note: Snapshot capacity is not shown for dedupe-enabled file systems. </div>
Tier 0 Meta-data and Tier 1 User-data	<div style="background-color: #e0f2f1; padding: 5px; margin-bottom: 10px;">  Note: These areas are displayed only for tiered file systems. </div> <p>These areas display information about the space allocation and usage for the tiers making up the file system. The Tier 0 Meta-data section describes information about the metadata tier. The Tier 1 User-data section describes information about the user data tier.</p> <ul style="list-style-type: none"> ▪ % Total Used Space: Percentage of the file system's total allocated space that has been used. This total reflects data and snapshots, if any. ▪ Capacity: Total amount of formatted space (free + used space). ▪ Free: Total amount of file system space unused (free), in GiB and as a percentage of the total. ▪ Total Used: Total amount of file system space in use, in GiB and as a percentage of the total. <p>If you change any of the auto-expansion limits, click apply to make the changes effective.</p>
Configuration	
Status	Current status of the file system, showing whether the file system is mounted or unmounted.
Syslock	<p>Indicates whether the file system is in Syslocked mode (System Lock enabled), or if the file system is not in Syslocked mode (System Lock disabled).</p> <p>When System Lock is enabled for a file system, NDMP has full access to the file system and can write to it during a backup or replication, but the file system remains in read-only mode to</p>

Field/Item	Description
	<p>clients using the file service protocols (NFS, CIFS, FTP, and iSCSI).</p> <p>To enable/disable the System Lock for a file system, click enable or disable. When viewing the details of a read cache, the System Lock's enable/disable button is not available.</p>
Object Replication Target	<p>Indicates whether or not the file system is formatted as an object replication target. A file system must be formatted as an object replication target in order to be the target of a replication.</p>
Transfer Access Points During Object Replication	<p>Indicates whether or not the file system is enabled to allow transfer access points (shares and/or exports) during an object replication. If disabled, click enable to allow the file system to transfer access points during an object replication. If enabled, click disable to prohibit the transfer of access points during an object replication.</p>
Transfer XVLs as Links During Object Replication	<p>Indicates whether or not the file system is enabled to transfer migrated files as links during an object replication, rather than transferring the file contents. If disabled, click enable to select this option. It can only be enabled when the file system is unmounted.</p> <div data-bbox="667 1010 1393 1094" style="background-color: #fff9c4; padding: 5px; border: 1px solid #ccc;"> <p> Caution: Once enabled, this option cannot be disabled for the file system.</p> </div>
Deduplication	<p>Indicates whether the file system is converted (supports deduplication) or not converted (does not support deduplication). Clicking on the status displays the Deduplication page.</p>
Thin Provisioning	<p>Indicates whether thin provisioning is enabled or disabled.</p>
EVS	<p>EVS to which the file system is assigned.</p> <p>If the file system is not currently assigned to an EVS, a list of EVSs (to which the file system can be assigned) appears. Select an EVS from the list and click assign.</p>
Security Mode	<p>Displays the file system security policy defined for the file system. Possible values are <code>Mixed (Windows and Unix)</code> or <code>Unix</code>. This may be followed by <code>(Inherited)</code>, where the security mode is inherited from the EVS security mode and has not been manually changed. Clicking on the status displays the File System Security page.</p>
Block Size	<p>File system block size: 32 KiB or 4 KiB, as defined when the file system was formatted.</p>

Field/Item	Description
Read Cache	Indicates whether this file system is a read cache (Yes) or a regular file system (No).
WFS Version	Indicates the file system format.
Usage Thresholds	
File System Usage	<p>Usage thresholds are expressed as a percentage of the space that has been allocated to the file system.</p> <p>When a threshold is reached, an event is logged and, depending on quota settings, an email may be sent. The Current line displays information on current usage for each of the following:</p> <ul style="list-style-type: none"> ▪ Live file system: Percentage threshold for space used by data. ▪ Snapshots: Threshold for file system snapshots. ▪ Entire file system: Threshold for the total of the live file system data and snapshots. <p>You can use the edit boxes to specify the Warning and Severe thresholds:</p> <ul style="list-style-type: none"> ▪ The Warning threshold should be set to indicate a high, but not critical, level of usage. ▪ The Severe threshold should be set to indicate a critical level of usage, a situation in which an out-of-space condition may be imminent. <p>You can define both Warning and Severe thresholds for any or all of the following:</p> <ul style="list-style-type: none"> ▪ Live file system (data). ▪ File system snapshots. ▪ Total of the live file system and snapshots. <p>Click apply to save the new settings.</p> <p>To ensure that the live file system does not expand beyond its Severe threshold setting, fill the Do not allow the live file system to expand above its Severe limit check box.</p>
Associations	
Storage Pool	The name of the storage pool in which the file system or read cache was created.

Field/Item	Description
	<p>This area also displays the following information:</p> <ul style="list-style-type: none"> ▪ Capacity: The total space allocated to the storage pool. ▪ Free: The total storage pool free space, in MiB, GiB or TiB, and as a percentage of the total. ▪ Used: The total storage pool used space, in MiB, GiB or TiB, and as a percentage of the total.
<p>Related File Systems</p>	<p>Displays the name of any related file systems. A related file system is one that is either the:</p> <ul style="list-style-type: none"> ▪ Source of a migration or replication operation where this file system was the target of the operation. ▪ Target of a migration or replication operation where this file system was the source of the operation. <p>If there are related file systems, the start date/time of the most recent associated migration/replication will be displayed in the Last: field.</p>
Check/Fix	
<p>Status</p>	<p>File system status messages may be any of the following:</p> <ul style="list-style-type: none"> ▪ File system is not being checked or fixed. ▪ Checking. ▪ Fixing. <p>You can start a check of the file system, or just part of the file system, using the Scope settings and the browse and check buttons.</p> <p>If one or more file system check/fix operations are in progress, click Active Tasks.</p>
<p>Scope</p>	<p>The scope controls allow you to set the scope of a check by the entire file system or the directory tree.</p> <p>To check the whole file system, click the Entire File System radio button. This option is only available when the file system is not mounted.</p> <p>To check a part of the file system, click the Directory Tree radio button, then use the browse button to navigate to the part of the file system you want to check.</p> <p>Once you have set the scope, click check to start the check.</p> <p>The Cancel button requests that checkfs/fixfs be aborted. This is not a forceful cancellation, so the check/fix operation may not be aborted immediately.</p>

Field/Item	Description
mount	Use to mount the file system. <div style="border: 1px solid black; background-color: #e0f2f1; padding: 5px; display: inline-block;">  Note: The server remembers which file systems were mounted when it shuts down, and mounts them automatically during system startup. </div>
unmount	Use to unmount the file system.
format	Advances to the Format File System page. Not available when the file system is mounted.
delete	Use to delete the file system. Not available when the file system is mounted.
expand	Use to allocate more space to the file system.
Data Migration Paths	Advances to the Data Migration Paths page.
File System Versions	Advances to the File System Versions page.

Formatting a file system

Formatting a file system prepares it for use by clients for data storage. File systems created through the NAS Manager are formatted and mounted automatically. Therefore, this procedure should rarely, if ever, be used. This procedure assumes that the file system has already been mounted.

Procedure

1. From the Home page, navigate to **Storage Management > File Systems** to display a list of all file systems.
2. For the file system to be formatted, click **details**.
3. If the file system is mounted, do the following to unmount it:
 - a. In the Label column, select the file system.
 - b. In the Actions section, click **unmount**.
 - c. In the confirmation dialog, click **OK**.
4. Click **format** to display the **Format File System** page.
5. Use the radio buttons to select **32 KiB** or **4 KiB** as the block size for the file system.

6. If the file system will be the target of an object replication, select the **Object Replication Target** checkbox. If the file system will not be the target of an object replication, make sure the **Object Replication Target** checkbox is cleared.

A file system must be formatted as an object replication target in order to be the target of a replication. When this check box is selected, the file system will be formatted to allow shares and exports.

7. Click **OK** to format the file system and return to the **File System** details page.

Standard bitmap support

NAS file systems were traditionally formatted with enhanced bitmap resiliency to provide a greater chance of recovering from defects in a storage subsystem. In general, a free space bitmap with enhanced resiliency provides a greater chance of recovery from system errors; however, some systems, such as those using HDP storage, do not benefit from enhanced bitmap resiliency. Formatting them with standard bitmap resiliency can offer a performance advantage—the average number of back-end write operations is reduced from ~6 to ~2, for each client write.

With version 12.6 and later, all file systems are formatted by default:

- With standard bitmap resiliency on HDP storage
- With enhanced bitmap resiliency on non-HDP storage



Note: Standard bitmap formatting is only supported on WFS-2 file systems. Additionally, versions 12.5 or earlier do not support standard bitmap resiliency; therefore, before downgrading to version 12.5 or earlier, all file systems with standard bitmap resiliency must be converted to enhanced bitmap resiliency, or they will not mount correctly after the downgrade.

There is no special configuration needed to use this feature. File systems on HDP storage are automatically formatted with standard bitmap resiliency. Files on non-HDP storage are formatted with enhanced bitmap resiliency as before.

Two commands are provided to perform bitmap resiliency conversion:

`fs-convert-to-standard-bitmap-resiliency`

`fs-convert-to-enhanced-bitmap-resiliency`

When converting a file system from a resilient to standard format (or vice versa), note that:

- The file system must be unmounted before converting bitmap resiliency.
- Bitmap resiliency is reversible. You can undo the conversion by running one of the commands mentioned above.
- If the bitmap resiliency conversion fails, such as due to a power failure, the file system reverts to its state before the command was initiated, as if the command were never run. The file system is unchanged if the conversion fails.
- To use the **`fs-convert-to-enhanced-bitmap-resiliency`** command, the free space needed is about 8 times the data length of the free space bitmap. For the **`fs-convert-to-standard-bitmap-resiliency`** command, it is about 4 times.

Mounting a file system

Use this procedure to manually mount a file system. Mounting a formatted file system makes it available to be shared or exported, and thus accessible for use by network clients. This procedure may also be used when an auto-mount of a file system has failed, which can occur in the following situations:

- The file system was not mounted when the server was shut down.
- The command line interface was used to disable auto-mounting.
- A storage system failure caused the server to restart.

You can use the NAS Manager to mount a file system.

Procedure

1. Navigate to **Home > Storage Management > File Systems** to display a list of all file systems.
2. Select the check box next to the label of the file system to be mounted.
3. If the file system is unmounted, click **mount**.

Cloning files and directory trees

Two types of cloning are available:

- File cloning
- Directory tree cloning

The **file clone** feature allows for the quick copying of a file, which is space efficient when the source file and its copies ("clones") do not diverge significantly from each other over time. This feature is intended for situations in which long-lived and mostly unchanging copies of a large master file are required (for example, deploying many copies of an OS disk image for virtualized desktops). When a file clone is created, a snapshot of the master file is created, preserving an image of the source file at that point in time. The snapshot is only of the file, it is accessible only by the NAS server, and it is read-only. This snapshot allows the source file and clone to share common data blocks. If additional clones of the master file are created, new snapshots are created if necessary.

The maximum number of clones is limited by the availability of space in the file system. File cloning is most space efficient when related clones do not diverge significantly from each other. File clones are less space efficient when some clones are short-lived and others are long-lived, because the data blocks that were once shared between the source (master) file and all related clones are not freed until all related clones are deleted.

The **directory tree cloning** feature provides a way to quickly create space-efficient writeable copies of a directory tree. Space efficiency is achieved by creating clones of the source files in the destination tree. Space efficiency is reduced over time, as the source files and clones diverge. The source directory's structure is recreated in the destination directory, and all of the contained files are cloned. Security information from the source files and directories is applied to the destination.



Note: Some types of files in a directory tree are not cloneable. For example, files that are not regular (such as sockets, FIFOs, block special devices, and character devices) are not cloneable, and links such as hard links or cross volume links are not cloneable.



Note: It is not possible to tree-clone a snapshot directory.

During a directory tree cloning operation, timing can have an effect on what is cloned, because the directory tree is not protected against modification during the clone operation. The cloned directory tree is not a point-in-time replica (like a snapshot), because the source tree is online and may be in use during the clone operation, so the directory tree (or the files in the tree) may be modified while the tree clone operation is in progress. When a directory tree is modified during a tree clone operation, none, some, or all of the modifications may be included in the clone. To ensure that the directory tree is cloned precisely as it is at the time the clone operation is initiated (a consistent copy), you must ensure that the directory tree is not modified until after the clone operation is complete.

Tree cloning uses the same mechanism as file cloning to clone individual files within the tree, so the same limitations apply, and a File Clone license is required to enable file or directory tree clones features.



Notes: File cloning is supported on all NAS Platform models.

After this feature has been used, certain downgrade restrictions will apply. If you downgrade to a release that does not support file clones, file systems on which file cloning has been used cannot be mounted. Contact customer support for more information about downgrade restrictions when using file clones.

Cloning symlinks in directory trees

When a symlink is encountered during the cloning of a directory tree, only the symlink itself is cloned, the target of the symlink is not cloned. The symlink is cloned without change, and the path contained by the symlink is not altered in any way.

When a symlink is cloned, if the symlink contains:

1. A relative path to a directory or a file within the directory being cloned, the cloned symlink will function just as it did before being cloned. The path resolves correctly because the target of the symlink is also cloned by the tree clone operation.
2. A relative path to a directory or a file outside the directory being cloned, then target path contained by the cloned symlink will resolve correctly only if a target of the same name exists in the same location relative to the cloned symlink location.
3. An absolute path to a directory or a file (a global symlink), the path contained in the copied symlink should resolve correctly.

File clone commands

Currently NAS Manager does not include support for file cloning; to use this feature, you must use CLI commands. See the following CLI man pages for detailed information on configuring and using the file clone feature:

- **file-clone**
Provides a description of the file cloning mechanism.
- **file-clone-create**
Creates a new file, and makes its data stream a clone of the source file's data stream.
- **file-clone-stat**
Displays information about a clone object; specifically, the handles of the snapshot file objects on which the clone object depends.
- **file-clone-report-block-sharing**
Reports on the amount of block sharing between a clone and each of its predecessor snapshot file objects.
- **file-clone-declone**
Removes the dependency of a cloned file from its snapshot object, converting the clone into a normal file.
- **file-clone-stat-snapshot-file**
Displays information about a snapshot file object based on the file handle or object number of the snapshot file object.

Tree clone commands

Currently NAS Manager does not include support for tree cloning; to use this feature, you must use CLI commands. See the following CLI man pages for detailed information on configuring and using the tree clone feature:

- **tree-clone-job-submit**
Submits a request to clone a directory tree.
- **tree-clone-job-status**
Gets the status of a tree cloning job specified by its job id.
- **tree-clone-job-list**
Lists the status of tree cloning jobs.
- **tree-clone-job-abort**
Aborts a tree cloning job specified by its job id.

Deleting a tree directory with tree-delete

The `tree-delete` feature provides a mechanism to immediately remove a directory tree from its position in the file system and to perform the deletion as a background job. A directory tree consists of a specified directory and the hierarchy of subdirectories and files below it. When a directory tree is targeted for deletion, a `tree-delete` job is created and added to the job queue. The targeted directory tree is immediately removed from the file system namespace, moved to the system trash directory, and scheduled for background deletion.

The `tree-delete` interfaces are provided in the form of management APIs and CLI commands. No license is required.

The `tree-delete` feature provides the following benefits, compared to deleting a directory tree via a network client:

- The instantaneous removal of a directory tree from the listing of the parent directory, allowing the client to proceed with further actions.
- The server-side delete eliminates the need for a client to recursively delete the directory tree over the network, therefore using less system resources.
- The multi-threaded implementation allows parallel deletion of the contents of the directory tree.

Important considerations when using tree-delete



Note: The `tree-delete` feature can destroy user data.

Although the targeted directory tree is immediately removed from the file system namespace and the listing of the parent directory, the client continues to have access to parts of the directory tree it had acquired access to prior to deletion, until these parts are actually deleted in the background. Users should be mindful of the fact that changes to a directory tree (creation and deletion of files/directories) after submission for deletion are detected by `tree-delete`, and all such new content is deleted before `tree-delete` considers its job done.



Note: Quotas only reflect the physical deletion happening in the background.

Unmounting a file system and tree-delete

When a file system is unmounted, all related tree-deletion activity is suspended and resumed on a subsequent re-mount.

Tree-deletion activity on other mounted file systems remains unaffected.

Undeleteable directories

The following directories cannot be deleted using `tree-delete`:

- Root directory and system directories.
- Virtualization root or its sub-directory.
- Regular directory containing virtualization root(s). See the warning on deletion attempts of such a directory tree in the `tree-delete-job-submit` man page
- Sub-directory of a virtual volume.



Note: You can run `logtrace dump tree-delete` for more details about deleted files.

Using `tree-delete`

`tree-delete` is implemented with the commands:

- `tree-delete-job-abort`
- `tree-delete-job-list`
- `tree-delete-job-reschedule`
- `tree-delete-job-status`
- `tree-delete-job-submit`

Please see the man pages for details.

Submitting a `tree-delete` job

Note the following factors when using the `tree-delete-job-submit` command:

- The main activities pertaining to a submitted `tree-delete` job, such as its start and end, are logged to the event log..
- A maximum of 160 jobs can be handled by the system at any given time.

Troubleshooting `tree-delete`

Run `logtrace dump tree-delete` for details about deleted files. Contact customer support if necessary.

Controlling file system space usage

The server can monitor space allocation on a file system and trigger alerts when pre-set thresholds are reached; optionally, users can be prevented from creating more files once a threshold has been reached. Alternatively, the file system can be expanded either manually or automatically while online. The command `fs-usage` controls the monitoring. See the man pages for details.

Two activities consume system space:

- **Live file system.** Refers to the space consumed when network users add files or increase the size of existing files.
- **Snapshots.** Refers to consistent file system images at specific points in time. Snapshots are not full copies of the live file system, and snapshot sizes change depending on the live file system. As the live file system uses more space, snapshots use more space, and as the data in the live file system is changed, snapshots require less space.



Note: Deleting files from the live file system may increase the space taken up by snapshots, so that no disk space is actually reclaimed as a result of the delete operation. The only sure way to reclaim space taken up by snapshots is to delete the oldest snapshot.

The server tracks space taken up by:

- The user data in the live file system
- The file system metadata (the data the server uses to manage the user data files)
- Snapshots
- Entire file system

For each of these slices, both a warning and a severe thresholds can be configured. Although they differ from system to system, the following settings should work in most cases:

	Warning	Severe
Live file system	70%	90%
Snapshots	20%	25%
Entire file system	90%	95%

When the storage space occupied by a volume crosses the warning threshold, a warning event is recorded in the event log. When the Entire File System Warning threshold has been reached, the space bar used to indicate disk usage turns yellow.

When the space reaches the severe threshold, a severe event is recorded in the event log, generating corresponding alerts. If the Entire File System Severe threshold has been reached, the space bar used to indicate disk usage turns amber.

If file system auto-expansion is disabled, you can limit the growth of the live file system to prevent it from crossing the severe threshold, effectively reserving the remaining space for use by snapshots. To limit the live file system to the percentage of available space defined as the severe threshold, fill the Do not allow the live file system to expand beyond its Severe limit check box on the **File System Details** page.

File system utilization recommendations

The recommendations are structured to take into consideration file systems of various sizes and uses. The recommendations are broken into the following components:

- Type of file system
- Recommended maximum file system utilization
- Recommended file system thresholds

Archive file systems

Archive file systems are defined as file systems that maintain a data set for an extended period of time that has little to no change during that life time. This type of access pattern allows it to be utilized at very high levels.

Recommendation: The file system should be maintained at a usage level no higher than 97%. The Entire File System usage thresholds are recommended to be set at the following levels:

Warning	90%
Severe	97%

High activity file systems

High activity or high churn file systems are defined as file systems that have a high rate of data being accessed, deleted and created. Due to the workload type and to maintain a high level of write performance, sufficient free space is required. These amounts can vary based on file system size. The following recommendations take into account file system size.

- File system size range < 1 TiB
 - Recommendation: The file system should be maintained at a usage level no higher than 80%. The Entire File System usage thresholds are recommended to be set at the following levels:

Warning	User Definable*
Severe	80%

* User Definable: Choose a value that provides sufficient time to increase file system capacity.

- File system size range 1 TiB < 10 TiB
 - Recommendation: The file system should be maintained at a usage level no higher than 85%. The Entire File System usage thresholds are recommended to be set at the following levels:

Warning	70%
Severe	85%

- File System size range > 10 TiB
 - Recommendation: The file system should be maintained at a usage level no higher than 90%. The Entire File System usage thresholds are recommended to be set at the following levels:

Warning	80%
Severe	90%

Dynamic Superblocks (DSB)

The file system maintains a history of file system checkpoints known as Dynamic Superblocks. If the end user requires fast reclamation of free space after data deletions, the DSB count could be reduced to 2 for file systems <10TiB and 16 for file systems >10TiB. The default number of DSBs is 128. You can specify the setting at format time or change it at a later time by issuing the following command:

```
fs-set-dsb-count <file system> <dsb count>
```

Example:

To change the DSB count of "fs1" to two DSBs:

```
fs-set-dsb-count fs1 2
```

Note that changing the number of DSBs requires that the file system be unmounted.

Increasing the size of a file system

There are two methods to expand the amount of storage allocated to a file system:

- Manual expansion

Manually expanding a file system allows you to add storage capacity to a file system (or a tier of a tiered file system) immediately. You specify the new size of a file system, and the storage is allocated immediately. The maximum size that a file system or tier can attain is specified, and the file system size can be set to the maximum size supported by the storage pool in which the file system was created.

- Automatic expansion

File system auto-expansion allows a file system to grow to by adding chunks of storage on an as-needed basis, as long as the confinement limit or the maximum file system size has not been reached. For tiered file systems, auto-expansion can be applied independently to one or to all tiers of the file system, allowing one tier to expand independently of another.

When auto-expansion is enabled, and the file system (or tier) reaches approximately 80 percent of its allocated capacity, one or more additional chunks are allocated (refer to the *Storage Subsystem Administration Guide* for a discussion of chunks). The maximum size that a file system can attain can be specified, or the file system size can be allowed to grow to the maximum size supported by the storage pool in which the file system was created.



Note: Once storage is allocated to a file system, that storage becomes dedicated to that file system, meaning that once a file system is expanded, its size may not be reduced. Unused space in the file system cannot be reclaimed, allocated to another file system, or removed. To reclaim the storage space, the file system must be relocated to different storage or deleted.

Increasing the amount of storage allocated to a file system (manually or automatically) does not require that the file system be taken offline.

Thin provisioning file systems

Thin provisioning is a method of controlling how a file system's free space is calculated and reported. Administrators use thin provisioning to optimize the utilization of storage and to plan resource acquisition in a way that helps minimize expenses, while ensuring that there is enough storage for all the system needs.

Thin provisioning allows you to oversubscribe the storage connected to the storage server. As long as the available storage is not completely allocated to file systems, the oversubscription cannot be noticed by storage system users.

When thin provisioning is enabled and storage is oversubscribed, if a client attempts a write operation and there is insufficient storage space, the client will receive an insufficient space error, even though a query for the amount of free space will show that space is still available. When storage is oversubscribed, the storage server's administrator must ensure that this situation does not occur; the storage server does not prevent this situation from occurring. To resolve this situation, the storage server's administrator must either disable thin provisioning or add storage.

When thin provisioning is enabled, the storage server reports the amount of free space for a file system based on the file system's expansion limit (its maximum configured capacity), rather than on the amount of free space based on the amount of storage actually allocated to the file system. Because file systems can be allowed to automatically expand up to a specified limit (the expansion limit), additional storage is allocated to the file system as needed, instead of all the storage being allocated to the file system when it is created.

For example, a file system has an expansion limit of 20 TB, with 6 TB already used and 8 TB currently allocated. If thin provisioning is enabled, the server will report that the file system has 14 TB of free space, regardless of how much free space is actually available in the storage pool. For more information about storage pools, refer to the *Storage Subsystem Administration Guide*. If thin provisioning is disabled, the server will report that the file system has 2 TB of free space.

By default, thin provisioning is disabled for existing file systems and for newly created file systems. Enable and disable thin provisioning using the `filesystem-thin` command (currently there is no way to enable or disable thin provisioning using NAS Manager).

Thin provisioning works on a per file system basis, and does not affect the capacity reported by the `span-list --filesystems` and `filesystem-list` commands. Also, NAS Manager displays the actual file system size. As a result, the administrator can perform proper capacity planning.

When enabled, thin provisioning information is returned by the following CLI commands:

- `cifs-share list`
- `df`
- `filesystem-limits`
- `filesystem-list v`
- `fs-stat`
- `nfs-export list`
- `query`

For more information about CLI commands, refer to the *Command Line Reference*.

If thin provisioning is enabled and you disable file system auto-expansion for a storage pool, the free space reported for each of the file systems in that storage pool is the same as if thin provisioning were not enabled. This means that the free space reported becomes equal to the difference between the file system's current usage and the amount of space in all storage pool chunks currently allocated to that file system. If you re-enable file system auto-expansion for file systems in the storage pool, free space is again reported as the difference between the file system's current usage and its expansion limit, if an expansion limit has been specified.

When thin provisioning is enabled, and the aggregated file system expansion limits of all file systems exceeds the amount of storage connected to the server/cluster, warnings are issued to indicate that storage is oversubscribed. These warnings are issued because there is an insufficient amount of actual storage space for all file systems to grow to their expansion limit.

Managing file system expansion

File system growth management strategies can be summarized as follows:

- **Auto-expansion enabled, but not confined.** The file system is created with a defined size limit, and a small amount of that space is actually allocated when the file system is created. The file system is then allowed to expand automatically (auto-expansion enabled) until the storage pool hosting the file system is full (auto-expansion is not confined), as long as the file system expansion will not cause the file system to exceed the maximum allowable number of chunks in a file system.
- **Auto-expansion enabled, and confined.** The file system is created with a defined size limit, and a small amount of that space is actually allocated when the file system is created. The file system is then allowed to expand automatically (auto-expansion enabled) to the defined size limit (auto-expansion is confined), as long as there is available space in the storage pool and the file system expansion will not cause the file system to exceed the maximum allowable number of chunks in a file system.
- **Auto-expansion disabled.** The file system is created with the full amount of the specified size, and is not allowed to expand automatically (auto-expansion disabled).



Note: The size of a file system cannot be reduced.

File System Type	Auto-Expansion Enabled	Auto-Expansion Disabled
Untiered	<p>If auto-expansion is not confined, the size limit is ignored. The file system will be allowed to expand until the storage pool is full.</p> <p>If auto-expansion is confined, the size limit defines the maximum size to which a file system will be allowed to expand.</p> <p>When the file system is created, it is initially allocated a certain amount of space (the initial capacity), and the file system is allowed to expand automatically, up to its size limit. When the file system uses approximately 80% of its currently allocated space, it is expanded automatically up to its size limit. This expansion occurs in increments specified by the <i>guideline chunk size</i> (which is calculated by the system).</p> <p>The file system can be manually expanded, increasing the file system size limit.</p>	<p>The size limit defines the amount of space that is immediately allocated to the file system.</p> <p>When the file system is created, it is allocated the total amount of space specified by the <i>size limit</i>.</p> <p>The file system can be manually expanded, increasing the file system size limit.</p>

File System Type	Auto-Expansion Enabled	Auto-Expansion Disabled
Tiered	<p>If auto-expansion is not confined, the size limit is ignored if defined. The tiers of the file system will be allowed to expand until the storage pool is full.</p> <p>If auto-expansion is confined, the size limit defines the maximum size to which the tier of a file system will be allowed to expand.</p> <p>When the file system is created, the user data tier is initially allocated a certain amount of space (the <i>initial capacity</i>), and the user data tier is allowed to expand automatically, up to its size limit. When the user data tier uses approximately 80% of its currently allocated space, it is expanded automatically up to its size limit. This expansion occurs in increments specified by the <i>guideline chunk size</i> (which is calculated by the system).</p> <p>Either tier can be manually expanded, increasing the file system size limit.</p> <p>When the file system is created, the user data tier is initially allocated the total amount of space specified by the size limit.</p> <p>Either tier can be manually expanded, increasing the file system size limit.</p>	<p>The size limit defines the amount of space that is immediately allocated to the user-data tier.</p> <p>When the file system is created, the user data tier is initially allocated the total amount of space specified by the size limit.</p> <p>Either tier can be manually expanded, increasing the file system size limit.</p>

By default, file system auto-expansion is enabled and, when auto-expansion is enabled, the file system expands, without interruption of service if the following conditions exist:

- **Confined limit not reached** (only for file systems that have auto-expansion confined). As long as the file system expansion would not exceed the confined auto-expansion limit.
- **Available space**. Sufficient available free space and chunks remain in the storage pool.
- **Chunk limit**. The file system expansion will not cause the file system to exceed the maximum allowable number of chunks in a file system.
- **Maximum supported file system size**. The file system expansion will not cause the file system to exceed the maximum supported file system size.

Whether auto-expansion is enabled or disabled, you can limit the size of the file system of an untiered file system or either tier of a tiered file system. If necessary, you can manually expand the file system or of an untiered file system or a tier of a tiered file system.

Note: File system auto-expansion may be enabled or disabled for all file systems in a particular storage pool. When enabled for a storage pool, file system auto-expansion is enabled by default, but if it has been disabled on an individual file system, you can re-enable it. When file system auto-expansion is disabled for a storage pool, you cannot enable it (you must expand the file system manually).

Enabling and disabling file system auto-expansion

When file system auto-expansion is enabled or disabled for a storage pool, you cannot change the setting for a single file system in the storage pool; you must enable or disable file system auto-expansion for all file systems in the storage pool.

When file system auto-expansion is disabled and a file system requires expansion, you must expand the file system manually.

The ability for file systems in a storage pool to automatically expand is enabled or disabled at the storage pool level.

- When file system auto-expansion is enabled for a storage pool, file systems in the storage pool may be allowed to auto-expand or they may be confined to a specified size limit. File system auto-expansion is enabled or disabled for each file system independently
- When file system auto-expansion is disabled for a storage pool, file systems in the storage pool are not allowed to auto-expand. You cannot change the setting for an individual file system in the storage pool to allow auto-expansion.

When file system auto-expansion is disabled (at the storage pool level, or for an individual file system) and a file system requires expansion, you must expand the file system manually.

Expanding a file system

Manual file system expansion is supported through NAS Manager and through the CLI.

Procedure

1. Navigate to **Home > Storage Management > File Systems**.
2. Select a file system and click **details** to display the **File System Details** page.
3. Click **expand** to display the **Expand File System** page.

For an untiered file system the **Expand File System** page looks like the following:

Storage Management [Home](#) > [Storage Management](#) > [File Systems](#) > [File System](#) > Expand File System

Expand File System

Storage Pool: SP_1
Free Capacity: 29.75 TiB
Guideline Chunk Size: 11.67 GiB

File System Label: protectedsiteA
Current File System Capacity: 489.31 GiB
Current Size Limit: 500 GiB

New Size Limit: GiB [Rounded to nearest chunk](#)

Allocate On Demand
Allow the file system to automatically expand to the *new size limit*.

Allocate Now
Expand the file system to the full size specified by the *new size limit*.

The Allocate On Demand option is not available for UVM-backed spans.

For a tiered file system, the **Expand File System** page looks like the following:

Storage Management [Home](#) > [Storage Management](#) > [File Systems](#) > [File System](#) > Expand File System

Expand File System

Storage Pool: mt-span
Free Capacity: 3.54 TiB
Guideline Chunk Size: 18 GiB

File System Label: mt-fs
Current File System Capacity: 35.75 GiB

Tier 0 Meta-data
Current Capacity: 17.91 GiB
Current Size Limit: -

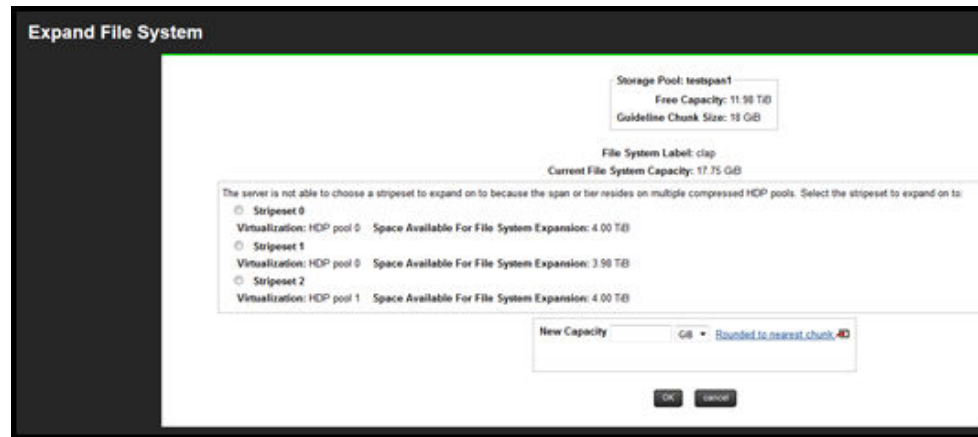
Tier 1 User-data
Current Capacity: 17.91 GiB
Current Size Limit: 1 TiB

New Capacity Limit: GiB [Rounded to nearest chunk](#)

Allocate On Demand
Allow the file system to automatically expand to the *new size limit*.

Allocate Now
Expand the file system to the full size specified by the *new size limit*.

In some circumstances, such as when the storage pool resides on a UVM span or in HDP compressed storage, a specific stripeset must be selected for expanding the file system. If the server cannot select the stripeset, the **Expand File System** page shows a list of stripesets from which to select.



4. To expand the file system manually, do one of the following:
 - For an untiered file system, specify the new file system capacity in the **New Capacity** field and use the list to select MiB, GiB, or TiB.
 - For a tiered file system, select the tier you want to expand, and specify the new file system capacity in the **New Capacity** field, then use the list to select MiB, GiB, or TiB.



Note: You can expand one tier per expansion operation. To expand both tiers, you must perform a manual expansion twice.

5. Click **OK**.



Note: Because space is always allocated in multiples of the chunk size set when the storage pool containing the file system was created, the final size of the file system may be slightly larger than you request.

Manual expansion of file systems is also supported through the command line interface. For detailed information on this process, run `man filesystem-expand` on the CLI.

Moving a file system

Moving a file system (or several file systems) may be necessary to improve performance or balance loads, to move data to different storage resources, to support changing network topography, or other reasons.

There are two basic methods of moving a file system:

- File System Relocation

File system relocation changes the EVS (virtual server) that hosts the file system, but it does not move file system data. Moving the file system from one EVS to another changes the IP address used to access the file system, and also changes CIFS shares and NFS Exports for that file system. For information on how to relocate a file system using File System Relocation, refer to the *Replication and Disaster Recovery Administration Guide*.

If the file system to be relocated is linked to from within a CNS, and clients access the CNS using a CIFS share or an NFS export, the relocation can be performed with no change to the configuration of network clients. In this case, clients will be able to access the file system through the same IP address and CIFS share/NFS export name after the relocation as they did before the relocation was initiated. For more information on CNS, refer to the *Server and Cluster Administration Guide*.



Caution: Whether or not the file system resides in a CNS, relocating a file system will disrupt CIFS communication with the server. If Windows clients require access to the file system, the file system relocation should be scheduled for a time when CIFS access can be interrupted.

- Transfer of primary access

A transfer of primary access is a replication-based method of copying data from a portion of a file system (or an entire file system) and relocating the access points for that data (copying the data and metadata). A transfer of primary access causes very little down time, and the file system is live and servicing file read requests during most of the relocation process. For a short period during the relocation process, access is limited to read-only. For more information on relocating a file system using transfer of primary access, refer to the *Replication and Disaster Recovery Administration Guide*.

The method you use to relocate a file system depends, in part, on what you want to move, and what you want to accomplish by relocating the file system.

- If you want to move the file system's access points, but not the actual data, using file system relocation is the most appropriate method.
- If you want to move the file system's data and access points, using a transfer of primary access is the most appropriate method.

File system relocation

Before it can be shared or exported, a file system must be associated with a Virtual Server (EVS), thereby making it available to network clients. The association between a file system and an EVS is established when the file system is created. Over time, evolving patterns of use and/or requirements for storage resources may make it desirable to relocate a file system to a different EVS.



Note: Read caches cannot be relocated.

A **file system hosted by an EVS on a cluster node** may be relocated to:

- An EVS on the same cluster node, or
- An EVS on a different node in the same cluster.

but may not be relocated to:

- An EVS on a stand-alone server, or
- An EVS on a node of a different cluster.

A file system hosted by an EVS on a stand-alone server may be relocated to

- An EVS on the same server

but may not be relocated to:

- An EVS on a different server, or
- An EVS on a node in a cluster.

Typically, File System Relocation is used to move a file system from an EVS on a cluster node to an EVS on a different cluster node in order to improve throughput by balancing the load between cluster nodes.

File system relocation performs the following operations:

- Re-associates the file system with the selected EVS.
- Transfers explicit CIFS shares of the file system to the new EVS.
- Transfers explicit NFS exports of the file system to the new EVS.
- Migrates FTP users to the new EVS.
- Migrates snapshot rules associated with the file system to the new EVS.
- Migrates the iSCSI LUs and targets.

File system relocation may require relocating more than just the specified file system. If the file system is a member of a data migration path, both the data migration source file system and the target file system will be relocated. It is possible for the target of a data migration path to be the target for more than one source file system. If a data migration target is relocated, all associated source file systems will be relocated as well.

If more than one file system must be relocated, a confirmation dialog will appear indicating the additional file systems that must be moved. Explicit confirmation must be acknowledged before the relocation will be performed.

File System Relocation will affect the way in which network clients access the file system in any of the following situations:

- The file system is linked to from the CNS tree, but is shared or exported outside of the context of the CNS.
- The cluster does not use a CNS.

In each of the above cases, access to the shares and exports will be changed. In order to access the shares and exports after the relocation, use an IP address of the new EVS to access the file service.

Relocating file systems that contain iSCSI Logical Units (LUs) will interrupt service to attached initiators, and manual reconfiguration of the IP addresses through which targets are accessed will be required once the relocation is complete. If relocating a file system with LUs is required, the following steps must be performed:

- Disconnect any iSCSI Initiators with connections to LUs on the file system to be relocated.
- Unmount the iSCSI LU.
- Relocate the file system as normal. This procedure is described in detail in the *Replication and Disaster Recovery Administration Guide*.
- Reconnect the new Targets with the iSCSI Initiators. Be aware that the Targets will be referenced by a new name corresponding to the new EVSs.



Note: All iSCSI LUs on a target must be associated with file systems hosted by the same EVS.

Using system lock on file systems

System Lock mode protects file systems during replication and transfer of primary access operations. Four important distinctions apply:

- **NDMP (Network Data Management Protocol) versus file service protocols.** When **System Lock** is enabled for a file system:
 - NDMP has full access (including writes) during backups, replication, and transfer of primary access.
 - The file system remains in read-only mode to clients using the file service protocols (NFS, CIFS, FTP, and iSCSI).
- **System Lock versus read only:**
 - When a file system is Syslocked, NDMP still has full access to that file system and can write to it.
 - When a file system is mounted as read-only, NDMP (like all other protocols) has read-only access to that file system, and cannot write to it. To ensure that a file system remains completely unchanged, you should mount it as read-only.
- **Replication versus transfer of primary access:**
 - *During replication operations*, the destination file system is put into System Lock mode.
 - *During transfer of primary access operations*, both the source file system and the destination file system are put into System Lock mode.
- **Read Cache Exception.** A read cache may not be put into System Lock mode.

Enabling and disabling system lock for a file system

Procedure

1. Navigate to **Home > Storage Management > File Systems**.
2. Select a file system and click **details** to display the **File System Details** page.

3. In the **Syslock** field, toggle the enable/disable button as appropriate.

When the file system is in System Lock mode, the Status changes to Syslocked, the System Lock status becomes enabled, and the **Enable** button becomes **Disable**.

When System Lock is enabled for a file system, NDMP has full access to the file system and can write to it during a backup or replication, but the file system remains in read-only mode to clients using the file service protocols (NFS, CIFS, FTP, and iSCSI).

When viewing the details of a read cache, the System Lock's enable/disable button is not available.

Recovering a file system

Following some system failures, a file system may require recovery before mounting. If required, such a recovery is performed automatically when you mount the file system. Performing recovery rolls the file system back to its last checkpoint and replays any data in NVRAM.

In extreme cases, when you mount a file system after a system failure, the automatic recovery procedures may not be sufficient to restore the file system to a mountable condition. In such a case, you must forcefully mount the file system, which discards the contents of NVRAM before mounting the file system.

Procedure

1. Navigate to **Home > Storage Management > File Systems**.
2. Select a file system and click **details** to display the **File System Details** page.
3. If a file system displays Not Mounted in the Status column, click **mount** to try to mount the file system.
 - If necessary, the automatic recovery processes will be invoked automatically. The file system was mounted successfully.
 - If the automatic recovery fails, the file system will not mount, and the **File Systems** page will reappear, indicating that the file system was not mounted. Navigate to the **File System Details** page.
4. For the file system that failed to mount, click **details** to display the **File System Details** page. In the Settings/Status area of the page, the file system label will be displayed, along with the reason the file system failed to mount (if known), and suggested methods to recover the file system, including the link for the **Forcefully mount** option.

5. Depending on the configuration of your system, and the reason the file system failed to mount, you may have several recovery options:
 - *If the server is part of a cluster*, you may be able to migrate the assigned EVS to another cluster node, then try to mount the file system. This can become necessary when another node in the cluster has the current available data in NVRAM that is necessary to replay write transactions to the file system following the last checkpoint. An EVS should be migrated to the cluster node that mirrors the failed node's NVRAM (for more information on NVRAM mirroring, refer to the *System Access Guide*. For more details on migrating EVSs, refer to the *Server and Cluster Administration Guide*.
 - *If the first recovery attempt fails*, click the **Forcefully mount** link. This will execute a file system recovery without replaying the contents of NVRAM.



Caution: Using the Forcefully mount option discards the contents of NVRAM, data which may have already been acknowledged to the client. Discarding the NVRAM contents means that all write operations in NVRAM (those write operations not yet committed to disk) are lost. The client will then have to resubmit the write request. Use the Forcefully mount option only upon the recommendation of customer support.

Restoring a file system from a checkpoint

Following a storage subsystem failure, it may be necessary to recover file systems.

File system corruption due to an event (such as RAID controller crash, storage system component failure, or power loss) often affects objects that were being modified around the time of the event.

The file system is configured to keep up to 128 checkpoints. The maximum number of checkpoints supported is 1024. The number of checkpoints preserved is configurable when the file system is formatted, but, once set, the number of checkpoints cannot be changed.

When a checkpoint completes, rather than immediately freeing the storage used for the previous checkpoint, the file system maintains a number of old checkpoints. As each new checkpoint completes, the oldest checkpoint is overwritten. This means that there can be multiple checkpoints on-disk, each of which is complete and internally consistent point-in-time view of the file system. If necessary, the file system can be restored to any of these checkpoints.

In the case of file system corruption, if there are enough checkpoints on disk, it may be possible to roll back to a previous checkpoint, pre-dating the event that caused the corruption and restoring the file system using the uncorrupted checkpoint. This may be possible even if this event occurred up to a few minutes before the file system was taken offline.

To restore a file system to a previous checkpoint, use the `fs-checkpoint-health` and the `fs-checkpoint-select` commands. Refer to the *Command Line Reference* for more information about these commands.

Note the following:

- Restoring a file system using a checkpoint does not affect snapshots taken prior to the checkpoint being restored, but, like any other file system update, snapshots taken after that checkpoint are lost.
- After restoring to a checkpoint, it is possible to restore again, to an older checkpoint and, if the file system has not been modified, restore again, to a more recent checkpoint. So, for example, it is possible to mount the file system in read only mode, check its status, and then decide whether to re-mount the file system in normal (read/write) mode or to restore to a different checkpoint.



Caution: Once you mount a restored file system in normal (read/write) mode, you cannot restore to a later checkpoint.

File system recovery from a snapshot

It is possible that, although corruption has occurred in the live file system, a good snapshot still exists. If so, it may be preferable to recover the file system from this snapshot, with some loss of data, rather than incur the downtime that might be required to fix the live file system. Recovering a file system from a snapshot restores the file system to the state that it was in when the snapshot was taken.

Recovering a file system from a snapshot makes it possible to roll back the file system to the state that it was in when a previous snapshot was taken.

File system recovery from a snapshot is a licensed feature, which requires a valid FSRS license on the server/cluster.



Note: You can recover a file system from a snapshot only when at least the configured number of preserved file system checkpoints have been taken since that snapshot was taken. For example, if a file system is configured to preserve 128 checkpoints (the default), then you can recover the file system from a snapshot only after a minimum of 128 checkpoints have been taken after the snapshot. If less than the configured number of checkpoints have been taken since the snapshot, you can either recover from an earlier snapshot or recover the file system from a checkpoint.

The following file system rollback considerations apply:

- File system rollback can be performed even if the live file system is corrupted.
- All snapshots are lost after the rollback.
- Even though the file system recovery happens very quickly, no new snapshots can be taken until all previous snapshots have been discarded. The time required before a new snapshot can be taken depends on the size of the file system, not on the number of files in the file system.



Note: Once you have recovered a file system from a snapshot, and mounted it in read-write mode, you cannot undo the recovery or recover again to a different snapshot or checkpoint.

To roll back a file system from a snapshot, use the `snapshot-recover-fs` command. Refer to the *Command Line Reference* for more information about this command.

An additional tool is available to kill all current snapshots, that is the `kill-snapshots` command (refer to the *Command Line Reference* for more information about this command). `snapshot-delete-all` is the preferred tool for deleting all snapshots as it does not require the file system to be unmounted, no space is leaked, and it does not affect checkpoint selection.

Automatic file system recovery

The `fixfs` utility is the main file system recovery tool, but it should only be used under the supervision of customer support personnel.

`fixfs` is capable of repairing a certain amount of non-critical metadata, for example performing orphan recovery. At all stages that have the potential to last longer than a few minutes, `fixfs` provides progress reporting, and the option to abort the fix. Note that progress reports are stage or operation based, for example Stage 3 of 7 complete. For some operations, `fixfs` will also provide an estimate of time until the completion of the operation.

The strategy used by `fixfs` to repair file systems can be summarized as:

- `fixfs` or `fs-checkpoint-health` are the recovery tools to be used if a file system is experiencing corruption. The default `fixfs` behavior may be modified by various command line switches, but often the required switch is suggested by `fixfs` during or at the end of a previous run.
- Where possible, `fixfs` will run with the file system in any state (there will be no need to perform file system recovery first, so that there's no need to worry about what happens if recovery cannot complete due to corruption). Where not possible (for example, if the file system is marked as "failed" or "requires expansion"), `fixfs` will not run. When `fixfs` does not run, it will give a clear indication of what needs to be done to get to the point where it can run.
- By default, `fixfs` will only modify those objects which are lost or corrupted.
- By default, `fixfs` will only bring the file system to the point where it can be mounted.
- Snapshots are considered expendable and are deleted.

Using deduplication file system

This section describes how to use deduplication on a file system.

Determining sufficient space for dedupe conversion

File systems that have not been formatted with dedupe support must be converted before they will support deduplication. However, the conversion process is an offline process which can sometimes fail because there is not sufficient scratch or free space available. In this scenario, you must expand the file system and then re-attempt the conversion.



Note: File systems formatted or expanded in releases newer than dedupe will have the extra scratch space required for conversion and the following procedure is unnecessary.

Before you start the dedupe conversion, use the following procedure to determine if a file system has sufficient scratch and free space available. It is not an offline procedure.

Procedure

1. Call customer support to obtain the instructions and dev password to execute the **fs-capacity-info** dev command.

This command can be run on a mounted file system.

2. Run the command with the name of the file system that is to be checked (for example, **fs-capacity-info f filesystem1**).

This command generates the following sample output:

```
fs-capacity-info/fs3: Underlying length : 0x017666400000 bytes,
0x00bb332000 sectors, 0x017666400 blocks
fs-capacity-info/fs3: Formatted length : 0x017636000000 bytes,
0x00bb1b0000 sectors, 0x017636000 blocks
fs-capacity-info/fs3: Scratch-space offset : 0x017636000000 bytes,
0x00bb1b0000 sectors, 0x017636000 blocks
fs-capacity-info/fs3: Scratch-space length : 0x000030400000 bytes,
0x0000182000 sectors, 0x000030400 blocks
fs-capacity-info/fs3: Reservable available space : 0x01748e51f000
bytes, 0x00ba4728f8 sectors, 0x01748e51f blocks
fs-capacity-info/fs3: Physical available space : 0x0175ff4bb000 bytes,
0x00baffa5d8 sectors, 0x0175ff4bb blocks
fs-capacity-info/fs3: Space required for RCB conversion :
0x0000e9e1c000 bytes, 0x000074f0e0 sectors, 0x0000e9e1c blocks
fs-capacity-info/fs3: Scratch-space required for RCB conversion:
0x00002eccc800 bytes, 0x0000176664 sectors, 0x00002eccc blocks
```

3. Locate the "**Scratch-space length**" value. In this example, **0x000030400000** bytes.
4. Locate the "**Scratch-space required for RCB conversion**" value. In this example, **0x00002eccc800** bytes.
5. If the **Scratch-space length** is greater than **Scratch-space required for RCB conversion**, then there is enough scratch space. In this example, **0x000030400000** is greater than **0x00002eccc800**, indicating that there is enough scratch space for the conversion.
6. Locate the "**Reservable available space**" value. In this example, **0x01748e51f000** bytes.
7. Multiply the **Space required for RCB conversion** value by 1.6. In this example, the value is **0x000176360000** bytes.
8. If **Reservable available space** is greater than **Space required for RCB conversion * 1.6**, then there is enough free space. **0x01748e51f000** is greater than **0x000176360000**, indicating there is enough free space for the conversion.

Preparing for dedupe conversion

In addition to the space required for the ondisk component (for example, 45 GB) of the dedupe index, a conversion process may also require additional scratch space that is normally carved out of the underlying partition when the file system is first formatted or expanded. It will not utilize storage from the free space available on the file system (for example, free space reported by `df`).

If you receive the following failure message during the dedupe conversion process, this indicates that additional scratch space needs to be setup.

```
2013-03-11 09:15:49.744-04:00 1 MMB1 fs-convert-to-support-dedupe request
for evs12-fs-d01-tgt: Checking if sufficient space is free on target file
system... 2013-03-11 09:15:49.744-04:00 1 MMB1 fs-convert-to-support-
dedupe request for evs12-fs-d01-tgt: Scratch space requires 4.21 GB. 2.31
GB is available. 2013-03-11 09:15:49.744-04:00 1 MMB1 fs-convert-to-
support-dedupe request for evs12-fs-d01-tgt: Expand the file system size
before re-submitting the conversion
```

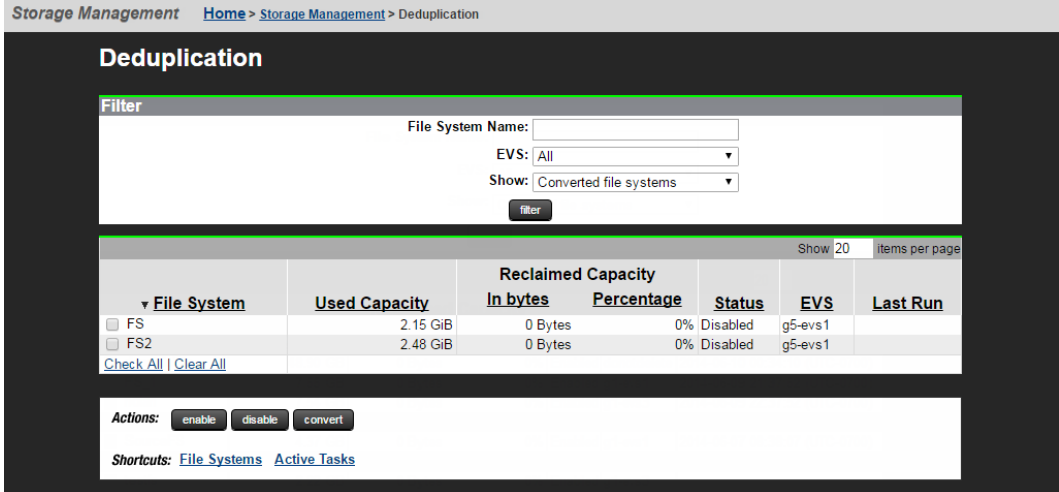
You can resolve this issue by expanding the file system. Usually one chunk is all that is required. Any storage beyond what is required for the scratch space from the chunk will be made available to the file system. Once this has been done, you can retry the conversion process.

Viewing the deduplication file system page

To view the Deduplication file system page:

Procedure

1. Navigate to **Home > Storage Management > File Systems > Dedupe File Systems**. The Deduplication page appears.



The screenshot shows the 'Deduplication' page. At the top, there is a breadcrumb trail: 'Storage Management > Home > Storage Management > Deduplication'. Below this is a 'Filter' section with a 'File System Name' input field, an 'EVS' dropdown menu set to 'All', and a 'Show' dropdown menu set to 'Converted file systems'. A 'filter' button is located below these fields. Below the filter is a table with the following data:

File System	Used Capacity	Reclaimed Capacity		Status	EVS	Last Run
		In bytes	Percentage			
<input type="checkbox"/> FS	2.15 GiB	0 Bytes	0%	Disabled	g5-evs1	
<input type="checkbox"/> FS2	2.48 GiB	0 Bytes	0%	Disabled	g5-evs1	

Below the table, there are 'Check All' and 'Clear All' links. At the bottom of the page, there are 'Actions' buttons: 'enable', 'disable', and 'convert'. There are also 'Shortcuts' for 'File Systems' and 'Active Tasks'.

The following table describes the fields on this page:

Field/Item	Description
File System Name	Use to search by file system name.
EVS	Use to search by a particular EVS.
Show	Use to search by type of file system version: <ul style="list-style-type: none"> ▪ Converted file systems ▪ Not converted file systems ▪ All file systems
Filter	Use to initiate your search criteria.
File Systems	Displays a list of file systems that match the search criteria.
Used Capacity	Displays the total amount of file system space in use, in GiB or TiB and as a percentage of the total.
Reclaimed Capacity	Displays the space savings based on deduplication, presented in bytes and in a percentage. The percentage is the ratio of the Deduped column to the sum of the Deduped and Use columns. Note that the dedupe process begins by taking a temporary snapshot of the file system and then dedupes it.
Status	Displays the status: <ul style="list-style-type: none"> ▪ Disabled: dedupe is disabled on the file system. ▪ Enabled: dedupe is enabled on the file system. ▪ Not converted: The file system needs to be converted in order to use the dedupe capability on the file system. ▪ Not assigned to an EVS: File system is not assigned to an EVS and cannot be converted until the file system is assigned to an EVS. ▪ EVS name (disabled): File system cannot be converted because the EVS is disabled. Enable the EVS first. ▪ Converting: Starting to convert a file system to support deduplication. ▪ Full run: Performing a full deduplication on the file system. ▪ Incremental run: Performing an incremental deduplication on the file system.
EVS	EVS to which the file system is assigned.
Last Run	The date and time of the last dedupe run.
enable	Dedupe-enables an existing file system.

Field/Item	Description
disable	Disables dedupe on an dedupe-enabled file system.
convert	Starts the conversion process on a file system that needs conversion to support and enable dedupe. See the <code>fs-convert-to-support-dedupe</code> man page for important information and considerations before starting the conversion process. For example: <ul style="list-style-type: none"> ▪ The file system should be in an unmounted state. ▪ The file system must have sufficient space.
File System	Returns to the File System page.
Active Tasks	Displays the current status of the conversion process.

Enabling dedupe for a file system

To enable dedupe for a file system:

Procedure

1. Navigate to **Home > Storage Management > File Systems > Dedupe File Systems**. The Deduplication page appears.
2. Enter the search criteria in the Filter section of the page and then click **Filter**.
3. Under the File System label, fill the check box next to the file system to be enabled. You can enable one or more file systems at a time.
4. Click **enable**.
The system immediately starts the dedupe-enable process. The Status column displays Enabled to reflect this action.

Converting a file system to enable dedupe

You can convert an existing file system to support and enable dedupe. After the conversion process takes place, the file system is dedupe-enabled. Converting a file system can take some time and it is recommended that you plan accordingly for the conversion time. It is recommended that you convert one file system at a time. See the `fs-convert-to-support-dedupe` man page for important information and considerations before starting the conversion process. For example:

- The file system should be in an unmounted state.
- The file system must have sufficient space.

Procedure

1. Navigate to **Home > Storage Management > File Systems > Dedupe File Systems**. The Deduplication page appears.

2. Select **Not converted file systems** from the **Show** list to display file systems do not have dedupe support enabled.
3. Click **Filter**.
4. The system displays the file systems that need conversion in order to be dedupe-enabled
5. Fill in the check box next to the file system to convert. It is recommended that you one file system at a time.
6. Click **Convert** and read the messages in the dialog that appears.
7. After you have read the message and ensure that you want to proceed with conversion, click **OK**.
8. Click **Active Tasks** to view the current conversion status.
After the conversion is done, the file system is dedupe capable and the file system is now queued for a full dedupe run. The dedupe process will start when the file system is queued for dedupe. The Status column displays Enabled.

If the status remains in the Needs Conversion status, check the Events page. Navigate to **Home > Status & Monitoring > Event Log**. This log reports any conversion errors. For example, an error may occur if there is not sufficient space for the file system to be converted or if the user-data-write-policy of the file system is set to anything other than never-overwrite. See the following CLI man page for more information:

`wfs-user-data-write-policy`

Dedupe support for object replication targets

- Deduplication of object replication target file systems is enabled in the same way as a normal file system.
- By default an incremental deduplication job will automatically be scheduled upon successful completion of the object replication.
- Further configuration details for this feature can be found in the best practices guide *Hitachi NAS Platform Deduplication Best Practice (MK-92HNAS031)*

Managing file system quotas

You can use a quota to allocate a maximum amount of disk space a user or group may use. It can be flexible in its adherence to the rules assigned and is applied per file system.


Quotas can be set for an entire virtual volume, and on individual users, and on groups of users. Default user and group quotas can be defined, and in the absence of explicit user or group quotas, the default quotas apply.

Managing usage quotas

The **Quotas by File System** page lists usage quotas for the selected file system.

Procedure

1. Navigate to **Home > Storage Management > Quotas by File System** to display the **Quotas by File System** page.



The following table describes the fields on this page:

Field/Item	Description
EVS/File System Label	The name of the selected EVS and file system.
change	Selects a different file system to display a different set of quotas.
Filter	Displays quotas based on specific users or groups.
User/Group Account	A quota name can consist of: <ul style="list-style-type: none"> ▪ A CIFS domain and user or group name, such as <code>bb\Smith</code> or <code>bb\my_group</code> (where <i>bb</i> is a domain, <i>Smith</i> is a user and <i>my_group</i> is a group). ▪ An NFS user or group such as <code>johns</code> or <code>finance</code> (where <i>johns</i> is an NFS user and <i>finance</i> is an NFS group).

Field/Item	Description
	A name may be '0' (if the quota was created for the owner of the directory at the root of the virtual volume).
File System	The file system on which the quota applies.
Quota Type	Type of file system activity. Possible values are User or Group.
Created By	Lists the method of quota creation: <ul style="list-style-type: none"> ▪ Automatically Created: A quota created using a quota default. ▪ User Defined: A uniquely defined quota.
Usage Limit	Overall limit set for the total size of all files in the file system owned by the target of the quota.
File Count Limit	Overall limit set for the total number of files in the file system owned by the target of the quota.
details	Displays the File System Quotas Details page for the selected file system.
add	Advances to the Add File System Quota page.
delete	Deletes a particular quota.
Delete All Quotas	Deletes all of the current quotas. This option is only visible if more than one quota is configured.
Refresh cache	Clears the NAS Manager cache and repopulates it with relevant objects. This is different from clicking the browser refresh button, which picks up any recent updates without clearing the cache.
User Defaults	Use to set, edit, or reset the user defaults.
Group Defaults	Use to set, edit, or reset the group defaults.
Modify Email Contacts	Use to add or delete email contacts who are notified when the file system exceeds its size threshold.
Download Quotas	Use to download the quotas (not file system quotas) for this virtual volume to a <code>.csv</code> file.

Setting user and group file system quota defaults

The quota default values define a template with which the system will automatically generate a quota in response to a file being saved on the file system. If a file is saved, and the respective defaults are set, a user quota will be created for the user, and a group quota created for the user's domain. In addition, as soon as the defaults are set, a user and group quota will be created for the owner of the directory at the root of the virtual volume.

Procedure

1. Navigate to **Home > Storage Management > Quotas by File System** to display the **Quotas by File System** page.
2. In the **Quotas by File System** page, click User Defaults or Group Defaults. User Defaults creates a user quota for the user; a Group quota creates a group quota for the user's domain.

Storage Management [Home](#) > [Storage Management](#) > [Virtual Volumes & Quotas](#) > [Quotas](#) > Quota


Quota

EVS / File System: g5-evs1 / g5-fs1
 Virtual Volume Name: fs1_vv1
 Created By: User Defined
 User/Group Account: n/a
 Quota Type: Virtual Volume

Usage	File Count
Used: 0 Bytes	Count: 1
Limit: <input type="text"/> GiB	Limit: <input type="text"/> 200000
Hard Limit: <input type="checkbox"/>	Hard Limit: <input type="checkbox"/>
Warning: <input type="text"/> 75 %	Warning: <input type="text"/> 75 %
Severe: <input type="text"/> 85 %	Severe: <input type="text"/> 85 %

Log Quota Events in the managed server's Event Log

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The EVS and file system on which the user file system quota applies.
Virtual Volume name	The name given to the virtual volume on creation.
Created By	Describes how the quota was created. The options are: <ul style="list-style-type: none"> User Defined - created by a user through the NAS Manager or CLI. System Defined - when an NFS or SMB client creates a file in the file system within the path of a virtual volume, and the user- or group-defaults for quotas in that virtual volume are set, a quota is created using the values in the defaults.
User/Group Account	A user or group has an SMB or NFS account. When using NFS, the ID number appears in this field. When using SMB, the AD account name appears in this field.
Quota Type	The type of quota. The options are: <ul style="list-style-type: none"> Virtual Volume - this quota type limits usage over the whole virtual volume, regardless of who is creating the files. User - this quota type limits the space usable in the virtual volume (or the number of files that can be created) by a single user. Group - this quota type limits the space usable in the virtual volume (or the number of files that can be created) by a whole group.
Automatically create quotas for Domain Users	<div style="border: 1px solid #add8e6; padding: 5px; margin-bottom: 10px;">  Note: This option only displays on the group file system quota page. </div> <p>This option allows the creation of default quotas for the group Domain Users. By default every NT user belongs to the group Domain Users. Enabling this option effectively includes every NT user in the quota, (unless each user's primary group has been explicitly set).</p>
Usage	
Used	The amount of used space in the virtual volume that is being tracked by this quota.
Limit	Amount of space to enable in Bytes: KiB, MiB, GiB, TiB, PiB, or EiB.

Field/Item	Description
Hard Limit	When enabled, the amount of space specified in the Limit field cannot be exceeded.
Warning	Percentage of the amount of space specified in the Limit field at which a Warning alert is sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert is sent.
File Count	
Count	The number of files in the virtual volume that are being tracked by this quota.
Limit	Maximum number of files to enable for this quota.
Hard Limit	When enabled, the number of files specified in the Limit field cannot be exceeded.
Warning	The percentage of the number of files specified in the Limit field at which a Warning alert is sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert is sent.
Log Quota Events in the managed server's Event log	Selecting this check box sets the default for all users or groups to have quota events logged in the server's event log.

3. For group file system quota defaults, select the **Automatically create quotas for Domain Users** check box to allow the creation of default quotas for the group domain users.
4. Under the Usage and File Count sections, enter the values as appropriate:
 - a. In the **Limit** field, enter the limit. Additionally, under Usage, select KiB, MiB, GiB, or TiB from the list.
 - b. Select the **Hard Limit** check box if the space specified in the **Limit** field cannot be exceeded.
 - c. In the **Warning** field, enter the warning.
 - d. In the **Severe** field, enter the value.
 - e. Select the **Log Quota Events in the managed server's Event Log** check box to set the default for all users or groups to have quota events logged in the server's event log.
5. Click **OK**.

Adding a quota

Describes how to allocate storage usage and file count by client.

Procedure

1. Navigate to **Home > Storage Management > Quotas by File System** to display the **Quotas by File System** page.
2. Click **add** in the **Add File System Quota** page.

Storage Management [Home](#) > [Storage Management](#) > [Virtual Volumes & Quotas](#) > [Quotas](#) > Add Quota

Add Quota

EVS / File System: g5-avs1 / g5-fs1
 Virtual Volume Name: fs1_vv1

Quota Type: User User Account i.e. domain/user (CIFS) or user (NFS)
 Group Group Account i.e. domain/group (CIFS) or group (NFS)
 Virtual Volume

Usage	File Count
Limit: <input type="text"/> GiB ▼	Limit: <input type="text"/>
Hard Limit: <input type="checkbox"/>	Hard Limit: <input type="checkbox"/>
Warning: <input type="text"/> 75 %	Warning: <input type="text"/> 75 %
Severe: <input type="text"/> 85 %	Severe: <input type="text"/> 85 %

Log Quota Events in the managed server's Event Log

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The EVS and file system on which the user file system quota applies.
Virtual volume name	The name of the virtual volume.
Quota Type	The Quota Types are: <ul style="list-style-type: none"> ▪ User - the user account name or ID. This quota type limits the space usable in the virtual volume, or the number of files that can be created, by a single user. ▪ Group - the group account name or ID. This quota type limits the space usable in the virtual volume, or the number of files that can be created, by a whole group. ▪ Virtual Volume - this quota type limits usage over the whole virtual volume, regardless of who is creating the files.
Usage	
Limit	Amount of space to enable in Bytes: KiB, MiB, GiB, TiB, PiB, or EiB.
Hard Limit	When enabled, the amount of space specified in the Limit field cannot be exceeded.
Warning	Percentage of the amount of space specified in the Limit field at which a Warning alert will be sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert will be sent.
File Count	
Limit	Maximum number of files to enable for this quota.
Hard Limit	When enabled, the number of files specified in the Limit field cannot be exceeded.
Warning	The percentage of the number of files specified in the Limit field at which a Warning alert will be sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert will be sent.
Log Quota Events in the managed server's Event log	Selecting this check box sets the default for all users or groups to have quota events logged in the server's event log.

3. Click **User** to add a user account, or **Group** to add a group account.
4. Under the Usage and File Count sections, enter the values as appropriate:

- a. In the **Limit** field, enter the limit. Additionally, under Usage, select KiB, MiB, GiB, or TiB from the list.
 - b. Select the **Hard Limit** check box if the space specified in the **Limit** field cannot be exceeded.
 - c. In the **Warning** field , enter the warning.
 - d. In the **Severe** field, enter the value.
 - e. Select the **Log Quota Events in the managed server's Event Log** check box to set the default for all users or groups to have quota events logged in the server's event log.
5. Click **OK**.

Modifying a file system quota

Procedure

1. Navigate to **Home > Storage Management > Quotas by File System** to display the **Quotas by File System** page.
2. Fill in the check box next to the quota, and click **details**.

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The EVS and file system on which the user file system quota applies.
Virtual volume name	The name of the virtual volume.
Quota Type	<p>The Quota Types are:</p> <ul style="list-style-type: none"> ▪ User - the user account name or ID. This quota type limits the space usable in the virtual volume, or the number of files that can be created, by a single user. ▪ Group - the group account name or ID. This quota type limits the space usable in the virtual volume, or the number of files that can be created, by a whole group. ▪ Virtual Volume - this quota type limits usage over the whole virtual volume, regardless of who is creating the files.
Usage	
Limit	Amount of space to enable in Bytes: KiB, MiB, GiB, TiB, PiB, or EiB.
Hard Limit	When enabled, the amount of space specified in the Limit field cannot be exceeded.
Warning	Percentage of the amount of space specified in the Limit field at which a Warning alert will be sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert will be sent.
File Count	
Limit	Maximum number of files to enable for this quota.
Hard Limit	When enabled, the number of files specified in the Limit field cannot be exceeded.
Warning	The percentage of the number of files specified in the Limit field at which a Warning alert will be sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert will be sent.
Log Quota Events in the managed server's Event log	Selecting this check box sets the default for all users or groups to have quota events logged in the server's event log.

3. Make any necessary changes, and click **OK**.

Deleting a file system quota

Procedure

1. Navigate to **Home > Storage Management > Quotas by File System** to display the **Quotas by File System** page.
2. Fill in the check box next to one or more quotas and then click **delete**.

Managing quotas on virtual volumes

Three types of quotas are maintained for each virtual volume:

- **Explicit User/Group Quotas.** A quota explicitly created to impose restrictions on an individual user or group, defining a unique set of thresholds.
- **Default User/Group Quotas.** A quota set automatically for all users and groups that do not have explicit quotas, set by defining a set of Quota Defaults (thresholds) for creating a quota automatically when a file is created or modified in the virtual volume. .

Default quotas for virtual volumes operate in the same way as those defined for file systems. User (Group) quota defaults define a set of thresholds for creating a quota for a user (or group) the first time that user (or group) saves a file in the virtual volume.

Initially, quota defaults are not set. When activity occurs in the virtual volume, it is tracked, but quotas are not automatically created. When at least one threshold is set to a non-zero value, a User or Group quota (as appropriate) will be created for the owner of the directory at the root of the virtual volume.

- **Virtual Volume & Quotas.** A virtual volume quota tracks the space used within a specific directory on the virtual volume . A quota can be explicitly created to define a set of thresholds restricting all operations in the virtual volume, unrelated to which user or group initiated them.



Note: Quotas track the number and total size of all files. At specified thresholds, emails alert the list of contacts associated with the virtual volume and, optionally, *Quota Threshold Exceeded* events are logged. Operations that would take the user or group beyond the configured limit can be disallowed by setting hard limits.

When *Usage* and *File Count* limits are combined, the server will enforce the first quota to be reached.

Important information about virtual volumes and quotas

The server treats the virtual volume 'root' directory, together with all its sub-directories, as a self-contained file system. The virtual volume tracks its usage of space and number of files, to provide a way of monitoring file system usage. This tracking allows quotas to be imposed on disk space usage, as well as the total number of files.

Quotas can be set for the entire virtual volume, and on individual users, and on groups of users. Default user and group quotas can be defined, and in the absence of explicit user or group quotas, the default quotas apply.

The following caveats apply in measuring the virtual volume status against quota thresholds:

- **Metadata and snapshot files.** Neither file system metadata nor snapshot files count towards the quota limits.
- **Symbolic link calculation.** Files with multiple hard links pointing to them are included only once in the quota calculation. A symbolic link adds the size of the symbolic link file to a virtual volume and not the size of the file to which it links.

Advertising NFS exports for Virtual Volumes

This feature enables or disables the reporting of 'fake' NFS exports at the roots of virtual volumes.

A file system can contain NFS exports at the root of the file system, for example, /home, and also virtual volumes for users' home directories, for example, /fred and /joe. When the automounter on an NFS client mounts one of these home directories, it sends a request to the server to obtain the list of exports. This is the same request that's used by the `showmount -e` command. On receiving this list, it selects the most suitable export to mount.

If, for example, the automounter has been told to mount /home/fred but the only suitable export it finds is /home, it first mounts /home, then changes directory to /fred. Although this is fine for general file system operation, and indeed many mounts are created this way, it does mean that quota information may not be returned correctly. The reason for this is that because the client has mounted the export /home, the free space that's reported for the mount is that of the entire file system, not of the virtual volume /fred.

In order to enable the server to report the correct quota information for the mount, the client must mount /home/fred directly, rather than mounting /home and then changing directory to /fred. An obvious way to achieve this is to add an export /home/fred. However, having to manually add an export at the root of every virtual volume is time consuming.

Therefore, the server has the facility to report a 'fake' NFS export at the root of each virtual volume. In the scenario above, with a single export /home at the root of the file system, the server reports the exports /home, /home/fred and /home/joe. On receiving this list, the client now sees that the most suitable export is /home/fred, so it mounts that path directly, therefore ensuring that the correct quota information is returned to the user fred.

Take the following information into account when using this feature:

- If UDP is used to transport MOUNT requests, only the first 64KiB of NFS export information is returned. So, TCP is recommended to ensure the complete list of NFS exports is always returned.
- The server returns a maximum of 1000 mounts in response to the `showmount -d` command.
- The Linux autofs4 client mounts all available exports, using a separate socket for each. So, if this feature causes a large number of exports to be advertised, Linux clients can end up with a large number of mounts. This can cause port exhaustion and can cause Linux clients to run very slowly.

For further information, see the `nfs-export-advertise-for-virtual-volumes` CLI command man page. This command also refreshes the list of exports manually and configures the interval at which automatic refreshes of the list occur. The list is always refreshed immediately when an export or virtual volume is added or removed, but it is also refreshed in the background at regular intervals to ensure other changes, such as directory renames, are reflected in the export list.

Viewing and modifying virtual volume quotas

This page lists all the current quotas for the specified virtual volume or the specified file systems, and allows a quota's details to be viewed.

Procedure

1. Navigate to **Home > Storage Management > Virtual Volume & Quotas**.
2. Fill in the check box next to the virtual volume to view or modify and then click **View Quotas** to display the **Quotas** page.

Storage Management [Home](#) > [Storage Management](#) > [Virtual Volumes & Quotas](#) > Quotas

Quotas

Virtual Volume		Filter	
EVS / File System: g1-eva3 / PHDS1 Virtual Volume Name: VirtualVolume203 Path: /PHDS1		Filter Quota Type: All Types	Where User/Group Account matches: <input type="text"/>
		<input type="button" value="filter"/>	

Show 20 items per page				
<input type="checkbox"/>	User/Group Account	Quota Type	Created By	Usage Limit
<input type="checkbox"/>		Virtual Volume	User Defined	100 GiB
				100

[Check All](#) | [Clear All](#)

Actions: | [User Defaults](#) [Group Defaults](#) | [Download Quotas for this Virtual Volume](#)

The following table describes the fields on this page:

Field/Item	Description
Virtual Volume	Identifies the virtual volume to which these quotas apply: <ul style="list-style-type: none"> ▪ EVS/File System: EVS and file system on which the virtual volume resides. ▪ Virtual Volume Name: Name of the virtual volume. ▪ Path: Directory on which the virtual volume has been created.
Filter	Because many user/group quotas can exist on a virtual volume, the server provides a way to filter the list. <ul style="list-style-type: none"> ▪ Filter Quota Type: You can filter on All Types (default), Users, Groups or Virtual Volumes. ▪ Where User/Group Account matches: The user or group name to be matched. The wildcard character * is permitted.
User/Group Account (also known as the target)	A quota name can consist of: <ul style="list-style-type: none"> ▪ A CIFS domain and user or group name, such as bb\Smith or bb\my_group (where <i>bb</i> is a domain, <i>Smith</i> is a user and <i>my_group</i> is a group). ▪ An NFS user or group such as richardb or finance (in which <i>richardb</i> is an NFS user and <i>finance</i> is an NFS group). <p>A name can be empty (if the quota is a virtual volume quota) or 0 (if the quota was created for the owner of the directory at the root of the virtual volume).</p>
Quota Type	Type of source of virtual volume activity. Possible values are User, Group, or Virtual Volume . The type applies to anyone initiating activity in the entire virtual volume, and only one quota with this target type may exist on each virtual volume.
Created By	Method of quota creation. Possible values are <i>Automatically Created</i> (created using a quota default) or <i>User Defined</i> (where the quota was set uniquely for one particular quota).
Usage Limit	Overall limit set for the total size of all files in the virtual volume owned by the target of the quota.
File Count Limit	Overall limit set for the total number of files in the virtual volume owned by the target of the quota.
details	Opens the Details page, in which you can view and edit the configuration of the selected quota.
add	Opens the Add Quota page.
delete	Deletes a selected quota.

Field/Item	Description
refresh cache	Clears the NAS Manager cache and repopulates the cache with the relevant objects. (This is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.)
User Defaults	Opens the User Quota Defaults page, in which you can set or change the defaults for users.
Group Defaults	Opens the User Quota Defaults page, in which you can set or change the defaults for groups.
Download Quotas	Downloads a comma-separated value (CSV) file containing all available quota information for the virtual volume.

Setting user/group defaults

This procedure illustrates the **User Default** page. The **Group Default** page is identical, except for the Automatically create quotas for Domain Users check box. This option allows default quotas for the group Domain Users to be created. By default, every NT user belongs to the group Domain Users, which includes every NT user in the quota unless each user's primary group has been explicitly set.

Procedure

1. Navigate to **Home > Storage Management > Virtual Volumes & Quotas** to display the **Quotas** page.
2. Click **User Defaults** to display the **User Quota Defaults** page.

The screenshot shows the 'User Quota Defaults' configuration page. At the top, it identifies the file system as 'EVS / File System: g5-evs1 / g5-fs1' and the virtual volume as 'fs1_vv1'. The page is split into two columns: 'Usage' and 'File Count'. Each column contains four rows of settings: 'Limit' (with a unit dropdown set to 'GiB'), 'Hard Limit' (with a checked checkbox), 'Warning' (with a percentage input field), and 'Severe' (with a percentage input field). Below these columns is a checkbox labeled 'Log Quota Events in the managed server's Event Log'. At the bottom of the page are three buttons: 'OK', 'clear defaults', and 'cancel'.

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The EVS and file system on which the user quota applies.
Virtual Volume Name	Name of the virtual volume to which a user quota created using these defaults is assigned. This option only appears when setting the quotas for a virtual volume.
Automatically create quotas for domain users	This option only appears for group quotas. It creates quotas for individual domain users.
Usage	
Limit	Amount of space to enable in Bytes: KiB, MiB, GiB, TiB, PiB, or EiB.
Hard Limit	When enabled, the amount of space specified in the Limit field cannot be exceeded.
Warning	Percentage of the amount of space specified in the Limit field at which a Warning alert will be sent.
Severe	Percentage of the amount of space specified in the Limit field at which a Severe alert will be sent.
File Count	
Limit	Maximum number of files to allow in this virtual volume.
Hard Limit	When enabled, the number of files specified in the Limit field cannot be exceeded.
Warning	Percentage of the number of files specified in the Limit field at which a Warning alert will be sent.
Severe	Percentage of the number of files specified in the Limit field at which a Severe alert will be sent.
Log Quota Events in the managed server's Event log	Selecting this check box sets the default for all users or groups to have quota events logged in the server's event log.
clear defaults	Prevents additional user quota defaults from being created in the virtual volume.

Exporting quotas for a specific virtual volume

You can export quotas for a specific virtual volume using the NAS Manager.

Procedure

1. Navigate to **Home > Storage Management > Virtual Volumes & Quotas** to display the **Virtual Volumes & Quotas** page.

2. Select a virtual volume and click **View Quotas** to display the **Quotas** page.

Storage Management [Home](#) > [Storage Management](#) > [Virtual Volumes & Quotas](#) > Quotas

Quotas

Virtual Volume
 EVS / File System: g1-eva3 / PHDS1
 Virtual Volume Name: VirtualVolume203
 Path: /PHDS1

Filter
 Filter Quota Type: All Types
 Where User/Group Account matches:
 filter

Show 20 items per page

User/Group Account	Quota Type	Created By	Usage Limit	File Count Limit
	Virtual Volume	User Defined	100 GiB	100

Check All | Clear All

Actions: [add](#) [delete](#) [refresh cache](#) | [User Defaults](#) [Group Defaults](#) | [Download Quotas for this Virtual Volume](#)

3. Click **Download Quotas for this Virtual Volume**.
4. Save the quota information to as a comma-separated value (CSV) file.
You also can choose to display the quota information in an application.

Managing virtual volumes

A file system can be divided into discrete areas of storage called virtual volumes. From a client's perspective, a virtual volume appears to be a normal file system. A virtual volume provides a simple method for allocating and controlling directories for projects, users, or groups. Capacity and number of files within a virtual volume can be controlled using quotas.

The terms user and group are used to indicate NFS or SMB users and groups.

Virtual volumes have the following characteristics:

- **Name:** A name or label by which the virtual volume is identified. This will often be the same as a SMB share or NFS export rooted at the volume's root directory.
- **File System:** The file system in which the virtual volume is created.
- **Path:** The directory at the root of the virtual volume.
- **Email Contacts:** A list of email addresses, to which information and alerts about virtual volume activity are sent. The list can also be used to send emails to individual users.
- **Moving Files:** You can move files in or out of a virtual volume.
- **Moving Directories:** Moving a directory in or out of a virtual volume will return a cross volume link error. However, most SMB or NFS clients will suppress this error and, instead, will copy the directory to the target location and then delete the original directory.
- **Linking Files** (hard links): You cannot link files across different virtual volumes.

Viewing virtual volumes

You can view the current virtual volumes for the selected EVS and file system using the **Virtual Volumes & Quotas** page in the NAS Manager. This page also allows the virtual volume's details to be viewed.

Procedure

1. Navigate to the **Home > Storage Management > Virtual Volumes & Quotas** to display the page.

Storage Management [Home](#) > [Storage Management](#) > Virtual Volumes & Quotas


Virtual Volumes & Quotas

EVS / File System Label		Filter		
g1-avs3 / protectedsiteA change...		Name: <input type="text"/>	<input type="text"/>	
		Path: <input type="text"/>	filter	

▼ Name	File System	Contact	Path	Usage
Check All Clear All				

Actions: [add](#) [delete](#) | [View Quotas](#) | [Download All Quotas](#)

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The name of the selected EVS and file system.
change	Select a file system associated with the chosen EVS.
Filter	Filters can be defined to reduce the number of virtual volumes displayed on the page and can be configured based on the name or the path.
Name	Name of the virtual volume.
File System	Name of the file system.
Contact	Contact email address for information and alerts about virtual volume activity.  Note: Only the first contact email address is shown; to view the full set of contacts or to modify the virtual volume, click details .
Path	Directory on which the virtual volume has been created.
Usage	Amount of data in the virtual volume.
details	Displays the Virtual Volume page for the selected virtual volume.
add	Displays the Add Virtual Volume page.
delete	Deletes the virtual volume and associated quota.
View Quotas	Displays quotas for the selected virtual volume.
Download All Quotas	Downloads a CSV (comma separated values) file listing all virtual volumes' configured quotas. The saved quota information includes: Quota Type, Created By, Usage, Usage Limit, Usage Hard Limit, Usage Reset (%), Usage Warning (%), Usage Severe (%), File Count, File Count Limit, File Count Hard Limit, File Count Reset (%), File Count Warning (%), and File Count Severe (%).

Adding a virtual volume


A file system can be divided into discrete areas of storage that are called virtual volumes. You can add a virtual volume in the NAS Manager.

Procedure

1. Navigate to **Home > Storage Management > Virtual Volumes & Quotas**
2. Click **add** to display the **Add Virtual Volume** page.

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The EVS and the file system to which to add this virtual volume. If the volume will be added to a different EVS/file system, click change and select an EVS/file system.
Virtual Volume Name	The name of the virtual volume.
Create a CIFS Share or NFS Export with the same name as the virtual volume	If a share or export with the same name as the virtual volume does not exist, selecting this check box ensures its creation. This is only intended for convenience in accessing the virtual volume through CIFS or NFS.
Allow exports to overlap	As overlapping exports can potentially expose security loopholes, the condition can be tested for and, if found, the export creation can be denied.
Path	Directory in the file system that will be the root of the virtual volume; for example, /company/sales. All subdirectories of this path will be a part of this volume. The path to the directory at the root of the virtual volume must be specified, or selected by browsing the file system. If the directory does not yet exist, then leaving the box checked will ensure it is created. It should be noted that if the system is left to create the directory in this way, the owner will be designated 'root', and the default quotas for this virtual volume will be named anonymously.
Email Contacts	Email contacts to receive information about virtual volume usage.

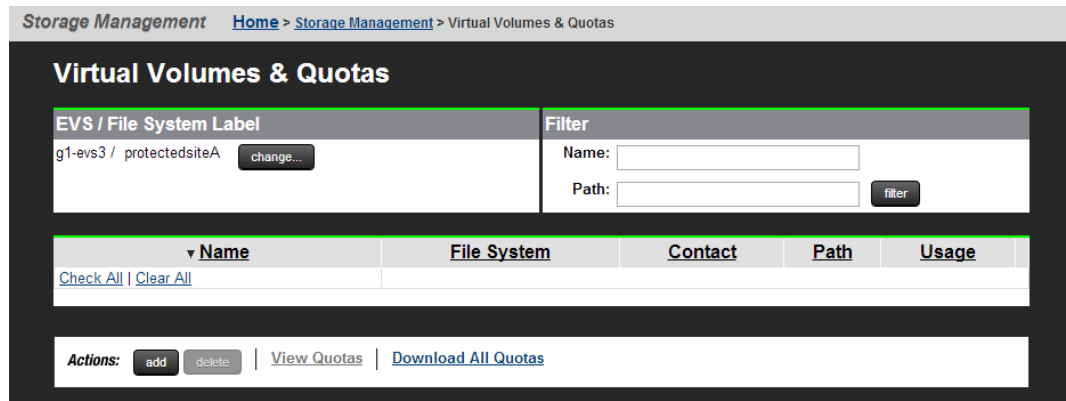
3. In the **Virtual Volume Name** field, enter the name. Note the the following:
 - The name can be up 128 characters.
 - Do not use the characters `?*+=+[]:./,<>\|` in the name.
 - The name `A$` is reserved for the Windows Event Viewer, and cannot be used.
 4. If a CIFS share of the same name as this virtual volume is required, fill the **Create a CIFS Share with the same name as the Virtual Volume** check box. Selecting this check box ensures its creation.
 5. if an NFS export with the same name as the virtual volume is required, fill in the **Create a NFS Export with the same name as the Virtual Volume** check box.
 6. If there is a possibility that this new NFS export will overlap an existing export, fill in the **Allow exports to overlap** check box.
 7. Enter the path to the directory at the root of the virtual volume or click **Browse** and navigate to the file system path.
 8. If the directory does not yet exist, fill in the **Create path if it does not exist** check box to ensure that it is created.
 9. Enter each email address in the **Email Contacts** box, and click **add** to append it to the list. Email lists are limited to a maximum of 512 characters.
 - *To configure email notification of threshold alerts*, designate explicit email recipients (for example, `admin@company.com`) to receive email notification any time a defined threshold has been reached.
 - *To send email to all affected user accounts when their user quota has been reached*, add an email address beginning with `*` to the Email Contacts list (for example, `*@example.com`).
-  **Note:** If no email contacts are specified for the virtual volume, the server generates events for quota warnings. To generate events in addition to email alerts, go to the server's command line interface and issue the command `quota-event--on`.
10. Click **OK**.

Modifying a virtual volume


You can modify the name and email address of an existing virtual volume in the NAS Manager.

Procedure

1. Navigate to **Home > Storage Management > Virtual Volumes & Quotas** to display the **Virtual Volumes & Quotas** page.



2. Click **details** to display the **Virtual Volume** page.



The following table describes the fields on this page:

Field/Item	Description
EVS/File System	The EVS and the file system to which to add this virtual volume. If the virtual volume will be added to a different EVS/file system, click change and select an EVS/file system.
Virtual Volume Name	The name of the virtual volume.
Path	Directory in the file system that is the root of the virtual volume; for example, /company/sales. All subdirectories of this path will be a part of this volume.
Security Mode	Displays the File System Security page.
Total Quota Limits Allocated	Displays the total quota limits allocated.
Total Usage	Displays the total usage excluding metadata. Use the <code>fs-analyze-data-usage</code> command to determine how much metadata exists on the file system. It is not possible to determine how much metadata exists for an individual virtual volume.
File Count	Displays the file count. For virtual volume quotas, the root of the virtual volume counts as belonging to the virtual volume - therefore an empty virtual volume displays a file count of one.
Email Contacts	Displays existing email contacts.
View Corresponding Share	Displays the corresponding shares.
View Corresponding Export	Displays the corresponding exports.
Data Migration Paths	Displays the Data Migration Paths page.

3. To modify the virtual volume name, enter the new name.
4. To add new email contacts, enter the email address and click **add**. To delete an existing email address, select the address from the list and click **x**.
5. Click **OK**.

Deleting a virtual volume

A virtual volume can only be removed when its associated directory is empty. For directories that are not empty, virtual volume removal commands provide a `--delete-contents` option to empty the directories before removing them.

- If the virtual volume is not empty, use the `--delete-contents` CLI option to remove its contents before removal:
 - `virtual-volume remove [--delete-contents] <file-system> <name>`
- To remove all virtual volumes of a given file system, issue the following command. If some volumes are not-empty, use the `--delete contents` option:
 - `virtual-volume removeall [--delete-contents] <file-system>`

Refer to CLI man pages for more information.

You can delete an empty virtual volume in the NAS Manager.

Procedure

1. Navigate to **Home > Storage Management > Virtual Volumes & Quotas**.
2. Select one or more virtual volumes.
3. Click **delete**.

A warning message displays asking for confirmation that this action is definitely required.

4. Click **OK**.

Enabling and disabling file system capacity and free space reporting based on virtual volume quotas

The file system capacity and free space reporting for virtual volume quotas option supports thin provisioning within a virtual volume. When this option is enabled and a virtual volume quota is created, capacity/free space counts returned to clients are derived solely from the virtual volume quota. This affects only those clients that have mounted an export or share within a virtual volume.

You may want to enable this option when data migration is configured. In this scenario, the primary file system could ingest more data than it has capacity itself for. You can define a quota for a virtual volume based on available capacity of migration target(s) and enable this feature so that the capacity defined by the quota is reported to protocol clients rather than the primary file system capacity/free space.



Note: This option is disabled by default.

Enabling file system capacity and free space reporting based on virtual volume quotas

- To enable this option for the virtual volume *vivol1* that resides on the file system *fs1*, issue the following CLI command:
 - `fs-space-reporting-based-solely-on-vivol-quota --on fs1 vivol1`

Displaying file system capacity and free space reporting based on virtual volume quotas

- To get the current setting for the virtual volume *vivol1* that resides on the file system *fs1*, issue the following CLI command:
 - `fs-space-reporting-based-solely-on-vivol-quota fs1 vivol1`

Disabling file system capacity and free space reporting based on virtual volume quotas

- To disable this option for the virtual volume *vivol1* that resides on the file system *fs1*, issue the following CLI command:
 - `fs-space-reporting-based-solely-on-vivol-quota --off fs1 vivol1`

See the *Command Line Reference* for more details.

Using the per-file system throttle feature

You can manage the file system performance by using per-file system throttling (PFST) commands. PFST allows you to place an upper limit on resource usage on a particular file system, which prevents a file system from using too many server resources. All PFST commands take effect on a cluster-wide basis.

To use the per-file system throttle feature, you must have Dev User access rights. The PFST feature allows you to:

- Enable and disable a PFST.
- List the PFST properties, including quota and number of queued operations.
- Create, delete, and modify PFST classes.
- Maintain mappings of file systems to PFST classes.



Note: Only NFSv2 and NFSv3 traffic is throttled. All other protocols, including SMB and NFSv4, are unaffected by throttling. For further information, refer to the `per-fs-throttle` concept CLI man page.

The following table lists the CLI commands to enable, disable, create, delete, and modify the per-file system throttle feature. For more information on a specific command, refer to the CLI man page.

Description	Command
To enable a per-fs-throttle bossock fiber quota	<code>per-fs-throttle-class-enable</code>
To disable a per-fs-throttle bossock fiber quota	<code>per-fs-throttle-class-disable</code>
To create a per-file system throttle class	<code>per-fs-throttle-class-create</code>
To delete a per-file system throttle class	<code>per-fs-throttle-class-delete</code>
To modify a per-file system throttle class	<code>per-fs-throttle-class-set</code>

Creating a read cache file system

A read cache is a special read-only file system that stores copies of individual files outside of their local file systems, enabling a server or a node to have a cached copy of the file.

Procedure

1. Navigate to **Home > Storage Management > File Systems** to display the **File System** page.
2. Click **Read Cache**.

Storage Management [Home](#) > [Storage Management](#) > [File Systems](#) > Create File System

Create File System

What do you want to create?

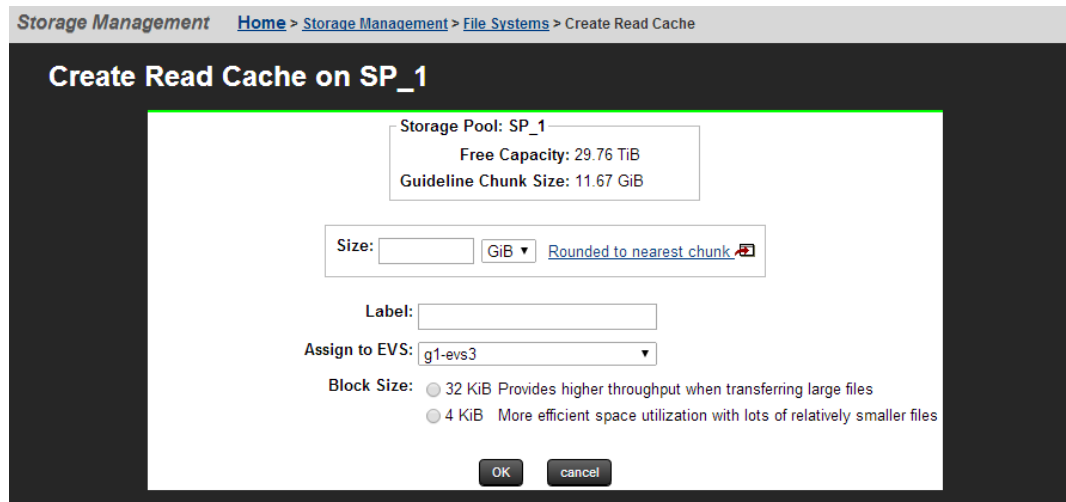
File System
 Read Cache

Select the storage pool on which to create a file system or read cache.


Label	Capacity	Free Space
<input checked="" type="radio"/> SP_1	42.72 TiB	29.75 TiB (70 %)

back next

3. Select a storage pool to contain the read cache, and then click **next** to display the **Create Read Cache** page.



The following table describes the fields on this page:

Field/Item	Description
Storage Pool	Displays the name of the storage pool in which the read cache file system is being created.
Free Capacity Amount	Displays the available space in the storage pool that can be used by read cache file systems.
Tier 0 Meta-data and Tier 1 User-data	Displays the size of the storage pool's metadata tier (Tier 0) and the user data tier (Tier 1). <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f7fa;"> <p> Note: This information applies only to tiered file systems, and is not displayed for untiered read cache file systems.</p> </div>
Guideline Chunk Size	Displays the approximate size of the chunks used in the selected storage pool.
Size	The size of the read cache being created. Enter a size for the read cache. <ul style="list-style-type: none"> For an untiered read cache, this defines the total size of the read cache. For a tiered read cache, this defines the size of the user-data tier of the read cache.


Field/Item	Description
	The size value can be changed on the File System Details page once the read cache has been created.
Rounded Size	Approximate size limit, based on the defined Size Limit and the chunk size defined for the storage pool.
Label	The label by which the read cache should be referenced.
Assign to EVS	The EVS to which the file system is assigned.
Block Size	Sets the optimal block size for the file system.

4. Enter a **Size Limit** for the file system.
This defines the maximum size to which the file system or user data tier can grow through Auto-Expansion. Once the file system has been created, this value can be changed on the **File System Details** page. This limit is not enforced for manual file system expansions performed through the CLI.
5. The value in the **Rounded Size Limit** field is calculated automatically, but can be changed. For more information, click **Rounded to nearest chunk**. If the specified size is not a multiple of the chunk size, the server rounds down to the nearest chunk boundary.
6. In the **Label** field, enter the name of the file system.
File system labels are not case sensitive, but they do preserve case (labels will be kept as entered, in any combination of upper and lowercase characters). Also, file system labels cannot contain spaces or any of the following special characters: "&*/<>?|. File system labels must be unique within a server or cluster. Also, a file system cannot have the same label as a storage pool.
7. From the EVS list, select the **EVS** to which the file system should be assigned.
8. In the **Block Size** field, enter the desired file system block size.
9. Click **OK**.


Monitoring file system operations

File system operations per second can be monitored from the File System Ops/Sec page in the Status and Monitoring section of the NAS Manager.

The graph on this page displays the number of operations received at the interface into the filesystem and includes internally-generated I/O such as replication as shown below:



Protocol operations can be monitored from the Performance Graphs page (Node Ops/Sec) in the Status and Monitoring section of the NAS Manager. Click on the Node Ops/Sec graph to open a page containing a separate graph for each node as shown below:



Note: Protocol operations do not map 1:1 to filesystem operations. Some protocols operations do not cause any filesystem operations, and some protocol operations require several filesystem operations.

In a typical configuration it is very common for the filesystem ops/s to be higher than that the protocols ops/s.

Chapter 3: Managing file system security

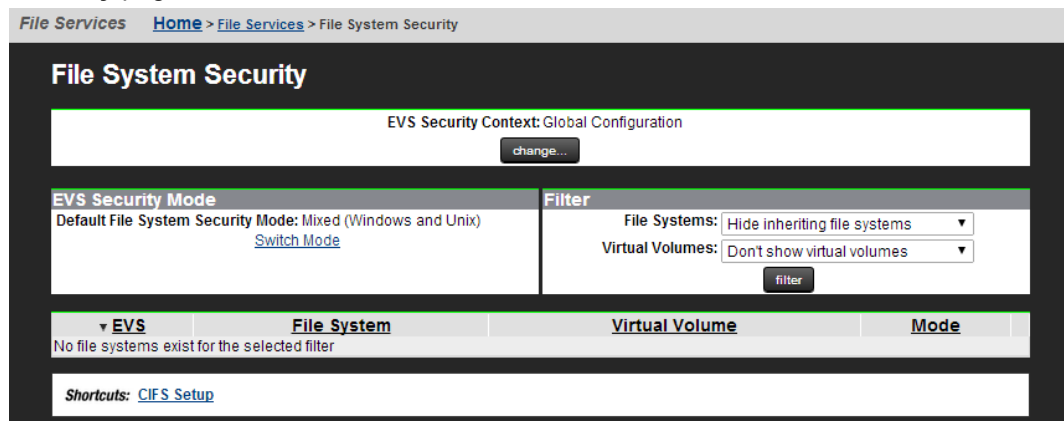
Security modes can be configured per-cluster/server, per-file system, or per-virtual volume. Selecting security modes on a tiered basis, rather than system-wide, enhances the granularity and convenience of managing system security.

Viewing file system security

Security modes can be configured per-EVS, per-file system, or per-Virtual Volume. You can view the EVSs, the file systems and the configured security modes in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > File System Security** to display the **File System Security** page.



The following table describes the fields on this page:

Field/Item	Description
EVS Security Context	Displays the currently selected EVS security context.
change	Selects a different EVS security context. You can select either Global Configuration which applies to all EVSs, or select a specific EVS.
EVS Security Mode	Displays current EVS security mode settings, and allows you to change those settings.
Default File System Security Mode	Indicates the default security mode that is in effect for the entire EVS. Click the Switch Mode link to switch the security mode for the entire EVS. You can switch between <code>Mixed mode</code> and <code>UNIX mode</code> .
Filter	Control the information displayed in this page. In the File Systems field, select whether to show file systems. In the Virtual Volumes field, select whether to show virtual volumes. Click filter to refresh the page based on the criteria selected in these two fields.
EVS	List of all virtual servers (EVSs) defined by the filter.
File System	<i>If this column is blank</i> , the displayed security mode is associated with the EVS. <i>If this column displays a file system label</i> , the displayed security mode is associated with this specific file system.
Virtual Volume	Lists the virtual volumes found on the file systems defined by the filter.
Mode	Security mode defined on the EVS or file system. File systems without an explicit security mode configuration inherit security mode from the EVS.
details	Advances to the Security Configuration page in which the security mode for the selected EVS can be modified.
CIFS Setup	Advances to a page in which CIFS setup can be performed.

NFS security and Kerberos

The NAS server supports Kerberos to provide authentication, integrity, and privacy when using NFS v2, v3, and v4. Kerberos provides a mechanism for entities (principals) to authenticate to each other and securely exchange session keys. The NAS server supports RPCSEC_GSS using Kerberos v5.

The Kerberos implementation has been updated with the Advanced Encryption Standard (AES). The Data Encryption Standard (DES) has been deprecated and is insufficiently secure.

Secure NFS requires configuration of the NFS server's Kerberos principal name, and secret keys. Kerberos related configuration settings are setup both globally and on a per-EVS basis. The NFS host name is configured on a per-EVS basis.

Kerberos principal formats

A Kerberos principal can take different forms, containing a varying number of components.

For the purposes of this feature, a principal can contain the following parts:

- **Primary** - this would typically be a username.
- **Instance** - this is an optional part which qualifies the primary. It can be a user role or a host name.
- **Realm** - this is the Kerberos realm which is usually a domain name.

In your environment, a principal could take the following form:

```
primary@REALM
```

For example:

```
user@EXAMPLE.COM
```

This user could operate on multiple clients.

Alternatively, your principal could take the following form:

```
primary/instance@REALM
```

For example:

```
user/machine1.example.com@EXAMPLE.COM
```

```
user/machine2.example.com@EXAMPLE.COM
```

Here, you could have different Kerberos principals that map to a single user.

To support these environments, the NAS server provides the **krb5-nfs-principal-format** command. By default, the Kerberos principal is unchanged before being mapped onto a user. For multiple Kerberos principals that are mapped to a single user, the **only-primary** option changes the Kerberos principal to `primary@REALM` before mapping it to a user. By using this setting, the principals in the second example above are interpreted as `user@EXAMPLE.COM` and so would require only a one-to-one mapping.



Important: It is expected that the **only-primary** configuration option is only selected when first configuring a security context and care should be taken when modifying the setting for an existing security context. If the setting is changed after Kerberos authentication for NFS has been used with the security context, this can result in a mixture of Kerberos Principal formats being stored in on-disk security. This is likely to result in users being unable to access files, which would need to be remedied by manually correcting the on-disk security.

Setting secure NFS

NFS supports three secure options:

- `krb5`: provides authentication only
- `krb5i`: provides authentication and integrity
- `krb5p`: provides authentication, integrity, and privacy (encryption of data)

NFS exports can be set to accept only secure connections. This is done by specifying the appropriate security options in the Access Configuration field of the **Add Export** page or the **NFS Export Details** page.

Setting the type of secure connections can be done using the CLI command `nfs-export` with the `mod -c` option. See the *Command Line Reference* for more information.

Mixed security mode

The server's mixed security mode supports both Windows and UNIX security definitions. Security is set up uniquely on each file (or directory), based on which user created, or last took ownership of, the file (or directory). If a Windows user, the security definition will be native CIFS and subject to Windows security rules; likewise, if a UNIX user, the security definition will be native NFS and subject to UNIX security rules.

AES support for SMB

Windows Vista and Windows Server 2008 introduced support for the Kerberos AES crypto profiles, in addition to the older crypto profiles (DES/DES3 and RC4) already implemented in earlier Windows versions.

The SMB implementation supports the new AES crypto profiles. The supported AES crypto profiles are:

- AES256: HMAC-SHA1-96 (the default if AES is supported)
- AES128: HMAC-SHA1-96. To force AES-128 encryption:
 - Configure the DC only: Set `msDS-SupportedEncryptionType 0x8 = (AES128_CTS_HMAC_SHA1_96)`.
 - Run `klist purge` on the client.



Note: Windows Server 2008 or higher is required.



Note: The `cifs-keytab-list` command can be used to display the encryption types supported by a CIFS name.

Configuration to Support AES with Existing CIFS Names (created on 12.2 or earlier)

- No configuration is required for existing CIFS names. AES is automatically enabled on upgrade to 12.3 or later.
- However, configuration is required on the DC for existing CIFS names. AES must be added to the supported encryption types list of existing CIFS names computer accounts.

Configuration to Support AES with New CIFS Names (create on 12.3 or later)

- No configuration is required for newly created CIFS names.



Note: The client credentials cache should be purged if the Kerberos configuration is changed on either the DC or clients (for example, to change the supported encryption types list). On windows clients this can be done using: `klist purge`.

Upgrades and downgrades

- For an upgrade (from a 12.2 or earlier to a 12.3 or later) , AES must be added to the supported encryption types of existing CIFS name DC computer accounts.
- For a downgrade (from a 12.3 or later to a 12.2 or earlier), AES must be removed from the supported encryption types of DC computer accounts for CIFS names that were created with 12.3 or later, or had AES explicitly enabled as per the above upgrade consideration. Otherwise, SMB authentication will fall back to NTLM.

SMB access to native SMB files

When an SMB client tries to access a native SMB file (that is, with Windows security information), the server checks the user information against the file's security information to determine whether an operation is permissible:

- **User Security.** This information is contained in an access token, which is made up of the user security identifier (SID), primary group SID, and other SIDs. The server receives the token from the domain controller and caches it for use throughout the user's session.
- **File Security.** This information is contained in a file's security descriptor, which is made up of the owner SID, group SID, and access control list (ACL). The ACL can contain several access control entries (ACEs), which specify the conditions for access.

ACE entries can be modified or deleted using a set of CLI commands called the `cacls` commands. This set of commands includes `cacls-add`, `cacls-del`, `cacls-fields`, `cacls-mask-in`, `cacls-mask-out`, and `cacls-set`. For more information on these commands, refer to the *Command Line Reference*.



Note: SMB can assign rights to machine (computer) accounts. A machine account is generated automatically by the operating system and registered in Active Directory. It can be used for authentication within a domain. A machine account authentication can be only done by an application which has built-in support. For example, Hyper-V server allows storing virtual machines on remote shares. Such shares should allow full access for the machine account of a computer running Hyper-V server. The feature acts the same way as authentication of a normal user for an SMB session. Authenticated connection using machine account will show up in "connection" command output as it was a normal user connection. Man pages for `cifs-saa` and `cacls-add` include an example of computer account use.

NFS access to native NFS files

When an NFS client tries to access a native NFS file (with UNIX security information), the server checks the user's UNIX credentials against the file's security information to determine whether or not an operation is permissible. The file security information is made up of a user ID, group ID, and read, write, and execute permissions.

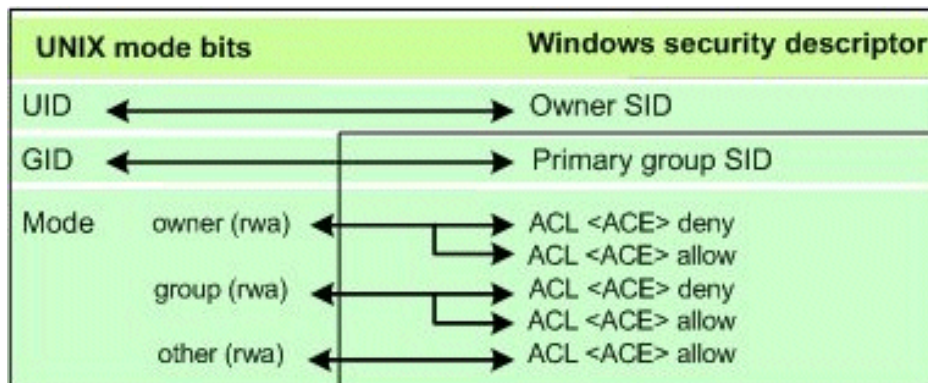
Client access to non-native files

SMB users may access files which have UNIX security information, and NFS users may access files which have Windows security information. The server supports this functionality with mapping tables, set up in NAS Manager, that associate the names of NFS users and groups with their Windows equivalents. For example, when an SMB user tries to access a file that has UNIX-only security information, the server automatically maps the user name to the corresponding NFS name in the mapping table.

- The server automatically translates user security information from UNIX to Windows format, or vice-versa, and caches it for the duration of the session:

UNIX credential		NT access token
UID	User mapping table	User SID
GID	Group mapping table	Primary group SID
Other groups	Group mapping table	Other groups

- The system automatically converts file security attributes from Windows to UNIX format and stores the result in file metadata, making the files native to both SMB and NFS clients. Although UNIX files are also converted to Windows format, the results are not stored in file metadata:



- Any changes that a user makes to a file's security attributes are applied equally to Windows and UNIX.

When an SMB user tries to access a file that has UNIX-only security information, the server maps the user to an NFS name and converts the user's access token to UNIX credentials. It then checks these credentials against the file's security attributes to determine whether or not the operation is permissible.

Similarly, when an NFS user tries to access a file that has Windows-only security information, the server maps the user to a Windows name and converts the user's UNIX credentials to a Windows access token. It then checks the token against the file's security attributes.

UNIX security mode

When the server is configured in UNIX security mode, it supports UNIX security for SMB and NFS clients. However, all security settings are saved with UNIX file attributes. As a result, NFS clients are always accessing files in native mode, while SMB clients are always accessing file non-native mode.



Note: With UNIX security mode, NFS users do not need to rely on the presence of a Windows domain controller (DC) in order to access files. As a result, they are fully isolated from potential DC failures.

Changing security modes

By default, the virtual volume or file system inherits the parent file system's security mode. In other words, when the parent file system has a UNIX security mode, the virtual volume associated with the file system will inherit the UNIX security mode. You can use the NAS Manager to modify the inherited security mode for the selected virtual volume.

Procedure

1. Navigate to **Home > File Services > File System Security** to display the **File System Security** page.
2. Click **details** to display the **Security Configuration** page.

The following table describes the fields on this page:

Field/Item	Description
EVS	Displays the selected virtual volume's parent EVS.
File System	Displays the selected virtual volume's parent file system (when selected from the File System Security page filter options).
Virtual Volume	Displays the selected virtual volume whose security mode can be changed (when selected from the File System Security page filter options).
Mode	The security mode defined on the virtual volume or file system, depending on the selection from the File System Security page filter options).

3. Select a security mode from the **Mode** list.
4. Click **OK**.

Mixed mode operation

The storage server allows network clients to share a common pool of storage on both Windows and UNIX clients. This is called mixed mode operation. Although the server does this as seamlessly as possible, the two protocols are considerably different, so mixed mode operation presents some challenges, discussed in the **File name representation** and **Symbolic links** sections.

File name representation

The maximum length of a file name is 255 characters, and file names may contain any Unicode character. Case-sensitivity in file names is significant to NFS and FTP clients, but not SMB clients.

Modern Windows (CIFS/SMB) clients (NT and newer) can make full use of UCS-2 (a two byte Unicode encoding). When communicating with Windows 9x, the server uses the Latin-1 version of extended ASCII.

When communicating with NFSv2 or NFSv3 clients, the server supports the Latin-1, UTF-8, EUC-KR, EUC-JP, and EUC-CN encodings. For NFSv2 or NFSv3 clients, the Latin-1 character set is the default. When communicating with NFSv4 clients, the server supports only the UTF-8 encoding.

If you plan to have file names that include non-ASCII characters, you should change the default encoding used by the server using the `protocol-character-set` command.

Symbolic links

Symbolic links (symlinks) are commonly used:

- To aggregate disparate parts of the file system.
- As a convenience, similar to a shortcut in the Windows environment.
- To access data outside of a cluster. For example, a symlink can point to data in another server in a server farm or a non- server.

There are two types of symlinks:

- **Relative symlinks** contain a path relative to the symlink itself. For example, `./dst` is a relative symlink.
- **Absolute symlinks** contain a path relative to the root of the file system on the NFS client that created the symlink (not relative to the root of the server's file system). For example, `/mnt/datadir/dst` is an absolute symlink.

When accessing the file system through NFS, the server fully supports symlinks. NFS/UNIX clients assume that files marked as symbolic links contain a text pathname that the client can read and interpret as an indirect reference to another file or directory. Any client can follow a symlink, but accessing the target file (or directory) still requires permission.

Clients using SMB1 cannot follow files marked as symlinks. For these files the server provides a server-side *symlink following capability*. When an SMB or FTP client accesses a server-side symlink, the server reads the path from the link and attempts to follow it automatically:

- **For relative symlinks**, the link can be followed because the server can follow the path from the link itself.
- **For absolute symlinks**, the server does not have access to the root of the file system on the NFS client that created the link, so it cannot follow the link automatically.

The server provides global symlinks, which allow clients using SMB1 to follow absolute symlinks:

- If an absolute symlink refers to a file or directory in the same SMB share as the symlink, the server follows the symlink (on behalf of the SMB client) internally.
- If an absolute symlink refers to an object in a different SMB share to the symlink, the SMB client is redirected to the link's destination via the Microsoft DFS mechanism.



Note: The Microsoft DFS mechanism supports redirection only to a directory (not a file). Therefore, absolute symlinks that refer to a file (in a different SMB share) will not be handled properly by an SMB client.

The link's destination may be on the same file system as the link, on a different file system within a server farm, or on a remote SMB server. To associate a global symlink with an absolute symlink, the server maintains a translation table between absolute symlink paths and global symlink paths.



Note: As of release 12.2 of the NAS Platform, clients using SMB2 or later can follow relative and absolute symlinks to files on storage without using server-side symlinks.

When accessing server-side symlinks, SMB clients cannot follow some symlinks which are perfectly valid for NFS, because the storage system follows the symlink on the SMB client's behalf and presents the linked-to file instead of the symlink. In this case, in line with the behavior of Samba, the server hides the existence of the symlink entirely from the SMB/FTP client. By default, a symlink that points outside of the scope of its own share (for example, to a different file system) is not followed.

Clients using SMB2 or later can follow relative and absolute symlinks to files on storage without using server-side symlinks. The following CLI commands manage this client-side SMB2 behavior:

- `smb2-client-side-symlink-handling-default`
- `smb2-client-side-symlink-handling-disable`
- `smb2-client-side-symlink-handling-enable`
- `smb2-client-side-symlink-handling-status`

For more information about the CLI commands, see the **Command Line Reference**.

Global symlinks (also called absolute symlinks) start with a slash character (/), and they allow you to set up links to data outside a cluster. NFS clients follow the global symlink directly and, for SMB clients, the server maintains a server-side translation table, that allows those clients to access the symlink destination. Both NFS and SMB clients can follow the same global symlink to the destination directory, when the global symlink, the exports, shares, and mount points are set up correctly. When a client encounters a global symlink:

- **For NFS clients**, the server returns the content of the global symlink, allowing the client to follow the link to the destination. This means that the NFS client's mount points and the NFS exports must be set up correctly.
- **For SMB clients**, the server causes the client to request a symlink lookup from the local EVS translation table. Once the client requests the lookup, the server returns the destination server name, share name, and path to the SMB client, allowing it to access the destination.



Caution: Symlink Destination Directory Alert! After the SMB client follows the path for the global symlink, it may not ask the server for another lookup for that symlink for an extended period of time. Because the symlink is not looked up every time the client follows the symlink, if the destination directory is changed or deleted, the SMB client may attempt to connect to the wrong destination, possibly causing the client to report an error.

Using global symlinks with SMB has a performance penalty. Therefore, global symlinks are disabled by default, but can be enabled by selecting the Follow Global Symbolic Links check box on the **Add Share** page (when creating the share) or **CIFS Share Details** page (after the share has been created).

Symlink translation tables are maintained on a per-EVS basis, meaning that:

- **Table entries do migrate with the EVS.** If an EVS is migrated, all of its table entries migrate along with the EVS.
- **Table entries do not replicate from the EVS.** When replicating data from one EVS to another, the mapping information for global symlinks is not automatically relocated, and it must be recreated in the translation table of the EVS into which the data was copied.
- **Table entries do not move with a file system.** If a file system is moved from one EVS to another, the mapping information for global symlinks is not automatically relocated and must be manually adjusted, except for those symlinks that are relative to a CNS tree (those symlinks do not require adjustment).
- **Table entries irrelevant for symlinks that are relative to a CNS.** When an EVS is migrated, no adjustment is necessary for symlinks that are relative to a CNS because, when the client follows the symbolic link, it is first referred to the CNS tree, then from the CNS tree to a real file system when the path crosses a CNS link.

The following CLI commands manage the symlink translation table:

- `global-symlink-add`
- `global-symlink-del`
- `global-symlink-delall`
- `global-symlink-list`

For more information about the CLI commands, see the **Command Line Reference**.

Mixed mode operation and LDAP servers

The storage server supports mixed mode access for file systems, meaning that a mapping is required between the file system permissions and owners in order to ensure consistent security and access. NIS/LDAP services allow the server to locate and map users and permissions based on an existing NIS/LDAP service on the network, instead of creating a local account on the storage server.

On an existing LDAP service, one of the following methods is typically used for allowing the server to locate and map users and permissions:

- RFC 2307 / RFC 2307bis schemas

RFC 2307 defines a standard convention for the storage and retrieval of user and group mapping information from an LDAP server. If your site uses the RFC 2307 (or RFC 2307bis) schema, and you configure your storage server/cluster to support both mixed mode operations and LDAP services, it is assumed that you have already loaded the RFC 2307 schema into your directory, and that you have already provisioned the user objects appropriately. This is the default method.

- Microsoft Active Directory schema

This setting configures your server to operate with Microsoft Active Directory 2012 and newer using the default Active Directory schema.

You can also configure the server to operate with two deprecated Microsoft LDAP services:

- Microsoft Windows Services for UNIX (SFU) schema
- Microsoft Identity Management for UNIX (IMU) schema

To ensure optimum performance when your server/cluster is configured to support both mixed mode operations and LDAP services, the most optimized configuration includes the creation of indexes in the LDAP service for attributes queried by the storage server. To ensure fastest responses to queries, exact-match indexes should be configured on the LDAP server for the attributes to be searched. The LDAP server on your network should index at least the following attributes:

Objects that:	RFC 2307 Class	Active Directory Class (also IMU and SFU)	Map to NIS Class
Describe user accounts	posixAccount	user	posixAccount
Describe the group identifier	posixGroup	group	posixGroup

Attributes for:	RFC 2307 Attribute	Active Directory Attribute	Services for UNIX Attribute	Identity Management for Unix Attribute	Map to NIS Attribute
User ID/login name	uid	sAMAccountName	sAMAccountName	uid	memberUid
User ID number	uidNumber	uidNumber	msSFU30UidNumber	uidNumber	uidNumber
Group name	cn	sAMAccountName	cn	cn	memberNisNetgroup
Group ID number	gidNumber	gidNumber	msSFU30GidNumber	gidNumber	gidNumber

To track indexing performance, you can use the `ldap-stats` command, which permits you to monitor response times for LDAP queries. It is necessary to first let the storage server complete some successful user lookups so that some statistical data can be gathered. In a short period of time, however, you should be able to determine whether any of the attributes are not indexed.

Mandatory and advisory byte-range file locks in mixed mode

The server supports both SMB and NFS locks. The server supports only one type of byte range locking for SMB clients. For NFS clients, however, additional byte range lock types are supported. When created by NFSv2 and NFSv3 clients, the server supports both monitored (NLM/NSM) and non-monitored advisory byte-range locks. When created by NFSv4 clients, the server supports both mandatory and advisory byte range locks.

The following tables describe the server behavior when a client locks a file using a byte range lock:

- Server default configuration (the server treats byte range locks by NFSv4 clients as mandatory):

Locked by client of type	Lock type when the client accessing the locked file is:		
	NFSv2 or NFSv3	NFSv4	SMB
NFSv2 or NFSv3	Advisory	Mandatory	Mandatory
NFSv4	Mandatory	Mandatory	Mandatory
SMB	Mandatory	Mandatory	Mandatory

- When the server is configured to treat byte range locks from NFSv4 clients as advisory:

Locked by client of type	Lock type when the client accessing the locked file is:		
	NFSv2 or NFSv3	NFSv4	SMB
NFSv2 or NFSv3	Advisory	Advisory	Mandatory
NFSv4	Advisory	Advisory	Mandatory
SMB	Mandatory	Mandatory	Mandatory

To change the server configuration so that it treats NFSv4 byte-range locks as advisory, use the command `set nfsv4-locking-is-advisory-only 1`.

To change the server configuration so that it treats NFSv4 byte-range locks as mandatory, use the command `set nfsv4-locking-is-advisory-only 0`.

Opportunistic locks (oplocks)

An oplock is a performance-enhancing technique used in Microsoft networking (SMB) environments. It enables applications to speed up file access and minimize network traffic by caching all or part of a file locally. As the data is kept on the client, read and write operations can be performed locally, without involving the server.

The server supports three categories of oplocks:

- **Exclusive.** An Exclusive oplock enables a single client to cache a file for both *reading and writing*. As the client that owns the oplock is the only client accessing the file, it can read and modify all or part of the file locally. The client does not need to post any changes to the server until it closes the file and releases the oplock.
- **Batch.** A Batch oplock enables a single client to cache a file for both *reading and writing*, as in the case of an exclusive oplock. In addition, the client can preserve the cached information even after closing the file; file open and close operations are also performed locally. The client does not need to post any changes back to the server until it releases the oplock.
- **Level II.** A Level II oplock enables multiple clients to cache a file for *reading only*. The clients owning the oplock can read file data and attributes from local information, cached or read-ahead. If one client makes any changes to the file, all the oplocks are broken.

When dealing with oplocks, the server acts in accordance with the SMB specification. Whether operating in a pure Windows environment or with a mix of SMB and NFS clients, the server allows applications to take advantage of local caches while preserving data integrity.

Exclusive and batch oplocks

An Exclusive or Batch oplock is an exclusive (read-write/deny-all) file lock that an SMB client may obtain at the time it opens a file. The server grants the oplock only if no other application is currently accessing the file.

When a client owns an Exclusive or Batch oplock on a file, it can cache part or all of the file locally. Any changes that the client makes to the file are also cached locally. Changes do not need to be written to the server until the client releases the oplock. In the case of an Exclusive oplock, the client releases the oplock when the server requests that it does so, or when it closes the file. In the case of a Batch oplock, the client may keep information (including changes) locally even after closing the file. While the client has an Exclusive or Batch oplock on a file, the server guarantees that no other client may access the file.

If a client requests access to a file that has an Exclusive or Batch oplock, the server asks the client with the oplock to release it. The client then writes the changes to the server and releases the oplock. Once this operation has finished, the server allows the second client to access the file. This happens regardless of the second client's network protocol.

In cases where an SMB client requests an oplock on a file that has an Exclusive or Batch oplock, the server breaks the existing oplock and grants both clients Level II oplocks instead.

Level II oplocks

A Level II oplock is a non-exclusive (read-only/deny-write) file lock that an SMB client may obtain at the time it opens a file. The server grants the oplock only if all other applications currently accessing the file also possess Level II oplocks:

- *If another client owns an Exclusive or Batch oplock*, the server breaks it and converts it to a Level II oplock before the new client is granted the oplock.
- *If a client owns a Level II oplock on a file*, it can cache part or all of the file locally. The clients owning the oplock can read file data and attributes from local information without involving the server, which guarantees that no other client may write to the file.
- *If a client wants to write to a file that has a Level II oplock*, the server asks the client that has the oplock to release it, then allows the second client to perform the write. This happens regardless of the network protocol that the second client uses.

User and group names in NFSv4

In NFSv4 users and groups are identified by UTF-8 strings of the form: `user@dns_domain` and `group@dns_domain`. The NAS Platform supports the following universal user/group identifiers:

OWNER@	The owner of a file.
GROUP@	A file's group.
EVERYONE@	The world.
NETWORK@	Accessed via the network.
ANONYMOUS@	Accessed without any authentication.
AUTHENTICATED@	Any authenticated user.

NFSv4 requires a mapping between Unix and NFSv4 names.

Configuring user and group mappings

When the server is operating in either mixed or UNIX security mode, it creates mappings between UNIX and Windows users and groups. For example, user John Doe could have a UNIX user account named `jdoe` and a Windows user account named `johnd`. These two user accounts are made equivalent by setting up a user mapping. Furthermore, the server assumes that equivalent user and group names are the same for both environments. For example, if no explicit mapping is found for user `janed`, the server assumes that the UNIX user account named `janed` is the same as the Windows user account with the same name.

There are two steps to follow when setting up user and group mappings on the server:

- Specify each NFS user and group's name and ID. Note that this step is not required for Windows users or groups, as the server obtains all of the information it needs from the domain controller (DC).
- Map the NFS user (group) names to Windows NT user (group) names.

Managing NFS user and group mapping

Windows access to a file created by a UNIX user (or vice-versa) is permitted when the UNIX name and Windows name are recognized as being the same user. However, NFS clients present an NFS operation to an NFS server with numerical UNIX User ID (UID) and UNIX Group ID (GID) as credentials. The server must map the UID and GID to a UNIX user or group name prior to verifying the UNIX to Windows name mapping.

The server uses the following methods to map from a numerical UNIX UID or GID to a UNIX user name or group name:

- If the server is configured to use the Network Information Service (NIS) no special configuration steps are needed; the server automatically retrieves the user (group) names and IDs from the NIS server.
- NFS user and group names can be added manually.
- NFS user and group names can be added by importing files. For example, the UNIX `/etc/passwd` file can be imported, providing the server with a mapping of user name to UID. The `/etc/groups` file should also be imported to provide the server with a mapping of Group name to GID.
- You can import the numerical ID to Name mappings directly from a NIS server or an LDAP server if one has been configured. Every time a UID is presented to the server, it will issue an NIS request to an NIS server to verify the mapping. This mapping can remain cached in the server for a configurable time. A cached ID to name binding for a User or Group will appear as Transient in the NFS Users or Groups list.



Note: When a Windows user creates a file and the UNIX user or group mapping fails, the server sets the UID or the GID to 0 (root). In previous releases, the server sets the UID or GID to 0 (root) or to 65534 (nobody).

Viewing NFS user mappings

Each UNIX user name and numerical UID can be manually entered, along with its corresponding Windows user and domain name. Users configured manually will appear as permanent in the NFS users list.

Procedure

1. Navigate to **Home > File Services > User Mapping** to display the **User Mapping** page.

File Services [Home](#) > [File Services](#) > User Mappings

User Mappings

EVS Security Context: Global Configuration [change](#)

Filter

Name:

UID: to

Include Discovered Information:

[filter](#)

Show: items per page

<input type="checkbox"/>	NFSv2/3 Name	UID	Windows Name	Windows ID	NFSv4 Name	Kerberos Name	
<input type="checkbox"/>			BUILTIN\Current Owner	S-1-5-32-21061	OWNER@		details
<input type="checkbox"/>	mccler	1045	TRAINING\mccler	S-1-5-21-3957838307...89-3186263376-1356			details
<input type="checkbox"/>	patrick	1134	TRAINING\patrick	S-1-5-21-3957838307...89-3186263376-1668			details


[Check All](#) | [Clear All](#) user mappings 1-20 of 21 : Page:

Actions: [add](#) [delete](#) [refresh cache](#) | [Refresh Mappings](#) [Import Users](#) [View Domain Mapping](#)

Shortcuts: [Group Mappings](#)

[Home](#) | [About](#) | [Sign Out](#)

The fields and items on this page are described in the table below:

Field/Item	Description
EVS Security Context	Displays the currently selected EVS security context. Click change to select a different EVS security context or to select the global configuration. Selecting a different EVS security context changes the EVS to which the mapping applies.
Filter	Filter the list of user mappings using any of the following criteria: <ul style="list-style-type: none"> ▪ Name, which applies to the NFSv2/v3, NFSv4 user names or the Windows user name. ▪ UID, which can be used to specify a range of UID values to display, or a minimum/maximum UID value to display. ▪ Fill the Include Discovered Information check box to also display information that has been discovered from NIS servers, LDAP servers, or domain controllers.
NFSv2/3 Name	User name configured in the UNIX environment.
UID	User ID configured in the UNIX environment.
Windows Name	User name configured in the Windows environment.
Windows ID	User ID configured in the Windows environment.
NFSv4 Name	Displays the NFSv4 user name.
Kerberos Name	Displays the Kerberos principal (of the form <code>user@realm</code>) for the user.
details	Changes the properties of a mapping, or to display more detailed information about a mapping.
add	Opens the Add User Mapping page.
delete	Deletes the selected mapping.
refresh cache	Clears the NAS Manager cache, and then repopulates the cache with the relevant objects. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;">  Note: This is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache. </div>
Refresh Mappings	Refreshes the mappings without clearing the cache.
Import Users	Opens the Import User Mappings page.
View Domain Mapping	Opens the Domain Mappings page.

Field/Item	Description
Group Mappings	Opens the Group Mappings page.

- If necessary, click **change** to select a different EVS security context or to select the global configuration.
The EVS Security Context displays the currently selected EVS security context. Changes made to mappings using this page apply only to the currently selected EVS security context.
 - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
 - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the mappings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

Creating an NFS user mapping

Each UNIX user name and numerical UID can be manually entered, along with its corresponding Windows user and domain name. Users configured manually will appear as permanent in the NFS users list.

Procedure

- Navigate to **Home > File Services > User Mapping** to display the **User Mapping** page.
- Click **add** to display the **Add User Mapping** page.

The following table describes the fields and items on this page:

Field/Item	Description
NFS v2/3 Name	The NFS user account name.
Unix ID	The Unix user identifier.
Windows Name	The Windows user account name.
Windows ID	The Windows group identifier.
NFSv4 Name	The NFSv4 user name.
Kerberos Name	The Kerberos principal (of the form user@realm) for the user.

3. Select the check box under each selection as appropriate:
 - **Save to NAS server** - The server relies on information you provide.
 - **Discover** - The server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.
 - **Ignore** - The server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers

Modifying user mappings

You can modify user mappings in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > User Mapping**.
2. Select the check box on the user mapping to modify and then click **details** to display the **User Mapping Details** page.

Field/Item	Description
NFS v2/3 Name	Displays the NFSv2/3 user account name.
Unix ID	Displays the UNIX user identifier.
Windows Name	Displays the Windows user account name.
Windows ID	Displays the Windows identifier
NFSv4 Name	Displays the NFSv4 user name.
Kerberos Name	Displays the Kerberos principal (of the form user@realm) for the user.

3. Make any changes as necessary to the fields described in the table.
 - **Save to NAS server** - The server relies on information you provide.
 - **Discover** - The server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.
 - **Ignore** - The server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers

Deleting a user mapping

You can delete a user mapping in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > User Mapping** to display the **User Mapping** page.
2. Select the check box next to the NFSv2/3 name of the user mapping to delete, and click **delete**.
3. Click **OK** to confirm the deletion.

Adding group mappings manually

You can add group mappings manually in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > Group Mappings** to display the **Group Mappings** page.
2. Click **add** to display the **Add Group Mapping** page.

The screenshot shows the 'Add Group Mapping' page with the following fields and options:

- NFSv2/3 Name:** [Text Input]
 - Save to NAS server
 - Discover
 - Ignore
- Group ID:** [Text Input]
 - Save to NAS server
 - Discover
 - Ignore
- Windows Name:** [Text Input]
 - Save to NAS server
 - Discover
 - Ignore
- Windows ID:** [Text Input]
 - Save to NAS server
 - Discover
 - Ignore
- NFSv4 Name:** [Text Input]
 - Save to NAS server
 - Discover
 - Ignore

Buttons: **OK** **cancel**

The following table describes the items and fields on this page:

Field/Item	Description
NFSv2/3 Name	Enter the NFS group account name.
Group ID	Enter the UNIX group identifier.
Windows Name	Enter the Windows group account name.
Windows ID	Enter the Windows group identifier.
NFSv4 Name	Enter the NFSv4 group account name.

3. Select the check box under each selection as appropriate:
 - **Save to NAS server** - The server relies on information you provide.
 - **Discover** - The server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.
 - **Ignore** - The server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers

Modifying group mappings

You can make changes to the group mapping in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > Group Mappings**.
2. Select the check box next to the group mapping to modify, and click **details** to display the **Group Mapping Details** page.

The following table describes the items and fields on this page:

Field/Item	Description
NFSv2/3 Name	NFS group account name
Save to NAS server	Requires that you manually enter the mapping name or ID.
Discover	The server uses information discovered from NIS servers, LDAP servers or domain controllers for the selected mapping.
Ignore	Grays out the name or ID field and the server does not use this information.
Group ID	UNIX group identifier
Windows Name	Windows group account name
Windows ID	Windows group identifier
NFSv4 Name	NFSv4 group account name

3. Make changes as necessary and click **OK**.

Deleting a group mapping

You can delete a group mapping in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > Group Mappings**.
2. Select the check box next the group mapping to delete, and click **delete**.
3. To confirm the deletion, click **OK**.

About importing user or group mappings from a file or an NIS LDAP server

You can specify user or group details by importing them from a file.

NFSv4 user and group names are distinct from the UNIX name associated with UNIX UIDs and GIDs. However, in many environments a user/group's NFSv4 name can be derived from their UNIX name by appending the NFSv4 domain. The storage server can perform this conversion automatically, based on the settings specified on the Domain Mappings page of NAS Manager or through the CLI command `domain-mappings-add`.

To display the Domain Mappings page, navigate to Home > File Services, select User Mapping or Group Mapping, and select the View Domain Mapping link. For more information on the `domain-mappings-add` command, refer to the *Command Line Reference*.

A UNIX `/etc/passwd` file can be imported, providing the server with a mapping of user name to UID. The `/etc/groups` file should also be imported to provide the server with a mapping of Group name to GID.

The server will ignore other fields from the `passwd` file, such as the encrypted password and the user's home directory. Users or Groups configured by importing from the `/etc/passwd` file will then appear in the appropriate list on the User Mappings page or the Group Mappings page.

Choose one of the three following formats and use it consistently throughout the file:

- NFSv2/3 user/group data only. The source of the user data can be a UNIX password file, such as `/etc/passwd`.

When using Network Information Service (NIS), use the following command to create the file:

```
yycat passwd > /tmp/x.pwd
```

The resulting file has the following format:

```
john:x:544:511:John Brown:/home/john:/bin/bash
keith:x:545:517:Keith Black:/home/keith:/bin/bash
miles:x:546:504:Miles Pink:/home/miles:/bin/bash
carla:x:548:504:Carla Blue:/home/carla:/bin/bash
```

- NFSv2/3-to-Windows user/group mappings only. Create a file with entries in the following format:

```
UNIXuser="NT User", "NT Domain"
```

with the following syntax rules:

- NT domain is optional.
- NFS user names cannot contain spaces.
- NT names must be enclosed in quotation marks.
- If the domain name is omitted, the server domain is assumed. If the empty domain name is required, it must be specified like this:

```
users="Everyone", ""
```

where the Everyone user is the only common account with an empty domain name.

- Both NFSv2/3 user/group data and NFSv2/3-to-Windows user mappings. Create a file with entries in the following format:

```
UNIXuser:UNIXid="NT User", "NT Domain"
```

with the same rules for NFS and NT names as for the NFSv2/3-to-Windows user mapping.

The resulting file has entries in the following format:


```
john:544="john", "Domain1"
keith:545="keith", "Domain1"
miles:546="miles", "Domain1"
carla:548="carla", "Domain1"
```

Importing a user mapping from a file or an NIS or LDSP server

You can import a mapping from a file stored on your machine or on a network drive.

Procedure

1. Navigate to **Home > File Services > User Mapping**.
2. Click **Import Users** to display the **Import User Mapping** page.



The following table describes the fields on this page:

Field/Item	Description
Import From File	Import a mapping from a file stored on your machine or on a network drive. Search for the file by clicking Choose File , and then click Import .
Import From NIS/LDAP	Import a mapping from the currently used NIS/LDAP server by clicking Import .


3. Complete the import process by choosing one of the following:
 - To import from a file, enter the file name in the **Filename** field or click **Browse** to locate the file, and then click **Import**.
 - To import from NIS/LDAP, click **Import**. The NIS or LDAP server displayed in the Current Domain will be contacted, and the mappings will be imported.

Importing a group mapping from a file or an NIS LDAP server

You can import group mappings from a file stored on your machine or on a network drive, or from an NIS/LDAP server.

Procedure

1. Navigate to **Home > File Services > Group Mappings**.
2. Click **Import Groups** to display the **Import Group Mappings** page.



The following table describes how the fields on this page:

Field/Item	Description
Import From File	Import a mapping from a file stored on your machine or on a network drive. Search for the file by clicking Choose File , and then click Import .
Import From NIS/LDAP	Import a mapping from the currently used NIS/LDAP server by clicking Import .

3. Complete the import process by choosing one of the following:
 - To import from a file, enter the file name in the **Filename** field or click **Browse** to locate the file, and then click **Import**.
 - To import from NIS/LDAP, click **Import**. The NIS or LDAP server displayed in the Current Domain will be contacted, and the mappings will be imported.

File system auditing

File system auditing monitors and records file access and modification operations performed through the SMB and NFSv3 protocols. Records are made using the Windows Eventlog format and can be stored to the file system's audit log or made available to third-party external tools.

File system audit logging is performed and controlled on a per file system basis.

Currently, file system auditing is only supported for operations using SMB and NFSv3. By default, when file system auditing is enabled, access to the audited file system is only allowed for these two protocols. However, access by clients using other protocols like NFSv2, can optionally be allowed. When such access is allowed, access to file system objects through these protocols is not audited.



Note: Auditing of SMB is based on the open and close operations; because NFSv3 is a stateless protocol and lacks equivalent operations, auditing checks must be performed on each I/O operation, which can be costly in terms of system performance. Therefore, if auditing was enabled for a file system in a previous release, on upgrade to this release NFSv3 auditing is disabled. The protocols which cause audit events to be generated can be controlled with the `--audit-protocol (-p)` option of the `filesystem-audit` command.

After a file has been externally migrated (migrated to an external server), for example to a Hitachi Content Platform (HCP) system, subsequent access to the file through the NAS server is audited as if the file were still local.

For known users (users with a Windows user mapping), the NAS server logs Object Access events 560, 562, 563 and 564. As with the Windows operating system, auditable events for objects are specified by SACLs (system access control lists). Auditing events are logged under the following conditions:

- 560 – open handle

This event is logged when a network client asks for access to an object. An access check is performed against the DACL (discretionary access control list) and an audit check is performed against the SACL. If the result of the access check matches the result of the audit check, an audit record is generated.

- 562 – close handle

This event is logged when an application closes (disposes of) an existing handle, and is logged in conjunction with event 560.

- 563 – open handle for delete

This event is logged when a network client asks for access to a file using the CreateFile call, and the delete-on-close flag is specified. An access check is performed against the DACL and an audit check is performed against the SACL. If the result of the access check matches the result of the audit check, an audit record is generated.

For successful deletions, the audit records the accesses that were granted, and for failures the audit records the accesses that were requested.

- 564 – delete

This event is logged when an application closes (disposes of) an existing handle, and is logged in conjunction with event 563.



Note: Events for any user who is a member of the Audit Service Accounts local group are excluded from the audit log. Adding the third party auditing software user to this group results in a small but measurable performance gain.

About file system audit logs

The file system audit log is buffered in memory, and may be permanently stored in a file in the file system being audited. Active audit log files are stored in Windows event log file format (.evt) so that standard tools can access them. The name, location, size of the active audit log file, log file retention, and active log file backup settings are defined when enabling auditing for a file system.



Note: File System Audit logs are saved in Windows XP format. An effect of this is that, depending upon how the saved .evt file is opened, a Windows Vista or Windows 2008 Server event viewer can report the file as corrupted, or might not be able to fully interpret the events. Note that the same situation occurs when a Windows Vista event viewer is used to display saved logs from an XP system. To display the logs correctly, use a Windows XP event viewer.

Audit log files are limited in size, and the retention behavior when a log fills is configurable. When an audit log reaches its maximum size, log entries (file system events) can be overwritten, or the full audit log can be saved, and a new log started

Note: All file system audit log parameters are specified on a per file system basis.

You can specify a backup policy, which backs up the active log at regular intervals, and starts a new active log file. Backup log files are created in the same directory as the active audit log file.

In the event of a server crash, active file system audit logs are recovered only if a rollback is performed on restart. Note that a rollback may reset the audit log file to a time when it can be recovered, thus saving some records that would otherwise be lost.

Controlling file system auditing

File system auditing requires that a file system audit policy be defined for the file system to be monitored, and that auditing is enabled for the specific file system. File system auditing is performed and controlled on a per file system basis.

Creating a file system audit policy

The file system audit policy specifies access restrictions for clients connecting through un-auditable protocols (if access is allowed or denied), and specifies audit log details. The audit log policy specifies naming, location in the file system, size, the log roll over policy, and the backup policy.

Procedure

1. Navigate to **Home > File Services > File System Audit Policies**, and click **add** to display the **Add File System Audit Policy** page.

The screenshot shows the 'Add File System Audit Policy' configuration page. The breadcrumb navigation is 'File Services > Home > File Services > File System Audit Policies > Add File System Audit Policy'. The main title is 'Add File System Audit Policy'. The configuration is for the file system 'EVS / File System: g5-evs1 / FS'. The 'Audited Protocols' section shows 'SMB: Always enabled' and 'NFSv3: [disabled]'. The 'Access via Unsupported Protocols' section has 'Deny Access' selected, with a note: 'Client access to the file system via un-auditable protocols will be denied; refer to Help for more information'. The 'Audit Log' section includes 'External: [disabled]', 'Active Log File Name: audit.evt' (with a note: '(File name entered must have .evt extension)'), 'Logging Directory: /.audit' (with a 'browse...' button and note: '(Directory will be created if it does not exist)'), 'Maximum Log File Size: 512 KiB', and 'Log roll over policy' with 'New' selected and 'Wrap' as an option. The 'Backup Policy' section shows 'Backup Interval: 0 minutes' and 'Number of files to retain: 10'. At the bottom are 'OK' and 'cancel' buttons.

Field/Item	Description
EVS/File System	Lists the currently selected EVS and file system, to which the audit policy will apply. Click change to go to the Select a File System page, where you can select a different EVS and file system.
Access via Unsupported Protocols	<p>When clients attempt to access the file system through a protocol that does not support auditing (such as NFSv2), this setting determines if those clients are permitted to access the file system. You can select either:</p> <ul style="list-style-type: none"> ▪ Deny Access. Client access to the file system using unauditible protocols (such as NFSv2) is denied. ▪ Allow Access. Allows client access to the file system using unauditible protocols (such as NFSv2), but does not create any auditing events.
Audited Protocols	<p>When clients attempt to access the file system through a protocol that does not support auditing (such as NFSv2), this setting determines if those clients are permitted to access the file system. You can select either:</p> <ul style="list-style-type: none"> ▪ smb. Only the SMB protocol is audited. Access to SMB is always allowed, and access via other protocols is determined via the Other Protocol Support option. ▪ smb, nfsv3. Both the SMB and NFSv3 protocols are audited. Access to SMB and NFSv3 is always allowed, and access via other protocols is determined via the Other Protocol Support option.
External	Stops the audit records from being stored locally (including audit log backups) and instead only makes them available to an external audit log server. To configure an external logging server, use the <code>audit-syslog</code> CLI command or for third-party audit logging applications, configure an audit log consolidated cache and then read the audit logs using the Windows EVENTLOG protocol.
Active Log File Name	Specify the file name for the file system audit log. The file name must have an <code>.evt</code> extension. The default file name is <code>audit.evt</code> .
Logging Directory	Specify the directory within the file system in which the file system audit log files are saved. You can use the browse button to search for an existing directory, or enter the name of a directory to be created.

Field/Item	Description
Maximum Log File Size	Specify the maximum size of the active audit log file in KiB or MiB. The default size is 512 KiB. The maximum log file size is 50 MiB.
Log roll over policy	Determines what the system does once the active audit log file is full (when it reaches the Maximum Log File Size). You can select either: <ul style="list-style-type: none"> ▪ Wrap, which causes the system to delete the oldest existing audit entry to allow room for a new entry. ▪ New, which causes the system to create a new active audit log file. The default is New.
Backup Interval	Specify the time (in minutes) between automatic backups of the active audit log. The backup interval must be between 5 and 14400 minutes (10 days). A value of 0 disables the automatic backups. The default is 0.
Number of files to retain	Specify the number of backup audit log files to retain. The default is 10 . The maximum number of files to retain is 50.

2. Specify the access settings for unsupported (unauditable) protocols

- **Deny Access.** Client access to the file system using unauditable protocols (such as NFS) is denied.

Specifying **Deny Access** generates an error if there is an NFS export mounted on an unauditable client or the file system has a FTP user that has a directory available. To ensure this error is not generated, you can remove the NFS export for the file system, remove the FTP user, or select the **Allow Access** option.

- **Allow Access.** Allows client access to the file system using unauditable protocols (such as NFS), but does not create any auditing events.

3. Specify the name for the active audit log file. The file type suffix must be `.evt`.

4. Click **browse** to specify an existing logging directory, or enter the name of a directory to create.

For ease of access to the audit log files, the logging directory should be within in a CIFS share that can be accessed by those who need to review the access log.

5. Specify the maximum log file size.
6. Specify the roll over (retention) policy.
7. Specify the backup interval.
8. Specify the number of files to retain.
9. Click **OK** to save the policy as specified.

Configuring auditing on the Windows client

After creating a file system audit policy, the next step is to configure which of these events gets audited from the Windows client. By default, no file system accesses will be audited.

Note: Only members of the Administrators local group have the right to edit the file system audit log policy from within Windows Explorer. A user that is not a member of Administrators Local Group cannot amend the audit settings of a file or directory.

Procedure

1. Right-click a folder that resides on a server file system that is configured for auditing and select **Properties**, and then the **Security** tab.
2. Click **Advanced** and select the **Auditing** tab.
3. Select **Add** and select which users get audited.
For example, select **Everyone** so that all users get audited.
4. A box pops up and allows you to specify which events are to be audited for the specified user.
5. You can choose to audit **Successful**, **Failed**, or **both** for each access type.

Enabling auditing for a file system

File system auditing can be enabled on a per-file system basis.

Note: By default, when file system auditing is enabled, access to the file system is limited to the SMB and NFSv3 protocols. Access by clients using other protocols, like iSCSI, can, however, be allowed. When such access is allowed, access to file system objects through these protocols is not audited.

To enable file system auditing for a particular file system, the file system must be added to the file system audit list.

Procedure

1. Navigate to **Home > File Services > File System Audit Policies**.

The screenshot shows the 'File System Audit Policies' management page. The breadcrumb navigation is 'File Services > Home > File Services > File System Audit Policies'. The main heading is 'File System Audit Policies'. At the top, there is a field 'EVS: g1-evs1' with a 'change...' button. Below this is a section for 'Audit Log Consolidated Cache' with 'Cache: Enabled - Unknown' and a 'modify' button. A table lists file systems with columns for 'File System', 'Status', and 'details'. The table contains one entry: 'FS_1' with status 'Enabled'. At the bottom, there are 'Check All' and 'Clear All' links, and an 'Actions' section with buttons for 'add', 'delete', 'enable', and 'disable'.

File System	Status	details
<input type="checkbox"/> FS_1	Enabled	details

Field/Item	Description
EVS	Lists the EVS to which host the file system is assigned. Click change to go to the Select an EVS page, where you can select a different EVS.
Audit Log Consolidated Cache	The server uses this cache for reporting file system audit events to Windows clients. Only one consolidated cache file can be configured per EVS.
modify	Enables the user to configure a file system, directory where the file is stored and file name for the audit log consolidated cache file.
File System	Lists all file systems in the specified EVS that have an audit policy.
Status	Indicates whether file system auditing is enabled or disabled. It also indicates whether auditing is external.
details	Displays the File System Audit Policy Details page, in which you can change the auditing options for a file system.
add	Displays the Add File System Audit Policy page, in which you can set the auditing options for a file system. Only one audit policy is allowed per file system.
delete	Deletes the audit policy for a selected file system.
enable	Enables file system auditing for the selected file system.
disable	Disables file system auditing for the selected file system.

2. If the file system on which you want to enable auditing is listed, an audit policy has already been defined for that file system.
 - If the Audit Policy Status is enabled, logging is already enabled for the file system, and no further actions are required.
 - If the Audit Policy Status is disabled, select the check box next to the file system name, and click **enable**.

If the file system on which you want to enable auditing is not displayed, a file system audit policy may not have been defined for that file system, or the file system may have an audit policy defined, but the file system is not in the currently selected EVS.

3. Click **change** to display the **Select an EVS** page, in which you can select a different EVS.
 - If, after selecting the EVS that hosts the file system, the file system on which you want to enable auditing is now listed on the **File System Audit Policies** page, select the check box next to the file system name, and click **enable**.
 - If, after selecting the EVS that hosts the file system, the file system on which you want to enable auditing is still not displayed, you must define a file system audit policy for that file system. Click **add** to display the **Add File System Audit Policy** page, in which you can set the auditing options for a file system.

Modifying a file system audit policy

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**.
If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, in which you can select a different EVS.
2. Click the **details** button on the file system with the audit policy you want to modify to display the **File System Audit Policy Details** page.

The screenshot shows the 'File System Audit Policy Details' page for 'Span19FS'. At the top, it indicates 'Auditing: Enabled' with a 'disable' button. Below this are sections for 'Audited Protocols' (SMB: Always enabled, NFSv3: checked), 'Access via Unsupported Protocols' (radio buttons for 'Deny Access' and 'Allow Access (without auditing)'), 'Audit Log' (External: unchecked, Active Log File Name: audit.evt, Logging Directory: /.audit, Maximum Log File Size: 512 KiB, Log roll over policy: New selected), and 'Backup Policy' (Backup Interval: 0 minutes, Number of files to retain: 10). At the bottom are 'OK' and 'cancel' buttons.

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Lists the currently selected EVS and file system, to which the audit policy will apply. Click change to go to the Select a File System page, where you can select a different EVS and file system.
Auditing	Indicates whether file system auditing is enabled or disabled. Click enable or disable to toggle the auditing mode.
Access via Unsupported Protocols	When clients attempt to access the file system through a protocol that does not support auditing (such as NFSv2),

Field/Item	Description
	this setting determines if those clients are permitted to access the file system. You can select either: <ul style="list-style-type: none"> ▪ Deny Access. Client access to the file system using unauditables protocols (such as iSCSI) is denied. ▪ Allow Access. Allows client access to the file system using unauditables protocols (such as iSCSI), but does not create any auditing events.
Audited Protocols	When clients attempt to access the file system through a protocol that does not support auditing (such as iSCS), this setting determines if those clients are permitted to access the file system. You can select either: <ul style="list-style-type: none"> ▪ smb. Only the SMB protocol is audited. Access to SMB is always allowed, and access through other protocols is determined through the Other Protocol Support option. ▪ smb,nfsv3. Both the SMB and NFSv3 protocols are audited. Access to SMB and NFSv3 is always allowed, and access through other protocols is determined via the Other Protocol Support option.
External	Stops the audit records from being stored locally (including audit log backups) and instead only makes them available to an external audit log server. To configure an external logging server, use the <code>audit-syslog</code> CLI command or for third-party audit logging applications, configure an audit log consolidated cache and then read the audit logs using the Windows EVENTLOG protocol.
Active Log File Name	Specifies the file name for the file system audit log. The file name must have an <code>.evt</code> extension. The default file name is <code>audit.evt</code> .
Logging Directory	Specifies the directory within the file system in which the file system audit log files are saved. You can use the browse button to search for an existing directory, or enter the name of a directory to be created.
Maximum Log File Size	Specifies the maximum size of the active audit log file in KiB or MiB. The default is 512 KiB. The maximum value is 50 MiB.
Log roll over policy	Determines what the system does once the active audit log file is full (when it reaches the Maximum Log File Size). You can select either: <ul style="list-style-type: none"> ▪ Wrap, which causes the system to delete the oldest existing audit entry to allow room for a new entry.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ New, which causes the system to create a new active audit log file. The default is New.
Backup Interval	Specifies the time (in minutes) between automatic backups of the active audit log. The backup interval must be between 5 and 14400 minutes (10 days). A value of 0 disables the automatic backups. The default is 0.
Number of files to retain	Specifies the number of backup audit log files to retain. The default is 10 .

3. Modify the policy as required.
4. Click **OK** to save the policy as specified.

Enabling or disabling auditing for a file system

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**.

If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, in which you can select a different EVS.

2. Select the check box next to the name of the file system with the audit policy you want to enable or disable.
3. Click **Enable** to allow a disabled policy to function again, or click **Disable** to stop the policy from functioning.
When disabled, file system access operations are not logged, and protocol restrictions are not enforced. Note that disabling a policy does not delete it.

Deleting a file system audit policy

Procedure

1. Navigate to **Home > Files Services > File System Audit Policies**.

If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, in which you can select a different EVS.

2. Select the check box next to the name of the file system with the audit policy you want to delete, and click **delete**.



Note: Existing log files are not deleted automatically when a policy is deleted. If you want to delete these logs, you must do so manually,

Displaying file system audit logs

The NAS server supports using a remote Windows Event Viewer to display file system audit log events. The audit log files are shown in the "FS" (file system) log, which can be displayed by the Windows Event Viewer, assuming that:

1. You have used the **audit-log-consolidated-cache** command to configure a single consolidated cache file (the audit-log-consolidated-cache).

If the cache file is not configured, the Windows Event Viewer cannot view file system events. The consolidated cache file has a default size of 10MB, and a maximum size of 50MB.



Note: Only one consolidated cache file can be configured per EVS. Audit events from all file systems assigned to that EVS are collected into this single consolidated cache file.

When you create the consolidated cache file, you must specify the name of the file system in which the file will be stored. The cache file is located in the `.audit` directory of the root of the named file system. The default name for the consolidated cache file is `audit_cache.evt` (audit log files for individual file systems have a default name of `audit.evt`).

2. The logging directory is within a CIFS share.

Using the Windows Event Viewer, you can display, save, and clear the local event logs, or those on a remote computer. Audit logs can be saved in several formats, including a `.evt` event format or a plain text file. The Windows Event Viewer can only save in `.evt` format to a file on the same computer as the event log, because it is the computer being viewed that does the copy (meaning the Event Viewer does not just read the event log and write it to a file). The Event Viewer can also be used to open and display saved audit log files.

Optionally, you can send file system audit logs to a remote syslog server using the **audit-syslog** command. Enter **man audit-syslog** at the CLI, or see the *Command Line Reference* for more information.

Chapter 4: Sharing resources with NFS clients

A fundamental component of most UNIX networks, the Network File System (NFS) protocol provides PC and UNIX workstations with transparent access to one another's files. This section describes how to set up NFS exports, and explains about NFS statistics, supported clients, and prerequisites.

The server implements the file-serving functions of an NFS server, providing normal file-serving functions, such as:

- Export manipulation
- File manipulation (read, write, link, create, and so forth)
- Directory manipulation (mkdir, readdir, lookup, and so forth)
- Byte-range file locking
- File access control (permissions)
- File and directory attributes (sizes, access times, and so forth)
- Hard links and symbolic (soft) links

Enabling NFS Protocol Support

To enable NFS access to the system, enter an NFS license key on the **License Keys** page in the NAS Manager.

Supported clients and protocols

The server supports all clients compliant with NFS version 2, version 3, or version 4 standards. NFS versions 2, 3 and 4 are supported over TCP. The following table summarizes the supported versions of NFS and other UNIX protocols:

Protocol	Supported versions
NFS	2, 3, and 4
Port Mapper	2
Mount	1 and 3
Network Lock Manager (NLM)	1, 3, and 4
Network Status Monitor (NSM)	1



Caution: While it is possible to use UDP with NFS on versions 2 and 3, it is not recommended due to inherent risks. On the NAS server, UDP is not automatically presented as a transport option for the NFS service by the Port Mapper service. To register NFS over UDP in the Port Mapper service, see the `rpc-service-nfs-udp` command.

Supported NFS versions

By default, the maximum version supported is version 3, meaning that the server supports versions 2 and 3 by default. To change the maximum supported version to NFSv4, use the CLI command `nfs-max-supported-version`. By setting the maximum supported version to 4, you allow the storage server to support NFS versions 2, 3, and 4.



Note: 'NFSv4' refers to NFSv4.0. There is currently no NAS support for NFSv4.1.

NFS statistics

NFS statistics for the server are available for activity since the previous reboot, or since the point when the statistics were last reset.

Statistics for NFS requests received by the server are broken down by NFS version and procedure, and are shown (in 10-second time slices) on the **NFS Statistics** page (refer to the Server and Cluster Administration Guide for more information). The statistics can also be sampled and viewed on-demand using the `nfs-stats` CLI command.

Unicode support

The storage server (or cluster) stores metadata about the files, directories, migration paths, CIFS shares, NFS exports, user names, group names, log entries, mount points and so on for the virtual servers, file systems, and namespaces served by the server/cluster.

When interacting with another network device, the metadata transmitted to or received by the storage server/cluster must use a character encoding supported by the other network device. Typically, clients/devices using the SMB/SMB2 protocol (Windows) encode data in the UCS-2 character encoding, and clients/devices that use the NFS protocol encode data in the UTF-8 character encoding.



Note: The data on storage subsystems attached to a storage server/cluster is not affected by changing the character encoding currently used by the server/cluster.

When using the FTP protocol to communicate with clients/devices, the storage server/cluster supports the UTF-8 character encoding for user names, passwords, and file/directory names.

NFS and NIS unicode support

Character sets supported by the NAS server to communicate with NFS clients and NIS servers include:

Communicating with	Character encodings supported	Default character encoding
NFSv2 and NFSv3 clients	Latin-1, UTF-8, EUC-KR, EUC-JP, and EUC-CN	Latin-1
NFSv4 clients	UTF-8	UTF-8
NIS servers	Latin-1 and UTF-8	Latin-1

You can specify the character encoding to be used when communicating with NFS clients and/or NIS servers using the `protocol-character-set` command.

The NFS character set controls:


- File, directory, and export names to/from NFSv2 and NFSv3 clients.
- Symlinks to/from NFS clients.

The NIS character set controls:

- NIS user and group names.
- LDAP user and group names.


Character set encoding may not be set for:

- Namespace links.
- Namespace directories.
- Communication with NFSv4 clients.

 **Note:** Communication with NFSv4 clients uses only the UTF-8 character set.

When multi-tenancy is not enabled, the configured character encoding applies on a cluster-wide basis (all clients communicate with the server using the same encoding). When multi-tenancy is enabled, the character encoding is configured on a per-EVS basis, allowing the use of multiple character sets on the same NAS server/cluster.


To correctly display the characters during communication between the client and the NAS server, both the program used by the the client to communicate with the NAS server and the NAS server itself must be configured with the same character encoding and locale.

 **Note:** When the EUC-KR, EUC-JP, or EUC-CN character encodings are enabled for NFSv2/NFSv3 clients, there is a performance penalty for operations that are not handle-based when compared to the UTF-8 or Latin-1 character encodings.

Changing the character set

By default, the storage server/cluster uses the ISO 8859-1 (Latin-1) character set when communicating with NFS clients and/or NIS servers. When NIS servers and NFS clients use different character sets, the administrator must specify which character set the NFS clients are using, and which character set the NIS servers are using. The `protocol-character-set` command allows an administrator to specify the character set to be used when communicating with NFS clients and/or NIS servers.

Refer to the *Command Line Reference* for more information on the `protocol-character-set` command.


 **Note:** After the `protocol-character-set` command is issued, the specified character set is put into use immediately, without the need to restart the server/cluster.

Enabling and disabling file services

You can enable and disable the required file services for the system in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > Enable File Services** to display the **Enable File Services** page.



The following table describes the fields on this page:

Fields/Item	Description
<ul style="list-style-type: none"> ▪ CIFS/Windows ▪ NFS/UNIX ▪ FTP ▪ iSCSI ▪ CNS ▪ ReadCache 	Select the check box for each service you want to enable.
apply	Click to apply the selections.

2. Select or deselect one or more services.
3. Click **apply**.



Notes:

- With the exception of FTP, all of these services require a valid license.
- If ReadCache is selected or deselected, a reboot may be required. If so, follow the on-screen instructions to restart the server.

Configuring NFS exports

NFS exports are configured on mounted file systems. NFS exports can be configured manually or export details can be imported from a file.

The NFSv4 pseudo file system

NFSv4 introduces the concept of the pseudo file system, where exports appear as directories. NFSv4 clients do not connect directly to NFS exports as in NFSv2/3. Instead all clients connect to the root of the pseudo file system, which is a virtual directory. The pseudo file system is generated automatically from the NFS exports, and is maintained automatically as exports are modified and removed. You can choose to present all the file systems in a single pseudo file system.

The server allows you to create views of many file systems from one point of contact, name spaces. These views are available on a per EVS basis or for the entire cluster.

This is an example of a pseudo file system:

A server named `numbers` has two exports: `/one` and `/two`. If a client wishes to get access to export `/one`, there are two ways to mount exports:

```
mount -t nfs4 numbers:/ /mnt
```

which mounts the pseudo file system at `/mnt`

```
mount -t nfs4 numbers:/one /mnt
```

which mounts the export `/one` at `/mnt`

The first method is only supported in NFSv4. The second method is supported in versions 2, 3, and 4. In the first method, the client can export `/one` with the command `cd /mnt/one`, and to export `/two` with `cd /mnt/two`.

Kerberos configuration

Before you begin



Note: The Kerberos implementation has been updated with the Advanced Encryption Standard (AES). The Data Encryption Standard (DES) has been deprecated and is insufficiently secure. AES pre-requisites are:

- Windows Server 2008 or higher is required to deploy a Microsoft Windows KDC that supports AES encryption.
- Configuration may be required on the clients. The configuration of the KDC and clients may vary depending on their operating systems.
- The Kerberos Principle accounts on the KDC may need to be configured to support AES.
- Supported AES encryption types are
 - AES256: HMAC-SHA1-96
 - AES128: HMAC-SHA1-96

Configuring the server requires the following steps:

Procedure

1. Create the principal and key of the service (the EVS) on the KDC (Key Distribution Center).

The `keytab` file must contain the service principal for the NFS service for the EVS. Once the NFS service principal for the EVS has been added, you can then create a keytab file specifically for the EVS. The type of key is critical.

- AES: To use AES, the keytab must contain an AES key to enable AES by default. If an AES only keytab is imported, DES is disabled. If an AES only keytab is imported, all clients must be configured to support AES and have an AES key in their keytabs.
- DES:
 - To use DES, the client must perform the Kerberos authentication with any of the supported encryption types except AES.
 - The server must have a key that corresponds to whatever encryption type the client used.
- AES and DES: The keytab must contain
 - An AES key and
 - Any old supported encryption type key (it does not have to be DES), provided that it is supported by the client as well.

For example, with an EVS named "man" in the Kerberos realm `AESIR.EXAMPLE.COM`, the keytab file for the NFS service on "man" should contain a principal `nfs/man.aesir.example.com@AESIR.EXAMPLE.COM`. The format of the principal starts with the service (nfs), followed by a slash, then the fully-qualified-domain name of the EVS, then the symbol @, and finally the Kerberos realm. Note that case is significant. Kerberos realms are always in uppercase. Also, there must be no trailing period after the Kerberos realm.

2. Export a keytab file from the KDC.
Typically you will use the `kadmin` utility run from the master KDC to export a `keytab` file. For details on creating an appropriate `keytab` file, refer to the documentation for the tools supplied with your version of Kerberos.
3. Import the `keytab` file into the server.

Transfer the keytab file to the flash of the server.

For example: securely move the `keytab` file to the NAS Manager and transfer it to the NAS server. Log on with `ssc`, and do the following:

```
SERVER:$ ssput man.nfs.keytab man.nfs.keytab
```

The first name is the local file name on the NAS Manager, the second name is the name to use on the server. Once the file has been placed on the server, import the keytab in the context of the EVS with:

```
SERVER:$ krb5-keytab import man.nfs.keytab
```

After the keytab has been imported, the uploaded `keytab` file can be safely removed with:

```
SERVER:$ ssrc man.nfs.keytab
```

4. Set the Kerberos realm for the server.

Set the realm by using the command `krb5-realm`. For example:

```
SERVER:$ krb5-realm AESIR.EXAMPLE.COM
```

The server's NFS hostname must be set, per EVS, using the command `nfs-hostname <hostname>`.

After performing these steps, the NAS server is able to complete the configuration. However, you may choose to create mappings between the Kerberos users/groups and the Active Directory users/groups.

Viewing NFS exports

You can view existing NFS exports and also add, modify and delete them on the NAS Manager NFS Exports page. This page can be configured to display all the exports associated with:


- The global cluster name space (CNS)
- An individual name space for an EVS
- A specific file system

Procedure

1. Navigate to **Home > File Services > NFS Exports** to display the **NFS Exports** page.

The screenshot displays the 'NFS Exports' page. At the top, there is a breadcrumb trail: 'File Services > Home > File Services > NFS Exports'. The main heading is 'NFS Exports'. Below this, there are two panels: 'EVS / File System Label' and 'Filter'. The 'EVS / File System Label' panel shows 'g1-evs3 / PHDS1' with a 'change...' button. The 'Filter' panel has input fields for 'Name:' and 'Path:', a dropdown menu for 'Transfer to Object Replication Target' set to 'None', and a 'filter' button. Below these panels is a table with columns 'Name', 'File System', and 'Path'. The table shows one entry: '/PHDS1', 'PHDS1', and '/'. There are 'Check All' and 'Clear All' links below the table. At the bottom, there are 'Actions' (add, delete, refresh cache) and links for 'Download Exports' and 'Backup & Restore'. 'Shortcuts' include 'Read Cache Options' and 'Read Cache Statistics'.

The following table describes the fields on this page:

Field/Item	Description
Cluster Name Space or EVS / File System	<p>Displays the currently selected name space or EVS/File System</p> <ul style="list-style-type: none"> When Cluster Name Space is displayed, the cluster (global) name space has been selected. When EVS / File System is displayed, a particular EVS (and optionally a particular file system) has been selected. <p>The currently selected name space controls which NFS exports are displayed on this page.</p>
change	Enables the user to select a different name space or EVS / File System.
Filter	<p>A subset of the exports on the EVS or file system can be viewed. Use the Name and Path text fields and the Transfer to Object Replication Target list to define the criteria for the selection. Valid selections for Transfer to Object Replication Target are None, Enable, Disable, or Use FS Default. To apply the filter to the list of exports, click filter.</p> <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note: If a field is left blank, it is ignored in the filtering process. When specifying a name to be matched, the wildcard character "*" may be used.</p> </div>
Name	The name of the NFS export.
File System	The name of the file system (or CNS link to a file system) to which the NFS exports is assigned.
Path	The path and directory to which the NFS export is directed.
details	Opens the NFS Export Details page in which you can display detailed information about the NFS export.
add	Advances to the Add Export page.
delete	Deletes the selected NFS export.
refresh cache	Clears the NAS Manager cache, and then repopulates it with the relevant objects. Note that this is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.
Download Exports	Downloads a CSV file containing a list of all configured NFS exports on the selected EVS and file system. Note that the downloaded file cannot be used to restore NFS exports (you must restore NFS exports from an NFS exports backup file). To download a list of exports from another file system, click change .

Field/Item	Description
Backup & Restore	Displays the NFS Export Backup & Restore page.
Read Cache Options	Advances to the Read Cache Options page.
Read Cache Statistics	Advances to the Read Cache Statistics page.


Adding an NFS export

You can add an NFS export in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > NFS Exports** to display the **NFS Exports** page.
2. Click **add** to display the **Add Export** page.

The following table describes the fields on the page:

Field/Item	Description
EVS/File System	Currently selected file system, to which the NFS Export will link.
Cluster Namespace	Currently selected cluster namespace, to which the NFS Export will link.
change / browse (depending on Web browser)	Enables the user to select a different file system or (on a cluster) a different cluster namespace.
Export Name	Name of the export.
Path / CNS Path	Path to the source directory for the export. To locate a source directory for the export, click the browse/change button.
Path Options	<p>Determines the path options:</p> <ul style="list-style-type: none"> ▪ Create path if it does not exist to create the path entered in the Path field (filesystems only). ▪ Allow this export to overlap other exports if nested NFS exports are allowed. <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note: If the file system is mounted read-only, for example it is an object replication target, it is not possible to create a new directory. Select a path to an existing directory.</p> </div>
Show snapshots	<p>Determines how to show snapshots:</p> <ul style="list-style-type: none"> ▪ Show and Allow Access, to display and allow access to snapshots. ▪ Hide and Allow Access, to hide snapshots, but still allow access to the hidden snapshots. ▪ Hide and Disable Access, to hide and disallow access to snapshots.
Local Read Cache (file systems only)	<p>Allows caching of files or cross file system links from the file system to which this export points:</p> <ul style="list-style-type: none"> ▪ Cache all files. Allows caching of files and cross file system links in the file system of the export. Cross file system links are local links that point to a data file in a remote file system. The remote file system may be on a remote server or storage device. ▪ Cache cross-file system links. Allows only cross file system links to be cached ▪ Do not cache files. Do not allow read caching of files and cross file system links.

Field/Item	Description
	Local read caching is not supported for NFSv4 clients.
Transfer to Object Replication Target (file systems only)	<p>When a file system is recovered from a snapshot, one of the final steps is to import the NFS exports found in the snapshot representing the selected version of the file system. Only those NFS exports marked as transferable will be imported.</p> <ul style="list-style-type: none"> ▪ Enable: NFS exports will be transferred to recovered file systems. ▪ Disable: NFS exports will not be transferred to recovered file systems. ▪ Use FS default: When the target file system is brought online, NFS exports will be transferred if Transfer Access Points During Object Replication option is enabled for the file system.
Access Configuration	<p>IP addresses, host names, or the NIS netgroups of the clients who are allowed to access the NFS export (up to 5957 characters). If the system has been set up to work with a name server, you can enter the NIS netgroup to which the clients belong, or the client's computer name rather than its IP address (not case sensitive).</p> <p>You can also specify the required flavors of NFS security in a colon-separated list using the option (sec=<list>).</p> <p>The supported flavors are:</p> <ul style="list-style-type: none"> ▪ none - Connect as a null user ▪ sys - The traditional security flavor used by NFS, users are not authenticated by the server ▪ krb5 - Kerberos authentication ▪ krb5i - Kerberos authentication with per-messaging integrity ▪ krb5p - Kerberos authentication with per-message privacy <p>For example: 10.1.*.*(sec=sys:krb5:krb5i)</p> <p>See the <code>mount-point-access-configuration</code> man page for further information.</p>

3. To add an export to a new EVS or file system, click **change** next to that line and make a selection from the **Select a File System** page.
4. Enter the `Export Name` through which clients will access the export.
5. Type the `path` to the directory being exported or click **browse...** to locate an existing directory.

6. Set Path Options as follows:

- To create the path automatically when it does not already exist, fill the **Create path if it does not exist** check box.



Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to rwxrwxrwx). It is recommended that such directories be created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created by this option.

- To allow this export path to overlap other exports, fill the **Allow this export path to overlap other exports** check box.

This option is useful if you expect to set up future, nested exports. For example, suppose you export the root directory of a volume and make it available to managerial staff only. By selecting this option, you can later export subdirectories of the root directory and make each of them available to different groups of users.

7. If snapshots are present, make them visible to clients by selecting from the list:

- **Show and Allow Access**, to display and allow access to snapshots.
- **Hide and Allow Access**, to hide snapshots, but still allow access to the hidden snapshots.
- **Hide and Disable Access**, to hide and disallow access to snapshots.

In order for this change to become effective on NFS clients, all NFS clients should unmount and then remount the export, or the administrator must run `' touch . '` from within the root directory of the export.

8. Select the Local Read Cache setting. To allow caching of files or cross file system links from the file system to which this export points, select one of the following:

- **Cache all files.** Allows caching of files and cross file system links in the file system of the export. Cross file system links are local links that point to a data file in a remote file system. The remote file system may be on a remote server or storage device.
- **Cache cross-file system links.** Allows only cross file system links to be cached.

Local read caching is not supported for NFSv4 clients.

9. Choose the **Transfer to Object Replication Target** option.

When a file system is recovered from a snapshot, one of the final steps is to import the NFS exports found in the snapshot representing the selected version of the file system. Only those NFS exports marked as transferable will be imported.

Use the list to specify one of the following:

- **Enable:** NFS exports will be transferred to recovered file systems.
- **Disable:** NFS exports will not be transferred to recovered file systems.
- **Use FS default:** When the target file system is brought online, NFS exports will be transferred if Transfer Access Points During Object Replication option is enabled for the file system.

10. In the Access Configuration field, type the IP addresses, host names, or the NIS netgroups of the clients who are allowed to access the NFS export (up to 5,957 characters). If the system has been set up to work with a name server, you can enter the NIS netgroup to which the clients belong, or the client's computer name rather than its IP address (not case sensitive). You can also specify the flavor of NFS security using the option (sec=<mode>). The table outlines what to type in this field.

What to Type	Means
Blank or *	All clients can access the export.
Specific address or name. Examples: 192.0.2.0, client.dept.example.com	Only clients with the specified names or addresses can access the export.
A range of addresses using Classless Inter-Domain Routing (CIDR) notation. Example: 192.0.2.0/24	Clients with addresses within the range can access the export.
Partial address or name using wildcards. Examples: 192.0.*.*, *.example.com	Clients with matching names or addresses can access the export.

11. Click **OK**.

IP address export qualifiers

The following table describes qualifiers that can be appended to IP addresses when specifying client access to an NFS export:

Qualifier	Description
<code>read_write, readwrite, rw</code>	Grants read/write access. This is the default setting.
<code>read_only, readonly, ro</code>	Grants read-only access.
<code>root_squash, rootsquash</code>	Maps user and group IDs of 0 (zero) to the anonymous user or group. This is the default setting.
<code>no_root_squash, norootsquash</code>	Turns off root squashing.
<code>all_squash, allsquash</code>	Maps all user IDs and group IDs to the anonymous user or group.
<code>no_all_squash, noallsquash</code>	Turns off all squashing. This is the default setting.
<code>secure</code>	Requires requests to originate from an IP port less than 1024. Access to such ports is normally restricted to administrators of the client machine. To turn it off, use the insecure option.
<code>insecure</code>	Turns off the secure option. This is the default setting.
<code>anon_uid, anonuid</code>	Explicitly sets an anonymous user ID.
<code>anon_gid, anongid</code>	Explicitly sets an anonymous group ID.
<code>noaccess, no_access</code>	Denies the specified clients access to the export.
<code>(sec=<mode>)</code>	Allows you to specify the flavor of NFS security, where <mode> is a colon delimited list of allowed security flavors (sys:krb5:krb5i:krb5p).

Here are some examples:

- `10.1.2.38 (ro)`
Grants read-only access to the client whose IP address is 10.1.2.38.
- `10.1.2.0/24 (ro)`
Grants read-only access to all clients whose IP address is within the range 10.1.2.0 to 10.1.2.255.
- `yourcompanydept (ro)`
Grants read-only access to all members of the NIS group yourcompanydept.
- `*.mycompany.com (ro, anonuid=20)`
Grants read-only access to all clients whose computer name ends.mycompany.com. All squashed requests are to be treated as if they originated from user ID 20.
- `10.1.*.* (readonly, allsquash, anonuid=10, anongid=10)`
Grants read-only access to all the matching clients. All requests are squashed to the anonymous user, which is explicitly set as user ID 10 and group ID 10.
- The order in which the entries are specified is important. Take the following two lines:
`*(ro)`
`10.1.2.38 (rw)`
The first grants read-only access to all clients, whereas the second grants read/write access to the specified client. The second line is redundant, however, as the first line matches all clients. These lines must be transposed to grant write access to 10.1.2.38.
- `10.1.1.*(sec=sys), 10.1.2.*(sec=krb5:krb5i:krb5p), *(sec=krb5p)`
 - Clients in the 10.1.1.* subnet use **sys** authentication.
 - Clients in the 10.1.2.* subnet to use **krb5**, **krb5i**, or **krb5p**.
 - All other clients use **krb5p**.



Note: To improve performance, when specifying clients that can access an export, it is recommended that you specify IP addresses or IP address ranges, including those that include wildcards, before specifying host names or NIS netgroups.

Specifying clients by name

The following list describes how to specify clients by name, and not an IP address.

- **Full Qualified Domain Name Required.**
Be sure to specify the fully qualified domain name of the client. For example, use `aclient.dept.example.com` rather than simply `aclient`.
- **Leading Wildcard Allowed.**
To specify a partial name, a single wildcard, located at the start of the name, may be used.

- **Export Options Change Requires Remount.**

When the client mounts the NFS export, it determines which export option to apply to a specific client. Subsequent changes to DNS, WINS, or NIS that would resolve the client's IP address to a different computer name are only applied to mounted exports when the client unmounts the exports and then remounts them.

- **Name Service Order is Significant.**

Application of export options to a client's mount request may be affected by the order in which the system applies DNS, WINS, and NIS information to resolve IP addresses. The first service in name order sequence that can resolve the client name supplies the name *and* searches configuration options for the export.

Modifying NFS Export Details


You can modify properties of the selected NFS export for either a cluster name space (CNS) or a file system (shown in two different tables) in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > NFS Exports** to display the **NFS Exports** page.
2. Select the check box next to the NFS export to display, and click **details** to display the **NFS Export Details** page.

The following table describes the fields and items on this page:

Field/Item	Description
EVS/File System	Currently selected file system, to which the NFS Export will link.
Cluster Namespace	Currently selected cluster namespace, to which the NFS Export will link.
change / browse (depending on Web browser)	Enables the user to select a different file system or (on a cluster) a different cluster namespace.
Export Name	Name of the export.
Path / CNS Path	Path to the source directory for the export. To locate a source directory for the export, click the browse/change button.
Path Options	Determines the path options: <ul style="list-style-type: none"> ▪ Create path if it does not exist to create the path entered in the Path field (filesystems only). ▪ Allow this export to overlap other exports if nested NFS exports are allowed.

Field/Item	Description
	 Note: If the file system is mounted read-only, for example it is an object replication target, it is not possible to create a new directory. Select a path to an existing directory.
Show snapshots	Determines how to show snapshots: <ul style="list-style-type: none"> ▪ Show and Allow Access, to display and allow access to snapshots. ▪ Hide and Allow Access, to hide snapshots, but still allow access to the hidden snapshots. ▪ Hide and Disable Access, to hide and disallow access to snapshots.
Local Read Cache (file systems only)	Allows caching of files or cross file system links from the file system to which this export points: <ul style="list-style-type: none"> ▪ Cache all files. Allows caching of files and cross file system links in the file system of the export. Cross file system links are local links that point to a data file in a remote file system. The remote file system may be on a remote server or storage device. ▪ Cache cross-file system links. Allows only cross file system links to be cached ▪ Do not cache files. Do not allow read caching of files and cross file system links. Local read caching is not supported for NFSv4 clients.
Transfer to Object Replication Target (file systems only)	When a file system is recovered from a snapshot, one of the final steps is to import the NFS exports found in the snapshot representing the selected version of the file system. Only those NFS exports marked as transferable will be imported. <ul style="list-style-type: none"> ▪ Enable: NFS exports will be transferred to recovered file systems. ▪ Disable: NFS exports will not be transferred to recovered file systems. ▪ Use FS default: When the target file system is brought online, NFS exports will be transferred if Transfer Access Points During Object Replication option is enabled for the file system.
Access Configuration	IP addresses, host names, or the NIS netgroups of the clients who are allowed to access the NFS export (up to 5957 characters). If the system has been set up to work with a name server, you can enter the NIS netgroup to which the clients belong, or the client's computer name rather than its IP address (not case sensitive). <p>You can also specify the required flavors of NFS security in a colon-separated list using the option (sec=<list>).</p>

Field/Item	Description
	<p>The supported flavors are:</p> <ul style="list-style-type: none"> ▪ none - Connect as a null user ▪ sys - The traditional security flavor used by NFS, users are not authenticated by the server ▪ krb5 - Kerberos authentication ▪ krb5i - Kerberos authentication with per-messaging integrity ▪ krb5p - Kerberos authentication with per-message privacy <p>For example: 10.1.*(sec=sys:krb5:krb5i)</p> <p>See the <code>mount-point-access-configuration</code> man page for further information.</p>

3. Make changes as necessary.

4. Click **OK**.

Deleting an NFS export

You can delete an NFS export in the NAS Manager.



Caution: Export Deletion Alert! Before carrying out the instructions that follow for deleting an export, verify that it is not currently being accessed. If an export is deleted while users are accessing it, their NFS sessions will be terminated and any unsaved data may be lost.

When replacing a storage enclosure, delete all the exports associated with it. Then, when the replacement enclosure is available, add new exports on the new system drives.

Procedure


1. Navigate to **Home > File Services > NFS Exports** to display the **NFS Exports** page.
2. Select the check box(es) next to the NFS export(s) to delete, and click **delete**.
3. To confirm the deletion, click **OK**.

Backing up or restoring NFS exports

You can back up and restore NFS exports in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > NFS Exports** to display the **NFS Exports** page.
2. Click **Backup & Restore** to display the **NFS Exports Backup & Restore** page.



3. Choose from the following options:

- To backup, click **backup**. In the browser, specify the name and location of the backup file, and click **OK** or **Save** (the buttons displayed and the method you use to save the backup file depend on the browser you use).

A backup file name is suggested, but you can customize it. The suggested file name uses the syntax:

```
NFS_EXPORTS_<YYYY>-<MM>-<DD>_<HH><MM><UTC-diff>.txt
```

For example,

```
NFS_EXPORTS_2015-11-04_1615+0000.txt
```

- To restore, navigate to the directory in which the backup file is stored, select the file, click **Open** and then click **restore**.

About the rquotad service

The rquotad protocol has been implemented as a service on the storage server. It functions as a read-only protocol and is only responsible for reporting information about user and group quotas. Quotas can be created, deleted, or modified through the Storage Management section of NAS Manager.

A UNIX/Linux NFS client can issue the **quota** command to retrieve information regarding quota usage of a user or group, based on their ID. The retrieved report contains *block count*, *file count*, *quota limits on both*, and other information (based on options invoked with the command). For accurate syntax, refer to the client's man pages, as implementation varies between client operating systems.

The server reports only **Hard Limit** quota information through rquotad. Three different quota limitations can be defined:

- User and group quotas to limit **space** and **file quantity** for **individuals** and **groups** *within a virtual volume*.
- User and group quotas to limit **space** and **file quantity** for **individuals** and **groups** *within an entire file system*.
- Virtual volume quotas to limit **space** and **file quantity** *by a virtual volume as a whole*.



Note: rquotad reports quota usage information on explicitly defined quotas and automatically created (default) quotas. Default quota information will be reported if an explicit quota has not been defined.

The rquotad service can be configured to report quota information using one of two modes:

- **Restrictive mode.** For the user or group specified in the client-side quota command, the rquotad service reports the quota information for the quota with the most constraints.
- **Matching mode.** For the user or group specified in the client-side quota command, the rquotad service reports the quota information for the first quota that meets the parameters defined by the client-side quota command.



Note: If the rquotad service is disabled, all requests are rejected with an error code of "EPERM".

Restrictive mode operation

When in Restrictive mode, the rquotad service picks the first applicable quota threshold crossed. It enables the user to determine the amount of data that can be safely recorded against this quota before reaching its Hard Limit. This is the default configuration option for rquotad on the server.



Note: The restrictive mode option returns quota information combined from the quota that most restricts usage and the quota that most restricts file count. For example:

If the user quota allowed 10 K of data and 100 files to be added, and the virtual volume quota allowed 100 K of data and 10 files to be added, rquota would return information stating that 10 K of data and 10 files could be added. Similarly, if the user quota is 10 K of data of which 5 K is used, and the virtual volume quota is 100 K of data of which 99 K is used, rquota would return information stating that 1 K of data could be added.

The console command `rquota` is provided to change between the two options, and also to disable access to quota information. For information on how to configure rquota, refer to the *Command Line Reference*.

Matching mode operation

When in Matching mode, the rquotad service follows a specific order to find a match for relevant quota information:

- If rquotad returns quota information for a user, it returns the *user's quota within the virtual volume* if it exists;
- Otherwise, it moves to the *user's file system quota* if that exists;
- If no file system quota exists for the user, then it will move to the *virtual volume quota*.

In this manner, rquotad keeps checking until a quota is found for the specified user or group. Once the **first** matching quota is found, rquotad stops searching and returns the quota information.

If a user does not have a specifically defined quota in a virtual volume, or in a file system, and the virtual volume quota allows all users 100 K of data and 10 files, rquotad would return information stating that user's quota is 100 K of data and 10 files. Similarly, if the user has a specified virtual volume quota that is 200 K of data and 20 files, and a file system quota that is 400 K of data and 40 files, rquotad would return information about only the first quota, stating that 200 K of data and 20 files could be added.

Chapter 5: Using SMB for Windows access

Windows networks use the SMB protocol for file sharing between workstations and servers. This section contains information on using SMB with the NAS server.

SMB protocol support

The server implements the SMB protocols as used by Microsoft Windows platforms. From the client perspective, the server is indistinguishable from a Windows file server.

It provides all of the normal file-serving functions, including:

- Share manipulation (for example, add, list, and delete).
- File manipulation (for example, read, write, create, delete, move, and copy).
- File locking and byte-range locking.
- File access control using standard Windows ACLs.
- File and directory attributes (for example, read-only, and archive).
- Automatic creation of user home directories.



Note:

The server does not support the following SMB features:

- Windows Extended Attributes (note that this should not be confused with NFS or POSIX xattr).
- BranchCache.
- Support for remote management from Server Manager (Windows Server 2012 or later).
- SMB2 large read/write MTU (NAS Server limited to 64KiB).
- SMB3 Directory Leasing.
- SMB Direct (SMB3 over RDMA).
- Offloaded Data Transfer (ODX).
- Library storage (for Hyper-V management tools).

Prerequisites

To enable SMB access to the server:

- Enter a CIFS license key.
- Enable the CIFS service.
- Configure the server.

Depending on the security model used on the SMB network, configure the server using one of the following methods:

Security Model	Client Authentication	Configuration Method
NT Domain security	NT 4 only	Add server to NT domain
Active Directory	NT 4 only	Add server to NT domain
	Kerberos and NT 4	Join Active Directory

When configured to join an Active Directory, the server functions the same way as a server added to an NT domain, except that after joining an Active Directory, the server can authenticate clients using the Kerberos protocol as well as NT 4-style authentication. Most modern Windows clients support both authentication methods, though a number of older Windows clients only support NT 4-style authentication.

Supported clients

The server supports platforms and clients that are compliant with SMB versions 1, 2, 2.1, and 3.

Domain controller interaction

The storage server relies on Windows domain controllers to authenticate users and to obtain user information (for example, group membership). The server automatically discovers and connects to the fastest and most reliable domain controllers. Because operating conditions can change over time, the server selects the best domain controller every 10 minutes.

By default, when authenticating clients in an Active Directory, the server uses the time maintained by the domain controller, automatically adjusting for any clock inconsistencies.

Dynamic DNS

The storage server supports DNS and DDNS. For more information, see the Network Administration Guide.

SMB (CIFS) Statistics

SMB statistics for the storage server (in 10-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Supported SMB versions

The NAS server supports the SMB file sharing protocols with the following versions: SMB1, SMB2.0, SMB2.1, SMB3.

Server version	Default max supported SMB version
11.2 and earlier	SMB1
11.3 and later	SMB2

The maximum supported SMB version advertised by the NAS server can be configured using the `smb-max-supported-version` CLI command (see below). The maximum supported SMB dialect is not server or cluster-wide - it is set on a per-EVS basis.

The NAS server supports UCS-2 character encoding when using the SMB protocols (the character set is not negotiable when using the SMB2 protocol).



Notes:

- A valid CIFS license is required in order to enable SMB2 or SMB3 support (CIFS is a dialect of SMB). For more information about license keys, refer to the Server and Cluster Administration Guide.
- One of the features of SMB is the ability to assign rights to machine (computer) accounts. The feature acts the same way as authentication of a normal user for an SMB session and can be used for authentication using machine accounts (SessionSetup SMB requests), and for management (add, delete, list) of rights for machine accounts. A machine account is generated automatically by the operating system and registered in Active Directory. It can be used for authentication within a domain. Machine account authentication can be only done by an application which has built-in support. For example, Hyper-V server allows storing virtual machines on remote shares. Such shares should allow full access for the machine account of a computer running Hyper-V server. Authenticated connections using machine accounts will show up in "connection" command output as if it was a normal user connection. The man pages for `cifs-saa` and `cacls-add` include an example of computer account use.

Specifying the SMB version for use by the EVS

To specify a version of the SMB protocol for use by the EVS, use the following commands:

- `smb-max-supported-version` - sets or displays the maximum supported version for both the NAS server and the client. The default is SMB2.
- `smb-min-supported-version` - limits the minimum supported version for both the NAS server and the client. The default is SMB1.



Note: SMB2 cannot be enabled if there are NT4 names and no ADS names configured on the server.

Disabling SMB1

To disable SMB1 on the NAS server, use the following command:

smb-min-supported-version 2

This command sets the minimum SMB version on the NAS server to SMB2, therefore preventing any new clients connecting using SMB1.

**Notes:**

- When a client initiates an SMB connection it advertises support for several versions/dialects. The server will choose the maximum version/dialect the client provides that is within its configured maximum/minimum. For example, a client that supports SMB1, SMB2 and SMB2.1 can establish an SMB2 connection if the max-supported version on the server is set to SMB2.
- Some SMB clients cache the connection type they last used with a server. If they last used SMB2/2.1/3, they may not offer SMB1 as an option until they are restarted.
- Existing SMB/SMB2/SMB3 client connections will continue to function after the minimum supported version has been raised.

Supported SMB3 functionality for Hyper-V

The NAS server supports SMB3 functionality for Microsoft Hyper-V over SMB shares, including transparent failover, continuous availability, and shadow copies.

- **Continuous Availability:** Enables files that are opened using SMB3 on a continuously available share to survive network failures or cluster node failures. For example, if one cluster node fails, the client transparently reconnects to another cluster node without interruption to the client applications.

If a continuously available share is changed from a cluster to a single server, and then back to a cluster, the server keeps the continuous availability of the share.



Note: Continuous Availability can impact SMB performance and should only be enabled where it is required, such as with Microsoft Hyper-V or Microsoft SQL Server. When this feature is in use, it is also recommended that the Administrator disables DDNS on the server.

- **Persistent file handles:** Enables clients to transparently reconnect to disconnected SMB sessions. A persistent handle is preserved after a disconnection and blocks any attempts to open files while it waits for the client to reconnect.
- **VSS for SMB file shares:** The File Server Remote VSS (Volume Shadow Copy Service) Protocol (FSRVP) is a protocol for Windows Server 2012 that creates shadow copies of file shares on a remote computer. This protocol is most commonly deployed with Hyper-V and enables backup applications to create application-consistent backup and restore of VSS-aware applications storing data on network file shares.
- **Service Witness Protocol:** Enables a registered client to receive notification of any state changes on a continuously available server, without needing to wait for the connection to time out. This ensures that there is a fast notification and recovery time from an unplanned failure, such as a network loss.
- **SMB3 Multichannel:** Enables file servers to use multiple network connections simultaneously. This increases the network performance and availability of the file servers, and improves data throughput and fault tolerance. With SMB3 Multichannel, applications can utilize all available network bandwidth and increase resilience during network failures.

SMB3 Multichannel support

SMB3 Multichannel enables file servers to use multiple network connections simultaneously. This feature increases the network performance and availability of the file servers.

SMB3 Multichannel benefits include:

- Automatic configuration.
- Client-side network processing on multiple CPU cores.
- Increased data throughput.
- Increased fault tolerance.
- Resilience during network failures.

SMB3 Multichannel is automatically enabled if the EVS is configured for version 3 of the SMB protocol. To set the version, use the `smb-max-supported-version 3` command.

CLI commands

All settings for Multichannel are per EVS. Use the following CLI commands to configure or view the maximum channels per session:

- **smb3-multichannel-max-channels-per-session-set**
Sets the maximum number of channels for all subsequent sessions.
 - Default: 32 channels
 - Minimum: 2 channels
 - Maximum: 64 channels
- **smb-multichannel-max-channels-per-session-show**
Shows the maximum number of channels per session.

For more information about the CLI commands, see the *Command Line Reference*.

SMB3 Encryption support

SMB Encryption provides end-to-end encryption of SMB data and protects against potential eavesdropping attacks on untrusted networks. Consider using SMB3 Encryption for any scenario in which sensitive data needs protection from man-in-the-middle (MITM) attacks.

SMB3 Encryption uses the Advanced Encryption Standard (AES)-CCM algorithm for both encryption and signing.

The main benefits of SMB3 Encryption are:

- No deployment requirements other than changing the SMB server settings.
- No dedicated hardware requirements unlike most storage area networks (SANs).
- Provides secure access to the server and shares.
- Protects data from eavesdropping attacks on untrusted networks.
- Provides end-to-end data encryption in-flight.

SMB3 Encryption is available only if the EVS is configured for version 3 of the SMB protocol. To set the version, use the **smb-max-supported-version 3** command.



Caution: SMB3 Encryption can severely impact SMB performance and should be enabled only where it is necessary.

CLI commands

To use SMB3 Encryption, the **cifs-auth** command must be set to **on**.

Use the following commands to enable or disable SMB3 Encryption on an EVS:

- **smb3-encryption-enable**
Enables encryption on the current EVS.
- **smb3-encryption-disable**
Disables encryption on the current EVS.

Use the following options on the `cifs-share` command to enable or disable encryption on a share:

- `--encrypt-data`
Enables encrypted client access to a share.
- `--no-encrypt-data`
Disables encrypted client access to a share.



Note: Always disable share-level encryption before you downgrade to a pre-feature build to prevent the share being deleted.

SMB2 clients cannot connect to a server or share that requires encryption. Use the following commands together with the `smb3-encryption` and `cifs-share` commands to allow or reject unencrypted access for SMB2 clients:

- `smb3-reject-unencrypted-access-enable`
Rejects unencrypted client access to the current EVS.
- `smb3-reject-unencrypted-access-disable`
Allows unencrypted client access to the current EVS.



Notes:

- SMB3 Encryption does not affect SMB1 clients. To prevent access by SMB1 clients, you must turn off the SMB1 server by using the `smb-min-supported-version 2` command.
- Some Remote Procedure Call (RPC) virus scanners are not compatible with SMB3 Encryption and will not work with `smb3-reject-unencrypted-access` enabled. Check with your virus scanner vendor for information about compatibility.

For more information about the CLI commands, see the *Command Line Reference*.

SMB3 Encryption client file access configurations

The following tables show how the EVS encryption, share encryption, and reject unencrypted access options affect SMB3 and SMB2 clients when they try to access files on a NAS Server.

File access for SMB3 clients

	EVS encryption ON <code>smb3-encryption-enable</code>	EVS encryption OFF <code>smb3-encryption-disable</code>
Share encryption ON <code>--encrypt-data</code>	Client must support encryption	Client must support encryption
Share encryption OFF <code>--no-encrypt-data</code>	Client must support encryption	Client not required to support encryption

File access for SMB2 clients

	Reject unencrypted access ON <code>smb3-reject-unencrypted-access-enable</code>		Reject unencrypted access OFF <code>smb3-reject-unencrypted-access-disable</code>	
	EVS encryption ON <code>smb3-encryption-enable</code>	EVS encryption OFF <code>smb3-encryption-disable</code>	EVS encryption ON <code>smb3-encryption-enable</code>	EVS encryption OFF <code>smb3-encryption-disable</code>
Share encryption ON <code>--encrypt-data</code>	Client must support encryption	Client must support encryption	Client not required to support encryption	Client not required to support encryption
Share encryption OFF <code>--no-encrypt-data</code>	Client must support encryption	Client not required to support encryption	Client not required to support encryption	Client not required to support encryption



Note: SMB3 Encryption does not affect SMB1 clients. To prevent unencrypted access by SMB1 clients, you must turn off the SMB1 server by using the `smb-min-supported-version 2` command.

Configuring SMB security

The server integrates seamlessly into the existing domain and simplifies access control by performing all authentications against the existing domain user accounts.



Note: Only accounts that have been created in the domain or in a trusted domain can access the server.

When a user attempts to access a share, the server verifies appropriate permissions; once access is granted at this level, standard file and directory access permissions apply.

The server operates on a specific domain and can, optionally, join an Active Directory. It interacts with a domain controller (DC) in its domain to validate user credentials. The server supports Kerberos-based authentication to an Active Directory, as well as NTLM authentication (using pre-Windows 2000 protocols). In addition to users belonging to its domain, the server allows connections from members of trusted domains.

The server automatically grants administrator privileges to domain administrators who have been authenticated by the DC. In addition, local administration privileges can be assigned, including backup operator privileges to selected groups (or users).



Note: SMB can assign rights to machine (computer) accounts. A machine account is generated automatically by the operating system and registered in Active Directory. It can be used for authentication within a domain. A machine account authentication can be only done by an application which has built-in support. For example, Hyper-V server allows storing virtual machines on remote shares. Such shares should allow full access for the machine account of a computer running Hyper-V server.

Assigning SMB names

Windows clients access the server through configured SMB names. Traditional Windows servers have a single host name. In environments where multiple Windows servers are being consolidated, the server can be configured with multiple SMB names.

In order to appear as a unique server on a Windows network, the server will do the following for each configured SMB name:

- Allow administration through the Microsoft Server Manager (NT 4) or Computer Management (Windows 2000 or later) administrative tools.
- If NetBIOS is enabled, register each SMB name with the domain Master Browser so each name appears as a unique server in Network Neighborhood.
- Register each SMB name with DDNS or WINS for proper host name resolution.
- Support up to 256 SMB names per EVS.

Viewing SMB setup

Windows clients access the server through configured SMB names. Traditional Windows servers have a single host name. In environments where multiple Windows servers are being consolidated, the server can be configured with multiple SMB names.

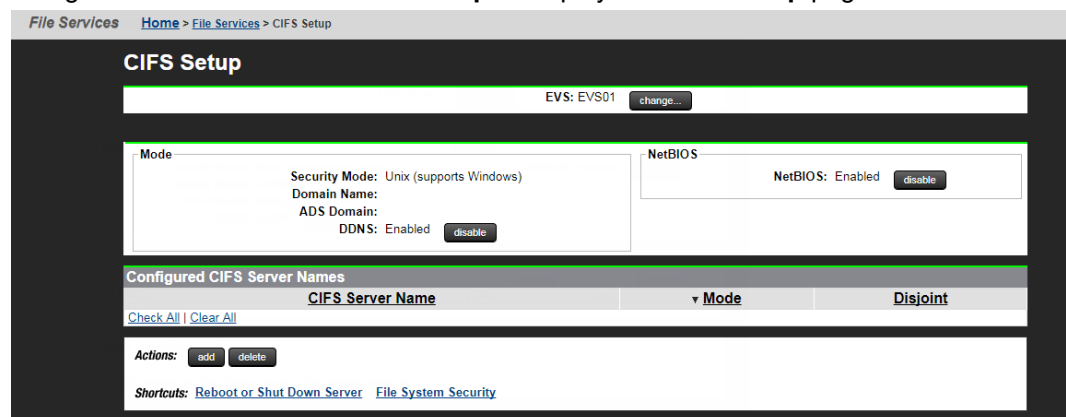
In order to appear as a unique server on a Windows network, the server will do the following for each configured SMB name:

- Allow administration through the Microsoft Server Manager (NT 4) or Computer Management (Windows 2000 or later) administrative tools.
- Register each SMB name as a server with the domain Master Browser so each name appears as a unique server in Network Neighborhood.
- Register each SMB name with DDNS and WINS for proper host name resolution.


You can view the SMB (CIFS) setup details in the NAS Manager.

Procedure

1. Navigate to **File Services > CIFS Setup** to display the **CIFS Setup** page.



The following table describes the fields on this page:

Field/Item	Description
EVS	Indicates the selected EVS. Click change to select another EVS.
Mode	
Security Mode	The currently configured security mode of the EVS.
Domain Name	The name of the NT domain in which the server resides. The domain is set when the first CIFS name is added.
ADS Domain	The domain where the server is located.
DDNS	Indicates whether DDNS is enabled or disabled.
NetBIOS	
NetBIOS	<p>When NetBIOS is enabled, it allows NetBIOS and WINS use on this server. If this server communicates by name with computers that use earlier Windows versions, this setting is required. By default, the server is configured to use NetBIOS. Click disable to disable NetBIOS.</p> <div style="background-color: #ffffcc; padding: 10px; border: 1px solid #ccc;"> <p> Caution: Before choosing to disable NetBIOS, verify that there is no need to use NetBIOS, WINS, or legacy NetBT-type applications for this network connection. If this server communicates only with computers that run Windows 2000, Windows XP, or Windows 2003, disabling NetBIOS will be transparent and may even result in a performance benefit.</p> </div>
Configured CIFS Server Names	
CIFS Server Name	A list of CIFS names added to the selected EVS.
Mode	<p>Displays the mode for each CIFS serving name. Mode defines the authentication protocol used to communicate with the Windows network clients and domain controllers. The mode can be:</p> <ul style="list-style-type: none"> ▪ ADS: The ADS-style communication protocol (Kerberos) is used to communicate with the Windows clients and domain controllers. ▪ NT4: The Windows NT 4-style communication protocol (NTLMSSP) is used to communicate with the Windows clients and domain controllers.
Disjoint	<p>Indicates whether the DNS suffix matches the Active Directory domain primary DNS suffix.</p> <ul style="list-style-type: none"> ▪ no: There is no disjoint namespace between the DNS and ADS. ▪ yes: There is a disjoint namespace between the DNS and ADS.

Field/Item	Description
add	Opens the Add CIFS Server Names page, in which you can add server names.
delete	Deletes the selected CIFS server.
Reboot or Shut Down Server	Opens the Reboot or Shutdown Server page, which enables you to shut down or reboot a server, a cluster node, or an entire cluster.
File System Security	Opens the File System Security page, which displays all EVSs and the configured security mode.


Joining an Active Directory

You can add an SMB (CIFS) server to an ADS domain in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > CIFS Setup** to display the **CIFS Setup** page.
2. Click **add** to display the **Add CIFS Server Names** page.

The following table describes the fields on this page:

Field/Item	Description
EVS	Displays the name of the EVS to which the new server name is added.
CIFS Server Name	The computer name through which CIFS clients will access file services on the server. In an ADS domain, the maximum number of characters for the CIFS server name is 63. In an NT4 domain, the maximum number of characters for the CIFS server name is 15.
Domain	<p>Indicates if the CIFS server is to be a part of an NT4 domain or an ADS domain, and allows you to specify the settings required to be a part of the domain.</p> <ul style="list-style-type: none"> Before an NT4 name is added to the server, a computer account must exist in the NT Domain. Use Server Manager on the Domain Controller to create the computer account. <p> Note: This must be done for each NT4 CIFS name added.</p> <ul style="list-style-type: none"> When adding an ADS name to the server, the server will automatically create a computer account in the Active Directory for each added name.
NT4	
NT4	Select the NT4 option to indicate that the CIFS server is to be a part of an NT4 domain.
Domain Name	Indicates that the CIFS server is to be part of an NT4 domain.
ADS	
ADS	Indicates that the CIFS server is to be a part of an ADS domain.
IP Address	The IP address of a domain controller in the Active Directory in which the server will be configured.
DC Admin User	A user account that is a member of the Domain Administrators group. This privilege is necessary to create a computer account in the Active Directory.
DC Admin Password	Password for the Domain Administrator user.
Folder	The folder in the Active Directory in which the computer account should be created. By default, the computer account will be created in the Computers folder.

Field/Item	Description
DNS Suffix	Use this option only if you need to set a DNS suffix other than the Active Directory domain's primary DNS suffix. (For example, set this if you have a disjoint domain.)

3. Enter the name corresponding with the newly created computer account into the field labeled **CIFS Server Names**.
 - In an ADS domain, the maximum number of characters for the CIFS server name is 63.
 - In an NT4 domain, the maximum number of characters for the CIFS server name is 15.
4. If you are adding a server to an NT domain, complete the following steps. If you are joining an ADS domain, see the next step.
 - a. Select the **NT4** option.
 - b. Enter the domain name.
 - c. Click **OK** to return to the **CIFS Server** page.
 - d. To create an NT 4 domain account, run Server Manager from a domain controller in the NT 4 Domain, and create a new Windows NT Workstation or Server account using the desired host name.
5. If you adding ADS domain, complete the following steps:
 - a. Select the **ADS** option.
 - b. In the **IP Address** field, specify the IP address of a domain controller in the Active Directory.
 - c. In the **DC Admin User** field, specify a user account that is a member of the Domain Administrators group. This privilege is necessary to create a computer account in the Active Directory. When specifying a user account from a trusted domain, the user account must be entered using the Kerberos format; that is, `administrator@ADdomain.mycompany.com`, not `ADdomain \administrator`.
 - d. In the **DC Admin Password** field, specify the password associated with the DC admin user name specified.
 - e. In the **Folder** field, specify the name of the folder in the Active Directory in which the computer account should be created. By default, the computer account will be created in the Computers folder.
 - f. Use the DNS Suffix option only if you need to set a domain name service suffix other than the AD domain's primary DNS suffix. For example, set this if you have a disjoint domain.
6. Click **OK**.

Removing SMB server names

SMB (CIFS) server names can be removed from the server's configuration in the NAS Manager.

When ADS SMB names are removed, the corresponding computer account in the Active Directory is also removed. Computer accounts in NT 4 Domains must be deleted manually through Server Manager.



Caution: SMB Name Deletion Alert! At least one SMB name must be configured on the server to support connections from Windows clients. As a result, if the last configured SMB name is removed, Windows clients are no longer able to access the server over SMB.



Note: DNS entries do not de-register automatically after removing a SMB server name, so the admin should delete the SMB server name entry from DNS manually.

Configuring local groups

In a Windows security domain, users and groups identify users (for example, `vsmith`) and groups of users (for example, `software`) on the network. Apart from the user-defined network group names (for example, `software`, `finance`, and `test`), Windows also supports a number of built-in or local groups with each providing various privileges and levels of access to the server on which they have been configured.

These groups exist on every Windows computer. They are not network groups, but are local to each computer. So, the user `vsmith` may be granted Administrator privileges on one computer and not on another.

On the server, the administrator can add users to any of the following local groups:


- **Root:** If a user is a member of the local Root group, the user bypasses all security checks, and can take ownership of any file in the file system.
- **Administrators:** If a user is a member of the local Administrators group, the user can take ownership of any file in the file system.
- **Audit Service Accounts:** If a user is a member of the Audit Service Accounts group, the server does not add any of their events to the audit log. However, the server does add events to the audit log for any user who is **not** a member of this group. These events consist of the Windows file access and deletion events which are recorded by the server. As an alternative to the NAS Manager, it is possible to use the `localgroup` CLI commands to add, remove or display the users for this group.
- **Backup Operators:** If a user is a member of the local Backup Operators group, the user bypasses all security checks, but cannot take ownership of a file in the file system. The privilege to bypass all security checks in the file system is required for accounts that run Backup Exec or perform virus scans. Virus scanner servers that are a part of the Backup Operators group can, however, take ownership of any file in the file system.
- **Forced Groups:** If a user is a member of the local Forced Groups group, when the user creates a file, the user's defined primary group is overridden and the user account will be used to indicate the file creator's name.

Adding a local group or local group members

You can add a local group or local group members in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > Local Groups** to display the **Local Groups** page.
2. If necessary, click **Change** to select a different EVS security context or to select the global configuration. Changes made to local groups using this page apply only to the currently selected EVS security context.
 - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
 - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To manage local groups for an EVS that uses an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.
3. Click **add** to display the **Add Local Group** page.



The following table describes the fields on this page:

Field/Item	Description
Group	<ul style="list-style-type: none"> ▪ Select Use existing local group and then select from the list to add from an existing local group. ▪ Select Add new local group and then enter the name to add a new local group.
Members	Enter the member's user name and then click add . To remove a member's user name, click on the X button.

4. To add a new member to an existing local group, complete the following.
 - a. Select the **Use existing local group** option.
 - b. Using the list of local groups, select the group to which you want to add a member.
 - c. Enter the new member's user name in the Members field.
 - d. Click **add**.
 - e. Repeat steps to add more members.
 - f. Click **OK**.
5. To add a new local group, complete the following:
 - a. Select the **Add new local group** option.
 - b. Enter the new local group name in the **Members** field.
 - c. If necessary, you can now enter group members for the new group. To enter members user names, enter each member's user name in the **Members** field.
 - d. Click **add**.
 - e. Repeat steps to add more members.

Deleting a local group or local group members

Once created, group names may not be changed. To change a group name, you must delete the group, then create a new group, and add members to the new group.

Procedure

1. Navigate to **Home > File Services > Local Groups** to display the **Local Groups** page.
2. If necessary, click **Change** to select a different EVS security context or to select the global configuration. Changes made to local groups using this page apply only to the currently selected EVS security context.

Deleting a local group is a two-stage process; you must delete all members of the group before you can delete the group itself.

3. Delete all members of the group:
 - a. Fill the check box next to all members of the group you want to delete.
 - b. Click **delete** to delete the selected group members.
 - c. Click **OK** to confirm the deletion return to the **Local Groups** page.
4. Delete the local group:
 - a. Fill the check box next to the group you want to delete.
 - b. Click **delete** to delete the selected group.
 - c. Click **OK** to confirm the deletion return to the **Local Groups** page.

Local user authentication for SMB and FTP users

Local User Authentication can be used by the NAS server to authenticate SMB and FTP users without reference, even indirectly, to an external source of authentication, like Kerberos or a Domain Controller. Users and passwords are configured and managed via the command line.

There is support for both NTLMv1 and all types of NTLMv2 authentication. The feature is available automatically and requires no license. It may be left un-configured with no effect. All configuration and management is performed via CLI, using a family of commands with the *local-password-* prefix. The commands share a common manual page.

There is no NAS Manager interface, and standard tools such as Microsoft Management Console are not supported. There is no way of importing passwords from another source, such as NIS, LDAP, Active Directory or */etc/passwd* files. No check is made of password strength and no password expiry is enforced. There is no supported, secure method for transferring local passwords between EVSs. When switching from a cluster-wide configuration to per-EVS, local password settings are not cloned.



Note: Although the NAS server supports both NTLM1 and NTLM2, it only supports the use of FTP with NTLM1. NTLM2 in FTP is not supported.

Using local user authentication

Many SMB clients require a user to be identified by a *username* and *domain*. When using local user authentication, the domain may be any string you choose and need not correspond to any other domain in use on the network, the IP of any EVS, or any SMB name.

Note that it is common to use ALL CAPS for domain names, but it is not a requirement.

Procedure

1. To add user *alice* in domain *EXAMPLE*, with password *Alligat0r*, enter:
`local-password-set EXAMPLE\alice Alligat0r`
 If preferred, the password may be entered interactively by omitting the password from the above command and then entering the password when prompted. Note that interactive entry limits password length to 127 characters.
 All three fields are case sensitive, with *EXAMPLE\Alice*, *example\alice* and *EXAMPLE\alice* considered three separate users.
2. To see which users are configured, enter:
`local-password-list`
 All configured users are shown, along with their stored password hashes.
3. To delete a single user, in this case, *bob* in domain *EXAMPLE*, enter:
`local-password-delete EXAMPLE\bob`
4. To irreversibly delete all configured users, enter:
`local-password-delete-all --confirm`
5. FTP: To use local user authentication with FTP, enable SMB-based authentication with:
`ftp-cfg --ntsecurity on`

SID mappings

The owner (typically the creator) of any file is identified by a Security Identifier (SID) associated with that user. The local users feature does not automatically create an SID for each user, so you must assign SIDs.

To assign SIDs:

Procedure

1. To assign SID *S-1-81-1* to user *EXAMPLE\david*, enter:
`user-mappings-add --nt-name EXAMPLE\david --nt-id S-1-81-1`
 If an SID is not assigned to a local user in this way, that user will still be able to authenticate but will be treated as the *Anonymous Logon* user. You should ensure that any SIDs are assigned before a user connects, as any changes will not take effect while they remain connected.
2. Typically, a user would also have a primary group SID (used to give a group to files created by that user) and may be a member of one or more additional groups. These may be configured using the `primary-group-set` and `localgroup` commands, respectively.
`$ primary-group-set "Unix user\1234" "Unix group\1234"`
`$ localgroup add "Backup Operators" "Unix user\521"`
 No programmatic assistance is provided for allocating SIDs. Instead, it is recommended that the first local user be given the SID *S-1-81-0*; the second, *S-1-81-1*, and so on. It is recommended that the first locally allocated group be given the SID *S-1-82-0*; the second, *S-1-82-1*, and so on.
3. Use `user-mappings-list` to see users with SIDs assigned.

Configuration

NTLMv2 authentication for local users using the NTOWF_V2 hash is not supported in versions prior to 12.3, is supported but is disabled by default in version 12.3, and is supported and enabled by default in versions 12.4 and later. The NAS server does not store user passwords in plain text. Instead, user passwords are passed through a variety of one-way functions to produce *hashes*; these do not permit retrieval of the originally entered password.

To perform NTLMv2 authentication using the NTOWF_V2 hash:

Procedure

1. Ensure the appropriate password hashes are available.
The server calculates up to four of these hashes: NTOWF_V1, LMOWF_V1, NTOWF_V2 and LMOWF_V2, with the set shown by the *local-password-list* command dependent on your configuration.
2. The server may also store passwords encoded using a two-way function. While passwords encoded this way are not stored as plain text, a skilled attacker could reverse the two-way function to obtain the original input. Consequently, it is recommended to disable the two-way function in version 12.3 and later. For version 12.3 and later, disable the two-way function using:
set shouldKeepObfuscatedLocalPasswords false
Note that you should only use this command once you are sure you do not wish to downgrade past version 12.3; software versions prior to this do not support one-way hashes and rely on the two-way function encoding to authenticate local users. Disabling the two-way function and downgrading will result in local users being unable to authenticate unless their passwords are reset.
3. In versions that support it, NTLMv2 authentication for local users using the NTOWF_V2 hash may be enabled with:
set ntlmV2-authentication-allowed true
This is a cluster-wide setting, is persistent across reboots, and will take immediate effect. NTLMv2 has long been supported for externally managed users.

Configuring SMB shares

SMB (CIFS) shares can be set up on mounted volumes. The server can support more than 1,000 shares. However, the exact limit of any share allocation depends on the server's configuration.

Adding an SMB share

You can add an SMB (CIFS) share in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > CIFS Shares**.
2. Click **add** to display the **Add Share** page.

File Services [Home](#) > [File Services](#) > [CIFS Shares](#) > Add Share

Add Share

EVS / File System: g1-eva3 / PHDS1 [change...](#)

Share Name:

Comment:

Path: [browse...](#)

Create path if it does not exist.
(See online help for this option).

Max Users:

Show Snapshots: ▼

Cache Options: ▼

Transfer to Object Replication Target: ▼ (The File System Default: Enabled)

Access Configuration:

(Enter IP-based values first, if possible)

Follow Symbolic Links: Follow Global Symbolic Links:

Force Filename to be Lowercase: Enable ABE:

Enable Virus Scanning: (Virus scanning is disabled on the EVS.)

Ensure Share Continuously Available:

[OK](#) [cancel](#)

Share Permissions

Everyone

Type:

Allow Deny

Full Control

Change

Read

After creating the share, go to Share details page to change the default permissions.

User Home Directory Mapping

Mode: ▼


Choose "ADS" mode to use the home directory path value provided by Active Directory.




Path: [browse...](#)


(Relative to Share path)

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Currently selected file system to which the CIFS share will link.
Cluster Namespace	Currently selected cluster namespace to which the CIFS share will link.
change / browse (depending on Web browser)	Enables the user to select a different file system or (on a cluster) a different cluster namespace.
Share Name	Name of the CIFS share.
Comment	Additional information associated with the CIFS share. This information is often displayed to clients along with the share name.
Path or CNS Path	<p>The directory to which the CIFS share points. Users accessing the share are able to access this directory, and any directories under it in the directory tree. To find a directory, click change / browse.</p> <p>On a file system only, select the Create path if it does not exist option to create the path if it does not already exist. If the file system is mounted read-only, for example it is an object replication target, it is not possible to create a new directory. Select a path to an existing directory.</p>
Max Users	The maximum number of users who can be associated with the CIFS share. The default is unlimited.
Show Snapshots	<ul style="list-style-type: none"> ▪ Show and Allow Access: Displays and allows access to snapshots. ▪ Hide and Allow Access: Hides snapshots, but still allows access to the hidden snapshots. ▪ Hide and Disable Access: Hides and disallows access to snapshots. <p>Changes to this setting become effective when a CIFS client refreshes its folder view.</p>
Cache Options	<ul style="list-style-type: none"> ▪ Manual Local Caching for Documents. The Manual mode permits the user to specify individual files required for offline access. This operation guarantees a user can obtain access to the specified files whether online or offline.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Automatic Local Caching for Documents. The Automatic mode is applied for all non-executable files on the entire share. When a user accesses any non-executable file in this share, it is made available to the user for offline access. This operation does not guarantee that a user can access all the non-executable files, because only those files that have been used at least once are cached. Automatic can also be defined for programs. ▪ Automatic Local Caching for Programs. The Automatic mode is applied for all executable files on the entire share. When a user accesses any executable file in this share, it is made available to the user for offline access. This operation does not guarantee that a user can access all the executable files, because only those executable files that have been used at least once are cached. Automatic can also be defined for documents. ▪ Local Caching Disabled. No caching of files or folders occurs. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;">  Note: The server supports Offline Files Access. This allows Windows 2000 and later clients to cache files that are commonly used from a network/file share. To use Offline Files, the client computer must be running Windows 2000 or later. </div>
Transfer to Object Replication Target	<p>When a file system is recovered from a snapshot, one of the final steps is to import the CIFS shares found in the snapshot representing the selected version of the file system. Only those CIFS shares marked as transferable will be imported.</p> <p>Use the list to specify one of the following:</p> <ul style="list-style-type: none"> ▪ Enable: CIFS shares will be transferred to recovered file systems. ▪ Disable: CIFS shares will not be transferred to recovered file systems. ▪ Use FS default (the default): When the target file system is brought online, CIFS shares will be transferred if Transfer Access Points During Object Replication is enabled for the file system.
Access Configuration	IP addresses of the clients who can access the share (up to 5,957 characters allowed in this field). Refer to IP Address Configuration** at the end of this table.
Follow Symbolic Links	Enables the following of symlinks for this share.

Field/Item	Description
	 Note: As of release 12.2 of the NAS Platform, clients using SMB2 or later are able to follow relative symlinks to files on storage without the use of global symlinks, if smb2-client-side-symlink-handling is enabled.
Follow Global Symbolic Links	Enables CIFS clients to follow global (absolute) symlinks via the Microsoft DFS mechanism for this share.
Force Filename to be Lowercase	Forces all filenames generated on this share to be lowercase. This is useful for interoperability of UNIX applications.
Enable ABE	<p>By default, ABE is disabled for shares and on the server/cluster as a whole. Before enabling ABE for a share, you must make sure ABE is enabled for the server/cluster as a whole (the CLI command to enable ABE support is <code>fsm set disable-ABE-support false</code>).</p> <p>When enabled, ABE filters the contents of a CIFS share so that only the files and directories to which a user has read access rights are visible to the user.</p>  Note: Enabling ABE can impact CIFS performance.
Enable Virus Scanning	<p>If virus scanning is enabled and configured for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is enabled by default. If virus scanning is not enabled for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is not enabled by default, but you can enable it a per-EVS basis.</p>  Note: Virus scanning is set up on a per-EVS basis, or for all EVSs using the global configuration context, but cannot be set up on a per-server or per-cluster basis.
Ensure Share Continuously Available	<p>Enables persistent file handles and transparent failover on the CIFS share. When enabled, Windows-based clients can continuously access the CIFS share if a network failure or a cluster node failure occurs. For example, if one cluster node fails, the client transparently migrates to another cluster node without any interruption to the client applications.</p> <p>This SMB3 option is available only in a clustered environment of more than one cluster node, and is disabled by default.</p>

Field/Item	Description
	 Note: Continuous Availability can impact CIFS performance and should only be enabled where it is required, such as with Microsoft Hyper-V or Microsoft SQL Server. When this feature is in use, it is also recommended that the Administrator disables DDNS on the server. If the file system is an object replication target, continuous availability is not effective until the file system is promoted.
Share Permissions	
Share Permissions List	By default, when a CIFS share is created, the group Everyone is added to the share permissions list.
User Home Directory Mapping	
Mode	<p>Used to specify how per-user home directories are created. The following options are available:</p> <ul style="list-style-type: none"> ▪ Off. Disable the home directory feature for this share. Do not automatically create home directories on this share for users. This is the default. ▪ ADS. Create the user home directories based on the home directory information supplied by the Active Directory server for each user. If you select ADS, do not specify a Path. ▪ User. Create the user's home directory by converting the user's Windows user name to lower case. (The user's Window's domain name, if any, is ignored.) For example, a user DOMAIN\John Smith would result in a home directory of john_smith. ▪ DomainAndUser. Create the user's home directory by creating a directory named for the user's Windows domain name, then converting the user's Windows user name to lower case and creating a sub-directory by that name. For example, a user DOMAIN\John Smith would result in a home directory of domain\john_smith. ▪ Unix. Create the user's home directory by converting the user's UNIX user name to lower case.
Path	Per-user home directories will be created in the specified Path , relative to the share root, which is specified without a leading \. If this field is left blank, user home directories will be created directly in the share root.

Field/Item	Description
	<p>By default, only one share per file system can be configured with home directories. The <code>cifs-home-directory</code> command can be used to relax this restriction, in which case great care must be taken not to configure conflicting home directories.</p> <p>For example, a share with a path of <code>\home1</code> and a share with a path of <code>\home2</code> would not cause a conflict, whatever home directory paths were configured. However, a share with a path of <code>\</code> and a default home directory path would conflict with a share with a path of <code>\dir</code> and a default home directory path.</p>

3. Click **change** to change the EVS/File System or Cluster Name Space (CNS) in which the CIFS share will reside.
4. Enter the Share Name. Clients will access the share through this name.
5. Type a comment that is meaningful to you or your users. This comment appears in Windows Explorer on client computers, and it is optional.
6. Type the Path to the directory being shared. Click **browse** to help find an existing directory (this button only exists if the path being created is the path in a file system, not a name space). To create the path automatically when it does not already exist, select the **Create path if it does not exist** check box.



Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users, that is, the permissions are set to `rxwxrwxrwx`. It is recommended that such directories are created via CIFS or NFS or that such directories are given the desired permissions explicitly after being created via this option.

7. To limit the number of users who can access the share simultaneously, enter the number of users in the **Max Users** field. By default, a share has unlimited access.



Note: This only limits the number of users that can concurrently access a share. It does not provide security restrictions.

8. If snapshots are present and you want them to be visible to clients, select the **Show snapshots** check box. If snapshots are not taken, or if you don't want clients to view snapshots, clear this check box.
9. To allow clients to traverse symbolic links, select the **Follow Symbolic Links** check box.
10. To enable CIFS clients to follow global (absolute) symlinks via the Microsoft DFS mechanism for this share, select the **Follow Global Symbolic Links** check box.
11. To force all characters to be lowercase when files and directories are created, select the **Force Filenames to be Lowercase** check box.
12. To disable Virus Scanning for the share, clear the **Enable Virus Scanning** check box. The default setting will add this share to the server-wide Virus Scan.



Note: Enable Virus Scanning is checked by default only if virus scanning is globally enabled.

13. To **enable ABE** (access based enumeration), select the check box.

ABE is disabled by default. When enabled, ABE filters the contents of a CIFS share so that only the files and directories to which a user has read access rights are visible to the user.



Note: Enabling ABE can impact CIFS performance.

14. To enable persistent file handles and transparent failover on the share, select the **Ensure Share Continuously Available** check box.



Note: Enabling Continuous Availability can impact CIFS performance.

15. To alter the caching option (Offline Files Access), select the desired new value from the Cache Options list.
16. To import the CIFS shares found in the snapshot representing the selected version of the file system, select the desired new value from the Transfer to Object Replication Target list. Only those CIFS shares marked as transferable will be imported.
17. In the **Access Configuration** field, specify the IP addresses of the clients who can access the share and the client's permissions for this share. The table outlines what to type in this field.

What to Type	Means
Blank or *	All clients can access the export.
Specific address or name. Examples: 192.0.2.0, client.dept.example.com	Only clients with the specified names or addresses can access the export.
A range of addresses using Classless Inter-Domain Routing (CIDR) notation. Example: 192.0.2.0/24	Clients with addresses within the range can access the export.
Partial address or name using wildcards. Examples: 192.0.*.*, *.example.com	Clients with matching names or addresses can access the export.

18. Click **OK**.

Viewing and modifying SMB shares details


You can view and modify SMB shares details in the NAS Manager.




Procedure


1. Navigate to **Home > File Services > CIFS Shares**.
2. Select the check box for the share to view or modify, and then click **details**.

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Currently selected file system to which the CIFS share will link.
Cluster Namespace	Currently selected cluster namespace to which the CIFS share will link.
change / browse (depending on Web browser)	Enables the user to select a different file system or (on a cluster) a different cluster namespace.
Share Name	Name of the CIFS share.
Comment	Additional information associated with the CIFS share. This information is often displayed to clients along with the share name.
Path or CNS Path	<p>The directory to which the CIFS share points. Users accessing the share are able to access this directory, and any directories under it in the directory tree. To find a directory, click change / browse.</p> <p>On a file system only, select the Create path if it does not exist option to create the path if it does not already exist. If the file system is mounted read-only, for example it is an object replication target, it is not possible to create a new directory. Select a path to an existing directory.</p>
Max Users	The maximum number of users who can be associated with the CIFS share. The default is unlimited.
Show Snapshots	<ul style="list-style-type: none"> ▪ Show and Allow Access: Displays and allows access to snapshots. ▪ Hide and Allow Access: Hides snapshots, but still allows access to the hidden snapshots. ▪ Hide and Disable Access: Hides and disallows access to snapshots. <p>Changes to this setting become effective when a CIFS client refreshes its folder view.</p>
Cache Options	<ul style="list-style-type: none"> ▪ Manual Local Caching for Documents. The Manual mode permits the user to specify individual files required for offline access. This operation guarantees a user can obtain access to the specified files whether online or offline.

Field/Item	Description
	<ul style="list-style-type: none"> <li data-bbox="678 254 1414 527">▪ Automatic Local Caching for Documents. The Automatic mode is applied for all non-executable files on the entire share. When a user accesses any non-executable file in this share, it is made available to the user for offline access. This operation does not guarantee that a user can access all the non-executable files, because only those files that have been used at least once are cached. Automatic can also be defined for programs. <li data-bbox="678 541 1414 814">▪ Automatic Local Caching for Programs. The Automatic mode is applied for all executable files on the entire share. When a user accesses any executable file in this share, it is made available to the user for offline access. This operation does not guarantee that a user can access all the executable files, because only those executable files that have been used at least once are cached. Automatic can also be defined for documents. <li data-bbox="678 829 1414 892">▪ Local Caching Disabled. No caching of files or folders occurs. <div data-bbox="695 919 1393 1108" style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;">  Note: The server supports Offline Files Access. This allows Windows 2000 and later clients to cache files that are commonly used from a network/file share. To use Offline Files, the client computer must be running Windows 2000 or later. </div>
Transfer to Object Replication Target	<p data-bbox="678 1136 1414 1268">When a file system is recovered from a snapshot, one of the final steps is to import the CIFS shares found in the snapshot representing the selected version of the file system. Only those CIFS shares marked as transferable will be imported.</p> <p data-bbox="678 1283 1170 1318">Use the list to specify one of the following:</p> <ul style="list-style-type: none"> <li data-bbox="678 1333 1414 1396">▪ Enable: CIFS shares will be transferred to recovered file systems. <li data-bbox="678 1411 1414 1474">▪ Disable: CIFS shares will not be transferred to recovered file systems. <li data-bbox="678 1488 1414 1621">▪ Use FS default (the default): When the target file system is brought online, CIFS shares will be transferred if Transfer Access Points During Object Replication is enabled for the file system.
Access Configuration	<p data-bbox="678 1669 1414 1759">IP addresses of the clients who can access the share (up to 5,957 characters allowed in this field). Refer to IP Address Configuration** at the end of this table.</p>
Follow Symbolic Links	<p data-bbox="678 1795 1235 1822">Enables the following of symlinks for this share.</p>

Field/Item	Description
	 Note: As of release 12.2 of the NAS Platform, clients using SMB2 or later are able to follow relative symlinks to files on storage without the use of global symlinks, if smb2-client-side-symlink-handling is enabled.
Follow Global Symbolic Links	Enables CIFS clients to follow global (absolute) symlinks via the Microsoft DFS mechanism for this share.
Force Filename to be Lowercase	Forces all filenames generated on this share to be lowercase. This is useful for interoperability of UNIX applications.
Enable ABE	<p>By default, ABE is disabled for shares and on the server/cluster as a whole. Before enabling ABE for a share, you must make sure ABE is enabled for the server/cluster as a whole (the CLI command to enable ABE support is <code>fsm set disable-ABE-support false</code>).</p> <p>When enabled, ABE filters the contents of a CIFS share so that only the files and directories to which a user has read access rights are visible to the user.</p>  Note: Enabling ABE can impact CIFS performance.
Enable Virus Scanning	<p>If virus scanning is enabled and configured for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is enabled by default. If virus scanning is not enabled for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is not enabled by default, but you can enable it a per-EVS basis.</p>  Note: Virus scanning is set up on a per-EVS basis, or for all EVSs using the global configuration context, but cannot be set up on a per-server or per-cluster basis.
Ensure Share Continuously Available	<p>Enables persistent file handles and transparent failover on the CIFS share. When enabled, Windows-based clients can continuously access the CIFS share if a network failure or a cluster node failure occurs. For example, if one cluster node fails, the client transparently migrates to another cluster node without any interruption to the client applications.</p> <p>This SMB3 option is available only in a clustered environment of more than one cluster node, and is disabled by default.</p>

Field/Item	Description
	 Note: Continuous Availability can impact CIFS performance and should only be enabled where it is required, such as with Microsoft Hyper-V or Microsoft SQL Server. When this feature is in use, it is also recommended that the Administrator disables DDNS on the server. If the file system is an object replication target, continuous availability is not effective until the file system is promoted.
Share Permissions	
Share Permissions List	By default, when a CIFS share is created, the group Everyone is added to the share permissions list.
User Home Directory Mapping	
Mode	<p>Used to specify how per-user home directories are created. The following options are available:</p> <ul style="list-style-type: none"> ▪ Off. Disable the home directory feature for this share. Do not automatically create home directories on this share for users. This is the default. ▪ ADS. Create the user home directories based on the home directory information supplied by the Active Directory server for each user. If you select ADS, do not specify a Path. ▪ User. Create the user's home directory by converting the user's Windows user name to lower case. (The user's Window's domain name, if any, is ignored.) For example, a user DOMAIN\John Smith would result in a home directory of john_smith. ▪ DomainAndUser. Create the user's home directory by creating a directory named for the user's Windows domain name, then converting the user's Windows user name to lower case and creating a sub-directory by that name. For example, a user DOMAIN\John Smith would result in a home directory of domain\john_smith. ▪ Unix. Create the user's home directory by converting the user's UNIX user name to lower case.
Path	Per-user home directories will be created in the specified Path , relative to the share root, which is specified without a leading \. If this field is left blank, user home directories will be created directly in the share root.

Field/Item	Description
	<p>By default, only one share per file system can be configured with home directories. The <code>cifs-home-directory</code> command can be used to relax this restriction, in which case great care must be taken not to configure conflicting home directories.</p> <p>For example, a share with a path of <code>\home1</code> and a share with a path of <code>\home2</code> would not cause a conflict, whatever home directory paths were configured. However, a share with a path of <code>\</code> and a default home directory path would conflict with a share with a path of <code>\dir</code> and a default home directory path.</p>

3. Modify the fields on this page as necessary.
4. Click **OK**.

Controlling access to shares using qualifiers

To specify which clients have access to an SMB share, qualifiers can be appended to the IP addresses:

Qualifier	Description
<code>read_write,</code> <code>readwrite,</code> <code>rw</code>	Grants read/write access. This is the default setting.
<code>read_only,</code> <code>readonly, ro</code>	Grants the specified client read-only access to the SMB share.
<code>no_access,</code> <code>noaccess</code>	Denies the specified client access to the SMB share.

Some SMB share qualifier examples are:

- `10.1.2.38 (ro)`
Grants read-only access to the client with an IP address of 10.1.2.38.
- `10.1.2.0/24 (ro)`
Grants read-only access to all clients whose IP address is within the range 10.1.2.0 to 10.1.2.255.
- `10.1.*.* (readonly)`
Grants read-only access to all clients with an IP address beginning with 10.1.

The order in which the entries are specified is important. For example,

```
* (ro)
10.1.2.38 (noaccess)
```

in which the first line grants read-only access to all clients, and the second denies access to the specified client. However, the second line is redundant, as the first line matches all clients. These lines must be transposed to ensure access is denied to 10.1.2.38

Controlling access to shares using permissions

Access to shares is restricted through a combination of share-level and file-level permissions. These permissions determine the extent to which users can view and modify the contents of the shared directory. When users request access to a share, their share-level permissions are checked first; if authorized to access the share, their file-level permissions are checked.

When the share-level permissions differ from the file-level permissions, the more restrictive permissions apply, as described in the following table, where [a] = “allowed” and [d] = “denied”:



Note: One of the features of SMB is the ability to assign rights to machine (computer) accounts. A machine account is generated automatically by the operating system and registered in Active Directory. It can be used for authentication within a domain. A machine account authentication can be only done by an application which has built-in support. For example, Hyper-V server allows storing virtual machines on remote shares. Such shares should allow full access for the machine account of a computer running Hyper-V server.

Activity	Read	Change	Full
View the names of files and subdirectories	a	a	a
Change to subdirectories of the shared directory	a	a	a
View data in files	a	a	a
Run applications	a	a	a
Add files and subdirectories	d	a	a
Change data in files	d	a	a
Delete files and subdirectories	d	a	a
Change permissions on files or subdirectories	d	d	a
Take ownership of files or subdirectories	d	d	a

When configuring access to a share, it is only possible to add users or groups that are:

- Known to domain controllers, and
- Seen by the server on the network.



Note: When a user is given access to a share, if the user has also a member of a group with a different access level, the more permissive level applies. For example, if a user is given *Read* access to a share, and that user also belongs to a group that has *Change* access to that same share, the user will have *Change* access to the share, because *Change* access is more permissive than *Read* access.

Adding or changing SMB share access permissions

You can add or modify SMB (CIFS) share access permissions in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > CIFS Shares** to display the **CIFS Shares** page.
2. Select the check box next to the share to modify, and then click **details**.
3. In the **Share Permissions** area of the **CIFS Share Details** page, click **change**.

The following table describes the fields on this page:

Field/Item	Description
New User/Group	Name for the new user or group.
Type	Displays a standardized identifier for the security group to which the user/ group being added belongs. The value is set automatically, based on the well known security identifier for the user/group being added.
modify	Saves any modifications made to the user or group settings.

4. To add a new user or group, follow these steps. To change permissions for an existing user or group, see step the next step.
 - a. Enter the name for the new user or group in the **New User/Group** field, and then click **add**.
 - b. Select the new user/group from the list.
 - c. Select the **Allow** or **Deny** check boxes to set the appropriate permissions. If the **Allow** check box is selected for full control, the user/group can perform all actions.
5. To change user/group permissions:
 - a. Select the user/group for which you want to change permissions.
 - b. Select the **Allow** or **Deny** check boxes to set the appropriate permissions. If the **Allow** check box is selected for full control, the user/group can perform all actions.
 - c. Save changes by clicking **modify**.

About Home Directories

The Home Directories feature simplifies the management of per-user home directories for larger environments:

- A per-user network directory is automatically generated when the user initiates an SMB connection to the EVS.
- If configured in the AD user profile, a Windows client will automatically map the drive letter from %HOMEDRIVE% to the network share %HOMESHARE% as a user logs in.
- These variables can be set automatically from Active Directory, or by a user login script.

Windows OS can be configured to automatically attach a remote CIFS share as a user's home directory when the user logs on. To do this, two environment variables are configured:

- %HOMEDRIVE% contains the drive letter to be used for the mapped drive.
- %HOMESHARE% contains the remote CIFS share to map

Using home directories with cluster EVS name spaces

The home directories feature is compatible with name spaces. However, note that home directories are not supported in a virtual file system.

Creating user home directories in a name space is considered a *lazy* process. When you first connect to the share in the name space, no home directory is created. If the user then browses or changes directory to the link from the name space to the regular file system, the server uses an SMB DFS referral to redirect them to a hidden share on the regular file system. When the DFS referral completes, and the user connects to the regular file system, their home directory is created.

Name space and file system layout example:

```
cns:
    \cnsdir
        \link --> Span0FS:
            \homes
```

When the user connects to the `cns` name space, no home directory is created. However, if the user later moves into `cns:\cnsdir\link`, their home directory is created, as that is the transition into a regular file system.

Offline file access modes

The server supports offline files access. This allows network clients to cache files that are commonly used from a network or file share. To use offline files, the client computer must be running Windows 2000 (or later). There are three different share caching modes (supporting all three modes of caching):

- **No Caching:** No caching of files or folders occurs.
- **Manual:** Allows user specification of individual files required for offline access. This operation guarantees a user can obtain access to the specified files whether online or offline.
- **Automatic:** Applies to the entire share. When a user accesses any file in this share, that file becomes available to that user for offline access. This operation does not guarantee a user can obtain access to the specified files, because only files that have been used at least once are cached. The Automatic mode can be defined for documents or programs.

Backing up and restoring SMB shares

When backing up and restoring SMB Shares:

- All SMB Shares in all EVSs are backed up (except those in the CNS tree).
- A SMB Share backup file is saved as a `.txt` file. The backup file contains the file system name and the share name, as well as most of the information about the share, including the settings for: *Ensure Path Exists*, *Show Snapshots*, *Follow Symbolic Links*, *Force Filename to Lowercase*, *Virus Scanning*, *Cache Options*, and *Max Users*.


When you restore SMB Shares from a backup file:

- The restore operation does not modify or overwrite currently existing shares that have the same name.
- With the exception noted above, all shares in the selected backup file are restored.

You can back up and restore SMB (CIFS) shares in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > CIFS Shares** to display the **CIFS Shares** page.
2. Click **Backup & Recovery** to display the **CIFS Shares Backup & Restore** page.



3. To back up: Click **backup**. In the browser, specify the name and location of the backup file, and then click **OK** or **Save** (the buttons displayed and the method you use to save the backup file depend on the browser you use).

A backup file name is suggested, but you can customize it. The suggested file name uses the syntax:

CIFS_SHARES_date_time.txt, where the following example illustrates the appropriate syntax: *CIFS_SHARES_Aug_4_2006_11_09_22_AM.txt*

4. To restore: Click **restore**. In the browser, the backup text file (*CIFS_SHARES_date_time.txt*) for the specific share(s) you want to restore, and then click **Open**. When the **CIFS Export Backup & Restore** page displays the name and location of the selected file, click **Restore**.

Considerations when using Hyper-V

In general, when using Hyper-V with an HNAS server:

- Store the Hyper-V virtual machine configuration locally on the Hyper-V server or server cluster. Do not store the Hyper-V virtual machine configuration on an HNAS CIFS/SMB3 share. Virtual hard disks (VHD and VHDX) can be stored on HNAS CIFS/SMB3 shares.



Note: Ensure that the virtual machine configuration remains consistent with the virtual hard disks. If an HNAS file system is reverted to an earlier snapshot, the Hyper-V virtual machine should be reverted to the most recent checkpoint in that snapshot, and the newer checkpoints deleted.

- Always make sure that the Hyper-V server is backed up.
- Enable SMB3 when using Hyper-V virtual machines with their virtual disks stored on SMB3 shares. Using SMB3 provides performance benefits and an improved experience. To set the protocol version, use the `cifs-max-supported-version` command.
- Enable Continuous Availability on SMB3 shares to ensure service continuity if an HNAS failover occurs.
- Do not scan virtual hard disk files for viruses. It is unlikely that a scan will detect any viruses on the virtual hard disks and it will impact virtual machine performance. Add VHD and VHDX extensions to the virus scanner exclusion list if required. For antivirus protection, run an antivirus product inside the virtual machine.
- Avoid using large fixed-sized virtual hard disks on SMB3 shares. Large fixed-sized disks can take a long time to create, so use smaller fixed-sized virtual disks or dynamically expanding virtual disks if large disk sizes are required.



Note: An HNAS server meets all requirements of Microsoft Scale-Out File Server (SOFS). Implementing a scale-out file server backed by an HNAS SMB3 share (using virtual machines with a shared virtual hard disk stored on an HNAS SMB3 share) is not supported. HNAS SMB3 shares are already equivalent to those provided by SOFS.

Configuring the Service Witness Protocol

For SMB3 transparent failover to perform efficiently, the Service Witness Protocol must be configured.

The Service Witness Protocol enables a registered client to receive notification of any state changes on a continuously available server. For example, if a continuously available SMB server provides SMB shares to a Hyper-V server, the protocol enables the Hyper-V server to receive notification of any state changes without needing to wait for the connection to time out. This ensures that there is a fast notification and recovery time from an unplanned failure, such as a network loss.

When a client connects to an SMB share, the client has to identify and then register with the witness EVS to receive witness event notifications. If the client cannot identify or register with the witness EVS, the NAS server cannot provide witness notifications to that client and standard timeouts will apply.

Configuring a witness EVS

To use the Service Witness Protocol, you need to create and configure a witness EVS to monitor a service EVS.

Before you begin

- A valid CIFS license is required. Review your licensed EVS count before making changes.
- The server to which the Witness protocol applies must be in a clustered environment of more than one cluster node.
- The service EVS must be running version 3 of the CIFS protocol. To set the version, use the `cifs-max-supported-version 3` command.
- Windows clients must be joined to the domain.



Note: You can use only CLI commands to configure and manage a witness EVS. For more information about the CLI commands, and for more CLI options for the Service Witness Protocol, see the *Command Line Reference*.

Procedure

1. Set up a share on the service EVS that you want the witness EVS to monitor.
2. Create a witness EVS for the service EVS that contains the share by using the following command:

```
evs create [-l <label>] -i [<ipaddr/prefix> | <ipaddr> -m <mask>]
-p <port> [-n <dst-nodeid>] [-w <witness-for>]
```

For example, a witness EVS on cluster node 2 for a service EVS on cluster node 1 would be:

```
evs create -l WITNESSEVS01 -i 192.0.2.1/24 -p ag1 n 2 -w EVS01
```

The witness EVS is created and bound to the service EVS.

3. Configure an ADS CIFS name for the witness EVS:
 - a. Put the witness EVS in context.

- b. Add an ADS CIFS name by using the following command:


```
cifs-name add -m ads -a <dc ipaddr> <my-ads-name>
```
 - c. Enter your user name and password.
4. Map the continuously available share in Windows 8, 8.1, Server 2012, or Server 2012 R2.



Note: When mapping the share, use a fully qualified domain name (FQDN) instead of the IP address of the service node.

5. To verify that the client is registered to the witness EVS, use the following command:


```
witness-registration-list-show
```

Using Windows server management

The Computer Management MMC tool, available for Windows 2000 or later, can perform share management tasks from any remote computer; for example:

- Viewing a list of all users currently connected to the system.
- Creating shares.
- Listing all shares on the system and the users connected to them.
- Disconnecting one or all of the users connected to the system or to a specific share.
- Closing one or all of the shared resources that are currently open.
- Viewing an event log.



Note: For older versions of Windows, the equivalent of this tool is provided by Server Manager.

Using the computer management tool



Note: The appearance of the screens depends on the operating system version.

To use the Computer Management tool:

Procedure

1. In the Windows interface, from Administrative Services, select **Computer Management**; then right-click on **Computer Management (Local)** to display a context menu, and select **Connect to another computer**:
2. Optionally, select the domain from the drop-down **Look in** field, then highlight a name or an IP address to use for file services on the server, and click **OK**.
Do not specify a server administration name or IP address for this purpose.
3. Click **Event Viewer** to display the server's event log:
4. On the event log window:
 - a. Click **Shares** to list all of the shares. Some or all of the users can be disconnected from specific shares.

- b. Click **Sessions** to list all users currently connected to the system. Some or all of the users can be disconnected.
- c. Click **Open Files** to list all the open shared resources. Some or all of the shared resources can be closed.

Restoring a previous version of a file

SMB clients can access previous versions of files stored on shares, as long as a snapshot containing the file/directory exists. A tab labeled Previous Versions is displayed on the **Properties** page for files and folders for which previous versions are available from snapshots. The tab displays the list of the available previous versions and their corresponding times.

The tab provides a means to access directly the previous versions listed, or to restore from them.

Chapter 6: Transferring files with FTP

This section explains how to set up File Transfer Protocol (FTP) so that users with FTP clients can access files and directories on the storage server.

FTP protocol support

The NAS server implements the file-serving functions of an FTP server. The server provides the file-serving functions required for:

- File manipulation
- Directory manipulation
- File access control (for example, permissions)

It is also possible for multiple clients to copy or read a file while it is being streamed over FTP to the server, using any protocol (FTP, SMB, or NFS). However, the file cannot be written/modified or locked. If a client copies a file while it is being written, the server only provides the data written up to that point.

Prior to allowing FTP access to the system, the FTP service must be enabled. No license key is required for this protocol.

FTP statistics for the storage server (in 10-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.



Note: The NAS server also supports FXP (inter-server data transfer).



Note: Although the NAS server supports both NTLM1 and NTLM2, it only supports the use of FTP with NTLM1. NTLM2 in FTP is not supported.

Configuring FTP preferences

As part of the process of setting up FTP, choose a service for authenticating the passwords of the FTP users. Also, a timeout must be set with which to end FTP sessions that have been inactive.

Procedure

1. Navigate to **Home > File Services > FTP Configuration** to display the **FTP Configuration** page.

File Services [Home](#) > [File Services](#) > FTP Configuration

FTP Configuration

Password Authentication Services

NT:

NIS:

Session Timeout

Timeout: minutes.

Anonymous User Permissions

ReadOnly:

Field/Item	Description
Password Authentication Services	<p>The configured security mode determines what password authentication service methods can be used.</p> <ul style="list-style-type: none"> ■ NT <p>If selected, FTP users must log in with an NT domain user name and password, which is authenticated via a domain controller.</p> ■ NIS <p>If selected, FTP users must log in with a UNIX user name and password, which is authenticated via a NIS server in the configured NIS domain.</p> <p>If both services are enabled, the FTP user will be authenticated against the configured NT domain first. If authentication fails, the server will attempt to authenticate the user against the configured NIS domain.</p> <p>If both NT and NIS are not selected, then authentication will not be performed using these methods and only the anonymous user will be able to log in.</p>
Anonymous User Permissions	<p>Specifies whether read-write is allowed for anonymous requests. Fill the ReadOnly check box to limit anonymous requests to read only.</p>
Session Timeout	<p>The number of minutes of inactivity after which to end an FTP session automatically (the Timeout). The valid range is from 15 to 14,400 minutes.</p>
apply	<p>Applies the configuration changes without closing the page.</p>

2. In the Password Authentication Services area, fill the check box for **NT** or **NIS**.
 If operating in UNIX or Mixed security mode, both NT and NIS password authentication are supported. If both services are enabled, the FTP user will be authenticated against the configured NT domain first. If authentication fails, the server will attempt to authenticate the user against the configured NIS domain.

3. Enter the Session Timeout value.
The valid range is between 15 and 35,000 minutes (35,000 minutes = 24 days).
4. Specify whether read-write is allowed for anonymous requests. Fill the **ReadOnly** check box to limit anonymous requests to read only.
5. Click **apply**.

Displaying FTP users

FTP users can be manually set up or display their details can be imported from a file.

Procedure

1. Navigate to **Home > File Services > FTP Users** to display the **FTP Users** page.

The screenshot shows the 'FTP Users' management interface. At the top, there is a breadcrumb trail: *File Services* > [Home](#) > *File Services* > *FTP Users*. The main heading is **FTP Users**. Below the heading, there are two main sections: 'EVS / File System Label' and 'Filter'. The 'EVS / File System Label' section shows 'g1-avs3 / PHDS1' with a 'change...' button. The 'Filter' section has 'Name:' and 'Path:' input fields and a 'filter' button. Below these sections is a table with columns: **Name**, **File System**, and **Path**. The table is currently empty, with 'Check All' and 'Clear All' links. At the bottom, there is an 'Actions' section with 'add', 'delete' buttons, and an 'Import Users' link.

The following table describes the fields on this page:

Field/Item	Description
EVS / File System Label	This field displays the EVS and File System where the FTP users listed on the page have been configured.
change	Click the change button to select a different file system.
Filter	The filter button allows you to filter the users based on user Name or Path.
Name	This column displays the existing FTP users. Up to 5000 users can be listed, but the FTP user list displays a maximum of 20 users per page. Use filter to control the display of users.
File System	Shows the file system containing the FTP user's initial directory path.
Path	Shows the path to the initial directory for the user after logging in over FTP.
details	Opens the FTP User Details page, allowing you to modify certain details about the selected user.
add	Opens the Add FTP User page, allowing you to set up a new user.
delete	Deletes the selected user.
Import Users	Opens the Import FTP Users page, allowing you to set up new users by importing them from a file.

Adding an FTP user

You can add an FTP user in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > FTP Users** to display the **FTP Users** page.
2. Click **add** to display the **Add User** page.

File Services [Home](#) > [File Services](#) > [FTP Users](#) > Add FTP User

Add FTP User


EVS / File System: g1-eva3 / PHDS1

User Name:

Initial Directory for the user:

Path Options
 These options only apply when 'path' or 'file system' values are changed.
 Create path if it does not exist. (See online help for security implications).

The following table describes the fields on this page:

Field/Item	Description
EVS / File System	Displays the selected file system.
change	Enables you to select a different EVS / File System.
User Name	The name with which the user is to log in. To allow anonymous logins to the mount point, specify the user name as <code>anonymous</code> or <code>ftp</code> .
Initial Directory for the user	The directory in which the user starts when logging in over FTP. Click browse to navigate to and insert the path. <div style="background-color: #e0f2f1; padding: 5px;"> <p> Note: Automatically created directories are owned by the root user and group (UID:0 / GID:0) and are accessible to all users (that is, the permissions are set to <code>rwxrwxrwx</code>). It is recommended that these directories are created through CIFS or NFS, or that they are given the desired permissions explicitly after being created using this option.</p> </div>
Path Options	The Create path if it does not exist option creates the path automatically when it does not already exist. If the file system is mounted read-only, for example it is an object replication target, it is not possible to create a new directory. Select a path to an existing directory.

- Enter the user name. To allow anonymous logins to the mount point, specify the user name as `anonymous` or `ftp`.
The password authentication service that you use determines whether users must log in with their NT domain name or UNIX user name.
- In the Initial Directory for the user field, type the path to the directory in which the user starts when he or she logs in over FTP.



Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to `rwxxrwxrwx`). It is recommended that such directories are created using SMB or NFS, or that such directories are given the desired permissions explicitly after being created with this option.

- To create the path automatically when it does not already exist, select the **Create path if it does not exist** check box.
- Click **OK**.

Importing an FTP user

You can import an FTP user in the NAS Manager.

Procedure

- Navigate to **Home > File Services > FTP Users**.
- Click **Import FTP Users** to display the **Import FTP Users** page.

The following table describes the fields on this page:

Field/Item	Description
Filename	The name of the file to import. Use the Choose File button to select the file.
Import	Click to import the file.

- In the **Filename** field, enter the file name that contains the user details, or click **Browse** to search for the file name.

The user details in the import file have the following syntax:

```
user_name file_system initial_directory
```

Each entry must be separated by at least one space. If either the *user_name* or *initial_directory* contains spaces, the entry must be within double-quotes. For example:

```
carla Sales /Sales/Documents
miles Sales "/Sales/Sales Presentations"
john Marketing /Marketing
```

If you cannot be certain that the initial directory exists, you can create it automatically by specifying the option `ENSURE_PATH_EXISTS` on a separate line in the file. For example:

```
ENSURE_PATH_EXISTS true
carla Sales /Sales/Documents
miles Sales "/Sales/Sales Presentations"
ENSURE_PATH_EXISTS false
john Marketing /Marketing
```

In the first instance of the `ENSURE_PATH_EXISTS` option, the `true` attribute turns on the option, and it applies to the two following entries until the option is turned off by the second instance of the option, with the attribute `false`. The default for the `ENSURE_PATH_EXISTS` option is `true` so that the initial directory is automatically created.

To insert a comment in the file, precede it with a hash character (`#`).



Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to `rwxrwxrwx`). We recommend that such directories are created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created via this option.

4. Click **Import**.


Modifying FTP users

You can modify FTP users in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > FTP Users**.
2. Fill the check box next to the user to display or modify, and click **details**.

The following table describes the fields on this page:

Field/Item	Description
File System	Displays the file system which owns the FTP user.
Change	Click change to select a different file system.
User Name	Displays the FTP user name.
Initial Directory for the user	<p>This directory is the location where the user starts after logging in over FTP. You can change the directory by typing the path to the new directory. You can click the browse button to find the required directory. This directory is the location where the user starts after logging in over FTP.</p> <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <p> Note: Automatically created directories are owned by the root user and group (UID:0 / GID:0) and are accessible to all users (that is, the permissions are set to rwxrwxrwx). It is recommended that such directories are created using CIFS or NFS, or that such directories are given the desired permissions explicitly after being created by this option.</p> </div>
Path Options	Click the Create path if it does not exist check box to create the path automatically when it does not already exist. If the file system is mounted read-only, for example it is an object replication target, it is not possible to create a new directory. Select a path to an existing directory.

3. Modify settings as needed:

- In the File System field, you can click **change** to select a different file system.
- In the Initial Directory for the user field, you can change the directory by typing the path to the new directory. You can click the **browse** button to find the desired directory. This directory is the location where the user starts after logging in over FTP.
- In the **Path Options** box, you can fill the check box **Create path if it does not exist** to create the path automatically when it does not already exist.

4. Click **OK**.

FTP statistics

FTP statistics for the storage server (in 10-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Configuring FTP audit logging

FTP generates an audit log to keep track of user activity. The system will record the event when each time a user takes any of the following actions:

- Logging in or out
- Renaming or deleting a file
- Retrieving, appending or storing a file
- Creating or removing a directory

The system also records when a session timeout occurs.

Each log file is a tab-delimited text file containing one line per FTP event. Besides logging the date and time at which an event occurs, the system logs the user name and IP address of the client and a description of the executed command. The newest log file is called `ftp.log`, and the older files are called `ftpn.log` (the larger the value of n , the older the file).

Chapter 7: Block-level access through iSCSI

The storage server supports iSCSI, the Internet Small Computer System Interface (iSCSI) protocol enables block-level data transfer between requesting applications and iSCSI target devices. Using Microsoft's iSCSI Software Initiator (version 1.06 or later), Windows servers can view iSCSI targets as locally attached hard disks. Windows can create file systems on iSCSI targets, reading and writing data as if it were on a local disk. Windows server applications, such as Microsoft Exchange and Microsoft SQL Server can operate using iSCSI targets as data repositories.


The server iSCSI implementation has attained the **Designed for Windows Server™ 2003** certification from Microsoft. The **Designed for Windows Server™ 2003** logo helps customers identify products that deliver a high quality computing experience with the Microsoft Windows Server 2003 operating system.

iSCSI support

To use iSCSI storage on the server, one or more iSCSI LUs (LUs) must be defined. iSCSI LUs are blocks of SCSI storage that are accessed through iSCSI targets. iSCSI targets can be found through an iSNS database or through a target portal. After an iSCSI target has been found, an Initiator running on a Windows server can access the LU as a "local disk" through its target. Security mechanisms can be used to prevent unauthorized access to iSCSI targets.

On the server, iSCSI LUs are just regular files residing on a file system. As a result, iSCSI benefits from file system management functions provided by the server, such as NVRAM logging, snapshots, and quotas.

The contents of the iSCSI LUs are managed on the Windows server. Where the server views the LUs as files containing raw data, Windows views each iSCSI target as a logical disk, and manages it as a file system volume (typically using NTFS). As a result, individual files inside of the iSCSI LUs can only be accessed from the Windows server. Server services, such as snapshots, only operate on entire NTFS volumes and not on individual files.



iSCSI MPIO

iSCSI MPIO (Multi-path Input/Output) uses redundant paths to create logical “paths” between the client and iSCSI storage. In the event that one or more of these components fails, causing the path to fail, multi-pathing logic uses an alternate path so that applications can still access their data.

For example, clients with more than one Ethernet connection can use logical paths to establish a multi-path connection to an iSCSI target on the server. Redundant paths mean that iSCSI sessions can continue uninterrupted in the event of the failure of a particular path. An iSCSI MPIO connection can also be used to load-balance communication to boost performance.

If you intend to use an offload engine, make sure it is compatible with Microsoft multi-path and load-balancing.

iSCSI MPIO is supported by Microsoft iSCSI Initiator 2.0.

iSCSI access statistics

Statistics are available to monitor iSCSI activity since the server was last started or its statistics were reset. The statistics are updated every 10 seconds.

iSCSI prerequisites

To enable iSCSI capability:

- Enter an iSCSI license key.
- Enable the iSCSI service.

Supported iSCSI initiators

The server currently supports the following iSCSI initiators:

- Microsoft iSCSI Initiator version 1.06 (or later).
- Microsoft iSCSI Initiator version 2.05 (provides MPIO support).
- Linux iSCSI initiator versions 3.4.2, 3.6.2, 3.6.3, and 4.0.188.13 (available from the Linux iSCSI project on SourceForge).
- Solaris 10 U2 (64 bit) native initiator, iscsiadm v1.0.
- Macintosh (OS X 10.4 Tiger) ATTO Xtend SAN v3.10.
- Open iSCSI version 2.0.865.



Note: Other iSCSI initiators and/or versions of the initiators listed above may also work with the server, but have not been tested. Check with your Hitachi representative for the latest list of supported iSCSI initiators.

Offload engines

The server currently supports the use of the Alacritech SES1001T and SES1001F offload engines when used with the Microsoft iSCSI initiator version 1.06 or later. Check with your Hitachi representative for the latest list of supported offload engines.

Configuring iSCSI

In order to configure iSCSI on the server, the following information must be specified:

- iSNS servers
- iSCSI LUs
- iSCSI targets (including iSCSI domain)
- iSCSI initiators (if using mutual authentication)

Configuring iSNS

The Internet Storage Name Service (iSNS) is a network database of iSCSI initiators and targets. If configured, the server can add its list of targets to iSNS, which allows Initiators to easily find them on the network.

The iSNS server list can be managed through the **iSNS Servers** NAS Manager page. The server registers its iSCSI targets with iSNS database when any of the following events occurs:


- A first iSNS server is added.
- An iSCSI target is added or deleted.
- The iSCSI service is started.
- The iSCSI domain is changed.
- A server IP address is added or removed.

Viewing iSNS servers

You can view the list of configured iSNS servers in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSNS Servers** to display the **iSNS Servers** page.



The following table describes the fields on this page:

Field/Item	Description
EVS	Selector for EVS and File System. To switch to a different EVS/File System, click change .
change	Click change to switch to a different EVS/Server.
IP Address	Lists the IP Address of the iSNS server.
Port	The port number the NAS server uses to communicate with the iSNS server.
Check All	Selects all iSNS servers.
Clear All	Deselects all iSNS servers.
add	Click this button to add an iSNS server to the list.
delete	Saves configuration changes, and closes the page.
iSCSI Targets	Opens the iSCSI Targets window.

Configuring iSCSI Logical Units

An iSCSI Logical Unit (LU) is a block of storage that can be accessed by iSCSI initiators as a locally attached hard disk. An LU is stored as a file on the server file system. Like any other data set on the file system, iSCSI LUs can be bound in size using the server's size management tools, including virtual volumes and quotas. LUs are created with a specific initial size but can be expanded over time, as demand requires.

After an LU has been created and the iSCSI domain name has been set, an iSCSI Target must be created to allow access to the LU. A maximum of 32 LUs can be configured for each iSCSI Target.

Logical unit management

An iSCSI LU is a file within one of the server's file systems. Such a file must have an `.iscsi` extension to identify it as an iSCSI LU. However, apart from this extension there is no other way to determine that a file does indeed represent an LU.



Note: customer support recommends that all iSCSI LUs are placed within a well-known directory, for example `/.iscsi/`. This provides a single repository for the LUs in a known location.

Logical unit security

As LUs are files, they can be accessed over other protocols, such as CIFS and NFS. This renders LUs vulnerable to malicious users who can modify, rename, delete or otherwise affect them.



Caution: customer support recommends setting sufficient security on either the LU file, the directory in which it resides, or both, to prevent unwanted accesses.

Concurrent access to logical units

The server's iSCSI implementation allows multiple initiators to connect to a single LU, which is necessary for applications and operating systems that support, or rely upon, concurrent file system access. However, concurrent access can be detrimental to a client machine when the client is unaware of other clients accessing the file system. For example:

- Simultaneous independent updates to the same files. Scenario: Two independent Microsoft Windows clients can connect to the same LU, containing an NTFS file system. Result: If allowed to simultaneously and independently modify data, metadata, and system files, conflicting disk updates will quickly corrupt the file system.
- Simultaneous access to separate partitions. Scenario: An LU contains two distinct NTFS partitions, with one Microsoft Windows client connected only to the first partition, and another connected only to the second partition. Result: Because a Microsoft iSCSI client will attempt to mount each partition it encounters on the LU, a Microsoft Windows client mounting an NTFS partition updates system files on all partitions; therefore, even though the two clients are accessing separate partitions within the LU, both will update system files on both partitions, causing conflicting system file updates, causing one or both of the clients to fail.

Taking snapshots of logical units

The contents of an iSCSI LU are controlled entirely by the client accessing it. The server cannot interpret the file systems or other data contained within an LU in any way. Therefore, the server has no knowledge of whether the data held within an iSCSI LU is in a consistent state. This introduces a potential problem when taking a snapshot of an LU.

For example, when a client creates a file, it must also insert the file name to the host directory. This means that more than one write is required to complete the operation. If the server takes a snapshot after the file object has been created, but before its name has been inserted into the directory, the file system contained within the snapshot will be inconsistent. If another client were to view the snapshot copy of the file system, it would see a file object without a name in a directory. This example provides only one possible scenario for snapshot inconsistency.



Caution: customer support recommends that prior to taking a snapshot of an iSCSI LU, all applications should be brought into a known state. A database, for example, should be quiesced. Disconnecting the iSCSI initiators from the LUs undergoing snapshot is also recommended. This guarantees that all pending writes are sent to the LU before the snapshot is taken.

Volume full conditions

Unexpected volume full conditions can occur with iSCSI LUs, as illustrated by the following two examples:

- **Directly Attached Disks.** When a client uses a directly attached disk, it can monitor the amount of available free space. If a partition contains no free space, the client can return a Volume Full condition. In this way, the client can ensure against file system corruption due to running out of disk space part way through an operation.
- **iSCSI LU.** By way of background, on iSCSI LUs with snapshots enabled, old data is preserved, not overwritten. Therefore, overwriting an area of an LU causes the server to allocate extra disk space, while using no extra disk space within the client's partition, causing a Volume Full condition to occur, even when partitions within the LU contain free space. Under this scenario, a client may receive a Volume Full condition part-way through an operation, causing file system corruption. Although this corruption should be fixable, this situation should be avoided.



Note: customer support recommends allocating sufficient disk space on the server to contain all iSCSI LUs and snapshots, as well as careful monitoring of free disk space.

Managing iSCSI logical units

This section describes how to manage iSCSI logical units.

Viewing the properties of iSCSI logical units

You can view the iSCSI logical units in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Logical Units** to display the **iSCSI Logical Units** page.

File Services [Home](#) > [File Services](#) > iSCSI Logical Units

iSCSI Logical Units

EVS / File System Label	Filter
g1-eva3 / PHDS1 change...	Alias: <input style="width: 100%;" type="text"/> File System: <input style="width: 100%;" type="text"/> Path: <input style="width: 100%;" type="text"/> Used in Target: <input style="width: 100%;" type="text"/> <div style="text-align: right; margin-top: 5px;">filter</div>

Alias	File System:Path	Size	Status	Used in Target	
<input type="checkbox"/> test	PHDS1/PHDS1.iscsi	1 MiB	Mounted		details



Check All
Clear All
Show 20 items per page

Actions: mount
unmount
delete
add

Shortcuts: [iSCSI Targets](#) [iSCSI Initiator Authentication](#)

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Selector for EVS and File System where LUs reside, or where LUs can be created. To switch to a different EVS/File System, click change .
change	Click change to switch to a different EVS/File System.
Filter	Use Filter to display a subset of the iSCSI Logical Units. The options are: <ul style="list-style-type: none"> ▪ Alias - name of the LU. ▪ File System - the file system for the LU.

Field/Item	Description
	<ul style="list-style-type: none"> ▪ Path - the path of the LU. ▪ Used in Target - the target of the LU.
Filter	Click Filter to display a subset of the iSCSI Logical Units.
Alias	Name of the LU.
File System:Path	The file system and path for the LU.  Note: LUs appear as regular files in server file systems.
Size	Size of the LU. This value cannot exceed the amount of available free space on the configured file system.
Status	Indicates of the LU status, usually whether the LU is mounted.  Note: The status will display Unmounted while an LU is being created asynchronously, and will then display Mounted once the creation has completed.
Used in Target	Displays the target.
mount	Select an unmounted LU and click mount to mount the LU.
unmount	Select a mounted LU and click unmount to unmount the LU.
delete	Deletes the selected iSCSI LUs.
add	Opens the Add iSCSI Logical Unit page where you can create a new iSCSI LU.
details	Opens the iSCSI Logical Unit Details page for the selected LU.
iSCSI Targets	Advances to the iSCSI Targets settings page.
iSCSI Initiator Authentication	Advances to the iSCSI Initiator Authentication settings page.

Adding iSCSI logical units

You can add iSCSi Logical Units in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Logical Units** and then click **add** to display the **Add iSCSI Logical Units** page:

File Services [Home](#) > [File Services](#) > [iSCSI Logical Units](#) > Add iSCSI Logical Unit

Add iSCSI Logical Unit

EVS / File System: g1-eva3 / PHDS1 change...

File System Free Capacity: 496.85 GiB

Alias:

Path to File: browse...

(Including name of file; the extension .iscsi may be appended)

File Already Exists

Create File

Size: GiB ▼

Create path to file if it does not exist

Comment:

OK cancel

Shortcuts: [iSCSI Targets](#)

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Displays the Virtual Server and file system on which to create the Logical Unit.
change	Click to open the Select a File System page where you can select a different file system.
Alias	The name of the logical unit. The maximum number of characters is 255.
Path to File	The path where the logical unit resides. browse can be used to assist in finding the desired path of a predefined Logical Unit.
File Already Exists	Choose this option if the file already exists.
Create File	Choose this option if the file does not exist.
Size	Size of the LU. This value cannot exceed the amount of available free space on the configured file system.
Create path to file if it does not exist	Creates the path specified in the Path to File field.
Comment	Allows you to provide descriptive information about the Logical Unit.
iSCSI Targets	Advances you to the iSCSI Targets page, which allows you to add, modify, and delete iSCSI Targets. iSCSI Targets appear to iSCSI initiators as different storage devices accessible over the network.

2. If necessary, change the EVS and/or file system.
The EVS name displayed indicates the EVS and file system to which the LU will be added. Click **change** to select a different EVS or file system.
3. Specify the Logical Unit alias.
In the **Alias** field, enter a name for the LU.
4. If the path to the file already exists, specify the path to the Logical Unit. When entering a path for an LU file that already exists, use the following procedure:
 - a. Choose the file.
Click **browse** to display a dialog that will allow you to select the file for the LU. Alternatively, you can enter the path name of the file (including the extension) and not use the browse button.
 - b. Select the **File Already Exists** radio button.
 - c. Optionally, add a comment.
Using the comment field, you can provide descriptive information about the LU.
 - d. Save the Logical Unit definition.
Click **OK** to add the LU.

5. If the path to the file does not already exist, specify to create the path to the Logical Unit. There are several steps to complete when creating a new file for use as an LU:
 - a. Choose the path for the file.
Click **browse** to display a dialog that will allow you to select the directory for the LU file. The name of the file as well as the directory need to be specified. The file does not need an extension; `iscsi` is appended automatically. Alternatively, you can enter the file name and path (including the extension) and not use the browse button.
 - b. Select the **Create File** radio button.
 - c. Specify the logical unit size.
Using the **Size** field and the drop-down list, specify the size of the LU file.
 - d. Fill the **Create path to file if it does not exist** check box.
 - e. Optionally, add a comment.
Using the **Comment** field, you can provide descriptive information about the LU.
 - f. Save the logical unit definition.
Click **OK** to add the LU.


Modifying an iSCSI logical unit

You can modify an iSCSI logical unit in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Logical Units**.
2. Select the check box next to the iSCSI logical unit to modify and then click **details**.

The following table describes the fields on this page:

Field/Item	Description
EVS/File System	Displays the EVS and file system hosting the LU.
File System Free Capacity	The amount of free space available in the file system.
Status	Indicates whether the LU is mounted or unmounted. It is possible to mount/unmount a LU while its underlying file system remains mounted. If the LU is not mounted, click mount to mount the LU. If the LU is mounted, click unmount to unmount the LU. Filling the Ensure the underlying file system exists when mounting check box will ensure that the underlying file system exists when the LU is mounted.
Alias	Name of the LU. You can change this name.
Path to File	The complete file system path to the LU file.  Note: LUs appear as regular files in server file systems.
Comment	Use to provide descriptive information about the LU.
Size	Size of the LU. This value cannot exceed the amount of available free space on the configured file system.
iSCSI Targets	Opens the iSCSI Targets page, in which you can add, modify, and delete iSCSI Targets.

3. Modify the fields as necessary.
4. Click OK.

Deleting an iSCSI logical unit

You can delete an iSCSI logical unit in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Logical Units** to display the **iSCSI Logical Units** page.
2. Select the check box next to the logical unit to delete and then click **delete**.
3. Click **OK** to confirm the deletion.

Backing up iSCSI logical units

Only a client connected to the LU through its Target can access and backup individual files and directories contained in the LU. LUs back up as normal files on a server file system.



Caution: If backing up the iSCSI LU from the server, ensure that the iSCSI initiators are disconnected, or make the backup from a snapshot.

To back up an iSCSI LU:

Procedure

1. Disconnect the iSCSI Initiator from the Target.
2. Unmount the iSCSI Logical Unit.

To unmount the iSCSI LU, you can use the following CLI command:

```
iscsi-lu unmount <name>
```

Where <name> is the name of the iSCSI LU.

3. Back up the Logical Unit to a snapshot or backup device.

For safety, you should either back up the iSCSI LU to a snapshot or to another backup device.

4. Mount the Logical Unit.

To mount the iSCSI LU, you can use the following CLI command:

```
iscsi-lu mount <name>
```

Where <name> is the name of the iSCSI LU.

5. Reconnect to the iSCSI Target using the iSCSI Initiator.
6. If necessary, rescan disks.

You may have to use Window's Computer Manager rescan disks to make the LU reappear to clients.

Restoring iSCSI logical units

To ensure consistency of data on an LU, it may be necessary to restore it from a snapshot or a backup. To restore an iSCSI LU, perform the following steps:

Procedure

1. Disconnect the iSCSI Initiator from the Target.
2. Unmount the iSCSI logical unit.
Use the following CLI command: `iscsi-lu unmount <name>`, where *name* is the name of the LU
3. Restore the logical unit from a snapshot or backup.
4. Mount the iSCSI logical unit.
Use the following CLI command: `iscsi-lu mount <name>`, where *name* is the name of the LU.

5. Reconnect to the Target using the iSCSI Initiator.
6. If necessary, rescan disks in Computer Management.

Setting up iSCSI targets


An iSCSI target is a storage element accessible to iSCSI initiators. These targets appear to iSCSI initiators as different storage devices accessible over the network. The server supports a maximum of 32 iSCSI Targets per EVS and a maximum of 32 iSCSI sessions per Target.

Viewing the properties of iSCSI targets

You can view the properties of iSCSI targets in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Targets** to display the **iSCSI Targets** page.



The following table describes the fields on this page:

Field/Item	Description
EVS	Select the EVS on which the target will be hosted.
change	Click to select a different EVS.
EVS iSCSI Domain	Displays the iSCSI domain, which is the DNS domain used when creating unique qualified names for iSCSI targets. To set an iSCSI domain, enter a domain name and click apply .
Alias	Identifies the name of the target.
Comment	Additional information related to the target.
Globally Unique Name	The target's name. The name is generated automatically by the server, and is unique across the globe.
details	Displays the details for the selected iSCSI target.
add	Advances to the Add iSCSI target page where you can add an iSCSI target.
delete	Deletes the selected iSCSI target.
iSCSI Logical Units	Advances to the iSCSI Logical Units settings page.
iSNS Servers	Advances to the iSNS Servers settings page.
iSCSI Initiator Authentication	Advances to the iSCSI Initiator Authentication settings page.

Adding iSCSI targets

You can add an iSCSI target in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Targets** to display the **iSCSI Targets** page.
2. Click **add** to display the **Add iSCSI Target** page.

File Services [Home](#) > [File Services](#) > [iSCSI Targets](#) > Add iSCSI Target

Add iSCSI Target

EVS: g1-eva3 change...

EVS iSCSI Domain: us_sj_114

Alias:

Comment:

Authentication

Enable Authentication

Secret:

Access Configuration

(Enter IP-based values first, if possible)

Logical Unit selection

Available

Logical Unit Number	Logical Unit Name
<input type="text"/>	test

▶

◀

Selected

Logical Unit Number	Logical Unit Name

OK
cancel

The following table describes the fields on this page:

Field/Item	Description
EVS	The EVS that will host the iSCSI target.
change	Click to select a different EVS.
iSCSI Domain	The DNS domain used when creating the Globally Unique Name of an iSCSI target.
Alias	The name of the iSCSI target. This can be a maximum of 255 characters long.
Comment	Additional descriptive information about the target.
Secret	The password used to secure the target from any unauthorized access. The initiator authenticates against this password when connecting to the target. The secret should be greater than or equal to 12 characters, but less than 17 characters, in length. Although the secret may be between 1-255 characters in length, some iSCSI initiators will refuse to connect if the secret contains less than 12 characters or more than 16 characters.
Enable Authentication	By default, the check box is filled. Filling or clearing the check box enables or disables authentication. When authentication is disabled, initiators are permitted to connect to the target and its logical units without needing to know the target's secret.
Access Configuration	Enter the desired access configuration parameters. Refer to the Access Configuration table for details on how to define the Access Configuration List.
Available Logical Units	The list of Logical Units available to assign an iSCSI target.
Selected Logical Units	The list of Logical Units added to the target.
Logical Unit Number	Enter a Logical Unit Number. The number can be any unique number between 0 and 255.

3. Specify the required information.
 - a. The iSCSI Domain, Alias, Available LUs, and Logical Unit Numbers are required, all other fields are optional.
 - b. Optionally, specify the Comment, Secret, and/or Access Configuration for the Target.

The following table provides syntax for the Access Configuration field:

What to type	Means
Blank or *	All clients can access the target.
Specific address or name. Examples: 10.168.20.2, client.dept.company.com To deny access to a specific host, use the no_access or noaccess qualifier. For example, 10.1.2.38 (no_access) will deny access to the host with the IP address 10.1.2.38.	Only clients with the specified names or addresses can access the target.
Partial address or name using wildcards. Examples: 10.168.*.*, *.company.com To deny access to a specific host, use the no_access or noaccess qualifier. For example, 10.1.2.38 (no_access) will deny access to the host with the IP address 10.1.2.38.	Clients with matching names or addresses can access the target.

4. Click **OK**.

Adding a logical unit to an iSCSI target

You can add a logical unit to an iSCSI target in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Targets** to display the **iSCSI Targets** page.
2. Select the check box next to the target and then click **details** to display the **iSCSI Target Details** page.

The following table describes the fields on this page:

Field/Item	Description
EVS	Selector for EVS where LUs reside.
EVS iSCSI Domain	Displays the iSCSI Domain, which is the DNS domain used when creating unique qualified names for iSCSI Targets.
Alias	Name of the iSCSI Target.
Comment	Additional information about the iSCSI Target.
Secret	Password used to secure the Target from any unauthorized access. The initiator authenticates against this password when connecting to the Target. The secret should be greater than or equal to 12 characters, but less than 17 characters, in length. Although the secret may be between 1-255 characters in length, some iSCSI initiators will refuse to connect if the secret contains less than 12 characters or more than 16 characters.
Enable Authentication	Enable authentication of the iSCSI Target. By default, the check box is not filled. Filling or clearing the check box enables or disables authentication. When authentication is disabled, initiators are permitted to connect to the target and its LUs without needing to know the target's secret.
Access Configuration	The access configuration parameters.
Available logical units	The list of LUs available for assignment to the iSCSI Target. This list includes all LUs on the EVS. Some of these LUs may already be assigned to other targets.
Selected LUN - LUN Name	The list of LUs selected to be part(s) of the iSCSI Target.
Logical Unit Number	The number assigned to the LU (the LUN). Enter a Logical Unit Number in the range of 0-255.
OK	Saves configuration changes, and closes the page.

Field/Item	Description
cancel	Closes the page without saving configuration changes.

3. Select an LU from the Available Logical Units list, specify a number (0-255) in the Logical Unit Number field, and then click the right arrow to move the LU to the Selected Logical Units list.



Note: You should make sure that the LU is not already assigned to a target.

4. Click **OK**.

Modifying the properties of an iSCSI target

You can modify the properties of an iSCSI target in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Targets** to display the **iSCSI Targets** page.
2. Select the check box next to the target to modify and then click **details**.

The following table describes the fields on this page:

Field/Item	Description
EVS	Selector for EVS where LUs reside.
EVS iSCSI Domain	Displays the iSCSI Domain, which is the DNS domain used when creating unique qualified names for iSCSI Targets.
Alias	Name of the iSCSI Target.
Regenerate GUN	Click to regenerate the Globally Unique Name (GUN).
Comment	Additional information about the iSCSI Target.
Secret	Password used to secure the Target from any unauthorized access. The initiator authenticates against this password when connecting to the Target. The secret should be greater than or equal to 12 characters, but less than 17 characters, in length. Although the secret may be between 1-255 characters in length, some iSCSI initiators will refuse to connect if the secret contains less than 12 characters or more than 16 characters.
Enable Authentication	Enable authentication of the iSCSI Target. By default, the check box is not filled. Filling or clearing the check box enables or disables authentication. When authentication is disabled, initiators are permitted to connect to the target and its LUs without needing to know the target's secret.
Access Configuration	The access configuration parameters.
Available logical units	The list of LUs available for assignment to the iSCSI Target. This list includes all LUs on the EVS. Some of these LUs may already be assigned to other targets.
Selected LUN - LUN Name	The list of LUs selected to be part(s) of the iSCSI Target.
Logical Unit Number	The number assigned to the LU (the LUN). Enter a Logical Unit Number in the range of 0-255.

- The iSCSI Domain, Alias, Available LUs, and Logical Unit Numbers are required. Optionally, you can specify the Comment, Secret, and/or Access Configuration for the Target.



Note: Once set, the iSCSI Domain cannot be changed, but it will be overridden/replaced if you later specify a new iSCSI Target with a different iSCSI Domain in the same EVS. The most recently specified iSCSI Domain overrides all previously-specified iSCSI Domains set for all previously added iSCSI Targets in the EVS.

The following table describes what you can type in the Access Configuration field.

What to type	Means
Blank or *	All clients can access the target.
<p>Specific address or name. Examples: 10.168.20.2, client.dept.company.com</p> <p>To deny access to a specific host, use the no_access or noaccess qualifier. For example, 10.1.2.38 (no_access) will deny access to the host with the IP address 10.1.2.38.</p>	Only clients with the specified names or addresses can access the target.
<p>Partial address or name using wildcards. Examples: 10.168.*.*, *.company.com</p> <p>To deny access to a specific host, use the no_access or noaccess qualifier. For example, 10.1.2.38 (no_access) will deny access to the host with the IP address 10.1.2.38.</p>	Clients with matching names or addresses can access the target.

4. Click **OK**.

Deleting an iSCSI target

You can delete an iSCSI target in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Targets**.
2. Select the check box next to the target to remove and then click **delete**.
3. To confirm the deletion, click **OK**.

Configuring iSCSI security (mutual authentication)

The storage server uses the Challenge Handshake Authentication Protocol (CHAP) to authenticate iSCSI initiators. CHAP requires a “shared secret” known by the initiator and the target. The server also supports mutual authentication where, in addition to the initiator authenticating against the target on the server, the server must also authenticate against the initiator.

To facilitate the mutual authentication process, the server must maintain a list of the initiators with which it can authenticate and the shared secret for each initiator.

Configuring the storage server for mutual authentication

You can configure the storage server for mutual authentication in the NAS Manager.

Procedure

1. Navigate to **Home > File Services > iSCSI Initiator Authentication**.

Field/Item	Description
EVS	Displays the EVS on which to configure Initiator Authentication. Click change to select a different EVS.
Initiator Name	Identifies the initiator with a globally unique name.
Secret	Password used to secure the Initiator from any unauthorized access. The secret should be from 12 to 17 characters in length, but may be between 1-255 characters in length.
details	Click to display the iSCSI Initiator Details page for the selected initiator.
Check All	Click to fill the check box of all initiators in the list.
Clear All	Click to empty the check box of all initiators in the list.
add	Click to display the Add iSCSI Initiator page.
delete	Click to delete the selected iSCSI initiator.
iSCSI Targets	Click to display the iSCSI Targets page.

2. Ensure that the required EVS is selected.
Check the EVS name listed at the top of the page. If this is not the EVS that you want the iSCSI initiator to connect with, click **change** and select the required EVS.
3. Click **add** to add an iSCSI initiator.

[File Services](#) [Home > File Services > iSCSI Initiator Authentication > Add iSCSI Initiator](#)

Add iSCSI Initiator

EVS: g1-eva3

Initiator Name:

Secret:

Shortcuts: [iSCSI Targets](#)

Field/Item	Description
EVS	The EVS on which to configure Initiator Authentication.
Initiator Name	Identifies the initiator with a globally unique name. This name is display in the Change Initiator node name dialog of the Microsoft iSCSI initiator.
Secret	The Secret for the Initiator. This is the secret which will be entered in the Chap Secret Setup dialog of the iSCSI Initiator. This secret is a password which is used to secure the Initiator from unauthorized access. The secret should be from 12 to 17 characters in length, but may be between 1-255 characters in length.
iSCSI Targets	Opens the iSCSI Targets page.

4. Enter the iSCSI initiator name.
5. Enter the initiator secret (password).
6. Save the configuration.

Verify your settings, then click **OK** to save or **cancel** to return to the **iSCSI Initiator Authentication** page without adding the initiator.

Changing the storage server's mutual authentication configuration

Once the storage server's iSCSI initiator configuration has been set up, you can either change an initiator's secret or delete it entirely.

Procedure

1. Navigate to **Home > File Services > iSCSI Initiators**.

Field/Item	Description
EVS	Displays the EVS on which to configure Initiator Authentication. Click change to select a different EVS.
Initiator Name	Identifies the initiator with a globally unique name.
Secret	Password used to secure the Initiator from any unauthorized access. The secret should be from 12 to 17 characters in length, but may be between 1-255 characters in length.
details	Click to display the iSCSI Initiator Details page for the selected initiator.
Check All	Click to fill the check box of all initiators in the list.
Clear All	Click to empty the check box of all initiators in the list.
add	Click to display the Add iSCSI Initiator page.
delete	Click to delete the selected iSCSI initiator.
iSCSI Targets	Click to display the iSCSI Targets page.

2. Ensure that the required EVS is selected.

Check the EVS name listed at the top of the page. If this is not the EVS that you want the iSCSI initiator to connect with, click **change** and select the required EVS.

3. You can now either delete the initiator or change the initiator's secret.

- To delete an iSCSI initiator:

- a. Select the check box for the initiator you want to remove.
- b. Click **delete** to delete the selected initiator.

A confirmation dialog appears, and you can click **OK** to delete the iSCSI Initiator, or cancel to return to the **iSCSI Initiators** page without deleting the Initiator.

- To change the initiator's secret:

- a. Select the check box for the initiator you want to change.
- b. Click **details** to display the **iSCSI Initiator Details** page for the selected initiator.

Field/Item	Description
Initiator Name	The initiator's globally unique name. This name is displayed in the Change Initiator node name dialog of the Microsoft iSCSI initiator or from a file located in the /etc/iscsi directory for a Linux iSCSI initiator.
Secret	The Secret for the Initiator. This is the secret which will be entered in the Chap Secret Setup dialog of the iSCSI Initiator. This secret is a password which is used to secure the Initiator from unauthorized access. The secret should be from 12 to 17 characters in length, but may be between 1-255 characters in length.
iSCSI Targets	Click to display the iSCSI Targets page.

- c. In the Secret field, type the new secret.
The secret should be between 12 and 17 characters in length, but can be between 1-255 characters in length.
- d. Click **OK** to save the changed secret, or click **cancel** to return to the **iSCSI Initiator Authentication** page.
4. Click **OK** to save the changed secret, or click **cancel** to return to the **iSCSI Initiator Authentication** page.
5. Click **details** to display the **iSCSI Initiator Details** page.

Configuring the Microsoft iSCSI initiator for mutual authentication



Note:

- For the latest version of Microsoft iSCSI Software Initiator, visit: <http://www.microsoft.com/>.
- The visible screens depend on the operating system version.

To configure the Microsoft iSCSI Initiator for mutual authentication:

Procedure

1. Navigate to the iSCSI Initiator Properties on your Windows system:
 - a. Start the Microsoft iSCSI Initiator.
 - b. Open the **iSCSI Initiator Properties** dialog.
 - c. Select the **General** tab.
 - d. Click **Secret** to display the **CHAP Secret Setup** dialog.



Note: The shared secret is used to authenticate an initiator with a server, and it should be different from the secret specified when setting up the target.

- e. Enter a secret.
In the field, enter the secret which allows the target to authenticate with initiators when performing mutual CHAP.
 - f. Click **OK** to save the secret and return to the **General** tab of the **iSCSI Initiator Properties** dialog.
2. If necessary, change the initiator node name.
If necessary, you can change the initiator node name.
 - a. Click **Change** to display the **Initiator Node Name Change** dialog.
 - b. Change the name as necessary.

The initiator node name is the name which should be used as the initiator name on the storage server's **iSCSI Initiator Authentication** page (**Home > File Services > iSCSI Initiators**).
 3. Verify the configuration settings.
 4. Click **OK** to save the changes.

Accessing iSCSI storage

iSCSI LUs can be accessed through their targets using the Microsoft iSCSI Initiator. Discovered through iSNS or through the target portal, all iSCSI targets that are available will be displayed as available targets by the Initiator.



Note: Microsoft currently only supports creating a Basic Disk on an iSCSI LU. To ensure data integrity, do not create a dynamic disk. For more information, refer to the *Microsoft iSCSI Initiator User Guide*.

If its underlying volume is mounted read-only by the storage server, or if it is a snapshot copy of another LU, an iSCSI LU will also be read-only. In turn, if an LU is read-only, then any file systems contained within it will also be read-only. Clients accessing such read-only file systems will not be able to change any part of them, including file data, metadata or system files.

If Microsoft Windows clients are required to access read-only NTFS file systems over iSCSI, Microsoft Windows 2008 or Windows 2012 must be used.

Using iSNS to find iSCSI targets

Using iSNS is the easiest way to find iSCSI targets on the network. If the network is configured with an iSNS server, configure the Microsoft iSCSI initiator to use iSNS. To add an iSNS server:



Note: The appearance of the screens depend on the operating system version.

Procedure

1. Navigate to the iSCSI Initiator Properties on your Windows system.
2. Open the **iSCSI Initiator Properties** dialog.

3. Select the **Discovery** tab.
4. In the iSNS Servers area, click **Add** to display the **Add iSNS Server** dialog.
5. Enter the IP address or DNS host name for the iSNS server.
6. Click **OK** to save the IP address or host name and return to the **Discovery** tab of the **iSCSI Initiator Properties** dialog.
7. If necessary, add another iSNS server.



Note: After the iSNS servers have been added, all available iSCSI targets that have been registered in iSNS will appear as available targets.

8. Save your changes.
Verify your settings, and then click **OK** to save the iSNS servers or **Cancel** to decline.

Using target portals to find iSCSI targets

If there are no iSNS servers on the network, iSCSI targets can be found through the use of target portals. Add the file services IP address of the EVS to the target portals list to find targets associated with that server or EVS.



Note: The appearance of the screens depend on the operating system version.

Procedure

1. Navigate to the iSCSI Initiator Properties on your Windows system.
2. Open the **iSCSI Initiator Properties** dialog.
3. Select the **Discovery** tab.
4. In the Target Portals area, click **Add** to display the **Add Target Portal** dialog.
5. Enter the file services IP address of the EVS.
6. Click **OK** to save the IP address and return to the **Discovery** tab of the **iSCSI Initiator Properties** dialog.
7. If necessary, add another target portal.
8. Save your changes.

Verify your settings, then click **OK** to save the list of target portals or **Cancel** to decline.

Accessing available iSCSI targets

To access an available iSCSI Target:



Note: The appearance of the screens depends on the operating system version.

Procedure

1. Navigate to the iSCSI Initiator Properties on your Windows system.
2. Open the **iSCSI Initiator Properties** dialog.
3. Select the **Targets** tab.
4. Select a target.

- Click **Log On** to display the **Log On to Target** dialog.



Note: Each logon starts an iSCSI session, and a maximum of 32 iSCSI sessions are allowed per Target.

- Click **OK** to log on to the target.
- If authentication is enabled on the target, click **Advanced...** to open the **General** tab in the **Advanced Settings** dialog.
 - Select the **CHAP logon information** checkbox and enter the **Target secret** (the password configured when the iSCSI target was created).
 - If mutual authentication has been configured, select the **Perform mutual authentication** checkbox.
 - Click **OK** to save the settings and return to the return to the **Log On to Target** dialog.
- Optionally, configure multi-pathing.

If multi-pathing is supported by the Microsoft iSCSI initiator, and you want to use multiple paths to the target, fill the **Enable multi-path** checkbox.

To create multiple paths to the target, you must later start another session to the target.

- Establish the connection.

Click **OK** to establish the connection to the selected target.

Verifying an active connection

After the connection has been established, you can view any details about the newly established connection.



Note: The appearance of the screen depends on the operating system version.

Procedure

- Navigate to the iSCSI Initiator Properties on your Windows system.
- Open the **iSCSI Initiator Properties** dialog.
- Select the **Targets** tab.
- Look at the Status column for the target.

The Status column for the target should display "Connected."

Terminating an active connection

After the connection has been established, you can terminate the connection if necessary. To end the connection:



Note: The appearance of the screens depends on the operating system version.

Procedure

- Navigate to the iSCSI Initiator Properties on your Windows system.

2. Open the **iSCSI Initiator Properties** dialog.
3. Select the **Targets** tab.
4. Select the target with the connection you want to end.
5. Click **Details** to display the **Target Properties** dialog.
6. Select the session to terminate.

In the list of sessions, select the identifier for the session you want to end.

7. Click **Log off...** to terminate the session.
The initiator will attempt to close the iSCSI session if there are no applications currently using the devices.

Using Computer Manager to configure iSCSI storage

The iSCSI “local disk” must be configured through Windows **Disk Management** tools, and Microsoft recommends that:

- If the LU is smaller than 2 TB, it should be configured as a Basic Disk.
- If the LU is larger than 2 TB, it should be configured as a GPT Disk.

Once the disk is configured, use the Windows **Disk Management** tools to create and format a partition on the disk. For information and instructions on using the Windows **Disk Management** tools, refer to the online help of the operating system on your Windows computer.

Chapter 8: Hitachi Dynamic Provisioning

You can use Hitachi Dynamic Provisioning (HDP) software to improve your storage utilization. The HDP software uses storage-based virtualization layered on top of RAID technology (RAID on RAID) to enable virtual LUNs (dynamically provisioned volumes, DP-Vols) to draw space from multiple pool volumes. This aggregated space widens the storage bottleneck by distributing the I/O to more disks. The greater distribution insulates the server from the realities of the pool volumes (small capacities of individual disks).

HDP with a NAS server provides the following benefits:

- Improves performance by striping I/O across all available disks
- Supports larger LUs, up to 64 TiB
- Reduces the need to use the `span-expand` command. When HDP thin provisioning is used, an HDP pool can be expanded in small increments any number of times. However, if you expand a storage pool, make the increments as large as the initial size of the storage pool to avoid performance problems.
- File system creation and expansion are safe, even on thinly provisioned HDP pools, because they check the amount of available space.

If you are using HDP, see the *Hitachi NAS Platform HDP Best Practices* (MK-92HNAS063) for recommendations.



Note: When using a storage system, there are commonly used Host Mode Options (HMOs) and System Option Modes (SOMs) which should be set correctly. For example, on Hitachi Enterprise RAID systems always enable HMO 7 and 68. When using an HNAS Gateway in a stretched cluster with GAD, enable HMO 78 on the host group containing the HNAS WWPN for the NAS that is considered remote to the system. Contact customer support for more information.

HDP high-level process

The following flow chart shows the high-level process for provisioning storage with HDP:





Figure 1 High-level process for HDP provisioning

Understanding HDP thin provisioning

Thin provisioning allows space to be allocated to an application without it being physically mapped on the storage system until it is actually used. Thin provisioning also decouples the logical provisioning of storage to an application from the physical addition of storage capacity to the storage system.

For example, given 30 TiB of physical storage, you can create an HDP pool with 80 TiB of DP-Vols and create an 80 TiB NAS storage pool on those DP-Vols. None of the available space is allocated until you create and expand file systems. Because only 30 TiB of real space is available, the NAS server will not create more than 30 TiB of file systems in the storage pool. If you later add more parity groups or pool volumes to the HDP pool, you can expand the file systems in the storage pool beyond 30 TiB without creating additional DP-Vols or expanding the NAS storage pool.

The ability to expand file systems instead of the storage pool is advantageous because it enhances performance by spreading the storage chunks used to expand a file system across all the SDs and physical disks in the DP pool. In contrast, a storage pool expansion limits performance by restricting the individual chunks to a small number of SDs and physical disks.


 **Note:** It is strongly recommended that you always use thin provisioning with HDP.

The NAS server reads the real space available in a DP pool. When you create or expand a file system, the NAS server checks for available space, then pre-allocates the space needed for that operation. If the DP pool has too little free space for the operation to complete successfully, the NAS server safely aborts the creation or expansion of the file system.

Every new storage pool should use a single stripeset that resides on a thinly provisioned HDP pool. This way, storage can be expanded in small increments without loss of performance, and all I/O will use all DP-Vols (and all their queue depth) and all physical storage media. For more about queue depth, see the `sd-queue-depth` man page. To use a single stripeset, follow the instructions below.


The process is as follows:

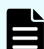
1. When provisioning a new NAS server storage pool, use just enough real disk space to meet your immediate needs for performance and capacity.
2. Place all your parity groups into a single HDP pool, then create DP-Vols whose total capacity roughly meets your expected needs for the next 18 to 24 months.

 **Note:** It does not matter if you over-estimate or under-estimate your capacity needs, because you can easily expand the storage pool beyond the total capacity of the original DP-Vols by adding another set of DP-Vols.

3. Create a NAS server storage pool on these DP-Vols, placing all the DP-Vols into a single stripeset.

Use enough DP-Vols to provide adequate queue depth in the future, after you have added enough parity groups to match the total capacity of the DP-Vols. Four DP-Vols is the bare minimum, but eight DP-Vols will provide better performance than four, and sixteen DP-Vols will be faster than 8 DP-Vols. In practice, a storage pool usually contains an even number of DP-Vols, and the capacity of each DP-Vol is 8 TiB.

 **Note:** If using the CLI `span-create` command, list all the SDs in the initial `span-create` command. Do not run a single `span-create` command, then a series of `span-expand` commands.

 **Note:** When using an application to create a storage pool, specify all the available SDs when creating the storage pool; do not create a single storage pool on a subset of the available SDs, then expand that storage pool onto the rest of the available SDs.

If there are more than 32 available DP-Vols, create the minimum possible number of NAS server stripesets consistent with making all stripesets identical, even if this means creating slightly more or slightly fewer DP-Vols than would otherwise have been created. For example, if you initially estimate that, in two years, you will need 50 8 TiB DP-Vols, you should now create 48 DP-Vols and make 2 stripesets of 24 DP-Vols each.

4. To expand the NAS server storage pool beyond the total capacity of the original DP-Vols, simply add another, identical set of DP-Vols (refer to the `span-expand` man page for more information).



Note: Every new storage pool contains one stripeset, and every expansion (other than by adding storage to the underlying HDP pool) adds a further stripeset.

Understanding how HDP works with HNAS

Using HDP with HNAS provides many benefits.

HDP with HNAS provides the following benefits:

- Improves performance by striping I/O across all available disks
- Supports scalability of larger LUs (typically up to 64 TiB)
- Reduces the need to use the *span-expand* command, and eliminates dynamic read balancing (DRB) limitations. When HDP thin provisioning is used, a pool can be expanded in small increments any number of times. However, if you expand a storage pool, make the increments as large as the initial size of the storage pool to avoid performance problems.
- File system creation or expansion still fails safely, even in the presence of thinly provisioned pools

To fully realize those benefits, see the HDP configuration guidelines in the *Hitachi NAS Platform Storage Pool and HDP Best Practices* (MK-92HNAS048).

As a general rule, you should make a rough forecast of how much data storage capacity will be needed in the next 12 to 24 months, then configure your DP-Vols to be just a little larger than your estimate. If you overestimated your data storage requirements, not too much space will have been wasted; if you underestimated your data storage requirements, you can always add a second, equally large stripeset using the **span-expand** command, then continue to expand the DP pool in increments as small as desired.

Some limitations with HDP thin provisioning and HNAS exist. Consider the following:

- Some storage arrays and systems do not over-commit by more than a factor of ten to one.
- The amount of memory the storage needs for HDP is proportional to the size of the (large, virtual) DP-Vols, not the (smaller, real) pool volumes. Therefore, massive over-commitment causes the storage to prematurely run out of memory.

Hitachi



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact