




Data ONTAP® 8.2

File Access and Protocols Management Guide

For 7-Mode



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1(408) 822-6000
Fax: +1(408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07984_A0
May 2013

Contents

Considerations before configuring file access	13
File protocols that Data ONTAP supports	13
How Data ONTAP controls access to files	13
Authentication-based restrictions	13
File-based restrictions	13
How Data ONTAP supports file access using NFS	15
Enabling or disabling NFS on the storage system	15
Exporting or unexporting file system paths	16
Editing the /etc/exports file	16
Using the exportfs command	17
Enabling and disabling fencing of one or more NFS clients from one or more file system paths	19
Displaying the actual file system path for an exported file system path	20
Displaying the export options for a file system path	21
How the access cache works	21
Adding entries to the access cache	22
Removing entries from the access cache	23
Viewing access cache statistics	23
Optimizing access cache performance	24
Setting access cache timeout values	24
Enabling Kerberos v5 security services for NFS	25
NFS clients that support Kerberos v5 security services	26
Configuring Kerberos v5 security services for NFS to use an Active- Directory-based KDC	26
Configuring Kerberos v5 security services for NFS to use a UNIX-based KDC	29
Debugging mounting problems	35
Displaying mount service statistics	35
Tracing mountd requests	35
Use of hard mounts	36
Displaying NFS statistics	36
Controlling NFS requests from nonreserved ports	36

Enabling or disabling NFSv2	37
Enabling or disabling NFSv3	37
Support for VMware vStorage over NFS	38
Enabling or disabling VMware vStorage over NFS	38
Displaying statistics for VMware vStorage over NFS	39
Differences in file system ID handling for NFSv3 and NFSv4	39
Supporting NFSv4 clients	40
NFSv4.0 functionality supported by Data ONTAP	40
Limitations of Data ONTAP support for NFSv4	40
How the pseudo-fs in NFSv4 affects mountpoints	41
Enabling or disabling NFSv4.0	42
Specifying the user ID domain for NFSv4	42
Managing NFSv4 ACLs	42
Managing NFSv4 file delegations	45
Configuring NFSv4 file and record locking	48
How the name server database cache works	50
Flushing the name server database cache	50
Differences in NFS user and group ID formats	51
Allowing or disallowing NFSv4 user and group IDs as numeric strings	51
Supporting PC-NFS clients	52
How the pcnfsd daemon works	52
Enabling or disabling the pcnfsd daemon	52
Creating PC-NFS user entries in the storage system's local files	53
How umask works with NFS file permissions	53
Defining the umask for files and directories that PC-NFS users create	54
Supporting WebNFS clients	54
Enabling or disabling the WebNFS protocol	54
Setting a WebNFS root directory	55
Support for NFS over IPv6	56
Enabling or disabling NFS over IPv6	56
Textual representation of IPv6 addresses	56
File access using CIFS	57
Unsupported Windows features	57
Configuring CIFS licenses	57
Connecting the MMC to the storage system	58
Configuring CIFS on your storage system	58

Supported CIFS clients and domain controllers	58
What the cifs setup command does	58
Requirements for the CIFS server name	59
Setting up your system initially	59
Specifying WINS servers	59
Changing the storage system domain	60
Changing protocol modes	61
Specifying Windows user account names	63
Support for the Windows “Owner Rights” security principal	64
Considerations when reconfiguring CIFS	64
Reconfiguring CIFS on your storage system	65
Configuring SMB on your storage system	66
Support for the SMB 1.0 protocol	66
Support for the SMB 2.0 protocol	66
Support for the SMB 2.1 protocol	66
When to enable the SMB 2.0 and SMB 2.1 protocols	67
Enabling or disabling SMB 2.x	67
SMB 2.x and durable handles	68
Monitoring durable handle status	68
Support for SMB signing	69
Managing shares	73
What you need to know when creating a share	73
Share naming conventions	74
Creating a CIFS share from the MMC on a Windows client	74
Creating a CIFS share from the Data ONTAP command line	75
Displaying and changing the properties of a share	76
Deleting a share	84
Managing access control lists	85
How Data ONTAP uses share-level ACLs	85
How ACL inheritance works	85
Displaying and changing a share-level ACL	86
Displaying and changing a file-level ACL	92
Specifying how group IDs work with share-level ACLs	94
Managing home directories	95
About home directories on the storage system	95
How Data ONTAP matches a directory with a user	96

How symbolic links work with home directories	97
Specifying home directory paths	97
Displaying the list of home directory paths	98
Specifying the naming style of home directories	99
Creating directories in a home directory path (domain-naming style)	99
Creating directories in a home directory path (non-domain-naming style)	100
Creating subdirectories in home directories when a home directory path extension is used	101
Syntax for specifying a home directory using a UNC name	101
Enabling users to access other users' home directories	102
Accessing your CIFS home directory using a share alias	102
Enabling or disabling widelinks from a share	103
Disabling home directories	103
Using BranchCache to cache CIFS share content at a branch office	103
Supported BranchCache caching modes	104
BranchCache terminology	105
How Data ONTAP implements BranchCache	106
How BranchCache clients retrieve cached content	106
Network protocol support requirements	108
Data ONTAP and Windows hosts version requirements	108
Configuring BranchCache	108
Enabling BranchCache caching on a CIFS share	110
How Data ONTAP manages BranchCache hashes	112
Displaying BranchCache hash statistics	113
Changing the BranchCache server secret key	114
What happens when you disable BranchCache on the storage system	115
Disabling BranchCache on the storage system	115
Disabling BranchCache on an existing CIFS share	116
Managing local users and groups	116
Managing local users	116
Managing local groups	118
Applying Group Policy Objects	120
Requirements for using GPOs with storage systems	121
Associating the storage system with a different OU	121
Enabling or disabling GPO support on a storage system	122

Managing GPOs on the storage system	122
Improving client performance with opportunistic and lease oplocks	129
Write cache data-loss considerations when using oplocks	129
How lease oplocks provide enhanced locking	129
Enabling or disabling oplocks and lease oplocks on the storage system	130
Enabling or disabling oplocks and lease oplocks on a qtree	130
Monitoring lease oplock status	131
Changing the delay time for sending oplock breaks	132
Managing authentication and network services	133
Understanding authentication issues	133
Setting the storage system's minimum security level	134
Preventing Kerberos passive replay attacks	135
Selecting domain controllers and LDAP servers	135
Using null sessions to access storage in non-Kerberos environments	140
Creating NetBIOS aliases for the storage system	142
Disabling NetBIOS over TCP	143
Monitoring CIFS activity	144
Different ways to specify a user	144
Displaying a summary of session information	145
Timing out idle sessions	145
Tracking statistics	145
Viewing specific statistics	146
Saving and reusing statistics queries	147
CIFS resource limitations	147
Managing CIFS services	147
Disconnecting clients using the MMC	147
Disconnecting a selected user from the command line	148
Disabling CIFS for the entire storage system	148
Specifying which users receive CIFS shutdown messages	149
Restarting CIFS service	150
Sending a message to users on a storage system	150
Displaying and changing the description of the storage system	151
Changing the storage system computer account password	151
About file management using Windows administrative tools	153
Troubleshooting access control problems	154
Adding permission tracing filters	154

Removing permission tracing filters	155
Displaying permission tracing filters	155
Finding out why Data ONTAP allowed or denied access	156
Using FPolicy	157
Introduction to FPolicy	157
Use of FPolicy within Data ONTAP	163
How to use native file blocking	164
How to work with FPolicy	168
FAQs, error messages, warning messages, and keywords	213
Support for CIFS over IPv6	227
Enabling or disabling CIFS over IPv6	228
Listing IPv4 or IPv6 CIFS sessions	228
Listing cumulative IPv4 or IPv6 CIFS sessions	230
File sharing between NFS and CIFS	232
NFS and CIFS file naming dependencies	232
Characters a file name can use	232
Case-sensitivity of a file name	232
Creating lowercase file names	233
How Data ONTAP creates file names	233
Controlling the display of dot files from CIFS clients	233
Enabling file name character translation between UNIX and Windows	234
Character restrictions	235
Clearing a character mapping from a volume	235
About file locking between protocols	236
Breaking locks	236
Breaking locks on a specified file	237
Breaking locks on a specified host	238
Breaking locks for a specified lock owner	239
Breaking locks based on a specified protocol	240
Breaking NLM locks on a specified IP network family	240
How Data ONTAP treats read-only bits	241
Deleting files with the read-only bit set	241
Managing UNIX credentials for CIFS clients	242
How CIFS users obtain UNIX credentials	242
Ensuring that only intended CIFS users receive UNIX credentials	243
Managing the SID-to-name map cache	254

Enabling or disabling the SID-to-name map cache	254
Changing the lifetime of SID-to-name mapping entries	255
Clearing all or part of the SID-to-name map cache	255
Using LDAP services	256
Configuring LDAP services	256
Managing client authentication and authorization	262
Managing LDAP user-mapping services	263
Specifying base and scope values for user-mapping	264
Managing Active Directory LDAP servers	265
Managing LDAP schema	267
Enabling Storage-Level Access Guard using the fsecurity command	268
About the fsecurity command	269
Generating and editing the job definition file	269
Specifying job definition file elements	270
Creating a security job and applying it to the storage object	271
Checking the status of or canceling a security job	272
Displaying the security settings on files and directories	273
Removing the Storage-Level Access Guard	273
Auditing system access events	274
About auditing	274
Events that Data ONTAP can audit	274
Configuring system event auditing	276
Viewing and understanding event detail displays	287
Controlling CIFS access to symbolic links	291
Enabling CIFS clients to follow symbolic links	291
Specifying how CIFS clients interact with symbolic links	292
Why you should avoid symbolic links to files	292
About Map entries	293
About widelink entries	293
About disabling share boundary checking for symbolic links	294
Redirecting absolute symbolic links	295
How the storage system uses Map and Widelink entries	296
Optimization of NFS directory access for CIFS clients	297
Creating new directories in Unicode format	297
Converting existing directories to Unicode format	298
Preventing CIFS clients from creating uppercase file names	299

Accessing CIFS files from NFS clients	299
Adding mapping entries to the WAFL credential cache	300
Deleting mapping entries from the WAFL credential cache	300
Setting how long mapping entries are valid	301
Monitoring WAFL credential cache statistics	302
Managing mapping inconsistencies	303
Tracing CIFS logins	304
Tracing domain controller connections	304
Allowing CIFS clients without UNIX "execute" permissions to run .dll and .exe files	305
How file editing with Windows applications affects UNIX permissions	305
File access using FTP	307
Managing FTP	307
Enabling or disabling the FTP server	307
Enabling or disabling the TFTP server	307
Enabling or disabling FTP file locking	308
Specifying the FTP authentication style	308
Enabling or disabling the bypassing of FTP traverse checking	310
Restricting FTP access	310
Managing FTP log files	312
Viewing SNMP traps that the FTP server generates	314
Viewing FTP statistics	315
Resetting FTP statistics	316
Specifying the maximum number of FTP connections	316
Specifying the maximum number of TFTP connections	316
Setting the FTP connection threshold	317
Specifying the TCP window size for FTP operations	317
Specifying the FTP idle timeout	317
Managing anonymous FTP access	318
Managing the Secure File Transfer Protocol (SFTP)	319
About SFTP	319
Limitations of Data ONTAP support for SFTP	320
Enabling or disabling SFTP	320
Enabling or disabling SFTP file locking	320
SFTP authentication styles	321
Specifying the SFTP authentication style	321

Enabling or disabling SFTP bypass traverse checking	322
Enabling or disabling SFTP user home directory restrictions	322
Specifying the SFTP override path for user home directories	323
Enabling or disabling the overriding of UNIX permissions	323
Managing SFTP log files	324
Viewing SFTP statistics	325
Resetting SFTP statistics	325
Specifying the maximum number of SFTP connections	325
Specifying the SFTP idle timeout value	326
Managing FTP over SSL (FTPS)	326
Differences between implicit and explicit FTPS	326
Enabling or disabling explicit FTPS	327
Allowing or preventing the opening of explicit FTPS data connections in secure mode	328
Enabling or disabling implicit FTPS	328
Managing FTP over IPv6	329
Enabling or disabling FTP over IPv6	329
Viewing FTP over IPv6 statistics	330
File access using HTTP	331
Managing the Data ONTAP HTTP server	331
Enabling or disabling the Data ONTAP HTTP server	331
Enabling or disabling the bypassing of HTTP traverse checking	331
Specifying the root directory for the Data ONTAP HTTP server	332
Specifying the maximum size of the log file for the Data ONTAP HTTP server	332
Testing the Data ONTAP HTTP server	332
Specifying how the Data ONTAP HTTP server maps MIME content types to file name extensions	333
Specifying how the Data ONTAP HTTP server translates HTTP requests	334
Configuring MIME Content-Type values	337
Maintaining security for the Data ONTAP HTTP server	338
Displaying HTTP server statistics	344
Resetting statistics for the Data ONTAP HTTP server	347
Viewing HTTP server connection information	347
Purchasing and connecting a third-party HTTP server to your storage system	348

Support for HTTP and HTTPS over IPv6	349
Enabling or disabling HTTP and HTTPS over IPv6	349
Listing HTTP connections over IPv4 or IPv6	350
File access using WebDAV	352
Understanding WebDAV	352
Managing the Data ONTAP WebDAV server	353
Enabling or disabling the Data ONTAP WebDAV server	353
Pointing a WebDAV client to a home directory	353
Purchasing and connecting a third-party WebDAV server to your storage system	354
Event log and audit policy mapping	355
Event Log mapping values	355
Audit mapping values	356
Glossary	358
Copyright information	366
Trademark information	367
How to send your comments	368
Index	369

Considerations before configuring file access

Data ONTAP allows you to manage access to files by clients using different protocols. There are certain concepts you should be familiar with before configuring file access.

File protocols that Data ONTAP supports

Data ONTAP supports file access using the NFS, CIFS, FTP, HTTP, and WebDAV protocols.

How Data ONTAP controls access to files

Data ONTAP controls access to files according to the authentication-based and file-based restrictions that you specify.

When a client connects to the storage system to access files, Data ONTAP has to perform two tasks:

- **Authentication**
Data ONTAP has to authenticate the client by verifying the identity with a trusted source.
- **Authorization**
Data ONTAP has to authorize the user by comparing the user's credentials with the permissions configured on the file or directory and determining what type of access, if any, to provide.

To properly manage file access control, Data ONTAP must communicate with external services such as NIS, LDAP, and Active Directory servers. Configuring a storage system for file access using CIFS or NFS requires setting up the appropriate services depending on your environment in Data ONTAP.

Authentication-based restrictions

With authentication-based restrictions, you can specify which client machines and which users can connect to the storage system.

Data ONTAP supports Kerberos authentication from both UNIX and Windows servers.

File-based restrictions

With file-based restrictions, you can specify which users can access which files.

When a user creates a file, Data ONTAP generates a list of access permissions for the file. Although the form of the permissions list varies with each protocol, it always includes common permissions, such as reading and writing permissions.

When a user tries to access a file, Data ONTAP uses the permissions list to determine whether to grant access. Data ONTAP grants or denies access according to the operation that the user is performing, such as reading or writing, and the following factors:

- User account

14 | File Access and Protocols Management Guide

- User groups or netgroups
- Client protocol
- File type

As part of the verification process, Data ONTAP maps host names to IP addresses using the lookup service you specify—Lightweight Directory Access Protocol (LDAP), Network Information Service (NIS), Domain Name Service (DNS), or local storage system information.

How Data ONTAP supports file access using NFS

You can export and unexport file system paths on your storage system, making them available or unavailable, respectively, for mounting by NFS clients, including PC-NFS and WebNFS clients.

Enabling or disabling NFS on the storage system

Enabling or disabling NFS on the storage system includes managing the NFS license and the NFS server. You can use the `license show` command to manage the NFS license and the `nfs` command to manage the NFS server.

Steps

1. Check whether NFS is currently licensed on the storage system by entering the following command:

```
license
```

2. Perform one of the following actions:

If NFS is...	Enter the following command...
Not licensed and you want to license it	license add <i>nfs_license_key</i>
Licensed and you want to remove the license	license delete nfs

3. Check whether the NFS server is currently running on the storage system by entering the following command:

```
nfs status
```

4. Perform one of the following actions:

If the NFS server is...	Enter the following command...
Not running and you want to enable it	nfs on
Running and you want to disable it	nfs off

After you finish

If you installed an NFS license and enabled the NFS server, you must now configure exports before NFS clients can access data on the storage system.

Exporting or unexporting file system paths

You can export or unexport a file system path, making it available or unavailable to NFS clients, by editing the `/etc/exports` file or running the `exportfs` command.

Before you begin

To support secure NFS access (through using the `sec=krb*` export option), you must first enable Kerberos v5 security services.

About this task

If you need to make permanent changes to several export entries at once, it is usually easiest to edit the `/etc/exports` file directly. However, if you need to make changes to a single export entry or you need to make temporary changes, it is usually easiest to run the `exportfs` command.

Editing the `/etc/exports` file

To specify which file system paths Data ONTAP exports automatically when NFS starts, you can edit the `/etc/exports` file. For more information, see the `na_exports(5)` manual page.

About this task

If the `nfs.export.auto-update` option is on, which it is by default, Data ONTAP automatically updates the `/etc/exports` file when you create, rename, or delete volumes. For more information, see the `na_options(1)` manual page.

If you edit the `/etc/exports` file using a text editor, your changes do not take effect until you export all file system paths in the `/etc/exports` file or synchronize the currently exported file system paths with those specified in the `/etc/exports` file.

Note: Running the `exportfs` command with the `-b`, `-p`, or `-z` option also changes the `/etc/exports` file.

Note: The maximum number of lines in the `/etc/exports` file is 10,240. This includes commented lines. The maximum number of characters in each export entry, including the end of line character, is 4,096.

An export entry has the following syntax:

```
path-option[, option...]
```

In the export entry syntax, *path* is a file system path (for example, a path to a volume, directory, or file) and *option* is an export option that specifies the following information:

- Which NFS clients have which access privileges (read-only, read-write, or root)
- The user ID (or name) of all anonymous or root NFS client users that access the file system path

- Whether NFS client users can create `setuid` and `setgid` executables and use the `mknod` command when accessing the file system path
- The security types that an NFS client must support to access the file system path
- The actual file system path corresponding to the exported file system path

Steps

1. Open the `/etc/exports` file in a text editor on an NFS client that has root access to the storage system.
2. Make your changes.
3. Save the file.

Using the `exportfs` command

To export or unexport file system paths from the command line, you can run the `exportfs` command. For more information, see the `na_exportfs(1)` man page.

Exporting file system paths

You can export a file system path with or without adding a corresponding entry to the `/etc/exports` file. In addition, you can export all file system paths specified in the `/etc/exports` file.

Exporting a file system path and adding a corresponding entry to the `/etc/exports` file

You can use the `exportfs -p` command to export a file system path and add a corresponding export entry to the `/etc/exports` file.

Step

1. Enter the following command:

```
exportfs -p [options] path
```

options is a comma-delimited list of export options. For more information, see the `na_exports(5)` man page.

path is a file system path (for example, a path to a volume, directory, or file).

Note: If you do not specify any export options, Data ONTAP automatically exports the file system path with the `rw` and `sec=sys` export options.

Exporting a file system path without adding a corresponding entry to the `/etc/exports` file

You can use the `exportfs -io` command to export a file system path without adding a corresponding export entry to the `/etc/exports` file.

Step

1. Enter the following command:

exportfs -io [options] path

options is a comma-delimited list of export options. For more information, see the `na_exports(5)` man page.

path is a file system path (for example, a path to a volume, directory, or file).

Note: If you do not specify any export options, Data ONTAP uses the export options specified for the file system path in the `/etc/exports` file, if any.

Exporting all file system paths specified in the `/etc/exports` file

You can use the `exportfs -a` command to export all file system paths specified in the `/etc/exports` file.

Step

- 1. Enter the following command:

exportfs -a

Unexporting file system paths

You can unexport one file system path and optionally remove its corresponding entry from the `/etc/exports` file. In addition, you can unexport all file system paths without removing their corresponding entries from the `/etc/exports` file.

Unexporting one file system path

You can use the `exportfs -u` command to unexport one file system path without removing its corresponding entry from the `/etc/exports` file. You can use the `exportfs -z` command to unexport one file system path and remove its corresponding entry from the `/etc/exports` file.

Step

- 1. Perform one of the following actions.

If you want to unexport one file system path...	Then...
Without removing its corresponding entry from the <code>/etc/exports</code> file	Enter the following command: exportfs -u path <i>path</i> is the file system path.
And remove its corresponding entry from the <code>/etc/exports</code> file	Enter the following command: exportfs -z path <i>path</i> is the file system path.

Unexporting all file system paths

You can use the `exportfs -ua` command to unexport all file system paths without removing their corresponding entries from the `/etc/exports` file.

Before you begin

Note: Be aware that running this command unmounts all file system paths, disconnecting all NFS clients from the storage system.

Step

1. Enter the following command:

```
exportfs -ua
```

`-u` specifies to unexport file system paths.

`-a` specifies all file system paths.

Synchronizing the currently exported file system paths with those specified in the `/etc/exports` file

You can use the `exportfs -r` command to export all file system paths specified in the `/etc/exports` file and unexport all file system paths not specified in the `/etc/exports` file.

Step

1. Enter the following command:

```
exportfs -r
```

Enabling and disabling fencing of one or more NFS clients from one or more file system paths

You can use fencing to give multiple NFS clients temporary or permanent read-only or read-write access to multiple file system paths.

About this task

When you enable or disable fencing, Data ONTAP moves the NFS clients you specify to the front of their new access lists (`rw=` or `ro=`). This reordering can change your original export rules.

Step

1. Enter the following command:

```
exportfs -b enable | disable save | nosave allhosts |  
clientid[:clientid...] allpaths | path[:path...]
```

If you want to...	Then...
Enable fencing	Specify the <code>enable</code> option.
Disable fencing	Specify the <code>disable</code> option.
Update the <code>/etc/exports</code> file	Specify the <code>save</code> option.
Prevent the updating of the <code>/etc/exports</code> file	Specify the <code>nosave</code> option.
Affect all NFS clients	Specify the <code>allhosts</code> option.
Affect all exported file system paths	Specify the <code>allpaths</code> option.
Affect a specific set of NFS clients	Specify a colon-delimited list of NFS client identifiers.
Affect a specific set of file system paths	Specify a colon-delimited list of file system paths

Data ONTAP processes all of the NFS requests in its queue before it enables or disables fencing, thereby ensuring that all file writes are complete.

Displaying the actual file system path for an exported file system path

You can use the `exportfs -s` command to display the actual file system path for an exported file system path.

About this task

A file system's actual path is the same as its exported path unless you export it with the `-actual` option. For more information, see the `na_exports(5)` man page.

Step

1. Enter the following command:

```
exportfs -s path
```

path specifies the exported file system path.

Displaying the export options for a file system path

You can use the `exportfs -q` command to display the export options for a file system path, which can help you in debugging an export problem.

Step

1. Enter the following command:

```
exportfs -q path
```

path specifies the file system path.

Result

Data ONTAP displays the export options for the path you specify.

Note: Data ONTAP also displays a rule identifier for each option, but you do not need the rule identifier unless you are using diagnostic commands. For more information, contact technical support.

How the access cache works

The Data ONTAP access cache reduces the likelihood of having to perform a reverse DNS lookup or parse netgroups when granting or denying an NFS client access to a file system path. This results in performance improvements due to less time used for DNS lookups.

Whenever an NFS client attempts to access a file system path, Data ONTAP must determine whether to grant or deny access. Except in the most simple cases (for example, when a file system path is exported with just the `ro` or `rw` option), Data ONTAP grants or denies access according to a value in the access cache that corresponds to the following things:

- The file system path
- The NFS client's IP address, access type, and security type

This value might not exist in the access cache entry if Data ONTAP has not made a previous access determination or you have not created an access cache entry using the `exportfs -c` command for this particular NFS client-file system path combination. In this case, Data ONTAP grants or denies access according to the result of a comparison between the following things:

- The NFS client's IP address (or host name, if necessary), access type, and security type
- The file system path export rules

Data ONTAP then stores the result of this comparison in the access cache.

To reduce the likelihood that of having to perform a reverse DNS lookup or parse netgroups, Data ONTAP breaks this comparison into three stages. It performs each successive stage of the comparison only if necessary to determine whether the NFS client has access to the file system path.

In the first stage, Data ONTAP compares the NFS client's IP address with all export rules that consist entirely of IP addresses, including single IP addresses, subnets, and host names that Data ONTAP has previously resolved to IP addresses.

In the second stage, Data ONTAP performs a reverse DNS lookup on the NFS client's IP address, and then compares the NFS client's host name with all of the export rules that contain subdomains and host names that Data ONTAP has not resolved into IP addresses.

In the third stage, Data ONTAP parses netgroups.

Data ONTAP backs up the entry cache onto disk every 15 minutes so that the information in the access cache is available after reboots and after takeover or giveback.

Adding entries to the access cache

You can use the `exportfs -c` command to check whether an NFS client has a specific type of access to a file system path and simultaneously add a corresponding entry to the access cache.

Step

1. To check NFS client access and add an entry to the access cache, enter the following command:

```
exportfs -c clientaddr[:clientaddr...] path [accesstype] [securitytype]
```

clientaddr specifies the NFS client IP address.

path specifies the file system path.

accesstype specifies one of the following access type options:

- `ro`—read-only access
- `rw`—read-write access
- `root`—root access

If you do not specify an access type, Data ONTAP simply checks whether the NFS client can mount the file system path.

securitytype specifies one of the following security type options:

- `sys`—Unix-style security
- `none`—no security
- `krb5`—Kerberos Version 5 authentication
- `krb5i`—Kerberos Version 5 integrity service
- `krb5p`—Kerberos Version 5 privacy service

If you do not specify a security type, Data ONTAP assumes the NFS client's security type is `sys`.

Removing entries from the access cache

Data ONTAP automatically removes entries from the access cache when you unexport a file system path or the entries time out. You can use the `exportfs -f` command to manually remove entries from the access cache.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Remove all entries from the access cache	<code>exportfs -f</code>
Remove all entries for a specific file system path from the access cache	<code>exportfs -f <i>path</i></code>
Remove all entries for a specific client from the access cache	<code>exportfs -f -n <i>clientIP_or_hostname</i></code>
Remove all entries for a specific file system path for a specific client from the access cache	<code>exportfs -f -n <i>clientIP_or_hostname path</i></code>

If you are flushing the access cache due to a changed reverse lookup DNS entry for a client, you must use the `-n` parameter to also flush the reverse lookup DNS cache. This prevents issues due to stale reverse lookup DNS cache entries.

Note: You can only run the `exportfs -f` command once every 60 seconds.

For more information, see the `na_exportfs(1)` man page.

Viewing access cache statistics

You can use the `nfstat -d` command to view access cache statistics. This enables you to view detailed information about access cache statistics, connections, requests, and more for troubleshooting purposes.

Step

1. To view access cache statistics, enter the following command:

```
nfstat -d
```

For more information about these access cache statistics, see the `na_nfstat(1)` man page.

Optimizing access cache performance

To optimize access cache performance, you should reuse identical export rules as often as possible.

About this task

Data ONTAP maintains a single access cache entry for all export entries that specify the same rule.

Step

1. Reuse identical export rules whenever possible.

Example

Even though the `ro,rw=@group1` rule exists in both of the following export entries, Data ONTAP maintains a single access cache entry for the rule:

```
/vol/a -sec=sys,ro,sec=sys,rw=@group1,sec=krb5,rw=@group2
/vol/b -sec=sys,ro,sec=sys,rw=@group1
```

Setting access cache timeout values

You can set several options to customize the access cache timeout behavior. This enables you to balance access cache performance with how recent the stored information is.

Steps

1. To specify how long Data ONTAP keeps an entry in the access cache, enter the following command:

```
options nfs.export.harvest.timeout integer
```

integer specifies the idle expiration time for entries in the export access cache in seconds. The default is 3600 seconds (one hour). The minimum value is 60 seconds. The maximum value is 604800 seconds (seven days).

2. To specify how long Data ONTAP uses an access cache entry which was denied access before refreshing it, enter the following command:

```
options nfs.export.neg.timeout integer
```

integer specifies the timeout period in seconds. The default is 1800 seconds (30 minutes). The minimum value is 60 seconds. The maximum value is 604800 seconds (seven days).

3. To specify how long Data ONTAP uses an access cache entry which was granted access before refreshing it, enter the following command:

```
options nfs.export.pos.timeout integer
```


integer specifies the idle expiration time for entries in the export access cache in seconds. The default is 36000 seconds (ten hours). The minimum value is 60 seconds. The maximum value is 604800 seconds (seven days). For more information, see the `na_options(1)` man page.

Enabling Kerberos v5 security services for NFS

To enable Kerberos v5 security services for NFS, you can use the `nfs setup` command.

About this task

Data ONTAP provides secure NFS access using the Kerberos v5 authentication protocol to ensure the security of data and the identity of users within a controlled domain.

The Data ONTAP Kerberos v5 implementation for NFS supports two Kerberos Key Distribution Center (KDC) types: Active Directory-based and UNIX-based, as described in the following table.

KDC type	Description
Active Directory-based	The Kerberos realm for NFS is an Active Directory-based KDC. You must configure CIFS with Microsoft Active Directory authentication (which is Kerberos-based); then NFS will use the CIFS domain controller as the KDC.
UNIX-based	The Kerberos realm for NFS is an MIT or Heimdal KDC.
Multirealm	Uses a UNIX-based KDC for NFS and an Active Directory-based KDC for CIFS. Available in Data ONTAP 7.3.1 and later releases.

Note: To support Kerberos multirealm configurations, Data ONTAP uses two sets of principal and keytab files. For Active Directory-based KDCs, the principal and keytab files are `/etc/krb5auto.conf` and `/etc/krb5.keytab`, respectively, just as in releases prior to Data ONTAP 7.3.1. For UNIX-based KDCs, however, the principal and keytab files are `/etc/krb5.conf` and `/etc/UNIX_krb5.keytab`, respectively. Starting with Data ONTAP 7.3.1, the keytab file for UNIX-based KDCs has changed from `/etc/krb5.keytab` to `/etc/UNIX_krb5.keytab`.

Data ONTAP continues to use the old keytab file `/etc/krb5.keytab`, however, if you upgrade from a release prior to Data ONTAP 7.3.1 in which Data ONTAP was configured to use a UNIX-based KDC for NFS. You need only use the new keytab file `/etc/UNIX_krb5.keytab` for UNIX-based KDCs if you are reconfiguring CIFS after upgrading from such a release or if you are configuring NFS for the first time after configuring an Active-Directory-based KDC for CIFS.

NFS clients that support Kerberos v5 security services

Before using Kerberos v5 security services with an NFS client, you should make sure that the NFS client supports RFC1964 and RFC2203.

For more information, see the Interoperability Matrix at support.netapp.com/NOW/products/interoperability.

Configuring Kerberos v5 security services for NFS to use an Active-Directory-based KDC

You can configure Kerberos v5 security services for NFS to use an Active-Directory-based KDC before or after running the `cifs setup` command. The security service setup procedure adds your storage system to an Active Directory-based KDC as a service principal called `nfs/hostname.domain@REALM`.

Configuring Kerberos v5 security services for NFS to use an Active-Directory-based KDC before configuring CIFS

If you have not run `cifs setup` to configure CIFS, you must provide configuration information that would otherwise have been taken from your CIFS configuration.

Before you begin

You must configure your storage system to use the Active Directory-based domain name service, modify the `/etc/resolv.conf` file as necessary to ensure that it lists only Active Directory servers.

For example, for a Kerberos realm in which the Active Directory servers are 172.16.1.180 and 172.16.1.181, you would change `/etc/resolv.conf` to include only the following Active Directory server entries:

```
nameserver 172.16.1.180
```

```
nameserver 172.16.1.181
```

Make sure you remove all other Active Directory server entries for that realm.

About this task

If you have already used `nfs setup` to enter configuration information, the prompts you receive may differ from those shown in the following procedure.

Steps

1. Enter the following command:

```
nfs setup
```

You receive the following message:

```
Enable Kerberos for NFS?
```

2. To continue, enter the following: **y**

You are asked to specify the type of KDC:

```
The filer supports these types of Kerberos Key Distribution Centers
(KDCs):
1 - UNIX KDC
2 - Microsoft Active Directory KDC
Enter the type of your KDC (1-2):
```

3. Enter the following: **2**

You are prompted to specify the storage system name:

```
The default name of this filer will be 'SERVER'
Do you want to modify this name? [no]:
```

4. Enter **yes** to be prompted for a storage system name or press **Enter** to accept the default storage system name “SERVER”.

You are prompted to specify the domain name for the storage system’s Active Directory server:

```
Enter the Windows Domain for the filer []:
```

5. Enter the domain name for the Active Directory server.

The domain name you enter is also used as the Kerberos realm name.

You are prompted to set up a local administrator account.

6. Enter the local administrator account information.

Note: This step has no effect on Kerberos configuration for an Active Directory KDC.

7. After you enter local administrator account information, verify the resulting message.

It should look similar to the following example:

```
ADKDC.LAB.DOCEXAMPLE.COM is a Windows 2000(tm) domain.
```

This message verifies that the storage system was able to find the Active Directory server, and that the storage system has determined this server can function as a KDC server.

If you do not receive a message such as this one, it indicates that there may be a problem with the Active Directory server, or that the DNS server for the storage system is not an Active Directory server. Check your network configuration, then run `nfs setup` again.

8. When you receive the following type of message, enter name and password information for the Active Directory domain administrator:

```
In order to create this filer's domain account, you must supply the
name and password of an administrator account with sufficient
privilege to add the filer to the ADKDC.LAB.DOCEXAMPLE.COM domain.
Please enter the Windows 2000 user
[Administrator@ADKDC.LAB.DOCEXAMPLE.COM] Password for Administrator:
```

If the password is correct and the specified account has the proper permissions within the storage system domain, you receive the following type of message:

```
CIFS - Logged in as administrator@ADKDC.LAB.DOCEXAMPLE.COM.
Welcome to the ADKDC (ADKDC.LAB.DOCEXAMPLE.COM) Windows 2000(tm)
domain.
Kerberos now enabled for NFS.
NFS setup complete.
```

Result

You might see the following message in the output text upon completion of NFS setup. This output is an artifact of the installation process, and can be ignored:

```
CIFS is not licensed.
(Use the "license" command to license it.)
```

Configuring Kerberos v5 security services for NFS to use an Active-Directory-based KDC after configuring CIFS

If you have already run `cifs setup` and configured Data ONTAP to use Active Directory for CIFS, `nfs setup` automatically uses some of the configuration information you specified for CIFS.

About this task

Note: If you have already used `nfs setup` to enter configuration information, the prompts you receive might differ from those shown in the following procedure.

Steps

1. Enter the following command:

```
nfs setup
```

You receive the following message from `nfs setup`:

```
Enable Kerberos for NFS?
```

2. Enter `y` to continue.

You are asked to specify the type of KDC:

```
The filer supports these types of Kerberos Key Distribution Centers
(KDCs):
1 - UNIX KDC
2 - Microsoft Active Directory KDC
Enter the type of your KDC (1-2):
```

3. Enter 2.

You receive the following message:

```
Kerberos now enabled for NFS.
NFS setup complete.
```

Result

Data ONTAP is now configured for Active Directory-based KDC Kerberos over NFS.

Configuring Kerberos v5 security services for NFS to use a UNIX-based KDC

To configure Kerberos v5 security services for NFS to use a UNIX-based KDC, you can create a principal (a realm user ID) and generate a keytab (key table file) for your storage system and configure Data ONTAP to use your UNIX-based KDC.

Rules for configuring Kerberos v5:

- An NFS client and a UNIX-based KDC must be set up, with client principals for root and at least one non-root client.
- NFS access must be verified for a client and an existing network server.
- The storage system, NFS client, and Kerberos server time must be synchronized with an NTP server, must be in the same time zone, and must have a clock skew of less than five minutes.
- DNS must be enabled on your storage system before setting up and using secure NFS.

If the host component is not already a fully qualified domain name and DNS has not been enabled, then you must change all your NFS server principal names to enable DNS later.

Note: You cannot authenticate CIFS clients with a UNIX-based KDC (that is, because of proprietary restrictions, there are no UNIX-based Kerberos implementations that support CIFS clients). However, in Data ONTAP 7.3.1 and later releases, which provide Kerberos multirealm functionality, you can configure CIFS to use a Microsoft Active Directory-based KDC for authentication of CIFS clients while simultaneously configuring NFS to use a UNIX-based KDC for authentication of NFS clients.

The steps to configure your environment depend on your system configuration, such as the version and type of client operating system, domain controller, Kerberos, and DNS. Documenting all these third-party components of your environment is beyond the scope of this document. For more information, see the respective documentation for each component.

Creating and copying Kerberos configuration files

As part of the Kerberos configuration process, you must create the Kerberos configuration files `krb5.conf` and `keytab`, and copy them to the appropriate locations on the client and storage system.

Steps

1. Configure the `/etc/krb5.conf` file of the storage system and the NFS client with the `krb5.conf` file contents of your Kerberos server.

When you specify the default `keytab` name, use the `FILE:` prefix for the path to ensure that Data ONTAP reads the `keytab` file during authentication and avoid Kerberos authentication failure.

See the documentation for your Kerberos server and clients for specifics about the syntax and options.

Example

```
[libdefaults]
    default_realm = MITKDC01.EXAMPLE.COM
    default_tgs_etypes = des-cbc-crc
    default_tkt_etypes = des-cbc-crc
    dns_lookup_realm = false
    dns_lookup_kdc = false
    default_keytab_name = FILE:/etc/UNIX_krb5.keytab

[realms]
    MITKDC01.EXAMPLE.COM = {
        kdc = sun217-19:88
        admin_server = sun217-19:749
    }

[logging]
    kdc = FILE:/var/log/krb5kdc.log
    admin_server = FILE:/var/log/kadmin.log
    default = FILE:/var/log/krb5lib.log

[domain_realm]
    .sim.example.com = MITKDC01.EXAMPLE.COM
    sim.example.com = MITKDC01.EXAMPLE.COM
    lab.example.com = MITKDC01.EXAMPLE.COM
    .lab.example.com = MITKDC01.EXAMPLE.COM
```

2. On the Kerberos server, create the `keytab` file for the storage system and NFS client.
3. Log in to the Kerberos server as a user that can edit Kerberos and export keys, and then enter the following command:

```
kadmin.local
```

Most Kerberos commands are run from `kadmin`.

Example

It is useful to determine whether the key you are attempting to create already exists. For some Kerberos issues, the key should be re-created. The following steps export the key and install it on the storage controller if a new key is created:

- a. Generate a listing of existing principals:

```
list_principals
```

- b. Delete any unneeded principals:

```
delete_principal
```

- c. Add a new principal for the NFS server:

```
add_principal -e "des-cbc-crc:normal" -randkey nfs/  
f3040-234-52.lab.example.com@MITKDC01.EXAMPLE.COM
```

- d. Create a key for the principal and add it to a temporary keytab file:

```
ktadd -k /tmp/f3040-234-52.keytab -e "des-cbc-crc:normal" nfs/  
f3040-234-52.lab.example.com@MITKDC01.EXAMPLE.COM
```

- e. Add a new principal for the NFS client:

```
add_principal -e "des-cbc-crc:normal" -randkey nfs/  
vmsun217-13.lab.example.com@MITKDC01.EXAMPLE.COM
```

- f. Create a key for the principal and add it to a temporary keytab file:

```
ktadd -k /tmp/vmsun217-13.keytab1 -e "des-cbc-crc:normal" nfs/  
vmsun217-13.lab.example.com@MITKDC01.EXAMPLE.COM
```

- g. Add a new principal for the client root user:

```
add_principal -e "des-cbc-crc:normal" -randkey root/  
vmsun217-13.lab.example.com@MITKDC01.EXAMPLE.COM
```

- h. Create a key for the principal and add it to a temporary keytab file:

```
ktadd -k /tmp/vmsun217-13.keytab2 -e "des-cbc-crc:normal" root/  
vmsun217-13.lab.example.com@MITKDC01.EXAMPLE.COM
```

4. Copy the keytab file from the Kerberos server to the storage system and NFS client.

Example

On the storage system, copy /tmp/f3040-234-52.keytab to /etc/UNIX_krb5.keytab.

On the NFS client, merge the two files and copy them to /etc/krb5.keytab.

```
ktutil: rkt /tmp/vmsun217-13.keytab2  
ktutil: rkt /tmp/vmsun217-13.keytab1  
ktutil: wkt /etc/krb5.keytab
```

5. On the NFS client, enable NFS security by performing the appropriate action:

If the client operating system is...	Then...
Linux	Edit the <code>/etc/sysconfig/nfs</code> file, and add the following line: <div><code>SECURE_NFS=yes</code></div>
Solaris	Uncomment all the <code>krb</code> lines in <code>/etc/nfssec.conf</code> .

6. On the NFS client, restart `rpcgssd`:

If the client operating system is...	Enter the command...
Linux	<code>bash# /etc/init.d/rpcgssd restart''</code>
Solaris	<code>bash# svcadm -v disable svc:/network/rpc/gss:default</code> <code>bash# svcadm -v enable svc:/network/rpc/gss:default</code>

Enabling Kerberos v5 security services for NFS

To enable Kerberos v5 security services for NFS, you can use the `nfs setup` command.

Before you begin

The `nfs setup` command enables you to configure your storage system for a UNIX-based KDC before creating the server principal and `keytab` file. However, you need to create the server principal and `keytab` file before you can use Kerberos.

Steps

1. Enter the following command:

```
nfs setup
```

You receive the following message from `nfs setup`:

```
Enable Kerberos for NFS?
```

2. Enter `y` to continue.

You are asked to specify the type of KDC:

```
The filer supports these types of Kerberos Key Distribution Centers (KDCs):
1 - UNIX KDC
2 - Microsoft Active Directory KDC
Enter the type of your KDC (1-2):
```


3. Enter 1.

If you have not yet set up your server principal file and keytab file, you receive one of several warnings, but the setup process continues.

If you are running `nfs setup` after a fresh installation, you receive the following warning message:

```
There is no /etc/krb5.conf file yet. You will need to establish one.
Unix KDC uses the keytab file /etc/UNIX_krb5.keytab. There is no /etc/
UNIX_krb5.keytab file yet. You will need to establish one.
```

If you are running `nfs setup` after running `cifs setup` (and you configured CIFS to use an Active-Directory-based KDC), you receive the following warning message:

```
There is no /etc/krb5.conf file yet. You will need to establish one.
You have an existing keytab file /etc/krb5.keytab. Your new keytab
file for Unix KDC would be /etc/UNIX_krb5.keytab.
NOTE: If CIFS Active Directory based authentication has been
configured on this filer at any point in the past, the /etc/
krb5.keytab might belong to CIFS. Do you want to rename your existing
keytab file /etc/krb5.keytab to the new keytab file /etc/
UNIX_krb5.keytab.
(Yes/No)? n
Unix KDC uses the keytab file /etc/UNIX_krb5.keytab. There is no /etc/
UNIX_krb5.keytab file yet. You will need to establish one.
```

If you are running `nfs setup` for the first time after upgrading Data ONTAP from a release prior to Data ONTAP 7.3.1, you receive the following warning message:

```
Your new keytab file for Unix KDC would be /etc/UNIX_krb5.keytab.
NOTE: If CIFS Active Directory based authentication has been
configured on this filer at any point in the past, the /etc/
krb5.keytab might belong to CIFS. Do you want to rename your existing
keytab file /etc/krb5.keytab to the new keytab file /etc/
UNIX_krb5.keytab. (Yes/No)? y
/etc/krb5.keytab renamed to /etc/UNIX_krb5.keytab
```

If you respond negatively to either of the last two prompts, `nfs setup` proceeds without renaming the keytab file.

You are prompted to enter the Kerberos realm name.

```
Enter the Kerberos realm name.
```

4. Enter the realm name for the UNIX-based KDC.

The realm name is the realm-specific part of the NFS server's Kerberos principal name (the name you specified for the NFS server principal). For example, `MY_COMPANY.COM`. The realm

name you enter can be verified or modified later by changing the value of the `nfs.kerberos.realm` option:

```
options nfs.kerberos.realm realm_name
```

Note: Data ONTAP supports lowercase realm names for UNIX-based KDCs but not for Active Directory KDCs.

Example

The following command specifies the Kerberos realm name `LAB.MY_COMPANY.COM`:

```
options nfs.kerberos.realm LAB.MY_COMPANY.COM
```

You are prompted to enter a host instance.

```
Enter the host instance of the NFS server principal name [default:
server.lab.my_company.com]:
```

5. Enter a host instance.

If DNS is enabled, it is used to verify that you have entered a fully qualified domain name for your host. If you entered a partial name and your host was entered in DNS, the missing domain information is appended to your entry.

The host instance you enter can be verified by using the `nfs.kerberos.principal` option:

```
options nfs.kerberos.principal
```

The `nfs setup` command uses your entries for the host instance and realm name to identify the Kerberos principal. The principal is derived from `nfs setup` entries as follows:

```
nfs/value from nfs.kerberos.principal@value from nfs.kerberos.realm
```

After you enter the host instance and exit `nfs setup`, the storage system is configured to use the keytab file you generated. You can modify this configuration later by running `nfs setup` again.

Supported client and user principal formats

Client and user principals must use certain formats supported by Data ONTAP to ensure proper Kerberos functionality.

Client principals should use the format `root/client_host_fqdn@Kerberos_realm`.

Note: If you encounter a client that does not function properly with this client principal format, you can use the format `nfs/client_host_fqdn@Kerberos_realm` and create an `nfs` user for mapping to a valid UID/GID.

User principals should use the format `user_name@Kerberos_realm` or `user_name/host_fqdn@Kerberos_realm`. Data ONTAP supports both formats. In both cases, Data ONTAP uses the `user_name` portion of the principal as the user name to associate with the NFS request. This user name is then mapped to a UID/GID.

Debugging mounting problems

To debug mounting problems, you can display mount service statistics and trace mountd requests.

Displaying mount service statistics

To display mount service statistics, you can enter the `nfsstat -d` command.

Step

1. Enter the following command:

```
nfsstat -d
```

Result

Data ONTAP displays the following mount service statistics:

```
v2 mount (requested, granted, denied, resolving)
v2 unmount (requested, granted, denied)
v2 unmount all (requested, granted, denied)
v3 mount (requested, granted, denied, resolving)
v3 unmount (requested, granted, denied)
v3 unmount all (requested, granted, denied)
mount service requests (curr, total, max, redriven)
```

For more information, see the `na_nfsstat(1)` man page.

Tracing mountd requests

To trace mountd requests, you can add a `*.debug` entry to the `/etc/syslog.conf` file and set the `nfs.mountd.trace` option to on.

About this task

Because there is a possibility that the syslog will get hit numerous times during DOS attacks, this option should be enabled only during a debug session.

By default, the `nfs.mountd.trace` option is off.

Steps

1. Edit the `/etc/syslog.conf` file and add a `*.debug` entry.

For more information about adding an entry to the `syslog.conf` file, see the `na_syslog.conf(5)` man page.

2. To enable the `nfs.mountd.trace` option, enter the following command:

```
options nfs.mountd.trace on
```

For more information about the `nfs.mountd.trace` option, see the `na_options(1)` man page.

Use of hard mounts

When troubleshooting mounting problems, you need to be sure that you are using the correct mount type. NFS supports two mount types: soft mounts and hard mounts. You should use only hard mounts for reliability reasons.

You should not use soft mounts, especially when there is a possibility of frequent NFS timeouts. Race conditions can occur as a result of these timeouts, which can lead to data corruption.

Displaying NFS statistics

You can display NFS statistics for the storage system to monitor performance and diagnose issues.

Step

1. To display NFS statistics, enter the following command:

```
nfsstat
```

You can use the `nfsstat` command to display NFS statistics for all clients. Or, if the `nfs.per_client_stats.enable` option is set to `on`, you can use the `nfsstat -h` or `nfsstat -l` commands to display NFS statistics on a per-client basis.

In addition to displaying NFS statistics, you can use the `nfsstat` command to reset NFS statistics.

For more information, see the `na_nfsstat(1)` man page and the following topics:

- “Displaying mount service statistics”
- “Displaying NFSv4 file delegation statistics”

Controlling NFS requests from nonreserved ports

You can reject NFS mount requests from nonreserved ports by enabling the `nfs.mount_rootonly` option. To reject all NFS requests from nonreserved ports, you can enable the `nfs.nfs_rootonly` option.

About this task

By default, the option `nfs.mount_rootonly` is `on`.

By default, the option `nfs.nfs_rootonly` is `off`.

These options do not apply to the NULL procedure.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Allow NFS mount requests from nonreserved ports	<code>options nfs.mount_rootonly off</code>
Reject NFS mount requests from nonreserved ports	<code>options nfs.mount_rootonly on</code>
Allow all NFS requests from nonreserved ports	<code>options nfs.nfs_rootonly off</code>
Reject all NFS requests from nonreserved ports	<code>options nfs.nfs_rootonly on</code>

Enabling or disabling NFSv2

You can enable or disable NFSv2 by modifying the `nfs.v2.enable` option. This allows file access for clients using the NFSv2 protocol. By default, NFSv2 is enabled.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Enable NFSv2	<code>options nfs.v2.enable on</code>
Disable NFSv2	<code>options nfs.v2.enable off</code>

Enabling or disabling NFSv3

You can enable or disable NFSv3 by modifying the `nfs.v3.enable` option. This allows file access for clients using the NFSv3 protocol. By default, NFSv3 is enabled.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Enable NFSv3	<code>options nfs.v3.enable on</code>
Disable NFSv3	<code>options nfs.v3.enable off</code>

Support for VMware vStorage over NFS

Data ONTAP supports certain VMware vStorage APIs for Array Integration (VAAI) features in an NFS environment.

Supported features

The following features are supported:

- **copy offload**
Enables an ESXi host to copy virtual machines or virtual machine disks (VMDKs) directly between the source and destination data store location without involving the host. This conserves ESXi host CPU cycles and network bandwidth. Copy offload preserves space efficiency if the source volume is sparse.
- **space reservation**
Guarantees storage space for a VMDK file by reserving space for it.

Limitations

VMware vStorage over NFS has the following limitations:

- vStorage is not supported with FlexCache.
- Copy offload is not supported across vFiler units and across storage systems.
- vStorage is not supported over IPv6.
- Copy offload operations can fail in the following scenarios:
 - While running wafliro on the source or destination volume because it temporarily takes the volume offline
 - While moving either the source or destination volume
 - While performing takeover or giveback operations

Enabling or disabling VMware vStorage over NFS

You can enable or disable support for VMware vStorage over NFS by modifying the `nfs.vstorage.enable` option.

Before you begin

NFSv3 must be enabled on the storage system.

About this task

By default, support for VMware vStorage over NFS is disabled.

Step

1. Perform one of the following actions:

If you want to...	Enter the following command...
Enable VMware vStorage support	<code>options nfs.vstorage.enable on</code>
Disable VMware vStorage support	<code>options nfs.vstorage.enable off</code>

After you finish

You must install the NFS Plug-in for VMware VAAI before you can use this functionality. For more information, see *Installing the NetApp NFS Plug-in for VMware VAAI*.

Displaying statistics for VMware vStorage over NFS

You can use the `nfs vstorage stats` command to display various counters about VMware vStorage over NFS activity on the storage system. This information can be useful for troubleshooting purposes.

Step

1. Enter the following command to display vStorage statistics:

```
nfs vstorage stats
```

Differences in file system ID handling for NFSv3 and NFSv4

You can configure Data ONTAP to either return the same or a different file system ID (FSID) for `.snapshot` subdirectories and files as for the active file system in NFSv3 and NFSv4.

When you mount an exported path and get a directory listing of the `.snapshot` directory and subdirectories, the returned file and directory attributes include the FSID. The FSID of the `.snapshot` subdirectories should be the same as the FSID of the active file system. The FSID of the `.snapshot` subdirectories vary depending on the following two options.

If you...	Then...
Enable the option <code>nfs.v3.snapshot.active.fsid.enable</code>	For NFSv3 requests, the FSID returned for directories and files within the <code>.snapshot</code> directory is the same as the FSID of the active file system.
Disable the option <code>nfs.v3.snapshot.active.fsid.enable</code>	For NFSv3 requests, the FSID returned for directories and files within the <code>.snapshot</code> directory are different from the FSID of the active file system.

If you...	Then...
Enable the option <code>nfs.v4.snapshot.active.fsid.enable</code>	For NFSv4 requests, the FSID returned for directories and files within the <code>.snapshot</code> directory is the same as the FSID of the active file system.
Disable the option <code>nfs.v4.snapshot.active.fsid.enable</code>	For NFSv4 requests, the FSID returned for directories and files within the <code>.snapshot</code> directory is different from the FSID of the active file system.

Supporting NFSv4 clients

Supporting NFSv4 clients involves enabling or disabling the NFSv4 protocol, specifying an NFSv4 user ID domain, managing NFSv4 ACLS and file delegation, and configuring file and record locking.

NFSv4.0 functionality supported by Data ONTAP

Data ONTAP supports all the mandatory functionality in NFSv4.0 except the SPKM3 and LIPKEY security mechanisms.

The following NFSV4 functionality is supported:

- COMPOUND** Allows a client to request multiple file operations in a single remote procedure call (RPC) request.
- File delegation** Allows the server to delegate file control to some types of clients for read and write access.
- Pseudo-fs** Used by NFSv4 servers to determine mount points on the storage system. There is no mount protocol in NFSv4.
- Locking** Lease-based. There are no separate Network Lock Manager (NLM) or Network Status Monitor (NSM) protocols in NFSv4.
- Named attributes** Similar to Windows NT streams.

For more information about the NFSv4.0 protocol, see RFC 3530.

Limitations of Data ONTAP support for NFSv4

You should be aware of several limitations of Data ONTAP support for NFSv4.

- The SPKM3 and LIPKEY security mechanisms are not supported.
- The delegation feature is not supported by every client type.
- Names with non-ASCII characters on volumes other than UTF8 volumes are rejected by the storage system.
- All file handles are persistent; the server does not give volatile file handles.

- Migration, referrals, and replication are not supported.
- NFSv4 clients are not supported with read-only load-sharing mirrors.
Data ONTAP routes NFSv4 clients to the source of the load-sharing mirror for direct read and write access.
- All recommended attributes are supported, except for the following:
 - `archive`
 - `hidden`
 - `homogeneous`
 - `mimetype`
 - `quota_avail_hard`
 - `quota_avail_soft`
 - `quota_used`
 - `system`
 - `time_backup`

Note: Although it does not support the `quota*` attributes, Data ONTAP does support user and group quotas through the RQUOTA side band protocol.

How the pseudo-fs in NFSv4 affects mountpoints

NFSv4 uses a pseudo-fs (file system) as an entry point into your storage system for determining mountpoints. A pseudo-fs allows you to use one port for security, rather than several. All NFSv4 servers support the use of a pseudo-fs.

Because of the pseudo-fs used in NFSv4, you might experience inconsistencies with mountpoints between NFSv3 and NFSv4,

In the examples that follow, you have these volumes:

- `/vol/vol0` (root)
- `/vol/vol1`
- `/vol/home`

Example 1

In NFSv3 if you do not use the complete path from `/vol/vol0`, and you mount `filer:/`, the mountpoint is `filer:/vol/vol0`. That is, if the path does not begin with `/vol` in NFSv3, Data ONTAP adds `/vol/vol0` to the beginning of the path.

In NFSv4, if you do not use the complete path from `/vol/vol0` and you mount `filer:/`, you mount the root of the pseudo-fs and not `/vol/vol0`. Data ONTAP does not add `/vol/vol0` to the beginning of the path.

Therefore, if you mount `filer:/n/filer` using NFSv3 and try the same mount using NFSv4, you would mount a different file system.

Example 2

In the Data ONTAP implementation of the NFSv4 pseudo-fs, the nodes `/` and `/vol` are always present and form the common prefix of any reference into the pseudo-fs. Any reference that does not begin with `/vol` is invalid.

In this example, there is a `/vol/vol0/home` directory. In NFSv3, if you mount `filer:/home/users`, `/home` is considered as the directory `/vol/vol0/home`. In NFSv4, if you mount `filer:/home/users`, `/home` is not interpreted as the volume `/vol/home`; it is considered an invalid path in the pseudo-fs tree.

Enabling or disabling NFSv4.0

You can enable or disable NFSv4.0 by modifying the `nfs.v4.enable` option. This allows file access for clients using the NFSv4.0 protocol. By default, NFSv4.0 is disabled.

Step

1. Perform one of the following actions:

If you want to...	Enter the following command...
Enable NFSv4.0	<code>options nfs.v4.enable on</code>
Disable NFSv4.0	<code>options nfs.v4.enable off</code>

Specifying the user ID domain for NFSv4

To specify the user ID domain, you can set the `nfs.v4.id_domain` option.

About this task

By default, Data ONTAP uses the NIS domain for NFSv4 user ID mapping, if one is set. If an NIS domain is not set, the DNS domain is used. You might need to set the user ID domain if, for example, you have multiple user ID domains. The domain name must match the domain configuration on the domain controller. It is not required for NFSv3.

Step

1. Enter the following command:

```
options nfs.v4.id_domain domain
```

Managing NFSv4 ACLs

You can enable, disable, set, modify, and view NFSv4 access control lists (ACLs).

How NFSv4 ACLs work

A client using NFSv4 ACLs can set and view ACLs on files and directories on the system. When a new file or subdirectory is created in a directory that has an ACL, the new file or subdirectory

inherits all ACL Entries (ACEs) in the ACL that have been tagged with the appropriate inheritance flags.

For access checking, CIFS users are mapped to UNIX users. The mapped UNIX user and that user's group membership are checked against the ACL.

If a file or directory has an ACL, that ACL is used to control access no matter what protocol—NFSv3, NFSv4, or CIFS—is used to access the file or directory and is used even if NFSv4 is no longer enabled on the system.

Files and directories inherit ACEs from NFSv4 ACLs on parent directories (possibly with appropriate modifications) as long as the ACEs have been tagged with the appropriate inheritance flags.

The default limit of ACEs for each ACL is 192.

When a file or directory is created as the result of an NFSv4 request, the ACL on the resulting file or directory depends on whether the file creation request includes an ACL or only standard UNIX file access permissions, and whether the parent directory has an ACL:

- If the request includes an ACL, that ACL is used.
- If the request includes only standard UNIX file access permissions but the parent directory has an ACL, the ACEs in the parent directory's ACL are inherited by the new file or directory as long as the ACEs have been tagged with the appropriate inheritance flags.

Note: A parent ACL is inherited even if `nfs.v4.acl.enable` is set to `off`.

- If the request includes only standard UNIX file access permissions and the parent directory does not have an ACL, the client file mode is used to set standard UNIX file access permissions.
- If the request includes only standard UNIX file access permissions and the parent directory has a non-inheritable ACL, the new object is created only with mode bits.

The security semantics of a qtree are determined by its security style and its ACL (NFSv4 or NTFS):

For a qtree with UNIX security style:

- NFSv4 ACLs and mode bits are effective.
- NTFS ACLs are not effective.
- Windows clients cannot set attributes.

For a qtree with NTFS security style:

- NFSv4 ACLs are not effective.
- NTFS ACLs and mode bits are effective.
- UNIX clients cannot set attributes.

For a qtree with mixed security style:

- NFSv4 ACLs and mode bits are effective.
- NTFS ACLs are effective.
- Both Windows and UNIX clients can set attributes.

Note: Files and directories in a qtree can have either an NFSv4 ACL or an NTFS ACL, but not both. Data ONTAP remaps one type to the other, as necessary.

Benefits of enabling NFSv4 ACLs

There are many benefits to enabling NFSv4 ACLs.

The benefits of enabling NFSv4 ACLs include the following:

- Finer-grained control of user access for files and directories
- Better NFS security
- Improved interoperability with CIFS
- Removal of the NFS limitation of 16 groups per user

Compatibility between NFSv4 ACLs and Windows (NTFS) ACLs

NFSv4 ACLs are different from Windows file-level ACLs (NTFS ACLs), but Data ONTAP can map NFSv4 ACLs to Windows ACLs for viewing on Windows platforms.

Permissions displayed to NFS clients for files that have Windows ACLs are "display" permissions, and the permissions used for checking file access are those of the Windows ACL.

Note: Data ONTAP does not support POSIX ACLs.

How Data ONTAP uses NFSv4 ACLs to determine whether it can delete a file

To determine whether it can delete a file, Data ONTAP uses a combination of the file's DELETE bit, and the containing directory's DELETE_CHILD bit. For more information, see the NFS 4.1 RFC 5661.

Enabling or disabling NFSv4 ACLs

To enable or disable NFSv4 ACLs, you can modify the `nfs.v4.acl.enable` option. This option is disabled by default.

About this task

The `nfs.v4.acl.enable` option controls the setting and viewing of NFSv4 ACLs; it does not control enforcement of these ACLs for access checking. For more information, see the `na_options(1)` man page.

Step

1. Perform one of the following actions:

If you want to...	Then...
Enable NFSv4.0 ACLs	Enter the following command: <code>options nfs.v4.acl.enable on</code>

If you want to...	Then...
Disable NFSv4.0 ACLs	Enter the following command: options nfs.v4.acl.enable off

Managing NFSv4 file delegations

You can enable and disable NFSv4 file delegations and retrieve NFSv4 file delegation statistics.

How NFSv4 file delegations work

Data ONTAP supports read and write file delegations in accordance with RFC 3530.

As specified in RFC 3530, when an NFSv4 client opens a file, Data ONTAP can delegate further handling of opening and writing requests to the opening client. There are two types of delegations: read and write. A read file delegation allows a client to handle requests to open a file for reading that do not deny read access to others. A write file delegation allows the client to handle all open requests.

Delegation works on files within any style of qtree, whether or not opportunistic locks (oplocks) have been enabled.

Delegation of file operations to a client can be recalled when the lease expires, or when the storage system receives the following requests from another client:

- Write to file, open file for writing, or open file for “deny read”
- Change file attributes
- Rename file
- Delete file

When a lease expires, the delegation state is revoked and all of the associated states are marked “soft”. This means that if the storage system receives a conflicting lock request for this same file from another client before the lease has been renewed by the client previously holding the delegation, the conflicting lock is granted. If there is no conflicting lock and the client holding the delegation renews the lease, the soft locks are changed to hard locks and are not removed in the case of a conflicting access. However, the delegation is not granted again upon a lease renewal.

When the server reboots, the delegation state is lost. Clients can reclaim the delegation state upon reconnection instead of going through the entire delegation request process again. When a client holding a read delegation reboots, all delegation state information is flushed from the storage system cache upon reconnection. The client must issue a delegation request to establish a new delegation.

Enabling or disabling NFSv4 read file delegations

To enable or disable NFSv4 read file delegations, you can modify the `nfs.v4.read_delegation` option. By enabling read file delegations, you can eliminate much of the message overhead associated with the opening and closing of files.

About this task

By default, read file delegations are disabled.

The disadvantage of enabling read file delegations is that the server and its clients must recover delegations after the server reboots or restarts, a client reboots or restarts, or a network partition occurs.

Step

- 1. Perform one of the following actions:

If you want to...	Then...
Enable NFSv4 read file delegations	Enter the following command: <code>options nfs.v4.read_delegation on</code>
Disable NFSv4 read file delegations	Enter the following command: <code>options nfs.v4.read_delegation off</code>

Result

The file delegation options take effect as soon as they are changed. There is no need to reboot or restart NFS.

Enabling or disabling NFSv4 write file delegations

To enable or disable write file delegations, you can modify the `nfs.v4.write_delegation` option. By enabling write file delegations, you can eliminate much of the message overhead associated with file and record locking in addition to opening and closing of files.

About this task

By default, write file delegations are disabled.

The disadvantage of enabling write file delegations is that the server and its clients must perform additional tasks to recover delegations after the server reboots or restarts, a client reboots or restarts, or a network partition occurs.

Step

- 1. Perform one of the following actions:

If you want to...	Then...
Enable NFSv4 write file delegations	Enter the following command: options nfs.v4.write_delegation on
Disable NFSv4 write file delegations	Enter the following command: options nfs.v4.write_delegation off

Result

The file delegation options take effect as soon as they are changed. There is no need to reboot or restart NFS.

Displaying NFSv4 open delegation statistics

To display information about NFSv4 open delegation requests, you can use the `nfsstat` command.

Step

1. Use the `nfsstat` command to display open delegation statistics.

Result

Results returned by the `nfsstat` command include open delegation requests that have been granted as well as requests that have been denied due to an error.

After you finish

For information about open delegation requests that your storage system has denied, view the system log file.

Displaying NFSv4 open delegation statistics for all clients

To display NFSv4 open delegation information for all clients, you can enter the `nfsstat -l` command.

Step

1. Enter the following command:

```
nfsstat -l count
```

Result

The storage system returns individual NFSv4 open delegation statistics for each client up to the count you specify. If you do not specify a count, the storage system returns statistics for the first 256 clients in order of the total NFS operations performed by each client.

Displaying NFSv4 open delegation statistics for a specific client

To display NFSv4 open delegation information for a specific client, you can use the `nfsstat -h` command.

Step

1. Enter the following command:

```
nfsstat -h hostname or ip_address
```

Result

The storage system returns individual NFSv4 open delegation statistics for the specified client.

Displaying NFSv4 open delegation statistics for a vFiler unit

To display NFSv4 open delegation statistics for a vFiler unit, you can run the `nfsstat -d` command in the vFiler unit's context.

Step

1. Enter the following command:

```
vfiler run filename nfsstat -d
```

Displaying NFSv4 open delegation statistics for a storage system

To display NFSv4 open delegation information for a storage system, you can enter the `nfsstat -d` command.

Step

1. Enter the following command:

```
nfsstat -d
```

Result

The storage system returns the total number of NFSv4 open delegations handled by the storage system, including current NFSv4 open delegations and any that have been recalled. To view only current NFSv4 open delegations handled by the storage system, use the `lock status` command.

Configuring NFSv4 file and record locking

You can configure NFSv4 file and record locking by specifying the locking lease period and grace period.

About NFSv4 file and record locking

For NFSv4 clients, Data ONTAP supports the NFSv4 file-locking mechanism, maintaining the state of all file locks under a lease-based model.

In accordance with RFC 3530, Data ONTAP “defines a single lease period for all state held by an NFS client. If the client does not renew its lease within the defined period, all states associated with the client's lease may be released by the server.” The client can renew its lease explicitly or implicitly by performing an operation, such as reading a file.

Furthermore, Data ONTAP defines a grace period, which is a period of special processing in which clients attempt to reclaim their locking state during a server recovery.

Term	Definition (see RFC 3530)
Lease	The time period in which Data ONTAP irrevocably grants a lock to a client.
Grace period	The time period in which clients attempt to reclaim their locking state from Data ONTAP during server recovery.
Lock	Refers to both record (byte-range) locks as well as file (share) locks unless specifically stated otherwise.

Specifying the NFSv4 locking lease period

To specify the NFSv4 locking lease period (that is, the time period in which Data ONTAP irrevocably grants a lock to a client), you can modify the `nfs.v4.lease_seconds` option. Shorter lease periods speed up server recovery while longer lease periods are beneficial for servers handling a very large amount of clients.

About this task

By default, this option is set to 30. The minimum value for this option is 10. The maximum value for this option is the locking grace period, which you can set with the `locking.lease_seconds` option.

Step

1. Enter the following command:

```
options nfs.v4.lease_seconds number_of_seconds
```

Specifying the NFSv4 locking grace period

To specify the NFSv4 locking grace period (that is, the time period in which clients attempt to reclaim their locking state from Data ONTAP during server recovery), you can modify the `locking.grace_lease_seconds` option.

About this task

By default, this option is set to 45.

Note that this option specifies both the locking lease period and the grace period.

Step

1. Enter the following command:

```
options locking.grace_lease_seconds number_of_seconds
```

How the name server database cache works

For performance optimization, the name server database cache (NSDB) stores the results of external name server lookup operations.

If the environment uses NFSv4 or RPCSEC_GSS with Kerberos for authentication to the storage system, string names are used as user/group identifiers. These string names have to be translated to their appropriate UNIX credentials in the form of UID/GIDs. This translation uses external name server lookup.

To save time on name server lookup requests, Data ONTAP stores the results of previous external name server lookup operations in the NSDB. If a name server lookup request matches a previous one, Data ONTAP can retrieve the information quickly from the NSDB instead of having to communicate with the external name server, which takes longer.

If you have modified user credential information in the name service, the change might not take effect immediately. This can happen if the NSDB has cached user information from before the change and the cached entry has not expired yet. Flushing the NSDB cache removes the outdated information.

Flushing the name server database cache

You can use the `nfs nsdb flush` command to clear specific entries or all entries from the name server database (NSDB) cache. This removes outdated information from the cache after you made changes.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Flush all entries	<code>nfs nsdb flush -a</code>
Flush entries specified by user name	<code>nfs nsdb flush -U <i>username1[,username2,...]</i></code>
Flush entries specified by user ID	<code>nfs nsdb flush -u <i>userID1[,userID2,...]</i></code>
Flush entries specified by group name	<code>nfs nsdb flush -G <i>groupname1[,groupname2,...]</i></code>
Flush entries specified by group ID	<code>nfs nsdb flush -g <i>groupID1[,groupID2,...]</i></code>

Differences in NFS user and group ID formats

NFSv3 handles user and group IDs as 32-bit numeric values. NFSv4 handles user and group IDs by default as strings in the form of *user_name@domain_name*. However, NFSv4 clients can also specify user and group IDs as numeric strings, which Data ONTAP can handle two different ways.

The Data ONTAP NFSv4 implementation expects clients to send the user and group ID as strings and handles them in accordance with the NFSv4 standards. However, if an NFSv4 client specifies a user or group ID as a numeric string, Data ONTAP can handle the numeric string two different ways:

- Accept the numeric string format and map it to a local numeric ID.
If there is no match, Data ONTAP maps it to the user *nobody* (user ID 65534).
- Reject the numeric string format.
In this case, Data ONTAP returns the error `NFSV4_BADOWNER` to the client and maps it to the user *nobody* (user ID 65534) instead.

The way Data ONTAP handles numeric strings depends on how you configure the option `nfs.v4.id.allow_numerics`.

Allowing or disallowing NFSv4 user and group IDs as numeric strings

You can choose whether you want Data ONTAP to allow NFSv4 clients to specify user and group IDs as numeric strings. You control this behavior by modifying the `nfs.v4.id.allow_numerics` option.

About this task

This only applies when the authentication type is `AUTH_SYS`.

By default, this option is `off`.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Allow NFSv4 clients to specify user and group IDs as numeric strings	<code>options nfs.v4.id.allow_numerics on</code>
Disallow NFSv4 clients to specify user and group IDs as numeric strings	<code>options nfs.v4.id.allow_numerics off</code>

Supporting PC-NFS clients

To support PC-NFS clients, you can enable the `pcnfsd` daemon, create PC-NFS user entries in the storage system's local files, and define the `umask` for files and directories that PC-NFS users create on the storage system.

How the `pcnfsd` daemon works

Data ONTAP's `pcnfsd` daemon provides authentication services for clients using PC-NFS version 1 or 2. Authenticated PC-NFS users can mount file system paths on your storage system just like NFS users. The `pcnfsd` daemon does not support printer service.

When the `pcnfsd` daemon receives an authentication request, it can use local files or NIS maps to validate the user's password. The local file used can be the `/etc/shadow` file or `/etc/passwd` file. The NIS maps used can be `passwd.adjunct` or `passwd.byname`. When the shadow source is available, Data ONTAP uses it. The shadow source contains encrypted user information, instead of the password database.

The following list describes how the `pcnfsd` daemon uses local files or NIS maps for authenticating both PC-NFS version 1 and version 2 users:

If a shadow source is available, Data ONTAP uses the `/etc/shadow` file or the `passwd.adjunct` NIS map to determine the user's password.

If a shadow source is not available, Data ONTAP uses the `/etc/passwd` file or the `passwd.byname` NIS map to determine the user's user ID (UID), primary group ID (GID), and password.

When the `pcnfsd` daemon receives a PC-NFS version 2 authentication request, it looks up the `/etc/group` file or the `group.byname` NIS map to determine all the groups to which the user belongs.

Enabling or disabling the `pcnfsd` daemon

To enable or disable the `pcnfsd` daemon, you can set the `pcnfsd.enable` option to `on` or `off`, respectively.

Before you begin

NFS must be enabled on the storage system before you can enable the `pcnfsd` daemon.

About this task

You must enable the `pcnfsd` daemon if you want the storage system to authenticate PC-NFS users when they try to mount file system paths on the storage system. If you want another computer to authenticate users, you do not need to enable the `pcnfsd` daemon. Users authenticated by other computers can access file system paths on the storage system just like users authenticated by the storage system.

Step

1. Perform one of the following actions.

If you want to...	Then...
Enable the <code>pcnfsd</code> daemon	Enter the following command: <code>options pcnfsd.enable on</code>
Disable the <code>pcnfsd</code> daemon	Enter the following command: <code>options pcnfsd.enable off</code>

Creating PC-NFS user entries in the storage system's local files

To create PC-NFS user entries in the storage system's local files, you can copy the `/etc/passwd`, `/etc/shadow`, and `/etc/group` files to the storage system from a UNIX host that properly authenticates all of the PC-NFS users.

About this task

You must create PC-NFS user entries in the storage system's local files if you want to use local files to authenticate PC-NFS users and determine group membership.

Step

1. Copy the `/etc/passwd`, `/etc/shadow`, and `/etc/group` files to the storage system from a UNIX host that properly authenticates all of the PC-NFS users.

How umask works with NFS file permissions

Before defining the `umask` for files and directories, you need to understand how the `umask` is used to calculate file permissions.

The permissions for each file are defined by three octal values, which apply to owner, group, and other. When a PC-NFS client creates a new file, Data ONTAP subtracts the `umask`, which is a three-digit octal number you define, from 666. The resulting octal digits are used as file permissions.

By default, the `umask` is 022, which means that the effective octal digits for permissions are 644. These permissions enable the file owner to read and write the file, and enable the group and others to read the file.

The following table provides the description for each digit in the umask.

Digit in the umask	Description
0	Read and write permissions
2	Write permission
4	Read-only permission
6	No permission

Defining the umask for files and directories that PC-NFS users create

Unlike NFS users, PC-NFS users cannot execute the UNIX `umask` command to set the file mode creation mask (umask), which determines the default file permissions. However, Data ONTAP enables you to define the umask for all PC-NFS users by setting the `pcnfsd.umask` option.

Step

1. Enter the following command:

```
options pcnfsd.umask umask
```

umask is a three-digit octal number.

Supporting WebNFS clients

To support WebNFS clients, you can enable the WebNFS protocol and, optionally, set the WebNFS root directory.

About this task

If you enable the WebNFS protocol, WebNFS client users can specify a URL starting with `nfs://` to transfer a file from the storage system.

Enabling or disabling the WebNFS protocol

To enable or disable the WebNFS protocol, you can set the `nfs.webnfs.enable` option to `on` or `off` respectively.

Step

1. Perform one of the following actions.

If you want the Web NFS protocol...	Then...
Enabled	Enter the following command: options nfs.webnfs.enable on
Disabled	Enter the following command: options nfs.webnfs.enable off

Setting a WebNFS root directory

To set a WebNFS root directory, you can specify the name of the root directory; then you can enable the root directory.

About this task

If you set a root directory for WebNFS lookup, a WebNFS user can specify only the path name relative to the root directory instead of the absolute path name. For example, if the WebNFS root directory is `/vol/vol1/web`, a WebNFS user can access the `/vol/vol1/web/specs` file by specifying `nfs://specs` as the URL.

Specifying the name of the WebNFS root directory

You can set the `nfs.webnfs.rootdir` option to specify the name of the WebNFS root directory.

Step

1. Enter the following command:

```
options nfs.webnfs.rootdir directory
```

directory specifies the full path to the root directory.

Enabling the WebNFS root directory

To enable the WebNFS root directory, you can set the `nfs.webnfs.rootdir.set` option to on.

Before you begin

You must specify the name of the WebNFS root directory before you enable it.

Step

1. Enter the following command:

```
options nfs.webnfs.rootdir.set on
```

Support for NFS over IPv6

Starting with Data ONTAP 7.3.1, NFS clients can access files on your storage system over an IPv6 network.

Enabling or disabling NFS over IPv6

You can enable or disable NFS over IPv6 by setting the `nfs.ipv6.enable` option to `on` or `off`, respectively.

Before you begin

You must enable IPv6 on the storage system by setting the `ip.v6.enable` option to `on`. For more information about enabling IPv6 on your storage system, see the *Data ONTAP Network Management Guide for 7-Mode*.

About this task

If you have enabled NFS over IPv6 and you then disable IPv6 on your storage system by setting the `ip.v6.enable` option to `off`, NFS is automatically disabled over IPv6.

Steps

1. Perform one of the following actions:

If you want NFS over IPv6 to be...	Enter the command...
Enabled	<code>options nfs.ipv6.enable on</code>
Disabled	<code>options nfs.ipv6.enable off</code>

2. Restart NFS by entering the following commands:

```
nfs off  
  
nfs on
```

Textual representation of IPv6 addresses

According to RFC 3513, the preferred form for representing IPv6 addresses as text strings is `x:x:x:x:x:x:x`, where the `x`'s are “the hexadecimal values of the eight 16-bit pieces of the addresses.”

Examples

- `FEDC:BA98:7654:3210:FEDC:BA98:7654:3210`
- `1080:0:0:0:8:800:200C:417A`

For more information, search the web for RFC 3513.

File access using CIFS

You can enable and configure a CIFS server to let CIFS clients access files on your storage system.

Unsupported Windows features

Before you use CIFS in your network, you need to be aware of certain Windows features that Data ONTAP does not support.

Data ONTAP does not support the following Windows features:

- Encrypted File System (EFS)
- Logging of NT File System (NTFS) events in the change journal
- Microsoft File Replication Service (FRS)
- Microsoft Windows Indexing Service
- Remote storage through Hierarchical Storage Management (HSM)
- Local user account creation from the User Manager or Microsoft Manager Console
- Quota management from Windows clients
- Windows quota semantics
- The LMHOSTS file
- NTFS native compression

Configuring CIFS licenses

To be able to set up and use CIFS functionality, your storage system must have a CIFS license installed. You can use the `license` command to configure licenses on the storage system.

Before you begin

You need to be sure that you have obtained a valid CIFS license from your sales or support representative.

Steps

1. Display the features that are currently licensed on the storage system by entering the following command:

`license`
2. Perform one of the following actions:

If the CIFS feature...	Then...
Displays a license code	Your storage system is already licensed for CIFS and you do not need to take any further action.
Displays <div>not licensed</div>	Install a CIFS license by entering the following command: <pre>license add license_code</pre>

Connecting the MMC to the storage system

Certain CIFS management tasks can be performed using the MMC. Before performing these tasks, you need to connect the MMC to the storage system. You can connect the MMC to the storage system using the MMC menu commands.

Steps

1. To open the MMC on your Windows server, in Windows Explorer, right-click the icon for the local computer and select **Manage**.
2. On the left panel, select **Computer Management**.
3. Select **Action > Connect to another computer**.
The Select Computer dialog box appears.
4. Type the name of the storage system or click **Browse** to locate the storage system.
5. Click **OK**.

Configuring CIFS on your storage system

You can use the `cifs setup` command to configure CIFS on your storage system. If you are setting up your storage system for the first time, you should see the *Data ONTAP Software Setup Guide*.

Supported CIFS clients and domain controllers

Before you can use CIFS with your storage system, you need to know which CIFS clients and domain controllers Data ONTAP supports.

For the latest information about which CIFS clients and domain controllers Data ONTAP supports, see the Interoperability Matrix at support.netapp.com/NOW/products/interoperability.

What the cifs setup command does

In addition to performing initial CIFS configuration, the `cifs setup` command enables you to perform several tasks.

With the `cifs setup` command, you can perform the following tasks:

- Create and name a CIFS server that your CIFS clients can access
- Join the CIFS server to a domain or workgroup, or move between them
- Create a default set of local CIFS users and groups

Requirements for the CIFS server name

By default, the `cifs setup` command assigns the storage system's host name as the CIFS server name. You can accept the default or specify a different name. If you specify a different CIFS server name, the name must conform to Microsoft Active Directory naming conventions.

Because the storage system's CIFS server name must conform to Microsoft Active Directory naming conventions, the name cannot be longer than 15 characters. If the storage system's host name is longer than 15 characters, you must enter a different name for the CIFS server name. If you specify a CIFS server name during CIFS setup, it must be restricted to a name that is up to 15 characters long to continue with CIFS setup.

You can ensure that the CIFS server name conforms to all other Microsoft Active Directory naming conventions by referring to the following Microsoft KB: [Microsoft KB 909264: support.microsoft.com/kb/909264](https://support.microsoft.com/kb/909264)

Setting up your system initially

When a valid CIFS license is present, Data ONTAP automatically invokes the `cifs setup` command during the initial setup of your storage system. The `cifs setup` command prompts you for information such as authentication type, lookup services to be used, and so forth.

To learn about using the `cifs setup` command for initial CIFS configuration, including a list of the information you need when running `cifs setup`, see the *Data ONTAP Software Setup Guide for 7-Mode*.

Specifying WINS servers

To specify WINS servers, you can use the `cifs.wins_servers` option, which is nondisruptive, or the `cifs setup` command, which requires you to halt CIFS services.

About this task

The WINS server list is not additive—if you are adding a third WINS server, you must enter all three IP addresses in a comma-separated list, or your existing two WINS servers are replaced by the server you intended to add.

Step

1. Perform one of the following actions.

If you want to...	Then...
Specify WINS servers using the <code>cifs.wins_servers</code> option	Enter the following command: <code>options cifs.wins_servers servers</code> <i>servers</i> is a comma-delimited list of WINS servers. For more information about the <code>cifs.wins_servers</code> option, see the <code>options(1)</code> man page.
Specify WINS servers using the <code>cifs setup</code> command	Enter the following command: <code>cifs setup</code> Then, when prompted, specify up to four IPv4 WINS servers. For more information about the <code>cifs setup</code> command, see the <code>cifs(1)</code> man page.

Changing the storage system domain

If you already configured your storage system for Windows Domain authentication and you want to move the storage system to a different domain, you need to rerun the `cifs setup` command.

Before you begin

You must have an administrative account with permissions to add any Windows server to the domain.

About this task

After you change the storage system's domain, Data ONTAP updates the membership of the BUILTIN\Administrators group to reflect the new domain. This change ensures that the new domain's Administrators group can manage the storage system even if the new domain is not a trusted domain of the old domain.

Note: Until you put the CIFS server into a new domain or workgroup, you can cancel the CIFS setup process and return to your old settings by pressing Ctrl-c and then entering the `cifs restart` command.

Steps

1. If CIFS is currently running, enter the following command:

```
cifs terminate
```

2. Enter the following command:

```
cifs setup
```

The following prompt appears:

```
Do you want to delete the existing filer account information? [no]
```

3. Delete your existing account information by entering **yes** at the prompt.

Note: You must delete your existing account information to reach the DNS server entry prompt.

After deleting your account information, you are given the opportunity to rename the storage system:

```
The default name of this filer will be 'system1'.
Do you want to modify this name? [no]:
```

4. Keep the current storage system name by pressing Enter; otherwise, enter **yes** and enter a new storage system name.

Data ONTAP displays a list of authentication methods:

```
Data ONTAP CIFS services support four styles of user authentication.
Choose the one from the list below that best suits your situation.
(1) Active Directory domain authentication (Active Directory domains
only)
(2) Windows NT 4 domain authentication (Windows NT or Active
Directory domains)
(3) Windows Workgroup authentication using the filer's local user
accounts
(4) /etc/passwd and/or NIS/LDAP authentication
Selection (1-4)? [1]:
```

5. Accept the default method for domain authentication (Active Directory) by pressing Enter; otherwise, choose a new authentication method.
6. Respond to the remainder of the `cifs setup` prompts; to accept a default value, press Enter.

Upon exiting, the `cifs setup` utility starts CIFS.

7. Confirm your changes by entering the following command:

```
cifs domaininfo
```

Data ONTAP displays the storage system's domain information.

Changing protocol modes

When you have a valid CIFS license and a valid NFS license, you can change your protocol setting from `unix` (the default) to `ntfs` or `mixed` (multiprotocol) mode.

About this task

The protocol mode determines whether NFS, CIFS, or both clients have access to the files on the storage system.

You can set the protocol mode by running the `cifs setup` utility or setting the `waf1.default_security_style` option.

If you use `cifs setup` to change to multiprotocol mode, files are not immediately available to NFS clients. To make files available to NFS clients after changing to multiprotocol mode using `cifs`

setup, you must also change the root volume qtree security style to `unix`; then use the `chmod` command to permit UNIX client access as desired.

Note: An NFS client can also get access to a file `a` with a Windows ACL if Data ONTAP successfully maps the user's Unix user ID to a CIFS credential and verifies (with the CIFS credential) that the user can access the file. For example, if Data ONTAP successfully maps the Unix `root` user to a user in the `BUILTIN\Administrators` group, then the Unix `root` user can access the same files that the Windows user can access regardless of the security style.

Step

- 1. Perform one of the following actions.

If you want to...	Then...
Change the protocol mode using the <code>cifs</code> setup utility	Enter the following commands: <code>cifs terminate</code> <code>cifs setup</code> Then follow the prompts to change the protocol mode.
Change the protocol mode by setting the <code>wafl.default_security_style</code> option	Enter the following command: <code>options wafl.default_security_style {unix ntfs mixed}</code>

Effects of changing an NTFS-only storage system to a multiprotocol storage system

Changing an NTFS-only storage system to a multiprotocol storage system has several effects.

These are the effects:

- When you create a volume, its default security is `unix`.
- The `wafl.default_security_style` option is set to `unix`.

Existing ACLs and the security style of all current volumes and qtrees remain unchanged.

Note: Because the security style of the root volume remains `ntfs` after you change the storage system to multiprotocol, you might be denied access to the root volume when you connect from UNIX as `root`. You can gain access if the ACL for the root volume allows full control for the Windows user that maps to `root`. You can also gain access by setting the `cifs.nfs_root_ignore_acl` option to `on`.

Effects of changing a multiprotocol storage system to an NTFS-only storage system

Changing a multiprotocol storage system to an NTFS-only storage system has several effects.

These are the effects:

- If ACLs already exist on the storage system root directory (`/etc`) and on files in the `/etc` directory, the ACLs remain unchanged.

Otherwise, these ACLs are created such that the BUILTIN\Administrators group has full control; any in the `/etc/http` directory are assigned “Everyone Read”.

- ACLs on other files and directories remain unchanged.
- The security style of all volumes, except read-only volumes, is changed to `ntfs`.
- If the `/etc` directory is a qtree, its security style is changed to `ntfs`.
- Security style for all other qtrees remains unchanged.
- When you create a volume or qtree, its default security style is `ntfs`.
- The `waf1.default_security_style` option is set to `ntfs`.

Specifying Windows user account names

You can specify Windows user account names in some Data ONTAP commands and configuration files.

About this task

You can specify a Windows user account name in the following places:

- As the argument to the `cifs sessions` command to display information about a Windows user
- In the `/etc/usermap.cfg` file to map Windows names to UNIX names
- In the `/etc/quotas` file to establish quotas for Windows users

If you specify a UNIX user name with a backslash (\) in a configuration file, Data ONTAP treats the name as a Windows user account name. For example, UNIX names such as `corp\john` in the `/etc/quotas` file are interpreted as Windows user account names.

Note: The only command in which you can specify Windows user account names using the `user@domain` format is the `cifs setup` command. There are also rules for specifying Windows user account names that are specific to particular configuration files. For additional information about those rules, see the sections in this guide that relate to the particular configuration files.

Step

1. Perform one of the following actions.

If you want to...	Then...
Specify a Windows name in the pre-Windows 2000 format	Append a backslash and user name to the domain name. For example, <code>corp\john_smith</code> .
Specify the name of a Windows 2000 user in the pre-Windows 2000 format	Use the NETBIOS form of the domain name and make sure that the user name does not exceed 20 characters. For example, if <code>john_smith@engineering.my_company.com</code> is a Windows 2000 user, you can refer to this user as <code>engineering\john_smith</code> in Data ONTAP commands and configuration files.

If you want to...	Then...
Specify a local user account	Replace the domain name with the storage system name in the pre-Windows 2000 format. For example: filer1\john_smith.

Support for the Windows “Owner Rights” security principal

Data ONTAP 8.1.2 and later releases support the “Owner Rights” security principal. You can add this security principal to a file or directory DACL (discretionary access-control list) to override the default behavior of owners of files or directories.

“Owner Rights” is a well-known security principal with the well-known security identifier (SID) S-1-3-4 that is available with Windows 2008 server and Windows Vista and later releases. By default, when a user creates a file or directory, that user is the owner of the file and directory, and the owner has certain default rights. The administrator can override the default rights by adding the “Owner Rights” principle to a file or directory's DACL, and then specifying the desired permissions by assigning ACEs (Access Control Entries) to the “Owner Rights” security principal.

You can use either of the following methods to manage the “Owner Rights” security principal for files or directories residing within volumes:

- By using the **Security** tab on the Windows Properties window
- By using the *secedit* utility, available as a software download from the support site

Considerations when reconfiguring CIFS

Before reconfiguring CIFS on your storage system, you need to know about prerequisites and additional important information.

Make sure to complete all of the prerequisite steps that are appropriate to your setup before you reconfigure CIFS:

- If you want to change the storage system's domain from a Windows NT domain to another domain as you reconfigure your storage system, the storage system must be able to communicate with the primary domain controller for the domain in which you want to install the storage system.
You cannot use the backup domain controller for installing the storage system.
- If you want to change the name of the storage system, you must create a new computer account on the domain controller.
(Required only for versions of Windows after Windows 2000.)
- The CIFS server name must conform to Microsoft Active Directory naming conventions.
If the storage controller host name does not conform to Microsoft Active Directory naming conventions, you can enter a different CIFS server name during setup. For more information about Active Directory naming conventions, see Microsoft KB 909264.
- Your storage system and the domain controllers in the same domain must be synchronized with the same time source.

If the time on the storage system and the time on the domain controllers are not synchronized, an error message is displayed.

For details about setting up time synchronization services, see the *Data ONTAP System Administration Guide for 7-Mode*.

If you reconfigure CIFS with the `cifs setup` command when a UNIX-based KDC is configured for NFS, Data ONTAP renames your UNIX keytab file to include the string `UNIX`. To rename the keytab file for UNIX-based KDCs, enter **yes** when Data ONTAP displays the following message prompt during CIFS reconfiguration:

```
*** Setup has detected that this filer is configured to support
Kerberos *** authentication with NFS clients using a non-Active
Directory KDC. If *** you choose option 1 below, to allow NFS to
use the non-Active *** Directory KDC, your existing keytab file
'/etc/krb5.keytab' will be *** renamed to '/etc/UNIX_krb5.keytab'.
NFS will be using the new keytab *** file '/etc/UNIX_krb5.keytab'.
Do you want to continue. (Yes/No)?
```

If you enter **yes**, Data ONTAP renames the keytab file for UNIX-based KDCs; if you enter **no** or press Enter, Data ONTAP terminates the CIFS reconfiguration process. This renaming is needed for Kerberos multirealm configurations.

Related information

[Microsoft KB 909264: support.microsoft.com/kb/909264](https://support.microsoft.com/kb/909264)

Reconfiguring CIFS on your storage system

You can reconfigure CIFS after your initial setup by running the `cifs setup` utility again.

About this task

The CIFS configuration settings that you can change by running `cifs setup` are as follows:

- WINS server addresses
- Whether your storage system is multiprotocol or NTFS-only
- Whether the storage system uses Windows domain authentication, Windows workgroup authentication, or UNIX password authentication
- The domain or workgroup to which the storage system belongs
- The storage system name

Note: If you need to terminate the `cifs setup` utility when it is in progress, press Ctrl-C. You then enter the `cifs restart` command to restart CIFS using your old configuration information.

Steps

1. Enter the following command:

```
cifs terminate
```

CIFS service is stopped for the storage system.

2. Enter the following command:

```
cifs setup
```

Data ONTAP runs the `cifs setup` program, which displays a list of prompts for you to reconfigure CIFS.

Configuring SMB on your storage system

In addition to the CIFS protocol, Data ONTAP supports the Server Message Block (SMB) 1.0 protocol and the SMB 2.0 and SMB 2.1 protocols.

Support for the SMB 1.0 protocol

Data ONTAP supports the SMB 1.0 protocol, which extends CIFS with security, file, and disk-management features.

Support for the SMB 2.0 protocol

In addition to the SMB 1.0 protocol, Data ONTAP supports the SMB 2.0 protocol, which provides several enhancements.

The SMB 2.0 protocol is a major revision of the SMB 1.0 protocol in that it uses completely different packet formats.

Note: Data ONTAP does not support symbolic links, which are an optional feature of the SMB 2.0 protocol.

If the SMB 2.0 protocol is disabled on the storage system, communication between the SMB 2.0 client and the storage system falls back to the SMB 1.0 protocol (assuming that the SMB 2.0 client includes the SMB 1.0 dialect in its negotiate request).

For more information, see the SMB 2.0 protocol specification.

Support for the SMB 2.1 protocol

In addition to supporting SMB 2.0, Data ONTAP 8.1.1 and later in the 8.1 release family supports SMB 2.1. You need to know which SMB 2.1 enhancements are supported.

SMB 2.1 is a minor revision of the SMB 2.0 protocol. The following SMB 2.1 enhancements are implemented with Data ONTAP 8.1.1 and later:

- Lease oplocks
Data ONTAP uses SMB 2.1 lease oplocks, which is a new oplock model that provides advantages over traditional oplocks.
- Retrieval of content hashes used by BranchCache
BranchCache enabled storage systems use content hashes to provide clients with information about cached content.

Data ONTAP does not support the following SMB 2.1 enhancements:

- Resilient handles
- Large MTU
- Write-through

Enabling SMB 2.0 enables both SMB 2.0 and SMB 2.1. If the client is SMB 2.1 capable and requests the SMB 2.1 dialect in the negotiate request, Data ONTAP uses SMB 2.1. Otherwise, Data ONTAP uses SMB 2.0. If SMB 2.0 is disabled on the storage system, communication between the SMB 2.0 client and the storage system falls back to SMB 1.0 (assuming that the SMB 2.0 client includes the SMB 1.0 dialect in its negotiate request).

For more information, see the SMB 2.1 specification.

When to enable the SMB 2.0 and SMB 2.1 protocols

There are several file-transferring and interprocess communication enhancements that make the SMB 2.0 and SMB 2.1 protocols a more suitable choice to use in certain scenarios instead of the original SMB 1.0 protocol.

According to the SMB 2.0 protocol specification, such scenarios might include those that have the following requirements:

- More scalability with regard to simultaneously open files, number of shares, and user sessions
- Quality of service guarantees with regard to number of requests that can be outstanding against the server
- Stronger data integrity protection through the use of the HMAC-SHA256 hash algorithm
- Better throughput across networks with nonhomogeneous characteristics
- Better handling of intermittent losses of network connectivity through the use of durable handles

In addition to the enhancements present in SMB 2.0, SMB 2.1 provides these additional enhancements:

- Enhanced locking using lease oplocks
- Support for BranchCache, a feature that provides WAN bandwidth optimization and improved file access performance using client-side caching at branch offices

For more information, see the SMB 2.0 and SMB 2.1 protocol specifications.

Enabling or disabling SMB 2.x

You can enable or disable SMB 2.x by using the `cifs.smb2.enable` option of the `options` command. This option enables or disables all revisions of SMB 2.x that are available on the installed Data ONTAP version. By default, this option is set to `on`.

Step

1. Perform the applicable action:

If you want SMB 2.x to be...	Enter the command...
Enabled	<code>options cifs.smb2.enable on</code>
Disabled	<code>options cifs.smb2.enable off</code>

The requesting client and storage controller negotiate the protocol version based on what protocol versions are supported on each and on what version the clients present during protocol negotiation.

Note: Because SMB 2.1 must be enabled when using BranchCache, you must enable this option if you are using BranchCache to cache content on Windows clients at branch offices. Disabling SMB 2.x also disables BranchCache functionality.

SMB 2.x and durable handles

Durable handles allow SMB 2.0 clients to open a file and survive a temporary connection loss.

A durable handle is an open file handle that is preserved during a short network outage, which allows a client to reconnect when connectivity is restored.

To view information about durable handles, you use the `lock status` command.

Note: In Data ONTAP 8.1.1 and later releases in the 8.1 release family, durable handles are always enabled.

For more information, see the SMB 2.x protocol specifications.

Monitoring durable handle status

Starting with Data ONTAP 8.1.1 and later releases in the 8.1 release family, you can monitor and display information about durable handles by using the `lock status` command.

About this task

The `lock status` command displays the following information about durable handles:

- The new `durable_state` field holds the current durable state of the lock.
Durable handle states seen in the `durable_state` field include the following:
 - `DH_GRANTED`
 - `DH_NONE`

Note: This field is empty if lease locks are present.
- If a lease oplock is present, the `oplock` field holds lease oplock and durable handle information. The following two lease oplock values seen in the `oplock` field indicate that the lease oplock uses durable handles (as indicated by the H designation):
 - `Lease-RWH`
 - `Lease-RH`

Step

1. Use the `lock status` command to display durable handle information.

The `durable_state` field displays information about durable handles and the `oplock` field displays information on durable handles if lease oplocks are present.

Example

The following command displays lock information about CIFS durable handle locks, including information about durable handles with lease oplock locks:

```
lock status -p cifs
CIFS path=\\(/vol/vol1/) host=10.0.0.2(WIN1) owner=administrator
state=GRANTED mode=Read-denyN
oplock=None durable_state=DH_NONE fsid=0x1eea09f9 fileid=0x00000040

CIFS path=\\sample.txt(/vol/vol1/sample.txt) host=10.0.0.2(WIN1)
owner=administrator state=GRANTED
mode=RdWr-denyN oplock=None durable_state=DH_GRANTED
fsid=0x1eea09f9 fileid=0x00000060

CIFS path= host=10.0.0.2(WIN1) owner= state=GRANTED mode=None-denyN
oplock=Lease-None
durable_state= sid=0x1eea09f9 fileid=0x00000060
```

Support for SMB signing

Data ONTAP supports SMB signing (over the SMB 1.0 protocol and over the SMB 2.x protocol) when requested by the client. SMB signing helps to ensure that network traffic between the storage system and the client has not been compromised; it does this by preventing replay attacks (also known as *man in the middle* attacks).

When SMB signing is enabled on the storage system, it is the equivalent of the Microsoft Network server policy *Digitally sign communications (if client agrees)*. For performance reasons, SMB signing is disabled by default on the storage system.

How client SMB signing policies affect communications with the storage system

There are two SMB signing policies on Windows clients that control the digital signing of communications between clients and the storage system: "always" and "if server agrees."

Client SMB policies are controlled through security settings using the Microsoft Management Console (MMC). For more information about client SMB signing and security issues, see the Microsoft Windows documentation.

Here are descriptions of the two SMB signing policies on Microsoft Network clients:

- "Digitally sign communications (if server agrees)"

This setting controls whether the client's SMB signing capability is enabled. It is enabled by default. When this setting is disabled on the client, the client communicates normally with the storage system without SMB signing, regardless of the SMB signing setting on the storage system. When this setting is enabled on the client, communications between the client and storage system proceed as follows:

- If SMB signing is enabled on the storage system, all communications between client and storage system use SMB signing.
- If SMB signing is not enabled on the storage system, communications proceed normally without SMB signing.
- "Digitally sign communications (always)"

This setting controls whether the client requires SMB signing to communicate with a server. It is disabled by default. When this setting is disabled on the client, SMB signing behavior is based on the policy setting for "Digitally sign communications (if server agrees)" and the setting on the storage system. When this setting is enabled on the client, communications between the client and storage system proceed as follows:

- If SMB signing is enabled on the storage system, all communications between client and storage system use SMB signing.
- If SMB signing is not enabled on the storage system, the client rejects communication with it.

Note: If your environment includes Windows clients configured to require SMB signing, you must enable SMB signing on the storage system. If you do not, the storage system cannot serve data to these systems.

Performance impact of SMB signing

When SMB signing is enabled, all CIFS communications to and from Windows clients experience a significant impact on performance, which affects both the clients and the server (that is, the storage system running Data ONTAP).

The performance degradation shows as increased CPU usage on both the clients and the server, although the amount of network traffic does not change.

Depending on your network and your storage system implementation, the performance impact of SMB signing can vary widely; you can verify it only through testing in your network environment.

Most Windows clients negotiate SMB signing by default if it is enabled on the server. If you require SMB protection for some of your Windows clients, and if SMB signing is causing performance issues, you can disable SMB signing on any of your Windows clients that do not require protection against replay attacks. For information about disabling SMB signing on Windows clients, see the Microsoft Windows documentation.

Enabling SMB signing

Data ONTAP supports Server Message Block (SMB) signing when requested by the client. You can enable SMB signing on the storage system. By default, SMB signing is disabled.

About this task

SMB signing helps to ensure that network traffic between the storage system and the client has not been compromised by ensuring that all SMB messages have valid signatures. When SMB signing is enabled on the storage system, it is the equivalent of the Microsoft Network server policy "Digitally sign communications (if client agrees)".

If this option is enabled, then the storage system does the following:

- If the client requests SMB signing, communications between the client and storage system use SMB signing.
- If the SMB signing is not enabled on the client or the client does not request SMB signing, communications proceed normally without SMB signing.
- If the `cifs.smb2.signing.required` option is also set to on and an SMB 2.x-capable client requests SMB signing, the storage controller uses SMB 2.x with signing.

If the client cannot establish an SMB 2.x session with signing, the client falls back to an SMB session with or without signing, and the storage system uses whichever the client requests.

Note: You must terminate CIFS services by using the `cifs terminate` command to ensure that all existing CIFS connections are terminated prior to enabling SMB signing. After you have enabled SMB signing, you can restart CIFS services.

Steps

1. Terminate CIFS services by using the `cifs terminate` command.

See the `cifs terminate` man page for more information.

2. Enter the following command:

```
options cifs.signing.enable on
```

3. Restart CIFS services by using the `cifs restart` command.

Result

All new connections use SMB signing if requested by the client.

Enforcing the requirement for clients to sign SMB 2.x messages

If you require SMB signing, Data ONTAP accepts SMB 2.x messages only if they have valid signatures. You can enforce the requirement for clients to sign SMB 2.x messages by enabling the `cifs.smb2.signing.required` option. By default, this option is set to `off`.

About this task

When the `cifs.smb2.signing.required` option is enabled on the storage system, it is the equivalent of the Microsoft Network server policy "Digitally sign communications (always)". If this option is enabled, then the storage system does the following when communicating with an SMB 2.x capable client:

- If the client requests SMB 2.x signing, communications between the client and storage system use SMB 2.x signing.
- If the SMB 2.x signing is not enabled on the client or the client does not request SMB 2.x signing, what happens next depends of whether the `cifs.signing.enable` option is set to `on` or to `off`:
 - If it is set to `off`, the storage system rejects communication with the client because SMB 2.x signing is required.
 - If it is set to `on`, the client falls back to an SMB session with or without signing, whichever the client requests.

Steps

1. Terminate CIFS services by using the `cifs terminate` command to ensure that all existing CIFS connections are terminated prior to enabling SMB signing.

See the `cifs terminate` man page for more information.

2. Perform one of the following actions:

If you want the requirement that clients sign SMB 2.x messages to be...	Enter the command...
Enforced by Data ONTAP	<code>options cifs.smb2.signing.required on</code>
Not enforced by Data ONTAP	<code>options cifs.smb2.signing.required off</code>

3. Restart CIFS services by using the `cifs restart` command.

Managing shares

As an administrator, you can share directories with users on the storage system (create "shares").

What you need to know when creating a share

A CIFS share is a named access point in a volume that enables CIFS clients to view, browse, and manipulate files on a file server. There are certain guidelines you should take into consideration when creating CIFS shares.

Information needed when creating the share

When you create a share, you must provide all of the following information:

- The complete path in a volume to the CIFS share
- The name of the share entered by users when they connect to the share

When you create a share, you can optionally specify a description for the share. The share description appears in the Comment field when you browse the shares on the network.

If you create the share from the Data ONTAP command line, you can also specify a number of share properties. You can change these properties at any time after you create a share. You can specify the following share properties:

- Group membership for files in the share
- The `umask` value for the share
- Whether CIFS clients can follow symbolic links in the share to destinations anywhere on the same storage system
- Whether support for widelinks in the share is allowed
- Whether the share is browsable
- Whether virus scanning when files in the share are opened is disabled
- Whether file caching in the share by Windows clients is disallowed
- Whether support for automatic caching of documents and programs in the share by Windows clients is allowed
- Whether controlling the display of shared resources with Windows Access-based Enumeration (ABE) is allowed

Information needed when modifying share ACLs and properties

When you create a CIFS share, Data ONTAP creates a default ACL for the share with full control permissions. To manage CIFS share access control lists, you can use the `cifs access` command family.

After you create a share, you can specify these share properties:

- Maximum number of users who can simultaneously access the share

Note: If you do not specify a number of users, additional users are blocked only if there is no more storage system memory.

- A share-level ACL

Recovering from issues when ending a share comment with a backslash

You should avoid creating a CIFS share with a comment that ends with a backslash (\) because the CIFS share might disappear after rebooting the storage system or restarting CIFS. You must take several steps to recover your CIFS shares when this issue occurs.

About this task

To prevent this issue from occurring, do not end comments for CIFS shares with a backslash. To recover CIFS shares that disappeared due to this issue, follow these steps:

Steps

1. Open the `/etc/cifsconfig_share.cfg` file.
2. Locate the command that originally created the CIFS shares that disappeared.
3. Execute the command again, but either remove the trailing backslash or add a character such as a space after the backslash.

Share naming conventions

Share naming conventions for Data ONTAP are the same as for Windows. You should keep Data ONTAP share naming conventions in mind when you create a share.

For example, share names ending with the \$ character are hidden shares, and certain share names, such as ADMIN\$ and IPC\$, are reserved.

Share names are not case-sensitive.

Creating a CIFS share from the MMC on a Windows client

You can create a CIFS share from the MMC on a Windows client by connecting the MMC to the storage system and then running the Share a Folder wizard.

Sharing a folder using the Share a Folder wizard

You can run the Share a Folder wizard by using the MMC.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Shared Folders > Shares > Action**.

The wording of these menu items might vary slightly, depending on your Windows version.

4. Double-click **New Share**.
5. Follow the instructions in the **Share a Folder** wizard.

Creating a CIFS share from the Data ONTAP command line

You can create a CIFS share from the Data ONTAP command line by using the `cifs shares -add` command.

Step

1. To create a CIFS share, enter the following command:

```
cifs shares -add shareName path
```

shareName specifies the name of the new CIFS share.

path specifies the path to the share. Path separators can be backward or forward slashes, although Data ONTAP displays them as forward slashes.

For more information, see the `na_cifs_shares(1)` man page.

Example

The following command creates a “webpages” share that is accessible on the web in the `/vol/vol1/companyinfo` directory with a maximum of 100 users and in which all files that CIFS users create are owned by all users:

```
cifs shares -add webpages /vol/vol1/companyinfo -comment "Product  
Information" -forcegroup webgroup1 -maxusers 100
```

About the forcegroup option

When you create a share from the Data ONTAP command line, you can use the `forcegroup` option to specify that all files created by CIFS users in that share belong to the same group (that is, the “forcegroup”), which must be a predefined group in the UNIX group database.

Specifying a forcegroup is meaningful only if the share is in a UNIX or mixed qtree. There is no need to use forcegroups for shares in an NTFS qtree because access to files in these shares is determined by Windows permissions, not GIDs.

If a forcegroup has been specified for a share, the following becomes true of the share:

- CIFS users in the forcegroup who access this share are temporarily changed to the GID of the forcegroup.
This GID enables them to access files in this share that are not accessible normally with their primary GID or UID.
- All files in this share created by CIFS users belong to the same forcegroup, regardless of the primary GID of the file owner.

When CIFS users try to access a file created by NFS, the CIFS users' primary GIDs determine access rights.

The forcegroup does not affect how NFS users access files in this share. A file created by NFS acquires the GID from the file owner. Determination of access permissions is based on the UID and primary GID of the NFS user who is trying to access the file.

Using a forcegroup makes it easier to ensure that files can be accessed by CIFS users belonging to various groups. For example, if you want to create a share to store the company's web pages and give write access to users in Engineering and Marketing, you can create a share and give write access to a forcegroup named "webgroup1". Because of the forcegroup, all files created by CIFS users in this share are owned by the web group. In addition, users are automatically assigned the GID of the web group when accessing the share. As a result, all the users can write to this share without your managing the access rights of the Engineering and Marketing groups.

Displaying and changing the properties of a share

You can display and change share properties from the MMC or at the Data ONTAP command line.

About this task

You can change the following share properties:

- The description for the share
- The maximum number of users who can simultaneously access the share
- The share-level permissions
- Whether Access-Based Enumeration is enabled or disabled


Displaying and changing the properties of a share from the MMC on a Windows client

You can display and change the properties of a share from the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Shared Folders**.
4. Double-click **Shares**.
5. In the right pane, right-click the share.
6. Select **Properties**.

Properties for the share you selected are displayed as shown in the following example.



7. Select the **Share Permissions** tab.

The share's ACL appears.

8. To change the share's ACL to include an additional group or user, select the group or user from the **Group or user names** dialog box.
9. Change the permissions in the Permissions for *group* or *user name* dialog box.

Displaying the properties of a share from the Data ONTAP command line

You can display the properties of a share from the Data ONTAP command line by using the `cifs shares` command.

Step

1. Enter the following command:

```
cifs shares sharename
```

sharename is the name of a single share. If you omit *sharename*, the properties of all shares are displayed.

Result

Data ONTAP displays the share name, the path name of the directory that is shared, the share description, and the share-level ACL.

Changing the properties of a share from the Data ONTAP command line

You can change the properties of a share from the Data ONTAP command line by using the `cifs shares` command.

Step

1. Enter the following command:

```
cifs shares -change sharename {-browse | -nobrowse} {-comment desc | -nocomment} {-maxusers userlimit | -nomaxusers} {-forcegroup groupname | -noforcegroup} {-widelink | -nowidelink} {-symlink_strict_security | -nosymlink_strict_security} {-vscan | -novscan} {-vscanread | -novscanread} {-umask mask | -noumask} {-no_caching | -manual_caching | -auto_document_caching | -auto_program_caching}
```

For more information, see the `na_cifs_shares(1)` man page.

Note: You can use the question mark and asterisk characters as wildcards in the *sharename* to change the properties of multiple shares simultaneously. For example, to disable virus scanning of any file that a client opens in any share, enter the following command:

```
cifs shares -change * -novscan
```

Specifying `-nocomment`, `-nomaxusers`, `-noforcegroup`, and `-noumask` clears the share's description, maximum number of users, forcegroup, and umask values, respectively.

Enabling or disabling boundary checking for symbolic links from a share

You can disable boundary checking for symbolic links from a share to allow CIFS clients to follow symbolic links in that share to destinations anywhere on the same storage system.

About this task

By default, boundary checking for symbolic links is enabled to prevent users from accessing files outside the share.

If boundary checking is disabled, the storage system checks the share permissions of only the share that has the symbolic link.

Step

1. Perform one of the following actions.

If you want boundary checking for symbolic links from a share to be...	Then...
Disabled	On the Data ONTAP command line, enter the following command: <code>cifs shares -change sharename -nosymlink_strict_security</code>
Enabled	On the Data ONTAP command line, enter the following command: <code>cifs shares -change sharename -symlink_strict_security</code>

Enabling or disabling widelinks from a share

You can enable widelinks from a share if you want to allow CIFS clients to follow absolute symbolic links to destinations outside the share or storage system. By default, this feature is disabled.

Step

1. Perform one of the following actions.

If you want to...	Enter the command...
Enable widelinks from a share	<code>cifs shares -change sharename -widelink</code>
Disable widelinks from a share	<code>cifs shares -change sharename -nowidelink</code>

You can also enable widelinks from a share by specifying the `-widelink` option when you create the share.

After you finish

After you enable widelinks from a share, you need to create widelink entries in the `/etc/symlink.translations` file to specify how the storage system determines the destination represented by each widelink.

Specifying permissions for newly created files and directories in a share

You can specify the permissions of newly created files and directories in a share having mixed or UNIX qtree security style by setting the share's `umask` option.

About this task

You must specify the share's `umask` option as an octal (base-8) value. The default `umask` value is 0.

Note: The value of a share's `umask` option does not affect NFS.

Step

- 1. Perform one of the following actions.

If you want to specify the permissions for newly created...	Then...
Files and directories	On the Data ONTAP command line, enter the following command: cifs shares -change sharename -umask mask <i>mask</i> is an octal value that specifies the default permissions for newly created files and directories. Alternatively, set the <code>umask</code> option when you create the share.
Files, overriding the value of the share's <code>umask</code> option	On the Data ONTAP command line, enter the following command: cifs shares -change sharename -file_umask mask <i>mask</i> is an octal value that specifies the default permissions for newly created files, overriding value of the <code>umask</code> share option. Alternatively set the <code>file_mask</code> option when you create the share.
Directories, overriding the value of the share's <code>umask</code> option	On the Data ONTAP command line, enter the following command: cifs shares -change sharename -dir_umask mask <i>mask</i> is an octal value that specifies the default permissions for newly created directories, overriding value of the <code>umask</code> share option. Alternatively, set the <code>dir_mask</code> option when you create the share.

Example

To turn off write access for "group" and "other" permissions when a file is created in a share, enter the following command:

```
dir_umask 022
```


Enabling or disabling browsing

You can enable or disable browsing to allow users to see or prevent users from seeing a specific share.

Before you begin

You must enable the `-browse` option on each share for which you want to enable browsing.

Step

1. Perform one of the following actions:

If you want to...	Enter the following command...
Enable browsing on a specific share (has no effect if the <code>cifs.enable_share_browsing</code> option is on)	<code>cifs shares -change sharename -browse</code>
Disable browsing on a specific share Note: Windows 7 clients might experience delays in this scenario.	<code>cifs shares -change sharename -nobrowse</code>
Enable browsing on all shares	<code>options cifs.enable_share_browsing on</code>
Disable browsing on all shares	<code>options cifs.enable_share_browsing off</code>

Note: If you globally disable browsing by setting the `cifs.enable_share_browsing` option to `on`, you cannot enable per share browsing on any share.

For more information, see the `na_options(1)` man page.

Enabling or disabling virus scanning

You can enable or disable virus scanning on one or more shares to increase security or performance, respectively.

About this task

By default, Data ONTAP scans any file that a client opens for viruses.

Step

1. Perform one of the following actions.

If you want to...	Enter the command...
Enable virus scanning of all files that a client opens	cifs shares -change <i>sharename</i> -vscan
Disable virus scanning of all files that a client opens	cifs shares -change <i>sharename</i> -novscan
Enable virus scanning of read-only files that a client opens	cifs shares -change <i>sharename</i> -vscanread
Disable virus scanning of read-only files that a client opens	cifs shares -change <i>sharename</i> -novscanread

You can also disable virus scanning for a share when you create the share by specifying the `-nvscan` or `-nvscanread` option.

For more information about specifying virus scanning for CIFS shares, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Enabling or disabling caching

You can enable or disable caching to allow or prevent clients from caching files on a share.

About this task

You can specify whether clients must manually select files for caching; and, if not, whether Data ONTAP automatically caches programs, user files, or both in accordance with client settings. By default, clients must manually select files for caching.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Disable caching	cifs shares -change <i>sharename</i> -nocaching
Enable manual caching	cifs shares -change <i>sharename</i> -manual_caching
Enable automatic caching of documents	cifs shares -change <i>sharename</i> -auto_document_caching
Enable automatic caching of programs	cifs shares -change <i>sharename</i> -auto_program_caching

When you create a share, you can override the default caching option (`-manual_caching`) by specifying the `-nocaching`, `-auto_document_caching`, or `-auto_program_caching` option.

Setting client-side caching properties for a share

You can set client-side caching properties for a share using the Computer Management application on Windows 2000, XP, and 2003 clients. For more information, see the Microsoft Windows online Help system.

About access-based enumeration

When access-based enumeration (ABE) is enabled on a CIFS share, users who do not have permission to access a shared folder or file underneath it (whether through individual or group permission restrictions) do not see that shared resource displayed in their environment.

Conventional share properties allow you to specify which users (individually or in groups) have permission to view or modify shared resources. However, they do not allow you to control whether shared folders or files are visible to users who do not have permission to access them. This could pose problems if the names of shared folders or files describe sensitive information, such as the names of customers or products under development.

Access-based enumeration (ABE) extends share properties to include the enumeration of shared resources. ABE therefore enables you to filter the display of shared resources based on user access rights. In addition to protecting sensitive information in your workplace, ABE enables you to simplify the display of large directory structures for the benefit of users who do not need access to your full range of content.

Enabling or disabling access-based enumeration

You can enable or disable access-based enumeration (ABE) to allow or prevent users from seeing shared resources that they do not have permission to access.

About this task

By default, ABE is disabled.

Step

1. Perform one of the following actions.

If you want to...	Enter the command...
Enable ABE	<code>cifs shares -change <i>sharename</i> -accessbasedenum</code>
Disable ABE	<code>cifs shares -change <i>sharename</i> -noaccessbasedenum</code>

You can also enable ABE when you create a share by specifying the `-accessbasedenum` option.

Executing access-based enumeration commands from a Windows client

You can execute access-based enumeration (ABE) commands from a Windows client using the `abecmd` command to enable or disable ABE for shares.

Before you begin

You must enable ABE in Data ONTAP.

Step

1. From a Windows client that supports ABE, enter the following command:

```
abecmd [/enable | /disable] [/server name] {/all | ShareName}
```

For more information about the `abecmd` command, see your Windows client documentation.

Deleting a share

You can delete a share from the MMC or the Data ONTAP command line.

Deleting a share from the MMC

You can delete a share from the MMC.

Steps

1. Connect the MMC to the storage system.
2. If it is not selected already, in the left pane, select **Computer Management**.
3. Select **System Tools > Shared Folders**.
4. Double-click **Shares**.
5. In the right pane, right-click the share; then select **Stop Sharing**.
6. In the confirmation box, select **Yes**.

Result

The MMC deletes the share.

Deleting a share from the Data ONTAP command line

You can use the `cifs shares` command to delete a share from the Data ONTAP command line.

Step

1. Enter the following command:

```
cifs shares -delete [-f] sharename
```

`-f` option forces all files closed on a share without prompting. This is useful when using the command in scripts.

`sharename` specifies the name of the share you want to delete.

Managing access control lists

You can display and change share-level ACLs from the MMC or the command line. You can change file-level ACLs from the command line only.

How Data ONTAP uses share-level ACLs

When an SMB user tries to access a share, Data ONTAP always checks the share-level ACL (access control list) to determine whether access should be granted, regardless of the security style of the volume or qtree containing the share.

A share-level ACL consists of a list of access control entries (ACEs). Each ACE contains a user or group name and a set of permissions that determines user or group access to the share.

A share-level ACL only restricts access to files in the share; it never grants more access than the file-level ACLs.

How ACL inheritance works

It is important to understand how Access Control Lists (ACLs) and Access Control Entries (ACEs) are inherited and react to modification.

ACEs contain additional information that specifies how objects located at deeper levels in the file system hierarchy inherit them. For example, consider a file in a folder that is a subfolder to a parent folder. The parent folder might have an ACL that is automatically inherited by objects below the parent folder.

Inheritance operations only take place when you create new objects. For example, when a file is created in a folder, the file inherits ACEs from the folder only at that moment. When you change the permissions of the file later, the system does not automatically initiate a new inheritance operation from the folder to the file. However, when you change the permissions of the folder and this setting is marked for inheritance, the change also applies to the file that initially inherited settings from the folder.

Because ACE settings can be inherited from folders to the objects they contain, the system has to keep track of which ACEs in an ACL are inherited and which are not. When you modify an inherited ACE of an object, both the modified ACE and the inherited ACE are maintained on the object.

This behavior is slightly different for NFSv4. You can only change ACEs that are not inherited. Non-inherited entries are called *explicit*. To change an inherited entry, you either change the entry at the parent level (where this inherited entry came from), or delete the ACL for this object (thereby breaking the inheritance), replacing the inherited entries by explicit entries.

If you want to edit permissions locally on the storage system, you can use the `fsecurity` command to modify NTFS permissions of folders within a share. If you modify NTFS permissions using the `fsecurity` command, you can avoid issues with ACL inheritance.

Displaying and changing a share-level ACL

You can change a share-level ACL to give users more or less access rights to the share.

About this task

After you create a share, by default, the share-level ACL gives read access to the standard group named Everyone. Read access in the ACL means that all users in the domain and all trusted domains have read-only access to the share.

You can change a share-level ACL using the MMC on a Windows client or the Data ONTAP command line.

If you use the MMC, remember these guidelines:

- You can specify only Windows permissions.
- The user and group names specified must be Windows names.
- The share-level ACL must not have UNIX-style permissions.

If you use the Data ONTAP command line, remember these guidelines:

- You can specify either Windows permissions or UNIX-style permissions.
- The user and group names can be Windows or UNIX names.
- If the storage system is authenticated by the `/etc/passwd` file, the user or group name in the ACL is assumed to be a UNIX name.

If the storage system is authenticated by a domain controller, the name is at first assumed to be a Windows name, but if the name is not found on the domain controller, the storage system tries to look up the name in the UNIX name database.

Adding a user or group to a share-level ACL from the MMC on a Windows client

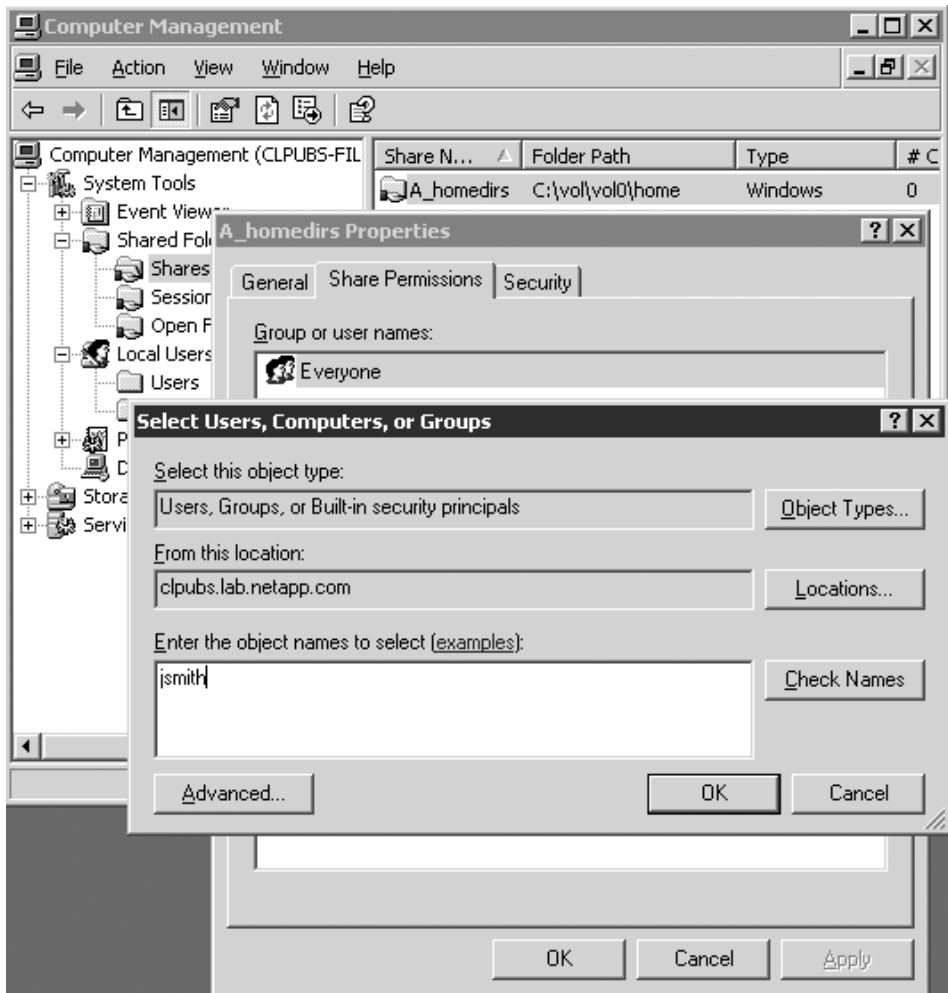
You can add a user or group to an ACL from the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Shared Folders**.
4. Double-click **Shares**.
5. In the right pane, right-click on the share.
6. Select **Properties**.
7. Select the **Share Permissions** tab.

The share's ACL appears.

8. Click **Add**.
9. In the "Select Users, Computers, or Groups" window, enter the name of the user in the "Enter the object names to select" box.



10. Click **OK**.

Result

The ACL now contains the new user or group.

Displaying and changing a share-level ACL from the MMC on a Windows client

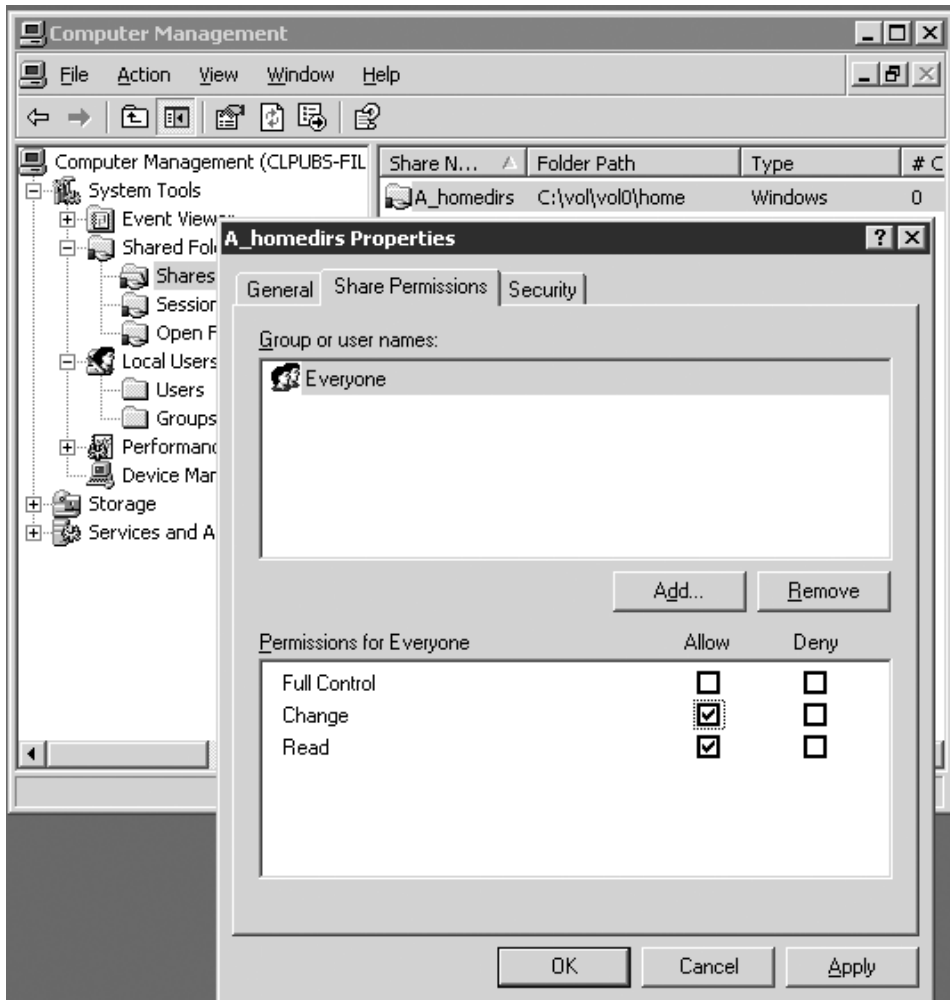
You can display and change a share-level ACL from the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Shared Folders**.
4. Double-click **Shares**.
5. In the right pane, right-click on the share.
6. Select **Properties**.
7. Select the **Share Permissions** tab.

The share's ACL appears.

8. To change the ACL for a group or user, select the group or user from the "Group or user names" box and change the permissions in the "Permissions for *group or user name*" box.



Removing a user or group from a share-level ACL using the MMC on a Windows client

You can remove a user or group from a share-level ACL using the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Shared Folders**.
4. Double-click **Shares**.
5. In the right pane, right-click on the share.

6. Select **Properties**.
7. Select the Share Permissions tab.
The share's ACL appears.
8. Select the user or group.
9. Click **Remove**.

Result

The ACL no longer contains the user or group.

Changing a share-level ACL from the Data ONTAP command line

You can change a share-level ACL from the Data ONTAP command line by using the `cifs access` command.

Step

1. Enter the following command:

```
cifs access share [-g] user rights
```

share is the name of the share (you can use the * and ? wildcards).

user is the name of the user or group (UNIX or Windows).

If *user* is a local group, specify the storage system name as the domain name (for example, toaster\writers).

rights are the access rights. For Windows users, you specify one of these choices of access rights: No Access, Read, Change, Full Control. For UNIX users, you specify one of these choices of access rights: r (read), w (write), x (execute).

Use the -g option to specify that *user* is the name of a UNIX group.

Examples

The following example grants Windows read access to the Windows user ENGINEERING\mary on the share releases:

```
cifs access releases ENGINEERING\mary Read
```

The following example grants UNIX read and execute access to the user *john* on the accounting share:

```
cifs access accounting john rx
```

The following example grants full access to the UNIX group wheel on the sysadmins share:

```
cifs access sysadmins -g wheel Full Control
```

Removing a user or group from a share-level ACL using the Data ONTAP command line

You can remove a user or group from an ACL using the Data ONTAP command line.

Step

1. Enter the following command:

```
cifs access -delete share [-g] user
```

share is the name of the share (you can use the * and ? wildcards).

user is the name of the user or group (UNIX or Windows).

If *user* is a local group, specify the storage system name as the domain name (for example, toaster\writers).

Use the -g option to specify that *user* is the name of a UNIX group (that is, that *user* is not a UNIX user, Windows user, or Windows group).

Example

The following example deletes an ACL entry for the user ENGINEERING\mary from a share named releases.

```
cifs access -delete releases ENGINEERING\mary
```

Specifying whether NFSv3 and NFSv4 clients display Windows ACL permissions based on minimum or maximum access

To specify that NFSv3 and NFSv4 clients should display Windows ACL permissions (not UNIX or NFSv4 ACL permissions) based on the minimum access granted by the Windows ACL, you can set the `nfs.ntacl_display_permissive_perms` option to `on`. Otherwise, you can set the option to `off`. By default, this option is `off`.

About this task

In versions of Data ONTAP earlier than 7.2.1, the permissions displayed to NFSv3 and NFSv4 clients on files were based on the maximum access granted by the Windows ACL. However, starting in Data ONTAP 7.2.1, the permissions displayed to NFSv3 and NFSv4 clients on files are based on the minimum access granted by the Windows ACL to any user.

Step

1. Enter the following command:

```
options nfs.ntacl_display_permissive_perms {on | off}
```

Displaying and changing a file-level ACL

You can change a file-level ACL to control whether specific users and groups have access to the file.

About this task

Permission settings for files and directories are stored in file-level ACLs. These ACLs follow the Windows 2000 NTFS security model. For files that have NTFS-style security, CIFS users can set and display file-level ACLs from their PC. All files in an NTFS-style qtree and some files in a mixed qtree might have NTFS-style security.

Files in a FAT (file allocation table) file system do not have ACLs; they use UNIX permissions. When viewed from a CIFS client, files without ACLs will not display the Security tab in the file Properties window.

The file system (FAT or NTFS) for a given resource depends upon the storage system authentication method and qtree style for that resource, as shown in the following table.

Qtree style and authentication method	File system
UNIX-style qtree and all authentication methods	FAT
Mixed or NTFS-style qtree and /etc/passwd authentication	FAT
Mixed or NTFS-style qtree and domain or workgroup authentication	NTFS

Steps

1. From the Windows desktop, right-click a file and select **Properties** from the pop-up menu.


Note: On an NT4 client, if you right-click a file that is located in a share that supports widelinks and select Properties, no Security tab is displayed. You can set security using a security tool such as cacs. Alternatively, you can either access files from a Windows 2000 client or access files using shares that don't support widelinks. You can have two different shares on the same directory, one that supports widelinks and one that does not, and use the share that does not support widelinks when setting security.

2. Click the **Security** tab.


Note: Depending on authentication method and qtree style, the Security tab might not be present.

3. Select the user or the group whose permissions you want to display from the "Group or user names" box.

The permissions for the group or the user you selected are displayed in the "Permissions for *user* or *group*" box.



- To add a user or a group to the file, click **Add**, then, in the "Select Users, Computers, or Groups" window, enter the name of the user or the group in the "Enter the object names to select" box.



A user or a group is added to the ACL.

Specifying how group IDs work with share-level ACLs

If a share contains files with UNIX-style security and if you want to use the share-level ACL to control access by UNIX groups, you must decide whether you want Data ONTAP to grant user access to files based on group ID.

About this task

If a share named `specs` exists in a UNIX-style qtree and you want two UNIX groups, engineering and marketing, to have full access to the share, you give `rx` permissions to these groups at the share level.

Suppose in this share, a file owned by the engineering group is named `draft` and it has the following permissions:

```
draft  rwxr-x---
```

When a member of engineering tries to access the `draft` file, the share-level ACL gives this user unrestricted access to the `specs` share, and access to the `draft` file is determined by the access rights assigned to the engineering group (`r-x`, in this example).

However, when a member of marketing tries to access the `draft` file, access is denied because the UNIX-style file permissions grant nonmembers of engineering no access to the file. To make the `draft` file readable by the marketing group, you need to change the file-level permissions to the following settings:

```
draft  rwxr-xr-x
```

The disadvantage of these permissions is that in addition to marketing, all UNIX users can read the file, which creates a security problem.

To solve this problem, you can configure Data ONTAP to disregard the GID when granting access.

If you configure Data ONTAP to disregard the user's GID when granting access, all users who are not the file's owner are considered members of the UNIX group that owns the file. In the preceding example, permissions that apply to the engineering group also apply to members of marketing who try to access the file. That is, both engineering members and marketing members have the r-x permissions to the draft file.

By default, Data ONTAP considers the user's GID before granting access. This default configuration is useful if either of the following statements is true:

- The share does not contain files with UNIX-style security.
- You do not use a share-level ACL to control any UNIX group's access.

Step

1. Perform one of the following actions.

If, when granting user access, you want to...	Then...
Disregard the user's GID	Enter the following command: options cifs.perm_check_use_gid off
Consider the user's GID	Enter the following command: options cifs.perm_check_use_gid on

Managing home directories

You can create user home directories on the storage system and configure Data ONTAP to automatically offer each user a home directory share.

About this task

From the CIFS client, the home directory works the same way as any other share to which the user can connect.

Each user can connect only to his or her home directories, not to home directories for other users.

About home directories on the storage system

Data ONTAP maps home directory names to user names, searches for home directories that you specify, and treats home directories slightly differently than regular shares

Data ONTAP offers the share to the user with a matching name. The user name for matching can be a Windows user name, a domain name followed by a Windows user name, or a UNIX user name. Home directory names are not case-sensitive.

When Data ONTAP tries to locate the directories named after the users, it searches only the paths that you specify. These paths are called home directory paths. They can exist in different volumes.

The following differences exist between a home directory and other shares:

- You cannot change the share-level ACL and the comment for a home directory.
- The `cifs shares` command does not display the home directories.
- The format of specifying the home directory using the Universal Naming Convention (UNC) is sometimes different from that for specifying other shares.

If you specify `/vol/vol1/enghome` and `/vol/vol2/mktghome` as the home directory paths, Data ONTAP searches these paths to locate user home directories. If you create a directory for `jdoe` in the `/vol/vol1/enghome` path and a directory for `jsmith` in the `/vol/vol2/mktghome` path, both users are offered a home directory. The home directory for `jdoe` corresponds to the `/vol/vol1/enghome/jdoe` directory, and the home directory for `jsmith` corresponds to the `/vol/vol2/mktghome/jsmith` directory.

How Data ONTAP matches a directory with a user

You can specify the naming style of home directories to determine how Data ONTAP matches a directory with a user.

These are the naming styles that you can choose from, and some information about each style:

- **Windows name**
Data ONTAP searches for the directory whose name matches the user's Windows name.
- **Hidden name**
If the naming style is hidden, users connect to their home directories using their Windows user name with a dollar sign appended to it (`name$`), and Data ONTAP searches for a directory that matches the Windows user name (`name`).
- **Windows domain name and Windows name**
If users from different domains have the same user name, they must be differentiated using the domain name.
In this naming style, Data ONTAP searches for a directory in the home directory path that matches the domain name. Then it searches the domain directory for the home directory that matches the user name.
Example: To create a directory for `engineering\jdoe` and a directory for `marketing\jdoe`, you create the two directories in the home directory paths. The directories have the same names as the domain names (`engineering` and `marketing`). Then you create user home directories in these domain directories.
- **Mapped UNIX name**
If the naming style is UNIX, Data ONTAP searches for the directory that matches the user's mapped UNIX name.
Example: If John Doe's Windows name `jdoe` maps to the UNIX name `johndoe`, Data ONTAP searches the home directory paths for the directory named `johndoe` (not `jdoe`) and offers it as the home directory to John Doe.

If you do not specify a home directory naming style, Data ONTAP uses the user's Windows name for directory matching. This is the same style used by versions of Data ONTAP prior to version 6.0.

How symbolic links work with home directories

The way symbolic links work depends on the home directory naming style.

If you do not specify a naming style, Data ONTAP follows any symbolic link that points to a directory outside the home directory path to locate a home directory.

Example: Suppose the home directory path is `/vol/vol0/eng_homes`. To locate the home directory for `jdoe`, Data ONTAP searches for `/vol/vol0/eng_homes/jdoe`, which can be a symbolic link pointing to a directory outside the home directory path, such as `/vol/vol1/homes/jdoe`.

If you specify a home directory naming style, by default a symbolic link works only if the symbolic link points to a directory in the home directory path.

Example: Suppose the home directory path is `/vol/vol0/eng_homes` and you use the Windows naming style. To locate the home directory for `jdoe`, Data ONTAP searches for `/vol/vol0/eng_homes/jdoe`. If the path is a symbolic link, the user can access the home directory only if the target of the symbolic link resides in the home directory path. For example, the symbolic link works if it points to the `/vol/vol0/eng_homes/john` directory; it does not work if it points to the `/vol/vol1/homes/john` directory.

Note: You can change the default storage system settings to allow CIFS clients to follow symbolic links to destinations outside the home directory path.

Because Data ONTAP now supports home directories in different volumes, you do not need to use symbolic links as home directory names. However, Data ONTAP continues to support symbolic links as home directory names for backward compatibility.

Specifying home directory paths

Data ONTAP searches the home directory paths in the order you specify for the directory that matches the user name. You can specify home directory paths by editing the `/etc/cifs_homedir.cfg` file.

About this task

You can specify multiple home directory paths. Data ONTAP stops searching when it finds the matching directory.

You can add an extension to the home directory path if you do not want users to access the top level of their home directories. The extension specifies a subdirectory that is automatically opened when users access their home directories.

You can change the home directory paths at any time by changing the entries in the `cifs_homedir.cfg` file. However, if a user has open files in a home directory path that you remove

from the list, Data ONTAP displays a warning message and requests a confirmation for the change. Changing a directory path that contains an open file terminates the connection to the home directory.

Data ONTAP creates a default `cifs_homedir.cfg` file in the `/etc` directory when CIFS starts, if the file does not already exist. Changes to this file are processed automatically whenever CIFS starts. You can also process changes to this file by using the `cifs homedir load` command.

Steps

1. Create directories to use as home directory paths.

For example, in the `/vol/vol0` volume, create a directory named `enghome`.

2. Open the `/etc/cifs_homedir.cfg` file for editing.
3. Enter the home directory path names created in Step 1 in the `/etc/cifs_homedir.cfg` file, one entry per line, to designate them as the paths where Data ONTAP searches for user home directories.

You can enter up to 1,000 path names.

4. Enter the following command to process the entries:

```
cifs homedir load [-f]
```

The `-f` option forces the use of new paths.

Displaying the list of home directory paths

You can use the `cifs homedir` command to display the current list of directory paths.

Step

1. Enter the following command:

```
cifs homedir
```

Note: If you are using the hidden naming style for home directories, when you display the list of home directory paths, Data ONTAP automatically appends a dollar sign to the home directory name (for example, `name$`).

Result

If you are using the hidden naming style for home directories, home directories are not displayed in the following cases:

- In DOS, when you use the `net view \\filer` command
- In Windows, when you use an Explorer application to access the storage system and display home directory folders

Specifying the naming style of home directories

You can specify the naming style used for home directories by setting the `cifs.home_dir_namestyle` option.

Step

1. Enter the following command:

```
options cifs.home_dir_namestyle {ntname | hidden | domain | mapped | ""}
```

Use `ntname` if the home directories have the same names as the Windows user names.

Use `hidden` if you want to use a Windows user name with a dollar sign (\$) appended to it to initiate a search for a home directory with the same name as the Windows user name.

Use `domain` if you want to use the domain name in addition to the Windows user name to search for the home directory.

Use `mapped` if the home directories have the UNIX user names as specified in the `usermap.cfg` file.

Use `" "` if you do not want to specify a name style and want Data ONTAP to match home directories to users by following any symbolic link that points to a directory outside the home directory path to locate a home directory.

By default, the `cifs.home_dir_namestyle` option is `" "`.

Creating directories in a home directory path (domain-naming style)

If the `cifs.home_dir_namestyle` option is `domain`, you can create home directories by editing the `/etc/cifs_homedir.cfg`, creating the directories, and setting the permissions on the directories.

Steps

1. Open the `/etc/cifs_homedir.cfg` file and add the path that represents where the home directories will exist.

The home directories will exist within folders named for the NetBIOS domains to which each user belongs. For example, add the path `/vol/vol1/homedir` to the `/etc/cifs_homedir.cfg` file.

2. In the directory that you added to the `/etc/cifs_homedir.cfg` file, create a directory for each domain.

For example, if there are two domains, HQ and UK, create a `/vol/vol1/homedir/hq/` directory and a `/vol/vol1/homedir/uk/` directory.

3. In each domain directory created in Step 2, create home directories for the users in that domain.

For example, if two users have the name jsmith and they are in the HQ domain and the UK domain, create the /vol/vol1/homedir/HQ/jsmith home directory and the /vol/vol1/homedir/UK/jsmith home directory.

4. Make each user the owner of his or her home directory.

For example, make HQ\jsmith the owner of the /vol/vol1/homedir/HQ/jsmith home directory and UK\jsmith the owner of the /vol/vol1/homedir/UK/jsmith home directory.

The user with the name HQ\jsmith can attach to the jsmith share corresponding to the /vol/vol1/homedir/HQ/jsmith home directory. The user with the name UK\jsmith can attach to the jsmith share corresponding to the /vol/vol1/homedir/UK/jsmith home directory.

5. Load the new CIFS homedir configuration into the storage system.

For example, enter the following command:

```
cifs homedir load -f
```

6. Make sure that the CIFS homedir domain name style is working by entering the following command:

```
cifs homedir showuser user_name
```

For example, enter one of the following commands:

```
cifs homedir showuser hq/jsmith
```

```
cifs homedir showuser uk/jsmith
```

Creating directories in a home directory path (non-domain-naming style)

If the `cifs.home_dir_namestyle` option is *not* domain, you can create home directories by creating the directories and making the users the owners of the directories.

Steps

1. In the specified home directory paths, create home directories.

Example

For example, if there are two users, jsmith and jdoe, create the /vol/vol0/enghome/jsmith and /vol/vol1/mktghome/jdoe home directories.

Users can attach to the share that has the same name as their user name and start using the share as their home directory.

2. Make each user the owner of his or her home directory.

Example

For example, make jsmith the owner of the /vol/vol0/enghome/jsmith home directory and jdoe the owner of the /vol/vol1/mktghome/jdoe home directory.

Note: If the naming style is hidden, users must enter their user name with a dollar sign appended to it (for example, name\$) to attach to their home directory.

The user with the name engineering\jsmith can attach to the share named jsmith, which corresponds to the /vol/vol0/enghome/engineering/jsmith home directory.

The user with the name marketing\jdoe can attach to the share named jdoe, which corresponds to the /vol/vol1/mktghome/marketing/jdoe home directory.

Creating subdirectories in home directories when a home directory path extension is used

You can create subdirectories that users can access in their home directories if you use a home directory path extension.

Step

1. For each home directory that resides in a home directory path with an extension, create a subdirectory that you want users to access.

For example, if the /etc/cifs_homedir.cfg file includes the /vol/vol0/enghome/%u%/data path, create a subdirectory named data in each home directory.

Users can attach to the share that has the same name as their user name. When they read or write to the share, they effectively access the data subdirectory.

Syntax for specifying a home directory using a UNC name

The convention for specifying a home directory when using UNC depends on the home directory naming style specified by the cifs.home_dir_namestyle option.

The following table lists UNC names, with examples, for each namestyle value.

Value of cifs.home_dir_namestyle	UNC name
ntname or " "	\\filer\Windows_NT_name Example: \\toaster\jdoe
hidden	\\filer\Windows_NT_name\$ Example: \\toaster\jdoe\$
domain	\\filer\~domain~Windows_NT_name Example: \\toaster\~engineering~jdoe
mapped	\\filer\~mapped_name Example: \\toaster\~jdoe

If cifs.home_dir_namestyle is domain but the UNC name in the access request does not specify a domain name, Data ONTAP assumes the domain to be the domain under which the request

is sent. If you omit the domain name in the access request, you can also leave out the tilde (~) before the user name.

Example

A user named jdoe is logged in as engineering\jdoe from a PC in the engineering domain. When he tries to access his home directory using his user name in the marketing domain, he can enter either of the following commands to request access:

```
net use * \\toaster\~jdoe /user:marketing\jdoe
net use * \\toaster\jdoe /user:marketing\jdoe
```

Enabling users to access other users' home directories

Although users can only attach to their own home directories, you can allow them to access other users' home directories.

Steps

1. Create a share that corresponds to the path name that is either one of the following:
 - A home directory path if `cifs.home_dir_name_style` is not domain
 - A domain directory in the home directory path if `cifs.home_dir_name_style` is domain

Example: If `/vol/vol0/enghome` is a home directory path, use the following command:

```
cifs shares -add eng_dirs /vol/vol0/enghome -comment "readable
engineering home directories"
```

2. Assign each user the appropriate access permissions to other users' home directories.

Example: Assign read-only permission to the engineering group for the `eng_dirs` share as follows:

```
cifs access eng_dirs engineering full
```

Members of the engineering group have read-only access to all home directories in the `eng_dirs` share.

Accessing your CIFS home directory using a share alias

For any CIFS home directory naming style, you can connect to your own CIFS home directory using either `cifs.homedir` or tilde (~) share aliases.

About this task

Connecting to your own CIFS home directory can be useful when you are writing scripts.

Step

1. Access your own CIFS home directory using either `cifs.homedir` or tilde (~) share aliases.

Examples

```
net use * \\toaster\cifs.homedir
net use * \\toaster\~
```

Enabling or disabling widelinks from a share

You can enable widelinks from a share if you want to allow CIFS clients to follow absolute symbolic links to destinations outside the share or storage system. By default, this feature is disabled.

Step

1. Perform one of the following actions.

If you want to...	Enter the command...
Enable widelinks from a share	<code>cifs shares -change <i>sharename</i> -widelink</code>
Disable widelinks from a share	<code>cifs shares -change <i>sharename</i> -nowidelink</code>

You can also enable widelinks from a share by specifying the `-widelink` option when you create the share.

After you finish

After you enable widelinks from a share, you need to create widelink entries in the `/etc/symlink.translations` file to specify how the storage system determines the destination represented by each widelink.

Disabling home directories

You can stop offering home directories by deleting the `/etc/cifs_homedir.cfg` file. You cannot use the `cifs shares -delete` command to delete home directories.

Step

1. Delete the `/etc/cifs_homedir.cfg` file on the storage system.

Using BranchCache to cache CIFS share content at a branch office

BranchCache was developed by Microsoft to enable caching of content on computers local to requesting clients. The Data ONTAP implementation of BranchCache can reduce wide-area network

(WAN) utilization and provide improved access response time when users in a branch office access content stored on a storage system using CIFS.

If you configure BranchCache, Windows BranchCache clients first retrieve content from the storage system and then cache the content on a computer within the branch office. If another BranchCache-enabled client in the branch office requests the same content, the storage system first authenticates and authorizes the requesting user. The storage system then determines whether the cached content is still up-to-date and, if it is, sends the client metadata about the cached content. The client then uses the metadata to retrieve content directly from the locally based cache.

Supported BranchCache caching modes

You can configure Microsoft BranchCache in one of two modes of operation: Hosted Cache mode and Distributed Cache mode. Both modes are supported with Data ONTAP BranchCache. Each mode has certain benefits to consider when configuring BranchCache.

You configure the mode of operation on the branch office Microsoft Windows clients that are participating in the BranchCache configuration. The mode of operation can be one of the following:

Hosted Cache mode The cached content is stored on a single branch office server. The clients retrieve cached content from the hosted cache server. If content requested by the client is not cached, the client retrieves the content from the storage system using the WAN, and then notifies the hosted cache server that it has content to cache. The hosted cache server then transfers the content to the hosted cache so that other clients requesting the same content can retrieve it from the hosted cache.

Distributed Cache mode Using a peer-to-peer model, the content is cached on a number of clients. Local branch office clients keep a copy of the content and make it available to other authenticated and authorized clients that request the same data. This differs from a type of local caching known as *offline files*, where the cached files are only available to the client on which the files are stored.

This mode only works across a single subnet. Also, clients that hibernate, or are otherwise unavailable from the network, cannot provide cached content to requesting clients.

The following outlines some key benefits to consider when choosing the BranchCache mode:

- Benefits of the Hosted Cache mode:
 - Increased cache availability
Content is available even if the client that originally requested the data is offline.
 - Supports access to cached content over more than a single LAN
Hosted caching can span multiple subnets, whereas distributed caching is limited to a single subnet.
- Benefits of the Distributed Cache mode:
 - Separate caching server not needed
All content is cached on local clients.

- Easy to configure and maintain

BranchCache terminology

To understand how Data ONTAP BranchCache works and whether you should implement BranchCache, you need to know some basic BranchCache terms.

Block	One of the fixed-size chunks of content that compose a segment. Each segment is divided into one or more blocks of equal size (64 kilobytes) except for the last block, which in general can be smaller. Every block belongs to a specific segment, and within a segment, blocks are identified by their progressive index.
Block range	A set of consecutive blocks within a segment described by a pair of integers, the first being the index of the first blocks in the range, and the second being the number of consecutive blocks in the range.
Block hash	A hash of a content block within a segment. Also known as a <i>block ID</i> .
Client	A peer that accesses certain content. With the distributed cache mode, the client also offers cached content to other peers in the branch office.
Content	A file that an application accesses. When cached, content is identified by segment and downloaded in blocks. Each content item consists of an ordered collection of one or more segments of the same size (32 MB), although they can be shorter if they represent the last segment of a content item the size of which is not a multiple of the segment size.
Content server	The storage system that the peer contacts to obtain either the hashes of the content or the actual content if the content is not cached.
Hosted cache	A centralized cache comprised of blocks added by peers, available when using hosted cache mode.
Peer	The nodes participating in content caching and retrieval system. A peer is a node that both accesses the content and, if using distributed caching mode, serves the content it caches for other peers.
Segment	A subdivision of content. Each segment has the same size (32 MB in the current version), except that the last segment can be smaller if the content size is not a multiple of the standard segment sizes.
Segment hash of data	The hash of the content block hashes of every block in the segment.
Segment ID	A hash that represents the content-specific label or public identifier that is used to discover content from other peers or from the hosted cache. This identifier is disclosed freely in broadcast messages. Knowledge of this identifier does not prove authorization to access the actual content.

Segment secret	The content-specific hash that is sent to authenticated and authorized clients along with the rest of the content information. It is generated by hashing the concatenation of the segment hash of data and the server secret.
Server secret	A SHA-256 hash of an arbitrary length binary string stored on the server. The server secret is created using the <code>cifs branchcache set key</code> command.

How Data ONTAP implements BranchCache

After you configure BranchCache and enable it on a share, BranchCache-enabled clients at a branch office initially retrieve content from the storage system. The retrieved content is then cached at the branch office. If another client needs the same content, it can download content directly from cache, providing that valid hashes are available.

For caching, Data ONTAP divides file content into smaller units known as *segments*, and then subdivides segments into *blocks*. In a BranchCache environment, Data ONTAP stores and retrieves segments and blocks rather than files.

Data ONTAP computes hashes on the block and on the segment. Segment hashes provide a unit of discovery. Block hashes provide a unit of download—a client retrieves cached content in block units. Block hashes and segment hashes primarily comprise the content metadata sent to a BranchCache enabled client that requests cached content.

When a client sends a request to the storage system for metadata about content that has been cached at the branch office, Data ONTAP authenticates and authorizes the requestor, exactly as it would without using BranchCache, then sends the metadata (hashes) about the content to clients. Data ONTAP always sends the content metadata to the client to ensure that the client always receives hashes for the most up-to-date content. If the content was not previously cached or the cached data is out-of-date, then Data ONTAP sends the content to the client and computes the hashes as it sends the data. Data ONTAP then stores the hash information for subsequent client requests.

If the client requests content made up of multiple blocks or multiple segments, the client must manage the retrieval and concatenation of blocks and segments to reconstruct the original content.

How BranchCache clients retrieve cached content

After the client receives metadata for cached content from the storage system, the method the client uses to retrieve cached content depends on whether Distributed Cache mode or Hosted Cache mode is configured at the branch office.

Distributed Cache mode

If content is not yet cached at the branch office, the storage system does the following:

1. Authenticates and authorizes the client.
2. Computes hashes for the requested content.
3. Sends the content and metadata to the client.

The client then caches the data locally and makes it available to other peers.

4. If another client requests the same content from the storage system, the storage system does the following:
 - a. Authenticates and authorizes the client.
 - b. Determines whether the hashes for the requested content are still valid.
 - c. If the hashes are still valid, sends metadata to the second client.

After the second client receives the metadata, it does the following:

1. Using the information from the metadata, the second client sends out a multicast message on the local network requesting information about the location of the cached data.
2. The client with the cached data responds, encrypts the requested content, and sends it to the second client.
3. The second client decrypts the data using information sent in the metadata from the storage system, and verifies that the received data is valid and has not been modified.

Hosted Cache mode

If content is not yet cached at the branch office, the storage system does the following:

1. Authenticates and authorizes the client.
2. Computes hashes for the requested content.
3. Sends the content and metadata to the client.
 The client does not cache this content locally. Participating branch office clients are configured with the fully qualified domain name of the hosted cache server. The client uses this information to contact the caching server and offers the content identifiers to the caching server. The caching server uses these identifiers to determine that this data is not yet cached and then retrieves the offered content blocks directly from the peer that offered them.
4. If another client requests the same content from the storage system, the storage system does the following:
 - a. Authenticates and authorizes the client.
 - b. Determines whether the hashes for the requested content are still valid.
 - c. If the hashes are still valid, sends metadata to the second client.

After the second client receives the metadata, it does the following:

1. The second client is configured with the address of the hosted cache server and queries that server directly for the block ranges of the requested segments.
2. The hosted cache server encrypts the data and returns it to the second client.
3. The second client decrypts the data using information sent in the metadata from the storage system, and verifies that the received data is valid and has not been modified.

Network protocol support requirements

You must be aware of the network protocol requirements for implementing Data ONTAP BranchCache.

You can implement the Data ONTAP BranchCache feature over IPv4 and IPv6 networks using SMB 2.1 or later.

All storage systems and branch office machines participating in the BranchCache implementation must have the SMB 2.1 or later protocol enabled. SMB 2.1 has protocol extensions that allow a client to participate in a BranchCache environment. This is the minimum SMB protocol version that offers BranchCache support.

Note: Even though the Microsoft BranchCache feature supports using both the HTTP/HTTPS and SMB protocols as file access protocols, Data ONTAP BranchCache only supports the use of SMB.

Data ONTAP and Windows hosts version requirements

Your storage system and branch office systems must meet certain version requirements before you can configure BranchCache.

Before configuring BranchCache, you must ensure that the storage system and participating branch office clients support SMB 2.1 or later and support the BranchCache feature. If you configure Hosted Cache mode, you must also ensure that you use a supported host for the cache server.

- Content server: Storage system with Data ONTAP 8.1.1 or later, operating in 7-Mode
- Cache server: Windows Server 2008 R2
- Peer: Windows 7 Enterprise, Windows 7 Ultimate, or Windows Server 2008 R2

For the latest information about which Windows CIFS clients support BranchCache, see the Interoperability Matrix at support.netapp.com/NOW/products/interoperability.

Configuring BranchCache

You configure BranchCache on your storage system using Data ONTAP commands. To implement BranchCache, you must also configure your clients, and optionally your caching server at the branch offices where you want to cache content. Next you must enable BranchCache on the CIFS shares for which you want to provide BranchCache caching services.

Configuring BranchCache on the storage system

You can configure the storage system to provide BranchCache services to BranchCache-enabled Windows clients at branch offices.

Before you begin

The following requirements must be met before configuring BranchCache on your storage system:

- Data ONTAP 8.1.1 operating in 7-Mode or later must be installed.
- CIFS must be licensed and configured.

- IPv4 or IPv6 network connectivity must be configured.

Steps

1. Enable SMB 2.1:

```
options cifs.smb2.enable on
```

The command enables both SMB 2.0 and SMB 2.1.

2. Enable BranchCache:

```
options cifs.smb2_1.branch_cache.enable on
```

3. Optional: Configure the time (in seconds) to keep an unused BranchCache hash in storage system memory:

```
options cifs.smb2_1.branch_cache.hash_time_out time
```

time can be 0 through 4000000. If you do not set this option, the default is to keep hashes in memory for 3600 seconds. In the example at the end of the procedure, the value was set at 7200.

4. Set the pass phrase that the storage system uses to generate hashes, enclosing the string in quotation marks if it contains spaces:

```
cifs branchcache set key "your_pass_phrase"
```

The pass phrase is a string that the storage system uses when generating BranchCache content hashes. You can reset the pass-phrase after you configure BranchCache. Changing the pass-phrase value results in all existing cached content being identified as stale. If data is identified as stale, content is retrieved from the content server and new hashes are generated and stored on the storage system.

```
options cifs.smb2.enable on
options cifs.smb2_1.branch_cache.enable on
options cifs.smb2_1.branch_cache.hash_time_out 7200
cifs branchcache set key "your_pass_phrase"
```

After you finish

Proceed to enable BranchCache on the shares on which you want to provide caching services.

Where to get information about configuring BranchCache at the branch office

In addition to configuring BranchCache on the storage system, you must install and configure BranchCache on client computers and, optionally, the cache server at your branch office.

Instructions for configuring branch office clients and, optionally, a caching server to use BranchCache are on the Microsoft BranchCache web site: [Microsoft BranchCache: *technet.microsoft.com/EN-US/NETWORK/DD425028*](http://technet.microsoft.com/EN-US/NETWORK/DD425028)

Enabling BranchCache caching on a CIFS share

After you configure BranchCache on the storage system and at the branch office, you can enable BranchCache on each CIFS share that contains content that you want to allow clients at branch offices to cache.

BranchCache caching is enabled on a share-by-share basis. You can use the `cifs shares -add` or `cifs shares -change` commands to enable BranchCache on the CIFS share. If you enable caching on an existing CIFS share, Data ONTAP begins computing hashes and sending metadata to clients requesting content as soon as you enable BranchCache on that share.

Note: If BranchCache on a CIFS share is subsequently disabled, Data ONTAP stops sending metadata to the requesting client. A client that needs data retrieves it directly from the content server (storage system).

Considerations when enabling BranchCache on a share

You need to keep certain limitation and support considerations in mind when you enable BranchCache on CIFS shares.

The following are limitations for BranchCache:

- BranchCache supports a maximum file size of 160 MB.
For larger file sizes, the client retrieves the content from the storage system.
- If a BranchCache client requests the hash of a file in a Snapshot copy, the hash information is not returned to the client.
The client must retrieve the content from the storage system.

Data ONTAP does not support BranchCache-enabled CIFS shares with the following configurations:

- When creating a CIFS share on a volume or qtree SnapMirror destination
- When modifying an existing CIFS share on a volume or qtree SnapMirror destination
- When a broken volume or qtree SnapMirror destination has CIFS shares with BranchCache enabled and is resynchronized to the SnapMirror source

Note: CIFS shares with BranchCache enabled are supported while the volume or qtree SnapMirror destination is in a broken state.

- When the Data ONTAP BranchCache feature is disabled
- When the Data ONTAP SMB 2.1 capability is disabled with the `options cifs.smb2.enable off` command

Creating a BranchCache enabled CIFS share

You can enable BranchCache on a CIFS share when you create the share.

About this task

The `-branchcache` option is one of several available client-side caching options. Enabling BranchCache on a share also enables manual client-side caching. If you enable BranchCache on a CIFS share, you cannot enable local client-side caching using the `-auto_document_caching` or `-auto_program_caching` options.

Step

1. To create a BranchCache enabled CIFS share, enter the following command:

```
cifs shares -add share_name path -branchcache
```

share_name specifies the name of the new CIFS share.

path specifies the path to the share. Path separators can be backward or forward slashes, although Data ONTAP displays them as forward slashes.

For more information, see the `na_cifs_shares(1)` man page.

Example

The following command creates a BranchCache enabled CIFS share named `projects` in the `/vol/vol1/projects` directory:

```
cifs shares -add projects /vol/vol1/projects -branchcache
```

Modifying an existing CIFS share to enable BranchCache

You can enable BranchCache by modifying an existing CIFS share. BranchCache caching is one of several client-side caching options that you can enable on a Data ONTAP CIFS share. You can only specify one client-side caching option when modifying a share. Therefore, if you specify the share properties `-branchcache` option, you cannot specify other client-side caching options.

About this task

Since a CIFS share can have only one of the client-side caching options enabled, specifying the `-branchcache` client-side caching option replaces any client-side caching option that you specified when you created the CIFS share. If you enable BranchCache on an existing CIFS share, you cannot also enable local client-side caching using the `-auto_document_caching` or `-auto_program_caching` options. Enabling BranchCache on a share also automatically enables manual client-side caching.

Step

1. Enable BranchCache on an existing CIFS share by entering the following command:

```
cifs shares -change share_name -branchcache
```

share_name specifies the name of the new CIFS share.

For more information, see the `na_cifs_shares(1)` man page.

Example

The following command enables BranchCache on an existing CIFS share named “projects”:

```
cifs shares -change projects -branchcache
```

How Data ONTAP manages BranchCache hashes

When deciding whether you want to enable BranchCache on a share, you need to understand the extent of turnover of your hash cache, and therefore, the turnover of cached content. Understanding when hashes are invalidated can help you in this evaluation.

Data ONTAP must manage BranchCache hashes to ensure that hashes are valid and not stale. If a hash is not valid or is stale, Data ONTAP invalidates the hash and computes a new hash the next time that content is requested, assuming that BranchCache is still enabled.

When a BranchCache-enabled client sends a request to the storage system for metadata (hashes) on content that is cached at the branch office, Data ONTAP does the following:

1. Data ONTAP checks to make sure that the hash for the requested content is valid and not stale. There are a number of reasons why the hash might be invalid or stale—for example, the content was modified after the hash was generated.
2. When Data ONTAP determines that a hash is invalid or stale, Data ONTAP returns an error message to the requesting client and at the same time generates a new hash and stores it in the hash cache.
3. The requesting client sends a second request to the storage system asking for the content directly (request for data).
4. Data ONTAP then sends content retrieved directly from the storage system, along with the freshly computed metadata, to the client.
5. At the branch office, the new data and metadata replaces the stale cached data.
6. The next time a BranchCache-enabled client sends a request to the storage system for metadata on this content, the storage system provides the new metadata information to the requesting BranchCache-enabled client.
7. The BranchCache-enabled client retrieves the newly cached content from the local cache.

Note: A non BranchCache-enabled client does not send requests for metadata to the storage system. Instead it sends a request for the actual content.

Data ONTAP invalidates hashes for the following reasons:

- A hash is flushed from the cache because the BranchCache hash timeout limit was reached for that hash.
This is a tunable parameter and can be modified to meet your business requirements.
- A file is modified either through CIFS or NFS access.
- A file for which there are computed hashes is deleted.
In this case, invalidation occurs when the hash timeout limit is reached.
- A volume that contains a BranchCache-enabled share is taken offline.
- A file for which there are computed hashes is restored using the `snap restore` command.
- A volume that contains CIFS shares that are BranchCache enabled is restored using the `snap restore` command.
- A volume that contains CIFS shares that are BranchCache enabled becomes a volume SnapMirror destination.
- The server secret key is modified.

Displaying BranchCache hash statistics

You can use the `cifs branchcache hash stat` command to display statistics about BranchCache hashes.

If you use the `cifs branchcache hash stat` command without specifying any options, it reports the number of hashes that were flushed in multiple 5-minute intervals and a list of file size ranges for which hashes were requested.

Displaying hash statistics by file size

You can generate a report that displays the number of files in various size ranges for which hashes were requested. You can use this report to determine whether you should adjust the client-side configuration to use BranchCache for smaller file sizes.

About this task

Reported file size ranges are as follows:

- Less than 10 KB
- 11 KB to 100 KB
- 101 KB to 250 KB
- 251 KB to 1 MB
- 1.1 MB to 10 MB
- 11 MB to 100 MB
- Greater than 100 MB

Step

1. Display the report on the number of files in various size ranges for requested hashes by entering the following command:

```
cifs branchcache hash stat -size
```

Example

```
system1>cifs branchcache hash stat -size

Number of files (size <= 10KB) for which hashes were asked for = 5
Number of files (10KB < size <= 100KB) for which hashes were asked for = 7
Number of files (100KB < size <= 250k) for which hashes were asked for = 120
Number of files (250KB < size <= 1MB) for which hashes were asked for = 9
Number of files (1MB < size <= 10MB) for which hashes were asked for = 1245
Number of files (10MB < size <= 100MB) for which hashes were asked for = 9
Number of files (size > 100MB) for which hashes were asked for = 27
```

Displaying number of hashes flushed

You can generate a report that displays the number of BranchCache hashes that were flushed in multiples of five-minute intervals. You can use this information to determine whether you should increase the hash timeout value.

Step

1. Display the report on the number of hashes flushed with the following command:

```
cifs branchcache hash stat -flush
```

If the report shows that hashes are flushed frequently, you can increase the hash timeout value with the `cifs.smb2_1.branch_cache.hash_time_out` option.

Example

```
cifs branchcache hash stat -flush

Number of hashes flushed in duration of 600 seconds = 5
Number of hashes flushed in duration of 900 seconds = 8
Number of hashes flushed in duration of 1200 seconds = 11
Number of hashes flushed in duration of 1500 seconds = 14
Number of hashes flushed in duration of more than 1800 seconds =
22
```

Changing the BranchCache server secret key

You can change the BranchCache server secret key by rerunning the `cifs branchcache set key` command.

About this task

Changing the pass-phrase value results in all existing cached content being identified as stale. If data is identified as stale, content is retrieved from the storage system and new hashes are generated and stored on the storage system.

Step

1. Change the server secret key by using the following command:

```
cifs branchcache set key pass-phrase
```

pass-phrase specifies the text string to use as the server secret key. If the pass phrase contains any spaces, enclose the pass phrase in quotation marks.

Example

The following example sets a new server secret key pass phrase that contains spaces:

```
cifs branchcache set key "my new secret server key phrase"
```

What happens when you disable BranchCache on the storage system

If you previously configured BranchCache but do not want the branch office clients to use cached content, you can disable caching on the storage system. You must be aware of what happens when you disable BranchCache.

When you disable BranchCache, Data ONTAP no longer computes hashes or sends the metadata to the requesting client. Thereafter, when BranchCache-enabled branch office clients request metadata information for content they want to access, Data ONTAP responds with an error. The client must then send a second request, requesting the actual content. In response to the request for content, the storage system sends the actual content that is stored on the storage system.

Disabling BranchCache on the storage system

You can disable BranchCache on the storage system either by disabling the BranchCache feature or by disabling SMB 2.1.

About this task

BranchCache requires that SMB 2.1 or later be enabled on the storage system; therefore, if SMB 2.1 is disabled, the BranchCache feature is also disabled.

Step

1. Run the appropriate command:

If you want to...	Then enter the following command...
Disable BranchCache but leave SMB 2.x enabled	options cifs.smb2_1.branch_cache.enable off
Disable BranchCache and SMB 2.x	options cifs.smb2.enable off

Disabling BranchCache on an existing CIFS share

You can disable BranchCache on an existing CIFS share by changing the client-side caching option using the `cifs shares -change` command. To disable BranchCache caching, you can change the client-side caching option to either `-no_caching` or `-manual_caching`.

Step

1. Disable BranchCache on an existing CIFS share, by entering the following command:

```
cifs shares -change share_name [-no_caching | -manual_caching]
```

share_name specifies the name of the CIFS share.

For more information, see the `na_cifs_shares(1)` man page.

Example

The following command disables BranchCache on an existing CIFS share named *projects*:

```
cifs shares -change projects -no_caching
```

Managing local users and groups

This section provides information about creating and managing local users and groups on the storage system.

Managing local users

Local users can be specified in user and group lists. For example, you can specify local users in file-level ACLs and share-level ACLs. You can also add local users to local groups.

When you should create local user accounts

There are several reasons for creating local user accounts on your storage system.

You should create one or more local user accounts if your system configuration meets the following criteria:

- If, during setup, you configured the storage system to be a member of a Windows workgroup. In this case, the storage system must use the information in local user accounts to authenticate users.
- If your storage system is a member of a domain:
 - Local user accounts enable the storage system to authenticate users who try to connect to the storage system from an untrusted domain.
 - Local users can access the storage system when the domain controller is down or when network problems prevent your storage system from contacting the domain controller.

For example, you can define a BUILTIN\Administrator account that you can use to access the storage system even when the storage system fails to contact the domain controller.

Note: If, during setup, you configured your storage system to use UNIX mode for authenticating users, you should not create local user accounts. In UNIX mode, the storage system always authenticates users using the UNIX password database.

Displaying the storage system's authentication method

You can display the storage system's authentication method, and thus determine whether you should create local users and groups, by entering the `cifs sessions` command.

Step

1. Enter the following command:

```
cifs sessions
```

For more information, see the `na_cifs_sessions(1)` man page.

Limitations of local user accounts

There are several limitations with local user accounts.

- You cannot use User Manager to manage local user accounts on your storage system.
- You can use User Manager in Windows NT 4.0 only to view local user accounts.
If you use User Manager in Windows 2000, however, you cannot use the Users menu to view local users. You must use the Groups menu to display local users.
- You can create a maximum of 96 local user accounts.

Adding, displaying, and removing local user accounts

You can add, display, and remove local user accounts by using the `useradmin` command.

About this task

You use the `useradmin` command for creating, displaying, and deleting administrative users on the storage system. (You can also use this command to manage non-local users through the `domainuser` subcommand.) For information about how to use the `useradmin` command, see the section about managing local user accounts in the introduction to storage system administration in the *Data ONTAP System Administration Guide for 7-Mode*.

Note: Data ONTAP keeps a single list of user accounts created by the `useradmin` command. The same type of information exists for local user accounts and administrative user accounts. CIFS users who have local user accounts with the appropriate Admin Roles can use Windows RPC calls to log in to the storage system. For more information, see the chapter on managing Administrator access in the *Data ONTAP System Administration Guide for 7-Mode*.

Managing local groups

You can manage local groups to control which users have access to which resources.

About this task

A local group can consist of users or global groups from any trusted domains. Members of a local group can be given access to files and resources.

Membership in certain well-known local groups confers special privileges on the storage system. For example, members of BUILTIN\Power Users can manipulate shares, but have no other administrative capabilities.

CIFS clients display the name of a local group in one of the following formats:

- FILERNAME\localgroup
- BUILTIN\localgroup

Adding, displaying, and removing local groups from the Data ONTAP command line

You can add, display, and remove local groups from the Data ONTAP command line by using the `useradmin` command.

About this task

For more information, see the *Data ONTAP System Administration Guide for 7-Mode*.

Adding a local group from the MMC on a Windows client

You can add a local group from the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Local Users and Groups**.
4. Right-click **Groups**.
5. Select **New Group**.
6. In the New Group box, enter the name and description of the group.
7. Click **Create**.

Result

A new group is created on the storage system.

Adding users to a local group from the MMC on a Windows client

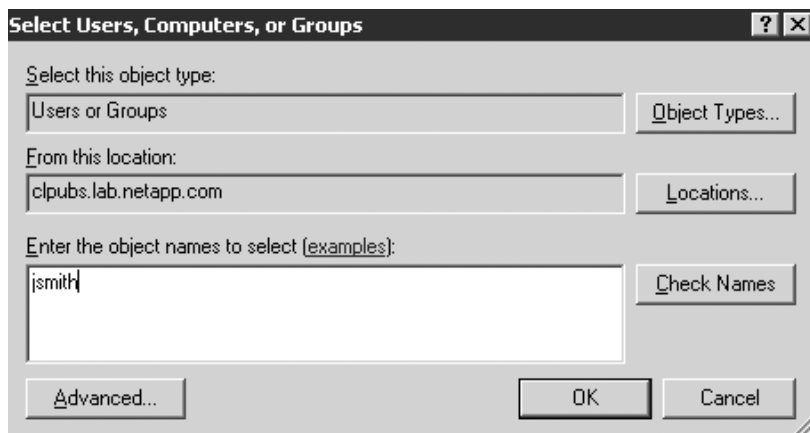
You can add users to a local group from the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Local Users and Groups**.
4. Double-click **Groups**.
5. In the right panel, right-click on the group to which you want to add a user.
6. Select **Add to Group**.

The MMC displays the Properties box.

7. In the Properties box, click **Add**.
8. In the Select Users, Computers, or Groups window, enter the name of the user in the "Enter the object names to select" box.



9. Click **OK**.

Result

The MMC adds the user to the group.

Removing a local group using the MMC on a Windows client

You can remove a local group using the MMC on a Windows client.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Select **System Tools > Local Users and Groups > Groups**.
4. In the right pane, right-click the local group that you want to remove.
5. Select **Remove**.
6. Click **OK**.

Result

The MMC removes the local group.

How SnapMirror works with local groups

Because the mirror is a read-only volume and you cannot change ACLs or permissions on it, do not use local groups in ACLs for files to be replicated by SnapMirror.

If you use the SnapMirror feature to copy a volume to another storage system and the volume has an ACL for a local group, the ACL does not apply on the mirror. This is because the group is local to the source storage system.

If you want to use local groups in ACLs for files to be replicated by SnapMirror, you can do this using the MultiStore product. For more information about the MultiStore product, see the MultiStore Management Guide.

Applying Group Policy Objects

Your storage system supports Group Policy Objects (GPOs), a collection of group policy settings that apply to computers in an Active Directory environment.

When CIFS and GPOs are enabled on your storage system, Data ONTAP sends LDAP queries to the Active Directory server requesting GPO information. If there are GPO definitions that are applicable to your storage system, the Active Directory server returns the following GPO information:

- GPO name
- Current GPO version
- Location of the GPO definition
- Lists of UUIDs (universally unique identifiers) for GPO policy sets

Note: For more information about Windows GPOs, see the Microsoft web site.

Although not all GPOs are applicable to your storage system, the storage system can recognize and process the relevant set of GPOs.

The following GPOs are currently supported for your storage system:

- Startup and shutdown scripts
- Group Policy refresh interval for computer (includes random offset)
- File System security policy
- Restricted Groups security policy
- Event Log policy
- Auditing policy
- Take Ownership user right
- Manage Auditing and Security Logs user right

Note: Event Log and Auditing policy settings are applied differently to storage systems than to Windows systems. Also, if you define a Take Ownership user list or group list that does not contain Windows built-in administrator accounts, these administrators lose Take Ownership privileges.

Requirements for using GPOs with storage systems

To use GPOs with your storage system, several requirements must be met.

Ensure that the following requirements are met:

- CIFS is licensed and enabled on the storage system.
- CIFS is configured.
- The storage system is joined to a Windows domain version 2000 or later
- GPOs are configured on a Windows Active Directory server by associating the storage system with an Organizational Unit (OU).
- GPO support is enabled on the storage system.

Associating the storage system with a different OU

The `cifs setup` process uses the default OU CN=Computers during setup unless a different OU is specified. You can change it to a different OU after setup.

Steps

1. On the Windows server, open the Active Directory Users and Computers tree.
2. Locate the storage system's Active Directory object.
3. Right-click the object and select **Move**.
4. Select the OU that you want to associate with the storage system.

Result

The storage system object is placed in the selected OU.

Enabling or disabling GPO support on a storage system

You can enable or disable GPO support on the storage system by setting the `cifs.gpo.enable` option.

Step

1. Perform one of the following actions.

If you want to...	Then...
Enable GPO	Enter the following command: <code>options cifs.gpo.enable on</code>
Disable GPO	Enter the following command: <code>options cifs.gpo.enable off</code>

Managing GPOs on the storage system

You can create, display, configure the updating of, and troubleshoot the Group Policy Objects on the storage system.

Creating File System security GPOs

You can specify GPO File System security settings directly on Data ONTAP file system objects (directories or files).

About this task

GPO File System security settings are propagated down the directory hierarchy; that is, when you set a GPO security setting on a directory, those settings are applied to objects within that directory.

Note: These File System security settings can only be applied in mixed or NTFS volumes or qtrees. They cannot be applied to a file or directory in a UNIX volume or qtree.

File System security ACL propagation is limited to about 280 levels of directory hierarchy.

Steps

1. On the Windows server, open the Active Directory Users and Computers tree.
2. Right-click the Organization Unit (OU) that contains the storage system.
3. Select the Group Policy tab, and select **New**.
4. Enter a name for the new GPO.
5. Highlight the new GPO and select **Edit**.

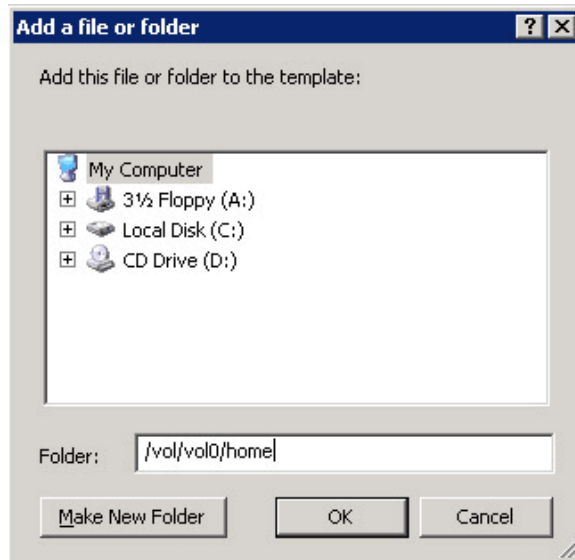
The Group Policy Object Editor appears.

6. Double-click **Computer Configuration > Windows Settings > Security Settings**.
7. Right-click File System and select **Add File**.

The "Add a file or folder" box appears.

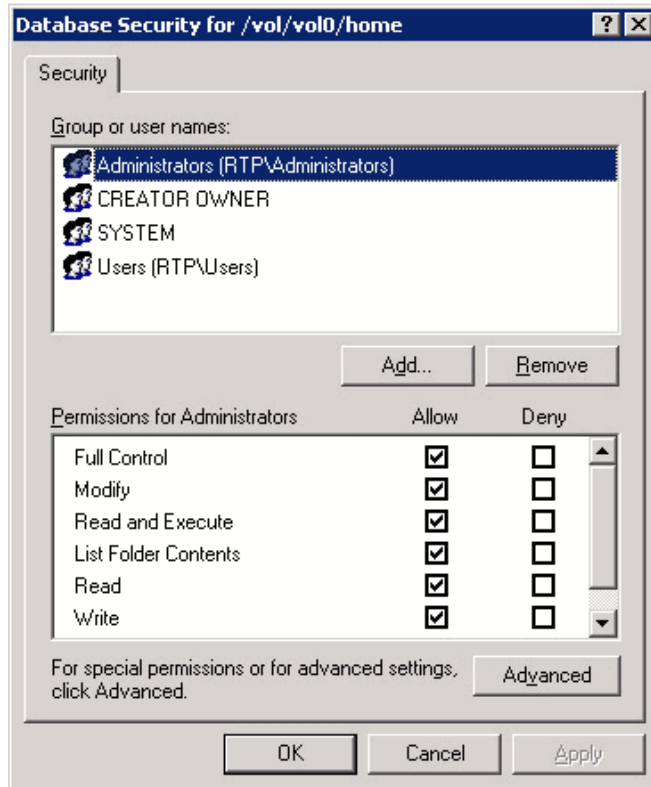
Note: Do not select the option to browse the local server's drives.

8. In the Folder field, enter the storage system path on which to apply the GPO; then click **OK**.



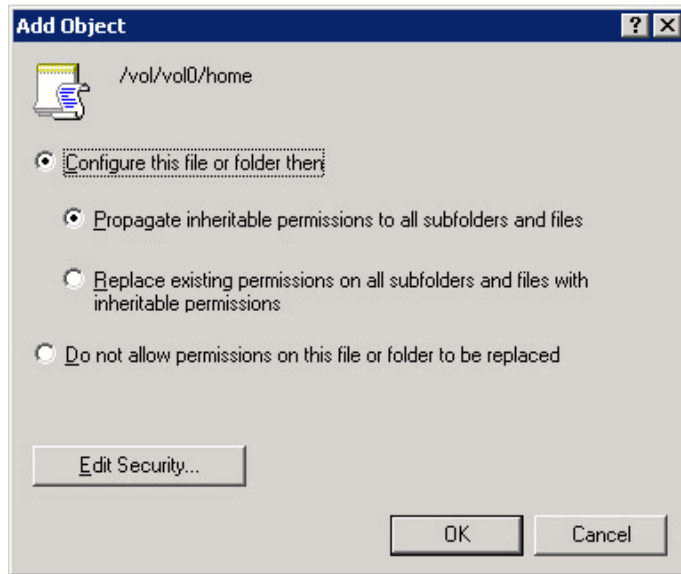
The Database Security window opens.

9. In the Database Security window, set the permissions you want; then click **OK**.




The Add Object window opens.

10. In the Add Object window, select the ACL inheritance you want; then click **OK**.



The Group Policy Editor displays the new object name.



11. Close the Group Policy Editor and the OU Properties dialog box.

12. On the storage system, enter the following command to retrieve and apply the new GPO:

```
cifs gpupdate
```

If you do not explicitly apply the new GPO with the `cifs gpupdate` command, the storage system applies the new GPO the next time it queries the Active Directory server (that is, within 90 minutes).

Displaying current GPOs and their effects

You can use the `cifs gpresult` command to display GPOs currently in effect for the storage system and the results of those GPOs.

About this task

The `cifs gpresult` command simulates the output of the Windows 2000/XP `gpresult.exe /force` command.

Note: The `cifs gpresult` command displays only those group policy settings that are relevant to your storage system and current Data ONTAP release.

Step

1. Enter the following command:

```
cifs gpresult [-r] [-d] [-v]
```

`-r` displays the results of applying current GPOs to the storage system.

`-v` generates a verbose display, including information about applicable GPOs and the results of applying them.

`-d` dumps the output from `cifs gpresult -v` to the `/etc/ad/gpresult_timestamp` file.

If you do not specify any options, the command displays information about the GPOs currently applicable to the storage system, including name, version and location.

Updating GPO settings

Data ONTAP retrieves and applies GPO changes every 90 minutes and refreshes security settings every 16 hours, but you can also force an update by entering the `cifs gpupdate` command.

About this task

Group Policy settings on the storage system can be updated in three ways:

- All GPOs are verified every 90 minutes.
By default, Data ONTAP queries Active Directory for changes to GPOs. If the GPO version numbers recorded in Active Directory are higher than those on the storage system, Data ONTAP retrieves and applies the new GPOs. If the version numbers are the same, GPOs on the storage system are not updated.
- Security Settings GPOs are refreshed every 16 hours.

Data ONTAP retrieves and applies Security Settings GPOs every 16 hours, whether or not these GPOs have changed.

Note: The 16 hour default value cannot be changed in the current Data ONTAP version. It is a Windows client default setting.

- All GPOs can be updated on demand with a Data ONTAP command.
This command simulates the Windows 2000/XP `gpupdate.exe /force` command.

Step

1. To update Group Policy settings manually, enter the following command:

```
cifs gpupdate
```

Troubleshooting GPO update problems

If Data ONTAP does not display messages on the console indicating that it has successfully applied GPO settings—for example, after you issue the `cifs gpupdate` command—you should check diagnostic information about storage system GPO connections by using the `cifs.gpo.trace.enable` option.

About this task

When updated Policy Settings have been applied on storage system GPOs, messages similar to one or both of the following appear on the storage system console:

```
CIFS GPO System: GPO processing is successfully completed.
CIFS GPO System: GPO Security processing is completed.
```

Steps

1. Enter the following command to enable GPO tracing:

```
options cifs.gpo.trace.enable on
```

2. Enter the following command to update GPO settings:

```
cifs gpupdate
```

GPO trace messages are written to the console and message logs until GPO tracing is turned off.

You see messages similar to the following that include Active Directory information about GPOs:

```
CIFS GPO Trace: Site DN: cn=Default-First-Site-Name,
cn=sites,CN=Configuration,DC=cifs,DC=lab,DC=company, DC=com.
CIFS GPO Trace: Domain DN: dc=CIFS,dc=LAB,dc=COMPANY, dc=COM.
CIFS GPO Trace: Filer DN: cn=user1,ou=gpo_ou,dc=cifs,
dc=lab,dc=company,dc=com.
```

```
CIFS GPO Trace: Processing GPO[0]: T_sub.
CIFS: Warning for server \\LAB-A0: Connection terminated.
```

3. Enter the following command to disable GPO tracing:

```
options cifs.gpo.trace.enable off
```

About startup and shutdown scripts on a storage system

When GPOs have been enabled on a storage system and specified in the Active Directory domain, Data ONTAP runs startup and shutdown scripts automatically whenever you start or shutdown CIFS.

The storage system accesses the scripts from the Domain Controller's sysvol directory and saves these files locally in the `/etc/ad` directory.

Note: Although the storage system periodically retrieves updates to the startup and shutdown scripts, startup scripts are not applied until the next time CIFS restarts.

About the `/etc/ad` directory

When GPO support is enabled on the storage system for the first time using the `cifs.gpo.enable` option, an `/etc/ad` directory is created.

This directory is used as a repository for the following files:

- GPO startup and shutdown scripts retrieved from the domain controller.
- Output for the `cifs gpresult -d` command.

Configuration requirements for Data ONTAP pathnames

The format of target file or directory names must be recognized by Data ONTAP and must be in absolute or relative form.

Here is more information about the path name forms:

- Absolute pathname—for example, `/vol/vol0/home`.
When an absolute pathname is supplied, Data ONTAP applies File System security settings to the specified target file or files within the target directories. In this example, the settings are applied to the `/home` directory in the storage system root volume.
- Relative pathname—for example, `/home`.
When a relative pathname is supplied (any pathname that does not begin with `/vol`), Data ONTAP applies File System security settings to any target file or directory containing the specified element. This is a convenient way to apply settings to multiple parallel targets in a single storage system; in this example, the settings are applied to all vFiler units with `/home` directories.

Improving client performance with opportunistic and lease oplocks

Oplocks (opportunistic locks) and lease oplocks enable a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then read from or write to a file without regularly reminding the server that it needs access to the file in question. This improves performance by reducing network traffic.

Lease oplocks are an enhanced form of oplocks available with the SMB 2.1 protocol and later. Lease oplocks allow a client to obtain and preserve client caching state across multiple file open calls originating from itself.

Write cache data-loss considerations when using oplocks

Under some circumstances, if a process has an exclusive oplock on a file and a second process attempts to open the file, the first process must invalidate cached data and flush writes and locks. The client must then relinquish the oplock and access to the file. If there is a network failure during this flush, cached write data might be lost.

- Data-loss possibilities

Any application that has write-cached data can lose that data under the following set of circumstances:

- The connection is made using SMB 1.0.
- It has an exclusive oplock on the file.
- It is told to either break that oplock or close the file.
- During the process of flushing the write cache, the network or target system generates an error.

- Error handling and write completion

The cache itself does not have any error handling—the applications do. When the application makes a write to the cache, the write is always completed. If the cache, in turn, makes a write to the target system over a network, it must assume that the write is completed because if it does not, the data is lost.

How lease oplocks provide enhanced locking

Data ONTAP avoids unnecessary oplock breaks by using lease oplocks. Available with SMB 2.1 as an enhancement over traditional opportunistic locks, the lease oplock model is based on a key called a *lease key* rather than a file handle that is used with traditional oplocks.

When a client makes multiple file open calls, lease oplocks allow the client to obtain and preserve client caching state across the file open calls that it originates. The client generates a lease key and uses it to link together multiple handles to the same file. Data ONTAP tracks the state using the client ID, which is a combination of the client's GUID and the lease key. Operations on handles with the same lease key do not affect the oplock state.

The client can upgrade the lease state on a subsequent file open call, or Data ONTAP can downgrade the lease state if another client makes a file open call. If the client and storage system connect using the SMB 2.1 dialect, lease oplocks are used instead of traditional oplocks.

For more information, see the SMB 2.1 protocol specification.

Enabling or disabling oplocks and lease oplocks on the storage system

You can enable or disable oplocks and lease oplocks on the storage system. Oplocks let clients lock files and locally cache information while preventing another user from changing the file. This increases performance for many file operations.

About this task

CIFS oplocks and lease oplocks are enabled on your storage system by default. You can leave CIFS oplocks and lease oplocks enabled in most cases, but you might want to disable them under the following circumstances:

- You are using a database application whose documentation recommends that oplocks be disabled.
- The CIFS clients are on an unreliable network.
- You are handling critical data and you cannot afford even the slightest data loss.

Enabling CIFS oplocks and lease oplocks does not override any client-specific settings. Disabling CIFS oplocks and lease oplocks disables all oplocks to or from the storage system. You can enable or disable CIFS opportunistic oplocks at individual clients using a Windows registry setting.

Step

1. Perform the applicable action:

If you want to...	Enter the command...
Enable oplocks and lease oplocks on the storage system	<code>options cifs.oplocks.enable on</code>
Disable oplocks and lease oplocks on the storage system	<code>options cifs.oplocks.enable off</code>

If the `cifs.oplocks.enable` option is set to `on`, the oplock setting per `qtree` takes effect. Otherwise, the oplocks for all `qtrees` are disabled regardless of the per-`qtree` oplock setting.

Enabling or disabling oplocks and lease oplocks on a qtree

If oplocks and lease oplocks are enabled at the storage system level, you can enable or disable oplocks and lease oplocks on an individual `qtree`.

Step

1. Perform the applicable action:

If you want to...	Then...
Enable oplocks and lease oplocks on a qtree	Enter the following command: qtree oplocks qtree_name enable
Disable oplocks and lease oplocks on a qtree	Enter the following command: qtree oplocks qtree_name disable

If the `cifs.oplocks.enable` option is set to on, the `qtree oplocks` command for a qtree takes effect immediately. If the `cifs.oplocks.enable` option is set to off, the `qtree oplocks` command does not take effect until you change the option to on.

Monitoring lease oplock status

Starting with Data ONTAP 8.1.1 and later in the 8.1 release family, you can monitor and display information about lease oplocks.

About this task

If a lease oplock is present, the `oplock` field from the output of the `lock status` command holds lease oplock information. If a durable handle is used on that lock, the `oplock` field also contains durable handle information.

The following lease oplock values are displayed in the `oplock` field:

- `Lease-RWH`
This is a lease oplock with a Read-Write-Batch lock.
- `Lease-RW`
This is a lease oplock with a Read-Write lock.
- `Lease-RH`
This is a lease oplock with a Read-Batch lock.
- `Lease-R`
This is a lease oplock with a Read lock.
- `Lease-None`
This is a lease oplock with no lease lock.

Note: The lease lock uses a durable handle if the value of the `oplock` field is either `Lease-RWH` or `Lease-RH`.

Step

1. Use the `lock status` command to display lease oplock information.

Example

The following command displays lock information about CIFS protocol locks, including information about lease locks:

```
lock status -p cifs
CIFS path= host=10.0.0.2(WIN1) owner= state=GRANTED mode=None-denyN
oplock=Lease-None
durable_state= sid=0x1eea09f9 fileid=0x00000060
```

Changing the delay time for sending oplock breaks

If a client that owns a file oplock sends a file open request, it is temporarily vulnerable to a “race condition” that can occur if the storage system requests an oplock break. To prevent this condition, the storage system delays sending an oplock break according to the delay time value (in milliseconds) specified by the `cifs.oplocks.opendelta` option.

About this task

By default, the default delay time is 0 milliseconds. If your storage system must support some older Microsoft Windows clients, including Microsoft Windows NT 4.0 without the latest Service Pack and Microsoft Windows NT 3.5.1, you should keep this default value to prevent the performance problem described in Microsoft Knowledge Base article 163525.

If you don't have clients running older version of Windows, you can set the delay time to another value, such as 8. This means that after the storage system receives or responds to a request to open a file, the storage system will make sure that 8 milliseconds have elapsed before sending an oplock break to that client.

You might want to increase the delay time if you issue the `cifs stat` command and the output shows a non-zero value for the `OpLkBkNoBreakAck` field.

You might also want to increase the delay time for sending oplock breaks if you see syslog messages similar to the following:

```
Mon Jan 21 15:18:38 PST [CIFSAdmin:warning]: oplock break timed out to
station JOHN-PC for file \\FILER\share\subdir\file.txt
```

Step

1. Enter the following command:

```
options cifs.oplocks.opendelta time
```

time is the delay in milliseconds.

Setting the `cifs.oplocks.opendelta` option postpones the sending of oplock break requests to clients that have just opened files. You must consult technical support if you are considering setting this value higher than 35.

Managing authentication and network services

This section provides information about storage system authentication, as well as procedures for managing the older NetBIOS protocol.

Understanding authentication issues

Your storage system supports three types of authentication: UNIX authentication, Windows workgroup authentication, and Kerberos authentication.

About UNIX authentication

Using UNIX mode, authentication is performed using entries in the `/etc/passwd` file and/or using NIS/LDAP-based authentication.

Using UNIX authentication:

- Passwords are sent “in the clear” (unencrypted).
- Authenticated users are given credentials with no unique, secure user identification (SID).
- The storage system verifies the received password against a “hash” (algorithmic variant) of the user password.

Passwords are not stored on the storage system.

In order to provide UNIX client authentication, the following items must be configured:

- Client information must be in the storage system `/etc/passwd` file.
- Client information must be entered in NIS and/or LDAP.
- Windows client registries must be modified to allow plain text passwords.

Because UNIX authentication transmits unencrypted passwords, Windows clients require a registry edit to enable them to send passwords without encryption. Clients that are not properly configured to send clear text passwords to the storage system might be denied access and display an error message similar to the following:

```
System error 1240 has occurred.
The account is not authorized to login from this station.
```

Refer to Microsoft support for information to enable plain text passwords, to allow clients to use UNIX authentication.

About Windows workgroup authentication

Workgroup authentication allows local Windows client access.

The following facts apply to workgroup authentication:

- Does not rely upon a domain controller

- Limits storage system access to 96 local clients
- Is managed using the storage system's `useradmin` command

Kerberos authentication for CIFS

With Kerberos authentication, upon connection to your storage system, the client negotiates the highest possible security level. However, if the client cannot use Kerberos authentication, Microsoft NTLM or NTLM V2 is used to authenticate with the storage system.

Note: Extended security features are only available to clients that are members of a Windows Active Directory domain.

Setting the storage system's minimum security level

To set the storage system's minimum security level (that is, the minimum level of the security tokens that the storage system accepts from clients), you can set the `cifs.LMCompatibilityLevel` option. By default, this option is set to 1.

Step

1. Enter the following command:

`options cifs.LMCompatibilityLevel minimum_level`

minimum_level is the minimum level of security tokens that the storage system accepts from clients, as defined in the following table.

Value	Description
1 (default)	The storage system accepts LM, NTLM, and NTLMv2 session security; it also accepts NTLMv2 and Kerberos authentication.
2	The storage system accepts NTLM and NTLMv2 session security; it also accepts NTLMv2 and Kerberos authentication. The storage system denies LM authentication.
3	The storage system accepts NTLMv2 session security; it also accepts NTLMv2 and Kerberos authentication. The storage system denies LM and NTLM authentication.
4	The storage system accepts NTLMv2 and Kerberos authentication. The storage system denies LM, NTLM, and NTLMv2 session security.
5	The storage system accepts Kerberos authentication only.

Preventing Kerberos passive replay attacks

Kerberos replay cache prevents passive replay attacks by storing user authenticators on the storage system for a short time, and by insuring that authenticators are not reused in subsequent Kerberos tickets.

About this task

Storing and comparing Kerberos authenticators can result in a substantial performance penalty for certain storage system workloads. For this reason, the `kerberos.replay_cache.enable` option is set to `off` by default.

Step

1. Perform one of the following actions.

If you want to...	Then...
Enable the Kerberos replay cache	Enter the following command: <code>option kerberos.replay_cache.enable on</code>
Disable the Kerberos replay cache	Enter the following command: <code>option kerberos.replay_cache.enable off</code>

Selecting domain controllers and LDAP servers

Upon startup and as listed below, your storage system searches for a Windows domain controller. This section describes how and when the storage system finds and selects domain controllers.

About this task

The storage system searches for domain controllers where any of the following is true:

- The storage system has been started or rebooted.
- A `cifs resetdc` command has been issued.
- Four hours have elapsed since the last search.

Note: Active Directory LDAP servers are searched for under the same conditions.

Understanding the domain controller discovery process

When you run CIFS in a domain environment, your storage system attempts to rediscover all of its domain controllers by sending Internet Control Message Protocol (ICMP) packets once every 4 hours. Doing so enables it to verify that the current domain controller is still accessible and to prioritize available domain controllers using the packets' round trip time.

If a storage system loses access to a domain controller with a very good connection rate and has to go to a backup domain controller with a slower rate, the storage system rediscovers domain controllers

every 2 minutes until it finds a better connection. After the storage system finds that connection, it connects to the new domain controller and returns to sending discovery packets every 4 hours.

The following table describes the domain controller discovery process and priority groups. The storage system only progresses to a lower priority group when it has failed to contact all domain controllers in the priority group above it.

Note: For Active Directory environments, site membership is one of the criteria by which the storage system selects domain controllers (when no preferred domain controllers are available). Therefore, it is important to have the Sites and Services configured properly (with the storage system's subnet information included in the same site as the storage system).

Domain controller category	Priority groups: Order in which domain controllers are selected
Preferred: Controllers in the <code>prefdc</code> list	Group 1: Preferred domain controllers are selected by the order in which the controllers appear in the <code>prefdc</code> list.
Favored: Controllers that share the same Active Directory site membership with the storage system (This category is empty for storage systems in Windows NT environments.)	Group 2: Domain controllers from which a response was received within one second of being pinged, in the order of fastest response time. Group 3: Domain controllers that did not respond within one second, but share the same subnet as the storage system. Group 4: All non-local domain controllers that did not respond within one second of being pinged
Other: Controllers that do not share site membership	Group 5: Domain controllers from which a response was received within one second of being pinged, in the order of fastest response time. Group 6: Domain controllers that did not respond within one second, but share the same subnet as the storage system. Group 7: All non-local domain controllers that did not respond within one second of being pinged.

Note: Because site membership is specific to Active Directory domains, there is no “favored” category for Windows NT4 domains, nor for mixed-mode domains in which your storage system is configured as an NT4 server. In these environments, all domain controllers found through discovery are assigned the category “other.”

Ensuring successful authentication with Windows Server 2008 R2 domain controllers

If your CIFS domain contains Windows Server 2008 R2 domain controllers, you need to take certain steps to ensure successful authentication.

About this task

Data ONTAP requires either a writable domain controller or a read-only domain controller that is configured to replicate passwords for the storage system.

Step

1. Perform one of the following actions:

If you want to...	Then...
Ensure that the storage system uses a writable domain controller	Use the <code>cifs prefdc add</code> command to add the writable domain controller to the list of preferred domain controllers on the storage system.
Ensure successful authentication with read-only domain controllers	Configure the writable domain controller to allow the read-only domain controller to replicate passwords for the storage system. For more information, see the article about Windows Server 2008 Password Replication Policy Administration at technet.microsoft.com .

Specifying a list of preferred domain controllers and LDAP servers

You can specify a list of preferred domain controllers and LDAP servers using the `cifs prefdc add` command.

Step

1. Enter the following command:

```
cifs prefdc add domain address [address ...]
```

domain specifies the domain for which you want to specify domain controllers or LDAP servers.

address specifies the IP address of the domain controller or LDAP server.

Example

The following command specifies two preferred domain controllers for the lab domain.

```
cifs prefdc add lab 10.10.10.10 10.10.10.11
```

Note: To force the storage system to use a revised list of preferred domain controllers, or LDAP servers, use the `cifs resetdc` command.

Deleting servers from the preferred domain controller list

You can use the `cifs prefcdc delete` command to delete entries from the preferred domain controller list. You should use this command for example to remove servers from the list that are not online anymore or no longer serving as domain controllers.

Steps

1. Enter the following command:

```
cifs prefcdc delete domain
```

domain is the domain where the preferred domain controller or LDAP server resides.

2. Enter the following command:

```
cifs resetdc [domain]
```

domain is the domain you specified in step one.

After you delete a domain from the `prefcdc` list, you should always enter the `cifs resetdc` command to update the storage system's available domain controller information. The storage system does not update the domain controller discovery information from network services when the preferred domain controller list is updated. Failure to reset the domain controller information can cause a connection failure, if the storage system tries to establish a connection with an unavailable domain controller (or LDAP server).

Note: Storage systems do not automatically perform domain controller discovery operations upon restart; restarting the storage system does not update the available domain controller and LDAP server list.

Result

The storage system disconnects and searches for a domain controller in the order specified in the revised `prefcdc` list.

Example

The following command deletes the server `lab` from the list of preferred domain controllers:

```
vs1::> cifs prefcdc delete lab
```

Displaying a list of preferred domain controllers and LDAP servers

You can use the `cifs prefdc print` command to display a list of preferred domain controllers and LDAP servers.

Step

1. Enter the following command:

```
cifs prefdc print [domain]
```

domain is the domain for which you want to display domain controllers. When a domain is not specified, this command displays preferred domain controllers for all domains.

Example

The following command displays the preferred controllers and LDAP servers for the lab domain:

```
cifs prefdc print lab
```

Reestablishing the storage system connection with a domain

You can use the `cifs resetdc` command to reestablish the storage system connection with a domain.

About this task

The following procedure disconnects your storage system from the current domain controller and establishes a connection between the storage system and a preferred domain controller. It also forces domain controller discovery, updating the list of available domain controllers.

Note: This procedure also reestablishes LDAP connections, and performs LDAP server discovery.

Step

1. Enter the following command:

```
cifs resetdc [domain]
```

domain is the domain from which the storage system disconnects. If it is omitted, the storage system disconnects from the domain in which the storage system is installed.

Example

The following command disconnects the storage system from the domain controllers for the lab domain:

```
cifs resetdc lab
```

Using null sessions to access storage in non-Kerberos environments

Null session access provides permissions for network resources, such as storage system data, to client-based services running under the local system. A null session occurs when a client process uses the “system” account to access a network resource. Null session configuration is specific to non-Kerberos authentication.

How the storage system provides null session access

Because null session shares do not require authentication, clients that require null session access must have their IP addresses mapped on the storage system.

By default, unmapped null session clients can access certain Data ONTAP system services, such as share enumeration, but they are restricted from accessing any storage system data.

Note: Data ONTAP supports Windows RestrictAnonymous registry setting values with the `cifs.restrict_anonymous` option. This allows you to control the extent to which unmapped null users can view or access system resources. For example, you can disable share enumeration and access to the IPC\$ share (the hidden named pipe share). For more information, see the `options(1)` man page.

Unless otherwise configured, a client running a local process that requests storage system access through a null session is a member only of nonrestrictive groups, such as “everyone.” To limit null session access to selected storage system resources, you might want to create a group to which all null session clients belong; creating this group enables you to restrict storage system access and to set storage system resource permissions that apply specifically to null session clients.

Data ONTAP provides a mapping syntax in the `/etc/usermap.cfg` file to specify the IP address of clients allowed access to storage system resources using a null user session. After you create a group for null users, you can specify access restrictions for storage system resources and resource permissions that apply only to null sessions.

Any null user accessing the storage system from a mapped IP address is granted mapped user permissions. Consider appropriate precautions to prevent unauthorized access to storage systems mapped with null users. For maximum protection, place the storage system and all clients requiring null user storage system access on a separate network, to eliminate the possibility of IP address “spoofing.”

Granting null users access to file system shares

You can allow access to your storage system resources by null session clients by assigning a group to be used by null session clients and recording the IP addresses of null session clients to add to the storage system's list of clients allowed to access data using null sessions

Steps

1. Open the `/etc/usermap.cfg` file.
2. Add an entry for each null user using the following format:
`IPqual:"" => unixacct`
 IPqual specifies either an IP address (hostname or numeric dot-format) or a subnet (IP address + network mask); "" indicates null user; => indicates the mapping direction; and *unixacct* is the UNIX account (from `/etc/passwd` or NIS) that the mapped null user will have.
3. Set the `cifs.mapped_null_user_extra_group` option to the group name you intend to use for null session clients.
4. Set permissions to allow appropriate access rights to null session clients.

Examples

```
10.10.20.19:"" => exchuser
```

```
192.168.78.0/255.255.255.0:"" => iisuser
```

The client at IP address 10.10.20.19 is allowed null session access to the storage system. The null user account is mapped to a UNIX account called *exchuser*, which must exist in the `/etc/passwd` or NIS database.

Also, any clients establishing a connection from the 192.168.78.0 class C subnet are allowed null session access and are mapped to the UNIX account *iisuser*. Other null user connections to the storage system are not allowed.

Using machine accounts to access storage in Kerberos environments

Machine accounts are subjected to the same Kerberos authentication as user accounts, so they do not need to be mapped on the storage system.

About this task

When authenticated using Kerberos, clients that run local processes using the “system” account assign those processes to the machine account when accessing remote resources. The machine account is assigned the computer name registered with the domain controller, followed by a dollar sign (\$).

Preventing machine accounts from accessing data

By default, machine accounts (like any other authenticated user) always have access to data shares. However, for security reasons, you might want to prevent services running on a Kerberos-enabled client from accessing data using CIFS.

About this task

Note: Disabling machine account access to data shares can cause a number of services to fail, such as offline folders and roaming profiles. Be sure to evaluate your storage system service needs before disabling machine account access.

Step

1. Enter the following command:

```
cifs access -delete share_name -m
```

Note: Access to the IPC\$ share (the hidden named pipe share) cannot be changed and is always permitted.

For more information, see the `cifs_access(1)` man page.

Creating NetBIOS aliases for the storage system

You can create NetBIOS aliases by setting the `cifs.netbios_aliases` option or by editing the `/etc/cifs_nbalias.cfg` file.

About this task

NetBIOS aliases are alternative names for your storage system. You can connect to the storage system using any of the names in the list.

With the `cifs.netbios_aliases` option, you can create NetBIOS aliases as a comma-separated list. This list allows up to 255 characters, including commas. The `/etc/cifs_nbalias.cfg` file allows up to 200 entries.

Creating NetBIOS aliases from the command line

You can create NetBIOS aliases from the command line by setting the `cifs.netbios_aliases` option.

Steps

1. Enter the following command:

```
options cifs.netbios_aliases name,...
```

You can enter up to 255 characters, including commas.

Example

```
options cifs.netbios_aliases alias1,alias2,alias3
```

2. Enter the following command to process the entries:

```
cifs nbalias load
```

Creating NetBIOS aliases in the /etc/cifs_nbalias.cfg file

You can create NetBIOS aliases in the `/etc/cifs_nbalias.cfg` file.

About this task

Data ONTAP creates a default `cifs_nbalias.cfg` file in the `/etc` directory when CIFS starts, if the file does not already exist. Changes to this file are processed automatically whenever CIFS starts. You can also process changes to this file using the command `cifs nbalias load`.

Steps

1. Open the `/etc/cifs_nbalias.cfg` file for editing.
2. Enter NetBIOS aliases in the `/etc/cifs_nbalias.cfg` file, one entry per line.

Note: You can enter up to 200 NetBIOS aliases in the file, using either ASCII or Unicode characters.

3. Enter the following command to process the entries:

```
cifs nbalias load
```

Displaying the list of NetBIOS aliases

You can display the list of NetBIOS aliases by entering the `cifs nbalias` command.

Step

1. Enter the following command:

```
cifs nbalias
```

Disabling NetBIOS over TCP

If NetBIOS over TCP has been disabled in your Windows 2000 network, you can use the `cifs.netbios_over_tcp.enable` option to disable NetBIOS over TCP on your storage system.

About this task

NetBIOS over TCP is the standard protocol used for CIFS prior to Windows 2000. The option to use this protocol, `cifs.netbios_over_tcp.enable`, is enabled on your storage system by default. It corresponds to the “Enable NetBIOS over TCP” setting in the Windows 2000 Advanced TCP/IP settings tab.

To verify the status of NetBIOS over TCP on your storage system, use the `nbtstat` command, as described in the `nbtstat(1)` man page.

In order to disable NetBIOS over TCP, all storage system clients must be running Windows 2000 or later. After you disable NetBIOS over TCP, you can use only Windows 2000 or later domain controllers and virus scanners.

Note: After you disable NetBIOS over TCP, clients no longer receive Data ONTAP notification messages, such as shutdown messages and vscan warnings.

Step

1. To disable NetBIOS over TCP, enter the following command:

```
options cifs.netbios_over_tcp.enable off
```

Monitoring CIFS activity

This section provides information about monitoring CIFS sessions activity and collecting storage system statistics.

About this task

You can display the following types of session information:

- A summary of session information, which includes storage system information and the number of open shares and files opened by each connected user.
- Share and file information about one connected user or all connected users, which includes
 - The names of shares opened by a specified connected user or all connected users
 - The access levels of opened files
 - Security information about a specified connected user or all connected users, which includes the UNIX UID and a list of UNIX groups and Windows groups to which the user belongs.

Note: The number of open shares shown in the session information includes the hidden IPC\$ share.

Different ways to specify a user

When displaying session information about a connected user, you can specify the user by the user name or the IP address of the workstation. In addition, if the user is connected to your storage system from a pre-Windows 2000 client, you can specify the name of the workstation.

Clients sometimes connect with an unauthenticated “null” session. Such sessions are sometimes used by clients to enumerate shares. If a client has only the null session connected to the storage system, you will see the following status message:

```
User (or PC) not logged in
```


Displaying a summary of session information

You can use the `cifs sessions` command to display a summary of session information.

Step

1. Enter the following command:

```
cifs sessions
```

Timing out idle sessions

You can specify the amount of time that elapses (in seconds) before Data ONTAP disconnects an idle session.

About this task

If a user does not have a file opened on your storage system, the session is considered idle. By default, Data ONTAP disconnects a session after it has been idle for 30 minutes.

If an idle session is disconnected, it will automatically reconnect the next time the client accesses the storage system.

Step

1. Enter the following command:

```
options cifs.idle_timeout time
```

time specifies the new idle time before disconnecting in seconds.

The new value for this option takes effect immediately.

Tracking statistics

Using the `stats` commands, you can view system statistics to track performance.

About this task

The `stats` command is not specific to CIFS-related statistics. The two `stats` commands that output statistics data are `stats show` (for real-time statistical data) and `stats stop` (when you are tracking statistics over a range of time). Note that the `cifs stats` command is still available.

The statistics displayed by the `stats` command are accumulated in counters. You reference a specific counter using a hierarchical name with components:

object_name:instance_name:counter_name. For example, a counter might be named `system:system:cifs_ops`. You can use the `stats list` command to determine the *object_names*, *instance_names*, and *counter_names* available on your storage system.

The output of the `stats show` command provides data describing the storage system at the moment you issued the command. To track statistics over time, use the `stats start` command to mark the

beginning of the time period you want to track, and the `stats stop` command to mark the end of the time period for which you want to collect statistical data. Data ONTAP outputs the collected data as soon as you enter the `stats stop` command.

Data ONTAP enables you to use the `stats start` and `stats stop` commands to track different statistics concurrently. To do this, you can enter an instance (`-i`) argument with the `stats start` and `stats stop` commands.

For more information about usage and syntax, see the `stats(1)` man page.

Steps

1. Enter the following command to view a list of objects that are tracked by the `stats` command:

```
stats list objects
```

Data ONTAP returns a list of objects you can view by using the `stats show object_name` command.

2. Enter the following command to view a list of statistics instances:

```
stats list instances
```

Data ONTAP returns a list of instances you can view by using the `stats show` command. You can use these instances to focus the output of the `stats show` command.

3. Enter the following command to view a list of statistics counters:

```
stats list counters
```

Data ONTAP returns a list of counters you can view by using the `stats show` command.

4. Enter the following command to receive a description of all counters, instances, or objects:

```
stats explain counters
```

Data ONTAP returns a description of all counters, instances, and objects you can use to focus the output of the `stats show` command.

Viewing specific statistics

When you know the objects, instances, and counters you can monitor to track individual statistics, you can use them as command-line arguments to focus the output of the `stats show` command.

About this task

For more information, see the `stats(1)` man page.

Step

1. Enter the following command:

```
stats show [[object_name][:instance_name][:counter_name]]
```

Data ONTAP returns the specific statistics you request.

Saving and reusing statistics queries

You can store and reuse “preset” statistics queries you commonly perform. Preset queries are stored in XML files, in the following location and naming format: `/etc/stats/preset/preset_name.xml`. For information about how to store and reuse queries, see the `stats_preset(5)` man page.

CIFS resource limitations

Access to some CIFS resources is limited by your storage system's memory and the maximum memory available for CIFS services.

These resources include:

- Connections
- Shares
- Share connections
- Open files
- Locked files
- Locks

Note: If your storage system is not able to obtain sufficient resources in these categories, contact technical support.

Managing CIFS services

This section provides information about managing CIFS services on the storage system.

Disconnecting clients using the MMC

You can disconnect clients using the MMC.

Steps

1. Connect the MMC to the storage system.
2. If it is not already selected, in the left pane, select **Computer Management**.
3. Double-click **System Tools > Shared Folders > Sessions**.
4. Perform one of the following actions:

If you want to...	Then...
Disconnect specific clients	<ol style="list-style-type: none"> a. Right-click the client's name. b. Select Close Session. c. Click OK.

If you want to...	Then...
Disconnect all clients	<ul style="list-style-type: none">a. Right-click on Sessions.b. Select Disconnect All Sessions.c. Click Yes.

Disconnecting a selected user from the command line

You can use the `cifs terminate` command to disconnect a selected user from the command line.

Steps

1. To display a list of connected clients, enter the following command:

```
cifs sessions *
```

2. To disconnect a client, enter the following command:

```
cifs terminate client_name_or_IP_address [-t time]
```

client_name_or_IP_address specifies the name or IP address of the workstation that you want to disconnect from the storage system.

time specifies the number of minutes before the client is disconnected from the storage system. Entering 0 disconnects the client immediately.

Note: If you do not specify time and Data ONTAP detects an open file with the client, Data ONTAP prompts you for the number of minutes it should wait before it disconnects the client.

Example

The following command sends a message to the workstation named `jsmith-pc`, notifying the user of the impending disconnection. Five minutes after you enter the command, `jsmith-pc` is disconnected from the storage system.

```
cifs terminate jsmith-pc -t 5
```

Disabling CIFS for the entire storage system

The disabling of CIFS service is not persistent across reboots. If you reboot the storage system after disabling CIFS service, Data ONTAP automatically restarts CIFS.

Steps

1. To disable CIFS service, enter the following command:

```
cifs terminate [-t time]
```

time is the number of minutes before the storage system disconnects all clients and terminates CIFS service. Entering 0 makes the command take effect immediately.

Note: If you enter the `cifs terminate` command without an argument and Data ONTAP detects an open file with any client, Data ONTAP prompts you for the number of minutes it should wait before it disconnects the client.

2. Perform one of the following actions:

If you want CIFS service to...	Then...
Restart automatically after the next storage system reboot	Do nothing.
Not restart automatically after the next storage system reboot	Rename the <code>/etc/cifsconfig.cfg</code> file.

Result

Data ONTAP sends a message to all connected clients, notifying the users of the impending disconnection. After the specified time has elapsed, the storage system disconnects all clients and stops providing CIFS service.

After you disable CIFS for the entire storage system, most `cifs` commands become unavailable. You can use the following `cifs` commands with CIFS disabled:

- `cifs prefcd`
- `cifs restart`
- `cifs setup`
- `cifs testdc`

Specifying which users receive CIFS shutdown messages

When you issue the `cifs terminate` command, by default Data ONTAP sends a message to all open client connections to notify users when CIFS service will be disconnected. You can change the default setting so that Data ONTAP never sends these messages or sends them only to connected clients that have open files.

Step

1. Enter the following command:

```
options cifs.shutdown_msg_level {0 | 1 | 2}
```

Use 0 to never send CIFS shutdown messages.

Use 1 to send messages only to connected clients that have open files.

Use 2 to send messages to all open connections, which is the default setting.

Restarting CIFS service

You can restart CIFS service by entering the `cifs restart` command.

Step

1. Enter the following command:

```
cifs restart
```

Result

The storage system connects to the domain controller and restarts CIFS service.

Sending a message to users on a storage system

You can send a message to all users on your storage system to tell them of important events. The message appears in an alert box on the users' computers.

About this task

Data ONTAP automatically sends a message to connected users after you enter the `cifs terminate` command. However, if you want to send a message without stopping CIFS service, for example, to tell users to close all files, you can use Server Manager or the Data ONTAP command line to send a message.

Some clients might not receive broadcast messages. The following limitations and prerequisites apply to this feature:

- Windows 95 and Windows for Workgroups clients must have the WinPopup program configured.
- Windows 2003 and Windows XP Service Pack 2 clients must have the messenger service enabled.
By default, it is disabled.
- Messages to users can only be seen by Windows clients connected using NetBIOS over TCP.

Note: Network configuration can also affect which clients receive broadcast messages.

Step

1. Perform one of the following actions.

If you want to...	Enter the command...
Send a message to all CIFS users connected to the storage system	cifs broadcast * "message"
Send a message to a specific CIFS user connected to the storage system	cifs broadcast <i>client_name</i> "message"

If you want to...	Enter the command...
Send a message to all CIFS users connected to a particular volume	<code>cifs broadcast -v volume "message"</code>

Displaying and changing the description of the storage system

Adding an informative description enables you to distinguish your storage system from other computers on the network.

About this task

The description of your storage system appears in the Comment field when you browse the network. Initially, the storage system has no description. The description can be up to 48 characters.

Steps

1. Enter the following command to display the current description:

```
cifs comment
```

2. Enter the following command to change the description:

```
cifs comment "description"
```

Changing the storage system computer account password

You can change the storage system computer account password manually or you can use options to configure the storage system to update it automatically.

Changing the storage system computer account password manually

You can change the storage system's Active Directory computer account password manually. The `cifs changefilerpwd` command instructs the storage system to change its domain account password immediately.

About this task

Note: For Active Directory domains with multiple domain controllers, a password change might inhibit CIFS connections for a short time while Active Directory propagates the new password among the domain controllers.

Step

1. Enter the following command:

```
cifs changefilerpwd
```

The storage system responds with the following message:

```
password change scheduled.
```

The password change is scheduled, and should take place within a minute.

Configuring the storage system for automatic computer account password changes

You can configure Data ONTAP to automatically schedule Windows Active Directory computer account password changes.

Before you begin

CIFS must be licensed and configured on your storage system. Additionally your storage system must be joined to an Active Directory domain.

About this task

There are three options to configure automatic computer account password changes:

- When set to on, the `cifs.W2K_password_change` option causes the storage system to change its domain password once in every time period defined by the `cifs.W2K_password_change_interval` option.
The duration is counted in weeks. The password change occurs randomly within the time period specified by the `cifs.W2K_password_change_within` option. The default for this option is off.
- The `cifs.W2K_password_change_interval` option specifies the time duration in weeks after which the domain computer account password change is triggered.
The default value is 4w, with a range of 1w through 8w. The password change is attempted at approximately 01:00 a.m. on the Sunday morning following the day when the configured time duration expires. The value in this option has no effect if the `cifs.W2K_password_change` option is set to off.
- The `cifs.W2K_password_change_within` option sets the time duration in hours within which domain account password change attempts are made.
The default value is 1h, with a range of 1h through 6h. Attempts are made after the expiration of the configured W2K password change interval. The password change is attempted at a random interval between 01:00 a.m. and the time set by the `cifs.W2K_password_change_within` option.

Steps

1. Enable automatic computer account password changes:
`cifs.W2K_password_change on`
2. Configure the `cifs.W2K_password_change_interval` option:
`cifs.W2K_password_change_interval interval`
3. Configure the `cifs.W2K_password_change_within` option:
`cifs.W2K_password_change_within hours_within`

Example

In the following example, the storage system is configured for automatic domain computer account password changes using a 2-week interval and a 3-hour "change within" window:

```
system1>options cifs.W2K_password_change on
system1>options cifs.W2K_password_change_interval 2w
system1>options cifs.W2K_password_change_within 3h
```

About file management using Windows administrative tools

You can accomplish some CIFS file access management tasks by using Windows administrative tools.

The following Windows administrative tools are compatible with Data ONTAP:

- Computer Management snap-ins for Microsoft Management Console (MMC) to manage users and groups
- Microsoft Active Directory Users MMC snap-ins
- Microsoft Event Viewer
- Microsoft Perfmon

The procedures for managing the storage system using the Microsoft administrative tools listed above are similar to those for managing a Windows server. The procedures in this chapter provide information for Data ONTAP administration tasks that differ from a Windows server.

Unlike text you enter through Windows server administration tools, the Data ONTAP command line is case-sensitive. For example, when you specify a volume name in Windows, you can type in either lowercase or uppercase letters. You cannot use Windows tools to create a qtree named Test at the same level as a qtree named TEST, because Windows tools do not make a distinction between these names. You can create and distinguish these two qtrees from the Data ONTAP command line.

The following limitations apply to NT User Manager when you use NT User Manager for your storage system:

- Although the storage system supports local users, you cannot use the New Users command on the User menu to create or delete local user accounts.
- The Policies menu is disabled, but some policies can be controlled through options or group membership.

The following NT Server Manager features are not supported because they are not applicable to Data ONTAP:

- Stopping and starting services
- Specifying the recipients of alerts

Troubleshooting access control problems

To troubleshoot access control problems (that is, to determine why a client or user is given or denied access to a file on the storage system when you think it should not be), you can use the `sectrace` command.

Adding permission tracing filters

You can add permission tracing filters to instruct Data ONTAP to log information in the system log about why the storage system allows or denies a client or user to perform an operation.

About this task

Adding permission tracing filters has a minor effect on storage system performance; therefore, you should add permission tracing filters for debugging purposes only. When you are done debugging, you should remove all permission tracing filters. Furthermore, the filtering criteria you specify should be as specific as possible so that Data ONTAP does not send a large number of EMS messages to the console.

Keep the following limitations in mind:

- You can add a maximum of 10 permission tracing filters per vFiler.
- You can add permission tracing filters for CIFS requests only.

Step

1. Enter the following command:

```
sectrace add [-ip ip_address] [-ntuser nt_username] [-unixuser  
unix_username] [-path path_prefix] [-a]
```

ip_address specifies the IP address of the client attempting access.

nt_username specifies the Windows NT user name of the user attempting access.

unix_username specifies the UNIX user name of the user attempting access. You cannot specify a UNIX user name if you specify an NT user name.

path_prefix specifies the prefix of the path name of the files to trace access to. For example, specify `/vol/vol0/home/file` to trace access to all files having names that start with "file" in the `/vol/vol0/home/` directory, such as `/vol/vol0/home/file100` and `/vol/vol0/home/file200`.

`-a` specifies that the storage system should trace requests that it allows as well as requests that it denies.

Examples

The following command adds a permission tracing filter to trace all access requests from a client with an IP address of 192.168.10.23 that Data ONTAP denies.

```
sectrace add -ip 192.168.10.23
```

The following command adds a permission tracing filter to trace all access requests from the UNIX user `foo` to the path `/vol/vol0/home4` that Data ONTAP allows or denies:

```
sectrace add -unixuser foo -path /vol/vol0/home4 -a
```

Removing permission tracing filters

Because permission tracing filters have a minor impact on storage system performance, you should remove them when you are done debugging access errors.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Remove all permission tracing filters	sectrace delete all
Remove one permission tracing filter	sectrace delete index When you add a permission tracing filter, Data ONTAP assigns it an index between 1 and 10. <i>index</i> specifies the index of the permission tracing filter to delete.

Example

The following command removes the permission tracing filter with an index of 1:

```
sectrace delete 1
```

Displaying permission tracing filters

You can use the `sectrace show` command to display the permission tracing filters on a storage system or vFiler.

Step

1. Enter the following command:

```
sectrace show [index]
```

When you add a permission tracing filter, Data ONTAP assigns it an index between 1 and 10. *index* specifies the index of the permission tracing filter to display. If you do not specify an index, Data ONTAP displays all of the permission tracing filters.

Example

The following command displays all permission tracing filters on a storage system:

sectrace show

Data ONTAP displays all of the permission tracing filters in output like this:

```
Sectrace filter: 1
Hits: 5
Path: /vol/vol1/unix1/file1.txt
NT User: CIFS-DOM\harry
Trace DENY and ALLOW events
Sectrace filter: 2
Hits: 7
IP Addr: 10.30.43.42
Path: /vol/vol1/mixed1/dir1/file1.txt
NT User: CIFS-DOM\chris
Trace DENY and ALLOW events
Sectrace filter: 3
Hits: 1
Path: /vol/vol1/mixed1/file2.txt
NT User: CIFS-DOM\chris
Trace DENY events
```

Finding out why Data ONTAP allowed or denied access

Data ONTAP logs an EMS message to the console whenever the criteria for a permission tracing filter are met. To get more information about why Data ONTAP allowed or denied access to a particular client or user, you can use the `sectrace print-status` command.

Step

1. Enter the following command:

sectrace print-status *status_code*

status_code corresponds to the value of the "Status: " tag in the storage system log for the request that the storage system allowed or denied.

Example

Suppose you added a permission tracing filter that caused Data ONTAP to log the following EMS message to the console:

```
Thu Dec 20 13:06:58 GMT [sectrace.filter.allowed:info]: [sectrace
index: 1] Access allowed because 'Read Control, Read Attributes,
```

```
Read EA, Read' permission (0x20089) is granted on file or directory
(Access allowed by unix permissions for group) - Status:
1:6047397839364:0:0 - 10.73.9.89 - NT user name: CIFS-DOM\harry -
UNIX user name: harry(4096) - Qtree security style is MIXED and
unix permissions are set on file/directory - Path: /vol/voll/mixed1/
file1.txt
```

To get more information about why Data ONTAP allowed this particular user to access this particular file, enter the following command:

```
sectrace print-status 1:6047397839364:0:0
```

Note: When invoking the `sectrace print-status` command, you must specify the status code from the "Status:" line of the corresponding error message.

In response, Data ONTAP provides more information, such as the following:

```
secAccess allowed because 'Traverse' permission is granted on
requested path.
- Access allowed by unix permissions for others.
- Access allowed because requested permission is granted on file or
directory.
- Access allowed by share-level ACL.
- Access allowed by unix permissions for group.
trace print-status 1:6047397839364:0:0
```

Using FPolicy

You can use FPolicy to allow partner applications connected to your storage systems to monitor and set file access permissions.

Introduction to FPolicy

An introduction to FPolicy includes the system architecture, information on how it works, FPolicy's common use cases, various FPolicy applications, and limitations of FPolicy.

What FPolicy is

FPolicy is an infrastructure component of Data ONTAP that enables partner applications connected to your storage systems to monitor and set file access permissions.

Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access. This enables partners to set restrictions on files that are created or accessed on the storage system.

FPolicy allows you to create file policies that specify file operation permissions according to file type. For example, you can restrict certain file types, such as JPEG and .mp3 files, from being stored on the storage system.

When FPolicy was first introduced in Data ONTAP 6.4, it only supported the CIFS protocol. Support for the NFS protocol was added in Data ONTAP 7.0. However, FPolicy requires CIFS to be licensed even for NFS specific events.

FPolicy determines how the storage system handles requests from individual client systems for operations such as create, open, rename, and delete. The storage system maintains a set of properties for FPolicy, including the policy name and whether that policy is active. You can set these properties for FPolicy using the storage system console commands.

The FPolicy interface is a Data ONTAP API (called ONTAPI) that runs on a Distributed Computing Environment (DCE) and uses Remote Procedure Calls (RPC). Using these tools, the external applications can register as FPolicy servers.

The FPolicy interface allows a programmer to implement sophisticated file screening functionality on a storage system or NearStore system from an external application running on a separate platform.

An application utilizing the FPolicy interface can perform the following actions:

- Register one or more FPolicy servers with one or more storage systems
- Receive notifications for file operations such as opening, creating, or renaming files
- Block access to any file it has received a notification for

The following protocols are supported by FPolicy:

- CIFS
- NFS (version 2, version 3, version 4)

The following filters can be used by an FPolicy server:

- Protocol
- Volume name
- File extension
- Offline bit
- Operations

File screening in Data ONTAP can be enabled in two ways.

- Using external file screening software

The file screening software runs on a client that functions as a file screening server. File screening software provides flexible control and filtering of file content.

Note: For optimal performance, you should configure the FPolicy server to be on the same subnet as the storage system.

- Using native file blocking

The file screening software runs natively on the storage system. Native file blocking provides simple denial of restricted file types.

How FPolicy works

An FPolicy server should be registered with a storage system before it can be configured to send notification for access by clients using NFS and CIFS.

After registering the FPolicy server with the storage system, when a client makes a request for access to a file, the storage system notifies the FPolicy server for events that are registered for notification. The storage system sends information about client access to the FPolicy server as part of the notification sent on the client request. The information sent to the FPolicy server includes the file name, path name, client information, protocol information, and operations requested by the client. Based on the information received and how the FPolicy server is configured, the FPolicy server responds to the client's request. The FPolicy server communicates to the storage system whether to allow or deny the request from the client.

You can use file policies to specify file or directory operations, and place restrictions on them. Upon receiving a file or directory operation request (such as open, write, create, or rename), Data ONTAP checks the file policies before permitting the operation.

If the policy specifies screening for that file based on its extension, file screening takes place either on a file screening server or on the storage system. The following list describes these methods of file screening:

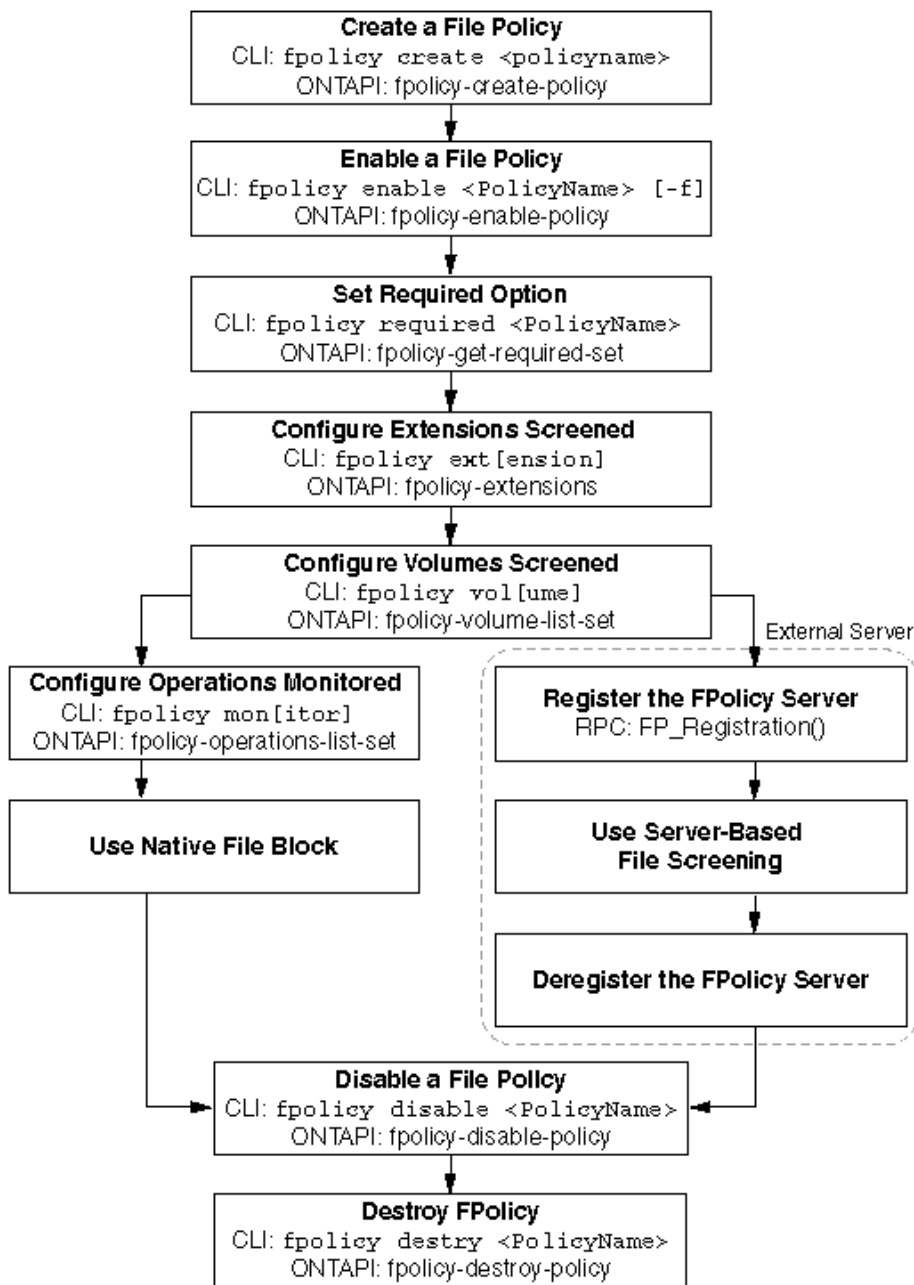
- On a file screening server (using external screening software):
The notification is sent to the file screening server to be screened and the file screening server, which applies rules to determine whether the storage system should allow the requested file operation. The file screening server then sends a response to the storage system to either allow or block the requested file operation.
- On the storage system (using native file blocking):
The request is denied and the file operation is blocked.

Related concepts

What native file blocking is on page 163

FPolicy work flowchart

The flowchart gives an overview of the usage model for FPolicy.



FPolicy in the storage environment


When a client requests a file, the request is sent to the Protocol Stack. If the FPolicy feature is enabled, the Protocol Stack identifies CIFS and NFS requests and marks them for FPolicy screening.

The request is then sent to the WAFL module. The WAFL module redirects the request from the storage system to the FPolicy server. The WAFL module sends the file request to the FPolicy engine.

The FPolicy engine consists of the FPolicy infrastructure, ONTAPIs and RPCs. It sends the request to the FPolicy server as an RPC call. When the FPolicy server returns the response, the FPolicy engine responds to the client request. This response is forwarded to the WAFL module which in turn forwards it to the Protocol Stack and then sends it to the client.

If the file access is allowed, the client is provided with the file. If file access is denied, an appropriate response is sent to the client. For CIFS clients, when file access is denied, the `STATUS_ACCESS_DENIED` error message is displayed.

The system architecture diagram provides an overview of the entire system architecture and indicates the FPolicy infrastructure in various layers of Data ONTAP.



What the multiple server configuration feature is

FPolicy supports load sharing among different servers registered for one policy. FPolicy allows more than one server to register for one policy. These servers can register as primary or secondary servers.

In a scenario where more than one FPolicy server registers to the same policy on the storage system, all FPolicy notifications for that policy are load-shared among the FPolicy servers. The storage

system performs load sharing by sending successive notifications to the FPolicy server that has the least number of outstanding requests. However, FPolicy gives priority to primary servers over secondary servers. If there is a mixed configuration of both primary and secondary servers registered to a given policy, the FPolicy notifications will be distributed among the primary servers.

If no primary server is available, the secondary server shares the notifications. If a primary server is available, the storage system sends the requests to the primary server and not to the secondary server.

If any one of the FPolicy servers hits the limit of maximum outstanding requests, which is 50, FPolicy redirects the notification to the other active servers. When all the registered servers reach this limit of maximum outstanding requests, all notifications are queued in the throttle queue.

The server configuration depends on the type of feature. For instance, features such as pass-through read, file size, and owner are server-based features. You need to enable these features on specific servers. However, features such as notification of permission changes, inode-to-file path, and offline bit are policy-wide features. That is, when you enable these features on one policy, the feature gets updated to all the FPolicy servers that use this policy.

Limitations of FPolicy

FPolicy limitations can be classified into protocol, screening and general limitations.

Following are the protocol limitations of FPolicy:

- FPolicy supports only CIFS and NFS protocols.
However, there are some operations for the CIFS and NFS protocols that FPolicy does not monitor, such as NFSv4 operations related to locking and delegation, session-related operations (SMB_COM_SESSION_SETUP_ANDX), operations not relevant to file system activity (print-related operations), and so on.
- FPolicy does not support other protocols such as FTP, HTTP, WebDAV, FileIO, and so on.
- You cannot configure CIFS and NFS operations separately on the same policy.

Following are the screening limitations of FPolicy:

- You must set up file screening on an entire volume.
You cannot screen individual qtrees and directories.
- FPolicy supports screening of CIFS operation on alternate data streams.
However, FPolicy does not support screening of NFS operations on alternate data streams.
- When you register multiple servers, the policy of all the servers connected changes based on the settings of the server that registers last.
- Multiple instances of FPolicy server from the same IP address cannot register to same policy.
- If the CIFS system resources used by FPolicy are exhausted, the CIFS screening by the FPolicy engine will stop.

Use of FPolicy within Data ONTAP


FPolicy can be used for native file blocking on a storage system.

What native file blocking is

Using native file blocking, you can deny any of the file or directory operations mentioned in the monitoring list.

The same set of filters and protocols that are supported by server-based file screening are also supported for native file blocking. You can configure native file blocking policy and FPolicy server-based file screening simultaneously, on different policies.

The following image displays the processing of a client request when native file blocking is enabled. The numbers depict the order of the request flow.



When a CIFS or NFS client makes a request, if native file blocking is enabled, the file is screened at the storage system. If the request matches the screening requirements the request is denied.

Native file blocking can be performed on any of the following operations:

- File open
- File create
- File rename
- File close
- File delete
- File read
- File write
- Directory delete
- Directory rename

- Directory create
- Getattr (NFS only)
- Setattr
- Create hard link (NFS only)
- Create symlink (NFS only)
- Lookup (NFS only)
- Notification of permission changes (CIFS only)
 - Change of owner
 - Change of group
 - Change of system ACL (SACL)
 - Change of discretionary ACL (DACL)

Related concepts

[How to use native file blocking](#) on page 164

Related references

[Events monitored through CIFS](#) on page 166

[Events monitored through NFS](#) on page 167

How to use native file blocking

To use native file blocking, you first create an FPolicy and then configure the FPolicy to provide notifications for certain operations. The native file blocking feature is available by default with Data ONTAP.

Native file blocking enables you to deny any of the file operations mentioned in the file screening section. Access to files with particular extensions can be blocked.

For example, to block files that contain .mp3 extensions, you configure a policy to provide notifications for certain operations with target file extensions of .mp3. The policy is configured to deny .mp3 file requests. When a client requests a file with an .mp3 extension, the storage system denies access to the file, based on the native file blocking configuration.

You can configure native file blocking and server-based file screening applications at the same time.

Note: The native file blocking feature only screens files based on the extensions and not on the content of the file.

Configuring native file blocking

To configure native file blocking, you create a policy and then configure it with a list of file extensions to block.

Before you begin

The CIFS protocol needs to be licensed and configured.

Steps

1. Create a file policy by using the following CLI command:

```
fpolicy create PolicyName Policytype
```

Example

To create a screening policy named “mp3blocker”, enter the following command:

```
fpolicy create mp3blocker screen
```

FPolicy creates the file policy with the specified policy name, using the `screen` policy type.

2. Configure the policy to monitor the mp3 extension by entering the following command:

```
fpolicy extensions include set PolicyName ext-list
```

Example

To configure the policy to monitor the mp3 extension, enter the following command:

```
fpolicy extensions include set mp3blocker mp3
```

3. Set the operations and protocols monitored by the policy by entering the following command:

```
fpolicy monitor {add|remove|set} PolicyName [-p protocols] [-f] op-spec
```

PolicyName is the name of the policy that you want to add operations to.

protocols is the set of protocols that you want to enable monitoring for. Use `cifs` to monitor CIFS requests, `nfs` to monitor NFS requests, or `cifs,nfs` to monitor both.

`-f` forces the policy to be enabled even if there are no servers available to enforce the policy.

op-spec is the list of operations you want to add.

Example

To replace the mp3blocker policy list of operations monitored for CIFS and NFS operations, enter the following command:

```
fpolicy monitor set mp3blocker -p cifs,nfs create,rename
```

Specify the `create` option to prevent creation of mp3 files. In addition, to ensure that an mp3 file is not copied onto the storage system with a different extension and renamed, also specify the `rename` option.

This CLI command sets specific operations to be monitored.

4. To enable mandatory screening, set the `required` option to `on`:

```
fpolicy options PolicyName required on
```

Example

To enable mandatory screening on the mp3blocker policy, enter the following command:

```
fpolicy options mp3blocker required on
```

This CLI command makes file screening mandatory before the files can be accessed.

5. Enable the FPolicy feature by entering the following command:

```
fpolicy enable PolicyName [-f]
```

Example

To enable the FPolicy mp3blocker, enter the following command:

```
fpolicy enable mp3blocker
```

This CLI command enables the file policy.

Result

After completing the steps, if a client tries to perform an operation that uses a blocked file, the operation fails and a `STATUS_ACCESS_DENIED` error code is sent.

Related concepts

[How to monitor operations using FPolicy](#) on page 209

Related tasks

[Creating a file policy](#) on page 169

[Specifying mandatory file screening](#) on page 170

[Enabling or disabling the FPolicy feature](#) on page 168

Events monitored through CIFS

FPolicy can monitor many CIFS events.

The following table lists the CIFS operations that FPolicy can monitor and a brief description of how FPolicy handles each operation.

Events	Description
File open	Notification sent when a file is opened
File create	Notification sent when a file is created
File rename	Notification sent when a file name is changed
File close	Notification sent when a file is closed
File delete	Notification sent when a file is deleted
File read	Notification sent when a file is read
File write	Notification sent when a file is changed
Directory delete	Notification sent when a directory is deleted
Directory rename	Notification sent when a directory name is changed
Directory create	Notification sent when a directory is created
Setattr	Notification sent when attribute information is set

Events monitored through NFS

FPolicy can monitor many NFS events.

The following table lists the NFS operations that FPolicy can monitor, and a brief description of each operation.

Events	Description
File open	Notification sent when a file is opened
File create	Notification sent when a file is created
File rename	Notification sent when a file name is changed
File close	Notification sent when a file is closed
File delete	Notification sent when a file is deleted
File read	Notification sent when a file is read
File write	Notification sent when a file is changed
Directory delete	Notification sent when a directory is deleted
Directory rename	Notification sent when a directory name is changed
Directory create	Notification sent when a directory is created
setattr	Notification sent when attribute information is set

Events	Description
getattr	Notification sent when attribute information is requested
link	Notification sent when a hard link is created
symlink	Notification sent when a symbolic link is created
Lookup	Notification sent when an NFS lookup occurs

How to work with FPolicy

Using CLI commands, you can create, enable, and configure FPolicy, monitor operations, and screen files based on volumes and extensions.

How to set up FPolicy

FPolicy can be set up using simple CLI commands.

For more information, you can see Technical Report TR 3640 for recommendations when setting up your FPolicy configuration.

Related information

Technical Report: FPolicy Safeguards in Data ONTAP: media.netapp.com/documents/tr-3640.pdf

Enabling or disabling the FPolicy feature

FPolicy is enabled by default when the CIFS protocol is licensed and configured. You can set the `fpolicy.enable` option to manually enable or disable the FPolicy feature.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Enable FPolicy	<code>options fpolicy.enable on</code>
Disable FPolicy	<code>options fpolicy.enable off</code>

Disabling the FPolicy feature overrides the enable or disable settings for individual policies and will disable all policies.

Creating a file policy

To set up a file policy, you first need to create it. To create a file policy, you use the `create` command.

About this task

To configure policies for notifications, create a file policy. A file policy can then be configured to send notifications, to the FPolicy server, for particular file operation requests or for native file blocking.

The `create` command creates a new file policy with a unique policy name.

After the new file policy is created, you can set the options and determine the requests that need to be screened for certain extensions.

Step

1. To create a file policy, enter the following command:

```
fpolicy create PolicyName policytype
```

PolicyName is the name of the file policy that you want to create. The policy name should be unique and not more than 80 characters long. The file policy name can consist of Unicode characters. The only special characters from the ASCII character set allowed in the policy name are the underscore (_) and the hyphen (-). In addition to not allowing most special characters in new policy names, FPolicy truncates the existing policy names that contains a "." (dot) in them by dropping the characters after and including the dot. Any options configured on this file policy will be lost after the upgrade.

policytype is the policy group to which this file policy should belong. Currently, the only policy type supported by FPolicy is `screen`.

Example

```
fpolicy create policy1 screen
```

Result

A file policy is created using the policy name *policy1* specified using the `screen` policy type.

Note: You can create and use up to 20 file policies for each vFiler unit at one time.

After you finish

For the file policy to work and take effect, enable the created file policy.

Related tasks

[Enabling or disabling the FPolicy feature](#) on page 168

Enabling the file policy

After you create a file policy, it is necessary that you enable it before notification policies can be configured. To enable a file policy, you can use the `enable` command.

Step

1. To enable the file policy, enter the following command:

```
fpolicy enable PolicyName
```

PolicyName is the name of the policy that you want to enable.

Example

```
fpolicy enable policy1
```

Result

The specified file policy is enabled.

Note: To activate the file policy, make sure that the `fpolicy.enable` option is set to `on`.

Specifying mandatory file screening

The `required` option determines if file screening should be mandatory.

About this task

When the `required` option is set to `on`, file screening becomes mandatory. If an FPolicy server is not available, since screening cannot be performed, the client request is denied. Use this option to enable native file blocking as well.

When the `required` option is set to `off`, file screening is not mandatory. If an FPolicy server is not connected, operations are permitted without screening.

Step

1. To make file screening mandatory, enter the following command:

```
fpolicy options PolicyName required on
```

PolicyName is the name of the policy for which you want to set the required option.

Result

This option is set to `off`, by default. If you turn on the `required` option for a policy when no file screening servers are available, the native file blocking feature blocks access to files specified in that policy.

Note: If you do not want to make file screening mandatory, set the same command to `off`.

Related concepts

What native file blocking is on page 163

Displaying information for a file policy

Important information on a particular file policy can be displayed using the `fpolicy show` command.

Step

1. Enter the following command:

```
fpolicy show PolicyName
```

PolicyName is the name of the file policy for which you want to view information.

Result

The `show` command displays the following information about a particular file policy:

- Status of the file policy
- List of operations monitored
- List of volumes screened
- List of extensions screened
- Total time that the server has been connected
- Number of requests screened
- Number of requests denied
- Number of requests blocked locally

Displaying information for all file policies

You can display important information about all the file policies by using the `fpolicy` command.

Step

1. Enter the following command:

```
fpolicy
```

Result

The `fpolicy` command displays the following information about all existing file policies:

- The list of FPolicy servers registered
- Status of all file policies
- List of operations monitored by each file policy
- List of volumes screened by each file policy
- List of extensions screened by each file policy
- Total time that the server has been connected

- Number of requests screened by each file policy
- Number of requests denied by each file policy
- Number of requests blocked locally

Disabling a file policy

When a file policy is disabled, the operations that are specified for that particular file policy will not be monitored. When a particular file policy is disabled, no file request notification is sent to the FPolicy server even if the FPolicy server is registered with the storage system.

Step

1. To disable a file policy, enter the following command:

```
fpolicy disable PolicyName
```

Example

```
fpolicy disable policy1
```

Destroying a file policy

Destroying a file policy immediately removes an existing file policy from the connected storage system.

About this task

To destroy or delete a particular file policy, use the `destroy` command. You must disable the file policy before destroying it. If an FPolicy server is connected to a file policy, the FPolicy server is deregistered.

Step

1. To destroy a file policy and remove it from a list of file policies, enter the following command:

```
fpolicy destroy PolicyName
```

Example

```
fpolicy destroy policy1
```

PolicyName is the name of the file policy you want to delete.

Result

When you enter this command, the specified file policy is destroyed or deleted from the list of policies.

Stopping server screening for disconnected CIFS requests

You can choose to stop the server from screening CIFS requests whose session is disconnected by enabling the `cifs_disconnect_check` option.

About this task

You can filter out redundant requests and reduce the load on the FPolicy server.

Step

1. To enable this feature on individual file policies, enter the following command:

```
fpolicy options PolicyName cifs_disconnect_check on
```

PolicyName is the name of the file policy for which you are enabling the check.

Result

Note: By default, this option is set to `off`.

Example

To enable `cifs_disconnect_check` for file policy `p1`, use the following command:

```
filer> fpolicy options p1 cifs_disconnect_check
fpolicy options p1 cifs_disconnect_check: off
filer> fpolicy options p1 cifs_disconnect_check on
```

Setting server timeout

You can set the limit on how long the system waits for the FPolicy server to respond to a request. You can set this limit individually for each file policy. This option ensures that the FPolicy server is making progress.

Step

1. To set the timeout value for individual file policies, enter the following command:

```
fpolicy options PolicyName serverprogress_timeout timeout-in-secs
```

PolicyName is the name of the file policy for which you want to set the FPolicy server timeout.

timeout-in-secs is the timeout value in seconds.

The minimum timeout value that you can provide is zero and the maximum timeout value is 4294967 seconds. Setting a timeout value of zero disables the `serverprogress_timeout` option.

Note: By default, this option is disabled and no timeout value is set.

Result

After the timeout value is set, if the FPolicy server does not respond before the set timeout value, it is disconnected.

Example

To set the timeout value for file policy p1, use the following command:

```
filer> fpolicy options p1 serverprogress_timeout
fpolicy options p1 serverprogress_timeout: 0 secs (disabled)
filer> fpolicy options p1 serverprogress_timeout 600
filer> fpolicy options fp1 serverprogress_timeout 4294967
```

Setting request screening timeout

You can set a limit on how long the system waits for the FPolicy server to screen a request. You can set this limit individually on each policy. This option improves the performance of the FPolicy server.

Step

1. To set the timeout value for individual file policies, enter the following command:

```
fpolicy options PolicyName reqcancel_timeout timeout-in-secs
```

PolicyName is the name of the file policy you want to set the screening timeout for.

timeout-in-secs is the timeout value in seconds.

Result

After the timeout value is set, if the screen request is not complete within the set timeout value, the screen request is cancelled.

Example

To set the timeout value for file policy p1, use the following command:

```
filer> fpolicy options p1 reqcancel_timeout
fpolicy options p1 reqcancel_timeout: 0 secs (disabled)
filer> fpolicy options p1 reqcancel_timeout 60
```

Enabling or disabling multiple open instances of the SMB named pipe

You can enable multiple open instances of the SMB named pipe on an FPolicy server by using the `fpolicy.multiple_pipes` option.

About this task

When you enable this option, the FPolicy engine can open up to 10 instances of the SMB named pipe simultaneously to an FPolicy server. However, when you disable this option, only one instance of the SMB named pipe is opened to an FPolicy server.

Step

1. To enable or disable multiple open instances of the SMB named pipe on an FPolicy server, enter the following command:

```
options fpolicy.multiple_pipes {on|off}
```

By default, this option is set to `on`.

Events screened for NFS and CIFS clients

The FPolicy server can screen a number of operations or events for file requests received from NFS and CIFS clients.

The following table lists the events screened in NFS and CIFS for both native file blocking and server-based screening.

Events	Protocols	Description
File open	CIFS and NFS(v4)	Notification sent when a file is opened
File create	CIFS and NFS	Notification sent when a file is created
File rename	CIFS and NFS	Notification sent when a file name is changed
File close	CIFS and NFS(v4)	Notification sent when a file is closed
File delete	CIFS and NFS	Notification sent when a file is deleted
File read	CIFS and NFS	Notification sent when a file is read
File write	CIFS and NFS	Notification sent when a file is worked upon
Directory delete	CIFS and NFS	Notification sent when a directory is deleted
Directory rename	CIFS and NFS	Notification sent when a directory name is changed
Directory create	CIFS and NFS	Notification sent when a directory is created

Events	Protocols	Description
Getattr	NFS	Notification sent of request for attribute information
Setattr	CIFS and NFS	Notification sent of setting attributes information
Create hard link	NFS	Notification sent when a hard link is created
Create symlink	NFS	Notification sent when a symbolic link is created
Lookup	NFS	Notification sent when an NFS lookup occurs

Note: Although the CIFS setattr event can perform a variety of functions, only setattr operations that change the Security Descriptor information are monitored by FPolicy. The security descriptor information includes owner, group, discretionary access control list (DACL), and system access control list (SACL) information.

FPolicy can be used to cover most events in the file system related NFS and CIFS operations. Some of the operations that FPolicy does not monitor are listed here.

- NFS (v2, v3, v4): ACCESS, COMMIT, FSINFO, FSTAT, PATHCONF, ROOT, READLINK, REaddir, REaddirPLUS, STATFS, MKNOD
- NFSv4: Operations related to locking and delegation
- CIFS:
 - Tree operations such as SMB_COM_TREE_CONNECT and SMB_COM_TREE_DISCONNECT
 - Session related operations such as SMB_COM_SESSION_SETUP_ANDX
 - Locking-related operations
 - Operations not relevant to file system activity, such as print-related operations

What a file or directory event is

A variety of file and directory operations are screened. Based on the configuration of the policy, notifications are sent to the FPolicy server for operation requests.

What file open request monitoring is

FPolicy receives a notification from the storage system for file open operations.

When a file open request is made by a CIFS or NFSv4 client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, file availability, and whether the file is being accessed by another client. After the file passes the checks, if the file extension is included in the file policy extension include list, the request is forwarded to the FPolicy server. The FPolicy server receives this request and allows or blocks the file open request, based on the configuration of the policies.

If the storage system reboots, NFSv4 clients can reclaim file handles for files that were open before shutdown. After the storage system is functional again, if the FPolicy server connects to the storage

system before the NFS clients, the storage system forwards the reclaim file as an open request to the FPolicy server.

If the FPolicy server connects to the storage system after the NFS clients, the storage system does not forward the open reclaim request as an open request to the FPolicy server. In this case, the NFS client gets the file handle using the NFSv4 reclaim operation.

To enable file extension-based screening for NFS operations, set the `no_i2p` option to `off` on the volume. This enables inode-to-path file name translation on the volume.

Previous releases of FPolicy do not support NFSv4 protocol and the `i2p` option.

Note: FPolicy supports the NFSv4 protocol and the `i2p` option on volumes beginning with the Data ONTAP 7.3 release.

If you are running an FPolicy for Data ONTAP based application in NFSv4 environments, you must upgrade the FPolicy application to support NFSv4.

NFSv4 adds support for file OPEN and CLOSE events. Therefore, in applications based on previous releases of FPolicy, these file operations might appear as UNKNOWN event errors to the FPolicy application.

The file open operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file open operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor file open operations through the CLI

You can use the `fpolicy monitor add` command to configure a file policy to monitor file open operations. This CLI command adds the file open operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file open operation, enter the following CLI command:

```
fpolicy monitor add PolicyName open
```

Configuring FPolicy to monitor file open operations through ONTAPI

You can use an ONTAPI call to configure a file policy to monitor file open operations.

Step

1. To set the monitoring options for file open operations, use the following ONTAPI call:

```
fpolicy-operations-list-set
```

In the monitored-operations input name field, the `monitored-operation-info[]` should contain the file-open operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a file open operation, both NFS and CIFS requests can be monitored.

Registering FPolicy for monitoring file open requests

You can monitor file open operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file open operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_OPEN 0x0001
```

Result

After the registration is complete, the FPolicy server monitors all file open requests.

What file create request monitoring is

The FPolicy server receives a notification from the storage system for file create operations.

When a file create request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in the FPolicy `extension include` list. The FPolicy server receives this request and allows or blocks the file create request, based on the configuration of the policies.

The file create operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file create operation can be monitored using the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor file create operations through the CLI

To configure a file policy to monitor file create operations, use the `fpolicy monitor add` command.

About this task

This CLI command adds the create file operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file create operation, use the following CLI command:

```
fpolicy monitor add PolicyName create
```

Configuring FPolicy to monitor file create operations through ONTAPI

You can use an ONTAPI call to configure a file policy to monitor file create operations.

Step

1. To set the monitoring options for file create operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the file-create operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a file create operation, both NFS and CIFS requests can be monitored.

Registering FPolicy for monitoring file create requests

You can monitor file create operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file create operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_CREATE 0x0002
```

Result

After the registration is complete, the FPolicy server monitors all file create requests.

What file close request monitoring is

The FPolicy server receives a notification from the storage system for file close operations.

When a file close request is made by a CIFS or NFSv4 client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permission, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server. After the file is closed, the storage system sends a notification to the FPolicy server that the file is closed.

The FPolicy server cannot block the file close operation.

The file close operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file close operation can be monitored using the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Open downgrade operations in NFSv4 are also considered close operations, and notifications are sent for such operations.

To enable file extension-based screening, for NFSv4 operations, set the `no_i2p` option to `off` on the volume. This enables the inode-to-path file name translation on the volume.

Note: Beginning with the Data ONTAP 7.3 release, FPolicy supports the NFSv4 protocol.

Configuring FPolicy to monitor file close operations through the CLI

You can use the `fpolicy monitor add` CLI command to configure a file policy to monitor file close operations.

About this task

This CLI command adds the close file operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file close operation, use the following CLI command:

```
fpolicy monitor add PolicyName close
```

Configuring FPolicy to monitor file close operations through ONTAPI

You can use an ONTAPI call to configure a file policy to monitor file close operations.

Step

1. To set the monitoring options for file close operations, use the following ONTAPI call:

```
fpolicy-operations-list-set
```

In the monitored-operations input name field, the monitored-operation-info[] should contain the file-close operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a file close operation, both NFS and CIFS requests can be monitored.

Registering FPolicy for monitoring file close requests

You can monitor file close operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file close operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_CLOSE 0x0008
```

Result

After the registration is complete, the FPolicy server monitors all file close requests.

What file rename request monitoring is

The FPolicy server receives a notification from the storage system for file rename operations.

When a file rename request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permission, checking if the file is available, checking if the file is being accessed by some other client, and so on.

After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in FPolicy `ext[ension] inc[lude]` list.

The rename request is sent to the FPolicy server only if either the old or the new extension is listed in the `ext[ension] inc[lude]` list. That is, if a file name is being changed from `test.txt` to `test.mp3`, either or both the extensions (`.txt` or `.mp3`) should be listed in the `extension include` list.

The FPolicy server receives this request and allows or blocks the file rename request, based on the configuration of the policies.

The file rename operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file rename operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor file rename operations through the CLI

Use the `fpolicy monitor add` CLI command to monitor file rename operations.

About this task

This CLI command adds the create file operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file rename operation, use the following CLI command:

```
fpolicy monitor add PolicyName rename
```

Configuring FPolicy to monitor file rename operations through ONTAPI

Use the `fpolicy-operations-list-set` ONTAPI call to configure a file policy to monitor file rename operations.

Step

1. To set the monitoring options for file rename operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the `monitored-operations` input name field, the `monitored-operation-info[]` should contain the `file-rename` operation. The `monitored-protocols` should contain the specific protocols that you want to monitor. In the case of file create, both NFS and CIFS requests can be monitored.

Registering FPolicy to monitor file rename requests

You can monitor file rename operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file rename operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_RENAME 0x0004
```

Result

After the registration is complete, the FPolicy server monitors all file rename requests.

What file delete request monitoring is

The FPolicy server receives a notification from the storage system for file delete operations.

When a file delete request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. When the checks are complete and the file passes the check, the request notification is sent to the FPolicy server. The FPolicy server receives this request and allows or blocks the file delete request, based on the configuration of the policies.

The file delete operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file delete operation can be monitored using the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

To enable file extension-based screening, for NFS operations, set the `no_i2p` option to off on the volume. This enables the inode-to-path file name translation on the volume.

Configuring FPolicy to monitor file delete operations through CLI

You can use the `fpolicy monitor` CLI command to monitor file delete operations.

About this task

This CLI command adds the delete file operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file delete operation, use the following CLI command:

```
fpolicy monitor add PolicyName delete
```

Configuring FPolicy to monitor file delete operations through ONTAPI

You can use the `fpolicy-operations-list-set ONTAPI` call to monitor file delete operations.

Step

1. To set the monitoring options for file delete operations, use the `fpolicy-operations-list-set ONTAPI` call.

In the `monitored-operations` input name field, the `monitored-operation-info[]` should contain the `file-delete` operation. The `monitored-protocols` should contain the specific protocols that you want to monitor. In the case of a file delete operation, both CIFS and NFS requests can be monitored.

Registering FPolicy for monitoring file delete requests

You can monitor file delete operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file delete operations, set the following bit in the `OpsToScreen bitmask` in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_DELETE 0x0010
```

Result

After the registration is complete, the FPolicy server monitors all file delete requests.

What file write request monitoring is

The FPolicy server receives a notification from the storage system for file write operations.

When a file write request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in the FPolicy `extension include list`.

The FPolicy server receives this request and allows or blocks the file write request, based on the configuration of the policies.

The file write operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file write operation can be monitored using the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

To enable file extension-based screening, for NFS operations, set the `no_i2p` option to off on the volume. This enables the translation of inode-to-path file name on the volume.

Configuring FPolicy to monitor file write operations through the CLI

You can use the `fpolicy monitor` CLI command to monitor file write operations.

About this task

This CLI command adds the write file operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file write operation, use the following CLI command:

```
fpolicy monitor add PolicyName write
```

Configuring FPolicy to monitor file write operations through ONTAPI

You can use the `fpolicy-operations-list-set` ONTAPI call to configure a file policy to monitor file write operations.

Step

1. To monitor the file write operation, use the `fpolicy-operations-list-set` ONTAPI call.

In the `monitored-operations` input name field, the `monitored-operation-info[]` should contain the write operation. The `monitored-protocols` should contain the specific protocols that you want to monitor. In the case of a file write operation, both CIFS and NFS requests can be monitored.

Registering FPolicy to monitor file write requests

You can monitor file write operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file write operations, set the following bit in the `OpsToScreen` bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_WRITE 0x4000
```

Result

After the registration is complete, the FPolicy server monitors all file write requests.

What file read request monitoring is

The FPolicy server receives a notification from the storage system for file read operations.

When a file read request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in FPolicy `ext[ension] include` list.

The FPolicy server receives this request and allows or blocks the file read request, based on the configuration of the policies.

The file read operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file read operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

To enable file extension-based screening, for NFS operations, set the `no_i2p` option to `off` on the volume. This enables the inode-to-path file name translation on the volume.

Configuring FPolicy to monitor file read operations through the CLI

You can use the `fpolicy monitor` CLI command to monitor file read operations.

About this task

This CLI command adds the read file operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the file read operation, use the following CLI command:

```
fpolicy monitor add PolicyName read
```

Configuring FPolicy to monitor file read operations through ONTAPI

You can use the `fpolicy-operations-list-set` ONTAPI call to monitor file read operations.

Step

1. To set the monitoring options for file read operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the `monitored-operations` input name field, the `monitored-operation-info[]` should contain the read operation. The `monitored-protocols` should contain the specific protocols that you wish to monitor. In the case of a file read operation, both CIFS and NFS requests can be monitored.

Registering FPolicy to monitor file read requests

You can monitor file read operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file read operations, set the following bit in the `OpsToScreen` bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_READ 0x2000
```

Result

After the registration is complete, the FPolicy server monitors all file read requests.

What link request monitoring is (for NFS only)

The FPolicy server receives a notification from the storage system for file link operations.

When a file link request is made by an NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in FPolicy extension include list. The FPolicy server receives this request and allows or blocks the file link request, based on the configuration of the policies.

The file link operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file link operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor file link operations through the CLI

You can use the `fpolicy monitor` CLI command to configure a file policy, to monitor file link operations.

Step

1. To monitor the file link operation, use the following CLI command:

```
fpolicy monitor add PolicyName link
```

Result

This CLI command can add the file link operations to the list of monitored events for NFS requests.

Configuring FPolicy to monitor file link operations through ONTAPI

You can use the `fpolicy-operations-list-set` ONTAPI call to monitor file link operations.

Step

1. To set the monitoring options for file link operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the link operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a file link operation, only NFS requests can be monitored.

Registering FPolicy to monitor file link requests

You can monitor file link operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file link operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

`FS_OP_LINK 0x0400`

Result

After the registration is complete, the FPolicy server monitors all file link requests.

What symlink (symbolic link) request monitoring is (for NFS only)

The FPolicy server receives a notification from the storage system for file symlink operations.

When a file symlink request is made by an NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in the FPolicy `extension include` list. The FPolicy server receives this request and allows or blocks the file symlink request, based on the configuration of the policies.

The file symlink operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file symlink operation can be monitored using the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

To enable file extension-based screening, for NFS operations, set the `no_i2p` option to off on the volume. This enables the inode-to-path file name translation on the volume.

Configuring FPolicy to monitor file symlink operations through the CLI

You can use a CLI command to configure a file policy, to monitor file symlink operations.

About this task

This CLI command adds the symlink file operations to the list of monitored events for NFS requests.

Step

1. To monitor the file symlink operation, use the following CLI command:

```
fpolicy mon[itor] add PolicyName symlink
```

Configuring FPolicy to monitor file symlink operations through ONTAPI

You can use an ONTAPI to configure a file policy, to monitor file symlink operations.

Step

1. To set the monitoring options for file symlink operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the symlink operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a file symlink operation, both CIFS and NFS requests can be monitored.

Registering FPolicy to monitor file symlink requests

You can monitor file symlink operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file symlink operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_SYMLINK 0x0800
```

Result

After the registration is complete, the FPolicy server monitors all file symlink requests.

What directory delete request monitoring is

The FPolicy server receives a notification from the storage system for directory delete operations.

When a directory delete request is made by a CIFS client using RMDIR operations or an NFS client using UNLINK operations to the storage system, the storage system conducts all the relevant checks on the directory. The relevant checks include checking permission, checking if the directory is available, checking if the directory is being accessed by some other client, and so on. After the directory passes the checks, the request is forwarded to the FPolicy server. If the `required` option is set to `on` in the file policy and a directory delete operation is requested, the request is denied.

The directory delete operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The directory delete operation can be monitored through CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor directory delete operations through the CLI

You can use a CLI command to configure a file policy to monitor directory delete operations.

About this task

This CLI command adds the directory delete operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the directory delete operation, use the following CLI command:

```
fpolicy mon[itor] add PolicyName delete_dir
```

Configuring FPolicy to monitor directory delete operations through ONTAPI

You can use an ONTAPI call to configure a file policy, to monitor directory delete operations.

Step

1. To set the monitoring options for directory delete operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the `monitored-operation-info[]` should contain the `directory-delete` operation. The `monitored-protocols` should contain the specific protocols that you want to monitor. In the case of a directory delete operation, both CIFS and NFS requests can be monitored.

Registering FPolicy to monitor directory delete requests

You can monitor directory delete operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of directory delete operations, set the following bit in the `OpsToScreen` bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_DELETE_DIR 0x0020
```

Result

After the registration is complete, the FPolicy server monitors all directory delete requests.

What directory rename request monitoring is

The FPolicy server receives a notification from the storage system for directory rename operations.

When a directory rename request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the directory. The relevant checks include checking permissions, checking if the directory is available, checking if the directory is being accessed by some other client, and so on. After the directory passes the checks, the request is forwarded to the FPolicy server. If the `required` option is set to `on` in the file policy and a directory rename operation is requested, the request is denied.

The directory rename operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The directory rename operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor directory rename operations through CLI

You can use a CLI command to configure a file policy to monitor directory rename operations.

About this task

This CLI command adds the directory rename operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the directory rename operation, use the following CLI command:

```
fpolicy mon[itor] add PolicyName rename_dir
```

Configuring FPolicy to monitor directory rename operations through ONTAPI

You can use an ONTAPI call to configure a file policy, to monitor directory rename operations.

Step

1. To set the monitoring options for directory rename operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the directory-rename operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a directory rename operation, both CIFS and NFS requests can be monitored.

Registering FPolicy to monitor directory rename requests

You can monitor directory rename operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of directory rename operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_RENAME_DIR 0x0040
```

Result

After the registration is complete, the FPolicy server monitors all directory rename requests.

What directory create request monitoring is

The FPolicy server receives a notification from the storage system for directory create operations.

When a directory create request is made by a CIFS or NFS client to the storage system, the storage system conducts all the relevant checks on the directory. The relevant checks include checking

permissions, checking if the directory is available, checking if the directory is being accessed by some other client, and so on. After the directory passes the checks, the request is forwarded to the FPolicy server. If the `required` option is set to `on` in the file policy and a directory create operation is requested, the request is denied.

The directory create operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The directory create operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor directory create operations through the CLI

You can use a CLI command to configure a file policy to monitor directory create operations.

About this task

This CLI command adds the directory create operations to the list of monitored events for CIFS and NFS requests.

Step

1. To monitor the directory create operation, use the following command:

```
fpolicy mon[itor] add PolicyName create_dir
```

Configuring FPolicy to monitor directory create operations through ONTAPI

You can use an ONTAPI call to configure a file policy, to monitor directory create operations.

Step

1. To set the monitoring options for directory create operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the `monitored-operations` input name field, the `monitored-operation-info[]` should contain the `directory-create` operation. The `monitored-protocols` should contain the specific protocols that you want to monitor. In the case of a directory create operation, both CIFS and NFS requests can be monitored.

Registering FPolicy to monitor directory create requests

You can monitor directory create operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of directory create operations, set the following bit in the `OpsToScreen` bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_CREATE_DIR 0x0080
```

Result

After the registration is complete, the FPolicy server monitors all directory create requests.

What file lookup request monitoring is (for NFS only)

The FPolicy server receives a notification from the storage system for file lookup operations.

When a file lookup request is made by an NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in the FPolicy `ext[ension] inc[lude]` list. The FPolicy server receives this request and allows or blocks the file lookup request, based on the configuration of the policies.

The file lookup operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The file lookup operation can be monitored using the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

Configuring FPolicy to monitor file lookup operations through the CLI

You can use a CLI command to configure a file policy, to monitor file lookup operations.

Step

1. To monitor the file lookup operation, use the following CLI command:

```
fpolicy mon[itor] add PolicyName lookup
```

Configuring FPolicy to monitor file lookup operations through ONTAPI

You can use an ONTAPI call to configure a file policy, to monitor file lookup operations.

Step

1. To set the monitoring options for file lookup operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the lookup operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a file lookup operation, only NFS requests can be monitored.

Registering FPolicy to monitor file lookup requests

You can monitor file lookup operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of file lookup operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_LOOKUP 0x1000
```


Result

After the registration is complete, the FPolicy server monitors all file lookup requests.

What getattr request monitoring is (for NFS only)

The FPolicy server receives a notification from the storage system for getattr operations.

When a get attributes (getattr) request is made by an NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in the FPolicy `ext[ension] inc[lude]` list.

The FPolicy server receives this request and allows or blocks the getattr request, based on the configuration of the policies.

The getattr operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The getattr operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

To enable file extension-based screening, for NFS operations, set the `no_i2p` option to off on the volume. This enables the inode-to-path file name translation on the volume.

Configuring FPolicy to monitor get attributes operations through CLI

You can use a CLI command to configure a file policy, to monitor getattr operations.

About this task

This CLI command adds the get attributes operations to the list of monitored events for NFS requests.

Step

1. To monitor the get attributes operation, use the following CLI command:

```
fpolicy mon[itor] add PolicyName getattr
```

Configuring FPolicy to monitor get attributes operations through ONTAPI

You can use an ONTAPI call to configure a file policy, to monitor getattr operations.

Step

1. To set the monitoring options for getattr operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the getattr operation. The monitored-protocols should contain the specific protocols that you want to monitor. In the case of a getattr operation, only NFS requests can be monitored.

Registering FPolicy to monitor get attributes requests

You can monitor getattr operations by registering for it when you register an FPolicy server.

Step

1. To enable the screening of getattr operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_GETATTR 0x0100
```

Result

After the registration is complete, the FPolicy server monitors all get attributes requests.

What setattr request monitoring is

The FPolicy server receives a notification from the storage system for setattr operations.

When a set attributes (setattr) request is made by an NFS client to the storage system, the storage system conducts all the relevant checks on the file. The relevant checks include checking permissions, checking if the file is available, checking if the file is being accessed by some other client, and so on. After the file passes the checks, the request is forwarded to the FPolicy server, if the file extension is included in the FPolicy `ext[ension] inc[lude]` list. The FPolicy server receives this request and allows or blocks the setattr request, based on the configuration of the policies.

When a set attributes (setattr) request is made by CIFS clients to the storage system using the `NT_TRANSACT_SET_SECURITY_DESC` operation, the storage system sends setattr notification if the CIFS client makes changes to the security descriptor. The security descriptor information includes owner, group, discretionary access control list (DACL), and system access control list (SACL) information. If the Windows-based CIFS client sends the `NT_TRANSACT_SET_SECURITY_DESC` operation to the storage system, without changing the security descriptor information, it does not forward the request to the FPolicy server.

The setattr operation should be added to the monitored operations list for the FPolicy server to receive a notification from the storage system. The setattr operation can be monitored through the CLI or ONTAPI. It can also be set by the FPolicy server using a bitmask.

To enable file extension-based screening, for NFS operations, set the `no_i2p` option to `off` on the volume. This enables the inode-to-path file name translation on the volume.

Configuring FPolicy to monitor set attributes operations through the CLI

You can use a CLI command to configure a file policy, to monitor setattr operations.

About this task

This CLI command adds the set attribute operations to the list of monitored events for NFS requests.

Step

1. To monitor the set attributes operation, use the following CLI command:

```
fpolicy mon[itor] add PolicyName setattr
```

Configuring FPolicy to monitor set attributes operations through ONTAPI

You can use an ONTAPI call to configure a file policy, to monitor setattr operations.

Step

1. To set the monitoring options for setattr operations, use the `fpolicy-operations-list-set` ONTAPI call.

In the monitored-operations input name field, the monitored-operation-info[] should contain the setattr operation. The monitored-protocols should contain the specific protocols that you wish to monitor. In the case of a setattr operation, only NFS requests can be monitored.

Registering FPolicy to monitor set attributes requests

You can monitor setattr operations using bitmasks when you register an FPolicy server.

Step

1. To enable the screening of setattr operations, set the following bit in the OpsToScreen bitmask in the `FP_registration()` call when you register the FPolicy server to the storage system:

```
FS_OP_SETATTR 0x0200
```

Result

After the registration is complete, the FPolicy server monitors all set attributes requests.

What screening by volume is

FPolicy enables you to restrict a policy to a certain list of volumes, by including or excluding volumes that need to be screened.

Using the include list, you can request notifications for the specified volume list. Using the exclude list, you can request notifications for all volumes except the specified volume list.

Note: If both an include list and an exclude list are set, the include list is ignored.

It is possible to set different include and exclude volumes for different policies.

The default volumes list for a file policy is:

- All volumes are listed in the include list.
- No volumes are listed in the exclude list.

You can perform the following operations on the exclude and include lists:

- Reset or restore the volume list to the default list.
- Show or display the volumes in an include or exclude list.

- Add a volume to the include or exclude list.
- Remove a volume from the include or exclude list.
- Set or replace the existing list with a new volume list.
- Display the list of volumes for a file policy with wildcard characters.

From the command line, you can display or change the list of included and excluded volumes.

The command syntax to reset or display the file volumes list is as follows:

```
fpolicy vol[ume] {inc[lude]|exc[lude]} {reset|show} PolicyName
```

The command syntax to work with file volumes is as follows:

```
fpolicy vol[ume] {inc[lude]|exc[lude]} {add| remove|set|eval} PolicyName  
vol-spec
```

include is used to make changes to the include list.

exclude is used to make changes to the exclude list.

reset is used to restore the file volume list to the default list.

show is used to display the exclude or include list as entered.

add is used to add a volume to the exclude or include list.

remove is used to remove a volume from the exclude or include list.

set is used to replace the existing list with a new volume list.

eval is used to display the list of volumes for a file policy with wildcard characters.

PolicyName is the name of the file policy.

vol-spec is the name of the volume list that you want to change.

Wildcard information for screening with volumes

You can use the question mark (?) or asterisk (*) wildcard characters, to specify the volume.

The question mark (?) wildcard character stands for a single character. For example, entering **vol?** in a list of volumes that contain **vol1**, **vol2**, **vol123**, **vol114** will match **vol1** and **vol2**.

The asterisk (*) wildcard character stands for any number of characters that contain the specified string. Entering ***test*** in a list of volumes to exclude from file screening excludes all volumes that contain the string such as **test_vol** and **vol_test**.

How to display the list of volumes

To display the list of volumes you have specified to include or exclude for a file policy, you can use the `show` or `eval` command.

Displaying volumes using the show command

You can display the list of specified volumes using the `show` command.

About this task

The `show` command of the `fpolicy volume` command displays the list of specified volumes as entered at the command line. If you specified a set of volumes using wildcard characters, the `show` command displays the wildcard character you entered. For example, `vol*`.

Step

1. To display the list of excluded volumes you specified for a file policy, enter the following command:

```
fpolicy vol[ume] exc[lude] show PolicyName
```

Result

When you enter this command, Data ONTAP responds with a list of entries from the exclude list for the file you specified. This might include volume names and wildcard characters that describe a set of volumes (for example, `vol*`).

Note: If you want to show volumes from the list of files to be included for file screening, use the `include (inc)` option in place of the `exclude (exc)` option.

Displaying volumes using the eval command

You can display the list of specified volumes using the `eval` command.

About this task

The `eval` command of the `fpolicy volume` command displays the specified volumes after evaluating any wildcard character included in the list you entered. For example, if your list includes `vol*`, the `eval` command lists all volumes including the string `vol`, such as `vol11`, `vol122`, or `vol_sales`.

Step

1. To display the list of excluded volumes for a file policy with the wildcard character evaluated, enter the following command:

```
fpolicy vol[ume] exc[lude] eval PolicyName
```

Result

When you enter this command, Data ONTAP responds with a list of volumes from the exclude list for the file you specified, with wildcard character evaluated. For example, if you entered `vol*`, the eval display includes all volumes including the string `vol`, such as `vol1`, `vol22`, or `vol_sales`.

Note: To use the `eval` command for the list of files to be included for file screening, use the `include (inc)` option instead of the `exclude (exc)` option.

How to add volumes to the list

You can add volumes to the include or exclude volume list.

Adding volumes to the include list

To add volumes to the include volumes list, you can use the `fpolicy volume include add` CLI command.

Step

1. To add volumes to the include list of volumes to be screened for a file policy, enter the following command:

```
fpolicy volume include add PolicyName vol-spec
```

Result

Files in the volumes you add to an include list will always be screened by the file screening server when the policy is enabled.

Example

To include `vol1`, `vol2`, `vol3` to the list of volumes screened, enter the following command:

```
fpolicy vol inc add imagescreen vol1,vol2,vol3
```

After the volumes are added, the policy `imagescreen` performs screening in the volumes `vol1`, `vol2`, and `vol3`.

Adding volumes to the exclude list

You can add volumes to the exclude volumes list using the `fpolicy volume exclude add` CLI command.

Step

1. To add volumes to the exclude list of volumes to be screened for a file policy, enter the following command:

```
fpolicy volume exclude add PolicyName vol-spec
```

Result

Files in the volumes you add to an exclude list will not be screened by the file screening server when that policy is enabled (unless contradicted by another enabled file screening policy).

Example

To exclude vol4, vol5, vol6 to the list of volumes screened, enter the following command:

```
fpolicy vol exc add default vol4,vol5,vol6
```

When the volumes are added to the list, the modified default policy will no longer perform file screening in the volumes vol4, vol5, and vol6.

How to remove volumes from the list

You can remove volumes from the include or exclude volume list.

Removing volumes from the include list

You can remove volumes from the include volumes list using the `fpolicy volume include remove` CLI command.

Step

1. To remove volumes from the include volumes list for a file screening policy, enter the following command:

```
fpolicy volume include remove PolicyName vol-spec
```

Example

```
fpolicy volume include remove default vol4
```

Files in the volume named vol4 are not screened.

Removing volumes from the exclude list

You can remove volumes from the exclude volumes list using the `fpolicy volume exclude remove` CLI command.

Step

1. To remove volumes from the exclude volumes list for a file screening policy, enter the following command:

```
fpolicy vol[ume] exc[lude] remove PolicyName vol-spec
```

Example

```
fpolicy volume exclude remove default vol4
```

Files in the volume `vol4` are screened if there are no volumes specified in the include list (for example, if the include list specifies a volume `vol1`, then even after removing `vol4` from the list the volume will not be screened).

How to specify or replace a list of volumes

Specify or replace an include list and an exclude list.

Setting the include volumes list

You can set the include volumes list using the `fpolicy volume include set` CLI command.

Step

1. To set or replace the entire volume include list for a file policy, enter the following command:

```
fpolicy volume include set PolicyName vol-spec
```

The new list of volumes you enter with this command replaces the existing list of included volumes so that only the new volumes are included in the screening.

Result

Note: Turn off the include list to no volumes by using the `set` option; for example,

```
fpolicy vol inc set PolicyName ""
```

However, this has the same effect as disabling the policy.

Setting the exclude volumes list

You can set the exclude volumes list using the `fpolicy volume exclude set` CLI command.

Step

1. To set or replace the entire volume exclude list for a file policy, enter the following command:

```
fpolicy volume exclude set PolicyName vol-spec
```

The new list of volumes you enter with this command replaces the existing list of excluded volumes so that only the new volumes are excluded from screening.

How to reset the volumes in a list

You can specify or replace volumes in the include or exclude volume list.

Resetting the include volumes list

You can reset the include volumes list using the `fpolicy volume include reset` CLI command.

Step

1. To reset all entries from the exclude or include list for a file policy to the default values, enter the following command:


```
fpolicy volume include reset PolicyName
```

This command resets all the entries in the include list. That is, all the volumes listed in the include list are removed.

Resetting the exclude volumes list

You can reset the exclude volumes list using a CLI command.

Step

1. To reset all entries from the exclude list for a file policy to the default values, enter the following command:

```
fpolicy vol[ume] exc[lude] reset PolicyName
```

Result

Here, all the volumes listed in the exclude list are removed.

What screening by extension is

FPolicy enables you to restrict a policy to a certain list of file extensions, by including or excluding extensions that need to be screened.

Using the include list, you can request notifications for the specified file extensions.

You can provide both an include list and an exclude list. The extensions are first checked in the exclude list. If the requested file's extension is not in the exclude list, the include list is checked. If the file extension is listed in the include list, the file is screened. If the file extension is not listed in the include list, the request is allowed without screening.

Note: The maximum length of file name extension supported for screening is 260 characters.

Screening by extensions is based only on the characters after the last period (.) in the file name.

For example, for a file named `file1.txt.name.jpg`, file access notification takes place only if a file policy is configured for the `.jpg` extension.

The screening by extension feature is policy-based. Therefore, you can specify different extensions for different policies.

The default extension lists for a file policy are as follows:

- All file extensions are listed in the include list.
- No file extensions are listed in the exclude list.

You can perform the following operations on the exclude and include lists:

- Reset or restore the extension list to the default list.
- Set or replace the existing list with a new extensions list.
- Add an extension to the include or exclude list.
- Remove an extension from the include or exclude list.

- Show or display the extension in an include or exclude list.
- Display the list of extensions for a file policy by using wildcard characters.

From the command line, you can display or change the list of included and excluded extensions.

The command syntax to reset or display the file extension list is as follows:

```
fpolicy extensions { include | exclude } { reset | show } PolicyName
```

The command syntax to work with file extension is as follows:

```
fpolicy extensions { include | exclude } { set | add | remove } PolicyName  
ext-list
```

include is used to make changes to the include list.

exclude is used to make changes to the exclude list.

reset is used to restore the file extension list to the default list.

show is used to display the exclude or include list as entered.

set is used to replace the existing list with a new list of extensions.

add is used to add an extension to the exclude or include list.

remove is used to remove an extension from the exclude or include list.

PolicyName is the name of the file policy.

ext-list is the list of extensions that you want to change.

Note: Extension-based screening is not performed for directory operations such as directory create, directory delete, and directory rename.

Wildcard information for screening with extensions

You can use the question mark (?) wildcard to specify the extension.

If the question mark (?) wildcard character is used in the beginning of the string, it stands for a single character. At the end of the string, it stands for any number of characters.

For example:

- Entering **?s** in a list of file extensions to include for file screening includes all file extensions that have two characters ending with s (such as **as** and **js** extensions).
- Entering **??m** in a list of file extensions to include for file screening includes all file extensions that have three characters ending with m (such as **htm** and **vtm** extensions).
- Entering **j?** in a list of file extensions to include for file screening includes all file extensions that begin with j (such as **js**, **jpg**, and **jpe** extensions).

How to display the list of extensions

You can display the list of included and excluded extensions using the `fpolicy extensions` CLI command.

Displaying the list of extensions in the include list

You can display the list of extensions in the include extensions list by using the `fpolicy extensions include show` CLI command.

Step

1. To display the list of included file extensions for a file policy, enter the following command:

```
fpolicy extensions include show PolicyName
```

Data ONTAP responds with a list of extensions from the include list for the file you specified.

Displaying the list of extensions in the exclude list

You can display the list of extensions in the exclude extensions list by using the `fpolicy extensions exclude show` CLI command.

Step

1. To display the list of excluded file extensions for a file policy, enter the following command:

```
fpolicy extensions exclude show PolicyName
```

Data ONTAP responds with a list of extensions from the exclude list for the file you specified.

How to add extensions to the list

You can add extensions to the list of included and excluded extensions using the `fpolicy extensions` CLI command.

Adding extensions to the include list

Add extensions to the include extensions list using the `fpolicy extensions include` CLI command.

Step

1. To add file extensions to the list of file extensions to be screened for a file policy, enter the following command:

```
fpolicy extensions include add PolicyName ext-list
```

Example

```
fpolicy ext inc add imagescreen jpg,gif,bmp
```

After the extensions are added to the list, the policy `imagescreen` performs screening for any files with file extension `.jpg`, `.gif`, or `.bmp`.

Result

The file extensions you add to an include list will always be screened by the file screening server when that policy is enabled.

Adding extensions to the exclude list

You can add extensions to the exclude extensions list using the `fpolicy extensions exclude` CLI command.

Step

1. To add file extensions to the list of file extensions to be excluded from file screening for a file policy, enter the following command:

```
fpolicy extensions exclude add PolicyName ext-list
```

Example

```
fpolicy ext exc add default txt,log,hlp
```

When the extensions are added to the list, the modified policy will no longer screen `.txt`, `.log`, and `.hlp` files to be screened by the file screening server.

Result

The file extensions you add to an exclude list will not be screened by the file screening server when that policy is enabled (unless contradicted by another enabled file screening policy).

How to remove extensions from the list

You can remove extensions from the list of included and excluded extensions using `fpolicy extensions` CLI command.

Removing extensions from the include list

You can remove extensions from the include extensions list using the `fpolicy extensions include remove` CLI command.

Step

1. To remove file extensions from the include extensions list for a file policy, enter the following command:

```
fpolicy extensions include remove PolicyName ext-list
```

Example

```
fpolicy ext inc remove default wav
```

Files with a `.wav` extension are not screened.

Result

This command removes entries from the current file extension list.

Removing extensions from an exclude list

You can remove extensions from the exclude extensions list using the `fpolicy extensions exclude remove` CLI command.

Step

1. To remove file extensions from the exclude extensions list for a file screening policy, enter the following command:

```
fpolicy extensions exclude remove PolicyName ext-list
```

Example

```
fpolicy ext exc remove default wav
```

Files with a `.wav` extension are screened.

Result

This command removes entries from the current file extension list.

How to set or replace a list of extensions

You can set or replace the list of included and excluded extensions using the `fpolicy extensions` CLI command.

Setting the include extensions list

You can set the include extensions list by using the `fpolicy extensions include set` CLI command.

Step

1. To replace the entire include list for FPolicy, enter the following command:

```
fpolicy extensions include set PolicyName ext-list
```

Result

The new list of extensions that you specified with this command replaces the existing list of included extensions so that only the new extensions are included for screening.

Note: You can also set the include list to not screen file extensions by using the `set` option. For example,

```
fpolicy ext inc set PolicyName ""
```

When this command is used, no files are screened.

Setting the exclude extensions list

You can set the exclude extensions list using the `fpolicy extensions exclude set` CLI command.

Step

1. To replace the entire exclude list for FPolicy, enter the following command:

```
fpolicy extensions exclude set PolicyName ext-list
```

Result

On entering this command, the new list of extensions you specified with this command replaces the existing list of excluded extensions so that only the new extensions are excluded from screening.

How to reset the extensions in the list

You can reset the list of included and excluded extensions using the `fpolicy extensions` CLI command.

Resetting the include extensions list

You can reset the include extensions list using `fpolicy extensions include reset` CLI command.

Step

1. To reset all entries from the include list for FPolicy to the default values, enter the following command:

```
fpolicy extensions include reset PolicyName
```

Result

This command restores the file extension include list to the default list.

Resetting the exclude extensions list

You can reset the exclude extensions list using the `fpolicy extensions exclude reset` CLI command.

Step

1. To reset all entries from the exclude list for FPolicy to the default values, enter the following command:

```
fpolicy extensions exclude reset PolicyName
```

Result

This command restores the file extension exclude list to the default list.

How to manage the file screening server

You can display important file screening server information using the CLI commands. You can also assign servers to the secondary server list, or remove them from the secondary server list.

Displaying the file screening server information

You can display important file screening server information using the `fpolicy servers show` CLI command. The information displayed includes the list of servers registered, the list of connected servers, and the features enabled.

About this task

The command displays the following information about a particular FPolicy:

- The list of FPolicy servers registered
- The list of FPolicy servers connected
- Total time for which the server has been connected
- The list of features enabled for the server supported in Data ONTAP 7.3
- The status of the primary server
- The status of the secondary server

Step

1. To display the status of file screening servers, enter the following command:

```
fpolicy servers show PolicyName
```

Result

When you enter this command, Data ONTAP returns the status of the file screening servers for the policy you specified.

Disabling the connection

When a server's connection is disabled, the connection between the FPolicy server and the storage system are terminated.

Step

1. To disable the connection to a file screening server, enter the following command:

```
fpolicy servers stop PolicyName server-IP-address
```

PolicyName is the name of the policy that you want to disable the connection for.

server-IP-address is the list of FPolicy server IP addresses that you want to disable from the storage system.

Result

The server's connection is disabled.

What secondary servers are

FPolicy servers can be used as both primary and secondary servers. You can designate a particular FPolicy server or a list of FPolicy servers as secondary servers using the `fpolicy options` command.

The storage system uses the secondary servers to enforce file policies only if no primary servers are available. That is, when an FPolicy server is designated as a secondary server, the storage system never uses it as long as a primary server is available. If all primary servers are unavailable, the storage system uses any secondary servers connected to the storage system until a primary server becomes available again.

Any FPolicy server not classified as secondary is considered a primary server.

Assigning secondary servers list

You can assign or designate a particular FPolicy server as a secondary server using the `fpolicy options secondary_servers` CLI command.

Step

1. To designate a list of secondary servers to be used when the primary file screening server is unavailable, enter the following command:

```
fpolicy options PolicyName secondary_servers [server_list]
```

PolicyName is the name of the policy that you want the secondary server to use.

server_list is the list of FPolicy server IP addresses that you want to designate as secondary servers. Use a comma (,) to separate the IP addresses. A connection from any of the IP addresses listed in this field is classified by the storage system as a secondary server.

Result

When you enter this command, the specified servers are designated as secondary servers for the specified FPolicy.

Note: When the comma-separated list of IP addresses is provided, any existing list is replaced with the new list. Therefore, to retain existing secondary servers, you must add their IP addresses to the new list.

Removing all secondary servers

You can convert all secondary servers to primary servers using the `fpolicy options` CLI command.

Step

1. To convert all secondary servers to primary servers, enter the following command:

```
fpolicy options PolicyName secondary_servers ""
```

PolicyName is the name of the policy that you want the secondary server to use.

Result

After running this command, all FPolicy servers assigned to be secondary FPolicy servers become primary FPolicy servers.

How to monitor operations using FPolicy

You use FPolicy to monitor file operations. Tasks to manage file operations monitoring include adding, removing, or setting the list of operations to be monitored.

Adding operations to the monitor list

For FPolicy to implement native file blocking, it first needs to monitor operations that need to be blocked natively. You can do that by adding the operations to the list of operations monitored.

Step

1. To add operations to the list of monitored operations to be screened for FPolicy, enter the following command:

```
fpolicy mon[itor] add PolicyName [-p {cifs|nfs|cifs,nfs} ] [-f] op-spec
```

PolicyName is the name of the policy you want to add operations to.

`-p {cifs|nfs|cifs,nfs}` specifies the protocols you want to enable monitoring for. Use `cifs` to monitor CIFS requests, `nfs` to monitor NFS requests, or `cifs,nfs` to monitor both. If the protocol information is not specified in the monitor command, the storage system sends notifications for both CIFS and NFS protocols. When a particular operation is set for CIFS operations and then later set for NFS operations, the operations are monitored for requests from both the protocols. However, when the operation is removed from one of the protocols, monitoring for that operation stops for both the protocols. When a particular operation is set only for CIFS and not on NFS, this operation is monitored for both the protocols. When this operation is removed from the list of monitored operations for NFS it also stops monitoring for CIFS.

`-f` forces the policy to be enabled even if there are no servers available to enforce the policy.

`op-spec` is the list of operations you want to add. You can also choose to set the monitoring options for all operations together, by replacing the list of operations with `all` option.

Result

The specified operation is added to the list of monitored operations.

Example

The following command adds read, write, and lookup operations to the list of monitored operations:

```
fpolicy mon add p1 read,write,lookup
```

After you enable the policy p1, it monitors read, write, and lookup operations along with any other operations that have been set previously.

Removing operations from the monitor list

You can remove operations from the list using `fpolicy monitor remove` CLI command. When you remove an operation from the list of monitored operations, the particular operation is not monitored by the FPolicy.

Step

1. To remove operations from the list of monitored operations to be screened for FPolicy, enter the following command:

```
fpolicy mon[itor] remove PolicyName [-p {cifs|nfs|cifs,nfs} ] [-f] op-spec
```

Result

The specified operation is removed from the list of monitored operations.

Example

To stop monitoring read and setattr operations and to remove them from the list of monitored operations, enter the following command:

```
fpolicy mon remove p1 read,setattr
```

After you enable the policy p1, it stops monitoring read and setattr operations and removes these two operations from the list of operations monitored.

Setting or replacing the list of monitored operations

You can replace the list of monitored operations using the `fpolicy monitor set` CLI command.

Step

1. To replace the list of operations monitored, enter the following command:

```
fpolicy mon[itor] set PolicyName [-p {cifs|nfs|cifs,nfs} ] [-f] op-spec
```

Result

The list of operations to be monitored is replaced with the new set of operations.

Example

To set or replace the list of operations monitored, enter the following command:

```
fpolicy mon set p1 read,setattr
```

After you enable the policy p1, it monitors only read operations and setattr operations. Any existing monitored lists are replaced by this one.

What the different CLI commands are

The following table lists the FPolicy CLI commands.

Input Name	Description
<code>fpolicy help [cmd]</code>	Used to show the CLI help
<code>fpolicy create PolicyName PolicyType</code>	Used to create a file policy
<code>fpolicy destroy PolicyName</code>	Used to delete a file policy
<code>fpolicy enable PolicyName [-f]</code>	Used to enable a file policy
<code>fpolicy disable PolicyName</code>	Used to disable a file policy
<code>fpolicy show PolicyName</code>	Used to display a file policy
<code>fpolicy servers show PolicyName</code>	Used to display the FPolicy server status information
<code>fpolicy servers stop PolicyName IP-address</code>	Used to disable the FPolicy server connection
<code>fpolicy options PolicyName required {on off}</code>	Used to turn on or turn off the required option for a file policy
<code>fpolicy options PolicyName secondary_servers [IP-address [,IP-address]*]</code>	Used to configure options for FPolicy server
<code>fpolicy extension {exclude include} show PolicyName</code>	Used to display extensions in the include or exclude lists
<code>fpolicy extension {exclude include} reset PolicyName</code>	Used to reset extensions in the include or exclude lists
<code>fpolicy extension {exclude include} add PolicyName ext-list</code>	Used to add extensions to the include or exclude lists

Input Name	Description
<code>fpolicy extension {exclude include} remove <i>PolicyName</i> ext-list</code>	Used to remove extensions from the include or exclude lists
<code>fpolicy extension {exclude include} set <i>PolicyName</i> ext-list</code>	Used to set or replace all the extensions in the include or exclude lists
<code>fpolicy volume {include exclude} show <i>PolicyName</i></code>	Used to display volumes in the include or exclude lists
<code>fpolicy volume {include exclude} reset <i>PolicyName</i></code>	Used to reset volumes in the include or exclude lists
<code>fpolicy volume {include exclude} add <i>PolicyName</i> vol_spec</code>	Used to add volumes to the include or exclude lists
<code>fpolicy volume {include exclude} remove <i>PolicyName</i> vol_spec</code>	Used to remove volumes from the include or exclude lists
<code>fpolicy volume {include exclude} set <i>PolicyName</i> vol_spec</code>	Used to set or replace all the volumes in the include or exclude lists
<code>fpolicy volume {include exclude} eval <i>PolicyName</i> vol_spec</code>	Used to display the volumes in the include or exclude lists evaluating volumes specified using the wildcard character
<code>fpolicy monitor add <i>PolicyName</i> [-p {nfs cifs cifs,nfs}] [-f] op_spec [,op_spec]</code>	Used to add operations to the list of operations that are being monitored
<code>fpolicy monitor remove <i>PolicyName</i> [-p {nfs cifs cifs,nfs}] [-f] op_spec [,op_spec]</code>	Used to remove files from the list of files that are being monitored
<code>fpolicy monitor set <i>PolicyName</i> [-p {nfs cifs cifs,nfs}] [-f] op_spec [,op_spec]</code>	Used to set or replace the list of files that are being monitored

FAQs, error messages, warning messages, and keywords

This section describes frequently asked questions and error and warning messages.

Frequently asked questions (FAQs)

General FAQs and FAQs about access rights and permissions and other specific subjects are covered in this section.

General FAQs

Is there a limit currently to the total number of active file policies?

Yes, currently the limit to the total number of active file policies is 20 per vFiler unit.

If two policies are created, will the storage system handle the requests in a sequential or parallel manner?

When two policies are created, the storage system handles the requests sequentially.

Can you prioritize the policies so that one is favored over the other?

The existing implementation of FPolicy does not support ordering of policies.

Can multiple policies be created for different FPolicy servers?

Yes. It is possible to create multiple policies and use individual policies for different FPolicy servers. For example, you can create two policies, one for the FLM and one for NTP, and point the two FPolicy servers to these two policies.

The order in which notifications will be sent is the same as the order in which policies are listed under the `fpolicy` command. This is the reverse of the order in which policies are created on the storage system. For example, if policy p1 was created followed by policy p2, notifications will be sent to p2 and subsequently to p1.

It is important to note the difference between "multiple file policies" and "multiple servers."

Some problems you might face are as follows:

- Currently the FPolicy engine sends requests sequentially (instead of sending them parallel) for the multiple policies so they might see double the performance degradation.

What licenses are needed to be enabled for FPolicy to work on your storage systems?

CIFS needs to be licensed and set up on the storage system for FPolicy to work.

Why do I need CIFS to be licensed and set up even on an NFS-only storage system?

An FPolicy server wields a lot of power, and it is authenticated using CIFS security to ensure that the server has Backup-Operator privileges (or more) on the storage system. Therefore, CIFS needs to be licensed even in an NFS exclusive environment. Also, to apply file policies to NFS files, you must also have NFS licensed and running.

Does FPolicy have any limitations?

Yes, the following are FPolicy limitations:

- FPolicy supports only CIFS and NFS protocols.
However, there are some operations for the CIFS and NFS protocols that FPolicy does not monitor, such as, NFSv4 operations related to locking and delegation, session-related operations (SMB_COM_SESSION_SETUP_ANDX), operations not relevant to file system activity (print-related operations), and so on.
- FPolicy does not support other protocols such as FTP, HTTP, WebDAV, FileIO, and so on.
- You cannot configure CIFS and NFS operations separately on the same policy.

Following are the screening limitations of FPolicy:

- You must set up file screening on an entire volume.
You cannot screen individual qtrees and directories.
- FPolicy supports screening of CIFS operations on alternate data streams.
However, FPolicy does not support screening of NFS operations on alternate data streams.
- When you register multiple servers, the policy of all the servers connected changes based on the settings of the server that registers last.
- Multiple instances of an FPolicy server from the same IP address cannot register to same policy.
- If the CIFS system resources used by FPolicy are exhausted, the CIFS screening by the FPolicy engine will stop.

Is FPolicy dependent upon Virus Scanning (vscan)?

FPolicy runs independently from vscan operations. FPolicy occurs before virus scanning operations, so that paths indicated in stub files (such as symlinks) can be traversed to load the actual file, instead of just scanning the stub file. Vscan operations are independent of file policies. That is, vscan can open and scan files that have been blocked by file policies. Therefore, there is no interdependence between FPolicy and vscan.

Where are FPolicy settings saved?

FPolicy settings are saved in the registry.

What happens when a user attempts to make changes to a migrated file that was accessed with read permission?

The FPolicy server has to do the following:

For CIFS and NFS version 4, it can recall the file at open time if the open request is for write (or read-write) access mode. Alternatively, it can do it when the write request is made. However, for this option the server has to be registered to monitor write operations.

Since NFSv2 and NFSv3 versions do not have an open call, the HSM server will need to register to monitor read and write operations. The HSM server will have to recall the file when it receives the write request. For read operations, the HSM server has an option of either using pass through read or write.

Access rights and permissions FAQs

What is the minimal access right for an account that connects to the storage systems, and registers as an FPolicy server listening to FPolicy events?

The FPolicy server needs backup privileges at least, to register to the storage system.

What is the minimal access right for an account that connects to the storage system and scans q-tree ACLs?

The right to scan ACLs is granted to CIFS logins using standard Windows methods. If you are connected to the storage system using an account that is a member of the Backup Operators or Administrators groups you can use the FILE_FLAG_BACKUP_SEMANTICS open mode, which allows you to access any file, regardless of security.

Performance FAQs

What factors does the performance of an FPolicy depend on?

The following are some of the factors that the performance of FPolicy depends on:

- Number of Operations (like read, open, close, and so on) being monitored
- Number of registered FPolicy servers (load sharing)
- Number of Policies screening the same operation
- Network bandwidth between storage system and FPolicy server (round-trip time of the screen request)
- Response time of the FPolicy server

How can we measure how FPolicy traffic is divided between CIFS and NFS traffic?

The output of the FPolicy command run at the storage system contains a counter for the total number of request screened by that particular file policy. However, currently there is no way to understand the division between CIFS and NFS traffic.

Every client request that goes through FPolicy screening generates some extra CIFS requests for internal FPolicy communication. This is true for both CIFS and NFS clients requests. Currently there is no way to measure this extra traffic.

If you switch on FPolicy before doing recalls, does that have an impact on performance?

Yes, switching on FPolicy before doing any recalls has an impact on the performance. The impact of the performance depends primarily on how FPolicy is configured. It is therefore recommended that you do not turn on FPolicy before doing any recalls.

When there are two FPolicy servers registered to a storage system with different performance levels, does the performance of the slower server affect the performance of the fast server?

Yes, the performance of the slower server does affect the performance of the faster server. It is therefore recommended that servers with same capabilities are used while connecting to a storage system.

Do we have a metric to determine the additional load on the CPU when FPolicy is enabled?

No, such data is not currently available for FPolicy.

File screening FAQs**How does file screening work?**

File screening policies are used to specify files or directories on which one wants to put some restrictions. Upon receiving a file operation request (such as open, write, create, or rename), Data ONTAP checks its file screening policies before permitting the operation.

If the policy specifies screening for that file based on its extension, the file name is sent to the file screening server to be screened. The file screening server applies policies to the file name to determine whether the storage system should allow the requested file operation. The file screening server then sends a response to the storage system to either allow or block the requested file operation.

Does the performance of the system go down while using file screening?

Yes, the performance of the system goes down while using file screening.

Can we use default options for setting file screening options?

There is a master setting for all file policies, the `fpolicy.enable` option, which is `on` by default. When an individual FPolicy is newly created, it is `off` by default. This allows the system administrator to fully configure the policy before activating it. Whether something is actually screened or not, depends upon whether or not there is a supported external file screening server running and accessible to the storage system. Remember that an external file screening server is a requirement in order to use FPolicy.

What happens if I create screening policies but do not have a screening server?

If you enable a policy when no file screening servers are available, nothing happens. However, if you have turned on the `fpolicy option required` for that policy, then access to files specified in that policy will be denied. The setting for 'required' on a policy is set to off by default.

How can I display the status of file screening servers?

You can display the status of the file screening server by using the following command:

```
fpolicy servers show PolicyName
```

Data ONTAP returns the status of the file screening server for the policy you specified.

Can I specify secondary screening servers? If yes, how can I do it?

Yes, you can designate a list of secondary servers to be used when the primary file screening server is unavailable. Use the following command:

```
fpolicy options PolicyName secondary_servers [ server_list ]
```

Any FPolicy server that connects to the storage system will be a primary server unless its IP address is in the secondary server list. Secondary servers will never be used by the storage system unless all primary servers are unavailable.

How can I disable the connection to a file screening server?

You can disable the connection to a file screening server by using the following command:

```
fpolicy servers stop PolicyName server-IP-address
```

Is FPolicy file screening applied at the volume level or at the qtree level?

FPolicy file screening is applied at the volume level, and not at the qtree level.

FPolicy server FAQs**What is the difference between primary and secondary servers?**

Primary servers are active servers that screen client requests. Secondary servers are registered for the fail safe mode. When all the primary servers are down, all the secondary servers start screening requests.

How can I register a secondary server?

To use a server as a secondary server, you have to add the server IP in the secondary server list.

When the server connects, it will be treated as secondary.

FPolicy error messages

Error messages might appear if FPolicy encounters errors.

fpolicy.fscreen.server.connectError severity="ERR"

Error message `fpolicy.fscreen.server.connectError severity="ERR"`

Cause This error occurs when the storage system encounters an error while attempting to communicate with an FPolicy (file policy) server. The communication failure causes the storage system to break its connection with this server.

The error can be due to network problems, security settings on the FPolicy server that deny access to the storage system, or hardware or software problems on the FPolicy server. The problem can also occur if a low memory situation on the storage system prevents the storage system from obtaining resources needed to perform the operation.

Recommended action Examine the error code to see whether it helps point to the cause of the problem.

Examine the event logs of the FPolicy server to determine whether it is disconnected from the storage system and why.

Examine the storage system syslog for error messages. Correct any problems that are found, such as network errors or hardware problems on the FPolicy server. Check whether a software patch was recently installed on the FPolicy server that might have changed security settings.

fpolicy.fscreen.server.closeError severity="ERR"

Error message `fpolicy.fscreen.server.closeError severity="ERR"`

Cause This error occurs when the storage system encounters an error while attempting to stop communication with an FPolicy (file policy) server. The error can be due to network problems or hardware or software problems on the FPolicy server.

Recommended action Examine the error code to see whether it helps point to the cause of the problem.

Examine the event logs of the FPolicy server to determine whether it is disconnected from the storage system and why. Correct any problems that you find, such as network errors or hardware problems on the FPolicy server.

This error might not be the underlying error. The storage system might be ending communication with the FPolicy server because an error occurred in an earlier attempt to communicate with the server. The error that occurs during the close of the connection might be a continuation of that previous error condition.

fpolicy.fscreen.server.requestError severity="ERR"

Error message `fpolicy.fscreen.server.requestError severity="ERR"`

Cause This error is generated when the storage system encounters an error while attempting to send a notification request to an FPolicy (file policy) server. The error can be due to network problems or hardware/software problems on the FPolicy server.

The storage system will retry this notification with another server of this policy if the policy has multiple servers. Otherwise, the storage system will proceed based on the policy's setting for the required option. If the required setting is on, the storage system will deny the request. If required is off, the storage system will allow the client request to proceed.

Recommended action Examine the error code to see if it helps point to the cause of the problem. Examine the event logs of the FPolicy (file policy) server to learn if it has disconnected from the storage system and why. Correct any problems that are found such as network errors or hardware problems on the FPolicy server.

fpolicy.fscreen.server.requestRejected severity="ERR"

Error message `fpolicy.fscreen.server.requestRejected severity="ERR"`

Cause This error is generated when a storage system's notification request to an FPolicy (file policy) server is rejected by the FPolicy server. The error can be due to software problems on the FPolicy server.

Recommended action Examine the error code to see if it helps point to the use of the problem. Examine the event logs of the FPolicy (file policy) server to learn if it has created an error to explain the problem. The FPolicy server may have detected an internal error, or may be unable to accept more requests.

fpolicy.fscreen.server.pingRejected severity="ERR"

Error message `fpolicy.fscreen.server.pingRejected severity="ERR"`

Cause From time to time if the FPolicy server connection is idle the storage system will send a status request to learn the status of the FPolicy server. This error is generated when a storage system's status request to an FPolicy server gets an error. This error can occur if the storage system is unable to contact the FPolicy server, or the error can occur if the server returns an error to the storage system's request. The error can be due to network problems or hardware/software problems on the FPolicy server. The storage system will break its connection with the server when this request fails.

Recommended action	Examine the error code to see if it helps point to the cause of the problem. Examine the event logs of the FPolicy (file policy) server to learn if it has disconnected from the storage system and why. Correct any problems that are found such as network errors or hardware problems on the FPolicy server.
---------------------------	---

fpolicy.fscreen.server.completionUnexpectedState severity="ERR"

Error message	<code>fpolicy.fscreen.server.completionUnexpectedState severity="ERR"</code>
Cause	This error occurs when the FPolicy server has completed a screen request and returned a completion. However, the internal storage system state for this request is not valid. This completion message is ignored by the storage system.
Recommended action	None

fpolicy.fscreen.server.requestStatusError severity="ERR"

Error message	<code>fpolicy.fscreen.server.requestStatusError severity="ERR"</code>
Cause	This error occurs when the FPolicy server has accepted a screen request but has not reported the completion of the request. The storage system will check on the status of incomplete requests. If the storage system is unable to send the request, or if the server does not support the request, this error occurs. The error can be due to network problems or hardware/software problems on the FPolicy server which have broken the connection of the server to the storage system.
Recommended action	Examine the error code to see if it helps point to the cause of the problem. Examine the event logs of the FPolicy (file policy) server to learn if it has disconnected from the storage system and why. Correct any problems that are found such as network errors or hardware problems on the server.

fpolicy.fscreen.server.connecting.internalError severity="ERR"

Error message	<code>fpolicy.fscreen.server.connecting.internalError severity="ERR"</code>
Cause	This error occurs when the file policy server registers with the storage system and offers to work as an FPolicy server or the storage system. The storage system has not been able to get memory it needs to hold information related to the FPolicy server.
Recommended action	Contact technical support.

fpolicy.fscreen.server.connecting.privError severity="ERR"

Error message `fpolicy.fscreen.server.connecting.privError severity="ERR"`

Cause This error occurs when the FPolicy server registers with the storage system and offers to work as an FPolicy server for the storage system. The server has connected as a user that has insufficient privileges to act as an FPolicy server. The storage system requires that the user under whose name the server connects to the storage system must be at least a member of the storage system's backup-operators group. This registration attempt will be rejected.

Recommended action Enable the option `cifs.trace_login` to see what user name the server is using to connect to the storage system. Remember to turn the option off after solving this problem because tracing can affect storage system performance. Check to see the user name under which the file policy service is running on the FPolicy server. Use the `wcc` command to learn which groups this user belongs to. Perhaps add this user to an appropriate group, or change the properties of the FPolicy server so that it runs under a different user name.

fpolicy.fscreen.cfg.pCreateErr severity="ERR"

Error message `fpolicy.fscreen.cfg.pCreateErr severity="ERR"`

Cause This error occurs when the storage system processes saved file screening configuration information from the registry. The storage system tried to create and initialize a new policy, but could not. This error suggests a problem with the consistency of the storage system registry. The storage system discards information related to the policy.

Recommended action You might be able to remove the policy by using the command `fpolicy destroy`. You can then re-create the policy by using the command `fpolicy create` and set the policy configuration as desired.

fpolicy.fscreen.request.pathError severity="ERR"

Error message `fpolicy.fscreen.request.pathError severity="ERR"`

Cause This error occurs when the storage system encounters an error as it builds a path for the fscreen server to use in accessing a file. Possible errors include a path that is too long or Unicode conversion problems. Users access a file with a path like this: `shareName\directories\fileName`. The storage system builds an absolute path for the server from the root of the storage system: `ontap_admin $ \vol\volName\sharePath\directories\FileName`. Typically, this is an internal storage system error (bug).

Recommended action None.

fpolicy.fscreen.server.requestTO severity="ERR"

Error message `fpolicy.fscreen.server.requestTO severity="ERR"`

Cause This error occurs when the server has accepted a file screen notification but has not reported the completion of the request. The storage system will check on the status of incomplete requests after a timeout has elapsed. If the server disavows all knowledge of a request which it has accepted but not completed, the request is considered to be timed out. Typically this indicates a server problem.

Recommended action Examine the event logs of the server to learn if it has noted any problems. Contact the server software vendor to learn if their product supports the request-status query. The storage system only times out requests which the server has accepted. It will not time out a request as long as the server affirms that it is still working on the request. The storage system sends request-status messages to the server to learn the status of requests which may have timed out.

fpolicy.server.fqdn.unavail severity="ERR"

Error message `fpolicy.server.fqdn.unavail severity="ERR"`

Cause This message occurs when the Reverse DNS lookup for the FPolicy server IP address fails and the storage system cannot determine the FPolicy server's Fully Qualified Domain Name (FQDN). If the FPolicy server is running on Microsoft Windows Server 2008 and later, the storage system requires the FPolicy server FQDN for authenticating itself to the FPolicy server.

Recommended action Verify the Reverse DNS lookup configuration on the DNS server.

FPolicy warning messages

Warning messages might appear if FPolicy encounters issues you might need to know about.

fpolicy.fscreen.server.connectedNone severity="WARNING"

Error message `fpolicy.fscreen.server.connectedNone severity="WARNING"`

Cause This warning occurs when no FPolicy (file screening) servers are connected to the storage system. This can be significant if the policy is required because the storage system will reject various operations on files and directories. This can be significant if the policy is not required because the storage system will allow

various operations on files and directories although no server has approved the operation.

Recommended action

Examine the event logs of the FPolicy server(s) to learn why they have disconnected from the storage system. Examine the storage system's syslog for error messages which could provide clues. Correct any problems that are found such as network errors or hardware problems on the FPolicy server.

fpolicy.fscreen.server.completionRequestLost severity="WARNING"

Error message `fpolicy.fscreen.server.completionRequestLost severity="WARNING"`

Cause

This warning occurs when the FPolicy (file policy) server has completed a screen request and returned a completion. However, the storage system cannot find the request which is being reported as complete.

This warning is an indication that the FPolicy server and storage system are out of synchronization. The problem can happen because of timing issues. For example, the completion can arrive shortly after file screening has been disabled on the storage system and all requests which had been waiting for completion have been allowed to continue. Or the FPolicy server has returned a completion prior to accepting the screen request. Or the request may have timed out, causing the storage system to ask for status on the request. If the FPolicy server was not able to find the request, the storage system allows the request to continue. Then, if the FPolicy server later completed the request, that request is not found.

Recommended action

If the problem is occurring repeatedly, a line trace can help in diagnosis. You can use the `pktt` command to get a line trace.

fpolicy.fscreen.server.completionInconsistent severity="WARNING"

Error message `fpolicy.fscreen.server.completionInconsistent severity="WARNING"`

Cause

This warning occurs when the FPolicy (file policy) server has completed a screen request and returned a screen completion. However, the file path in the completion message does not match the file path for which a screen request was made by the storage system.

When the storage system makes a screen request it provides the FPolicy server with both a file path and a request ID. The FPolicy server has sent a screen completion message to the storage system which has a valid request ID, however, the file path does not match the one given by the storage system in the screen request.

This problem might be a software defect in the FPolicy server or in the storage system. Or, if the FPolicy server is serving a group of storage systems, it is possible that it is completing requests but sending the completions to the wrong storage system, such that the request ID is valid in the storage system but the file path is not associated with the request ID.

Recommended action If the problem is occurring repeatedly, a line trace can help in diagnosis. You can use the `pktt` command to get a line trace.

fpolicy.fscreen.server.connecting.badOperationList severity="WARNING"

Error message `fpolicy.fscreen.server.connecting.badOperationList severity="WARNING"`

Cause This warning occurs when the FPolicy server provides the storage system with a list of operations for which it wishes to receive notifications. Examples of operations includes renaming a file or directory, creating a file, opening or closing a file and so on. The list can be provided to the storage system when the server registers, or at a later time. This warning has occurred because the list of operations provided by the FPolicy server includes operations for which the storage system does not provide notifications. The storage system ignores the list provided by the server.

Recommended action Check the versions of the storage system and the file policy software to ensure that they are compatible with each other.

fpolicy.fscreen.server.connecting.badParameter severity="WARNING"

Error message `fpolicy.fscreen.server.connecting.badParameter severity="WARNING"`

Cause This warning occurs when the FPolicy server registers with the storage system and offers to work as a file policy server for the storage system. A parameter provided by the server was not set to a value that the storage system understands. The storage system ignores this parameter provided by the server.

Recommended action Check the versions of the storage system and the file policy software to ensure that they are compatible with each other. The storage system ignores the invalid parameter and allows the server to register with the storage system.

fpolicy.srv.conn.badOptParam severity="WARNING"

Error message `fpolicy.srv.conn.badOptParam severity="WARNING"`

Cause	This message occurs when the FPolicy server registers with the system and makes itself available as a file policy server for the system. However, the Optional Parameter provided by the FPolicy server is either not set to a value that the system understands or the format of the Optional Parameter followed is not correct. The system ignores the invalid parameter and allows the FPolicy server to register.
Recommended action	Check the versions of the system and file policy software to ensure that they are compatible with each other.

fpolicy.fscreen.cfg.pCreateInfo severity="WARNING"

Error message	<code>fpolicy.fscreen.cfg.pCreateInfo severity="WARNING"</code>
Cause	This warning occurs when the storage system processes saved file screening configuration information from the registry. The storage system wishes to create and initialize a new policy. However, the storage system encountered a problem. This warning suggests a problem with the consistency of the storage system registry and should not happen.
Recommended action	It might be possible to remove the policy with the <code>fpolicy destroy</code> command. Then recreate the policy using the <code>fpolicy create</code> command and set the policy configuration as desired.

fpolicy.fscreen.server.droppedConn severity="WARNING"

Error message	<code>fpolicy.fscreen.server.droppedConn severity="WARNING"</code>
Cause	This warning is generated when connection to a file policy (fscreen) server is lost. A connection can be lost when the server voluntarily disconnects from the storage system. Other possible reasons for dropped connections include network errors, termination of CIFS services on the storage system, and hardware/software failures on the server.
Recommended action	Check to see if your server is still functioning. Check network connectivity between the storage system and the server. Check the server's event log to see if there are errors that explain the disconnect. Make sure that CIFS is running on your storage system.

fpolicy.fscreen.vol.i2p.off severity="WARNING"

Error message	<code>fpolicy.fscreen.vol.i2p.off severity="WARNING"</code>
----------------------	---

Cause	This message occurs when an FPolicy server registers with the system for a file policy and required inode-to- path name translation for file policy notifications against NFS requests. However, the volume monitored by the file policy has inode-to-pathname translation disabled. The FPolicy fails to generate a file path for NFS requests in case inode-topathname translation is disabled on the volume.
Recommended action	<p>Disable the <code>no_i2p</code> option by entering the following command:</p> <pre>vol options volumeName no_i2p off</pre> <p>This command enables the inode-to-path name translation for files on the volume. This operation might fail if a nondisruptive volume movement is being performed on the target volume.</p>

fpolicy.fscreen.server.unexpectedFileDataResponse severity="WARNING"

Error message	<code>fpolicy.fscreen.server.unexpectedFileDataResponse severity="WARNING"</code>
Cause	<p>This message occurs when the FPolicy server sends file data to the system for a RRD (Read Redirect) request but the system cannot find the request for which file data sent. This warning is an indication that the FPolicy server and system are out of synchronization.</p> <p>This problem can happen because of timing issues. For example, the file data might arrive shortly after file screening is disabled on the system and all requests that are waiting for completion have been allowed to continue; or the FPolicy server has returned a file data before accepting the screen request; or the request might have timed out, causing the system to ask for status on the request. Or, if the FPolicy server was not able to find the request, the system would allow the request to continue. Then, if the FPolicy server later sends file data for the request, that request would not be found by the system.</p>
Recommended action	If the problem occurs repeatedly, use a packet trace to help in diagnosis. You can use the <code>pktt</code> command to get a packet trace.

Keywords list for screening operations

All operations can be monitored in three ways. They can be set using a bitmask while registering an FPolicy server, they can be configured by using an ONTAPI call, and they can be configured using a CLI.

You can use different keywords to configure monitoring of the different operations supported by FPolicy. The following table lists the keywords used for each of the operations, when you attempt to configure them with the three options.

Operation name	CLI Keyword	ONTAPI keyword	Registration key words	
			Bit	Value
File open	open	file-open	FS_OP_OPEN	0x0001
File create	create	file-create	FS_OP_CREATE	0x0002
File close	close	file-close	FS_OP_CLOSE	0x0008
File rename	rename	file-rename	FS_OP_RENAME	0x0004
File delete	delete	file-delete	FS_OP_DELETE	0x0010
File write	write	file-write	FS_OP_WRITE	0x4000
File read	read	file-read	FS_OP_READ	0x2000
File link	link	link	FS_OP_LINK	0x0400
File symlink	symlink	symlink	FS_OP_SYMLINK	0x0800
Directory delete	directory-delete	directory-delete	FS_OP_DELETE_DIR	0x0020
Directory rename	directory-rename	directory-rename	FS_OP_RENAME_DIR	0x0040
Directory create	directory-create	directory-create	FS_OP_CREATE_DIR	0x0080
File lookup	lookup	lookup	FS_OP_LOOKUP	0x1000
Get attribute	getattr	getattr	FS_OP_GETATTR	0x0100
Set attribute	setattr	setattr	FS_OP_SETATTR	0x0200

Support for CIFS over IPv6

Starting with Data ONTAP 7.3.1, you can allow the CIFS clients to access the files in your storage system over IPv6.

Data ONTAP uses a dual stack mechanism to transition from IPv4 to IPv6. Therefore, your storage system can allow CIFS clients to access the files in the following three modes:

- Only IPv4 mode
- Only IPv6 mode
- IPv6/IPv4 mode

For more information about the dual stack mechanism, see the *Data ONTAP Network Management Guide for 7-Mode*.

Your storage system sends and receives data only on port 445 for providing CIFS service over IPv6.

Note: NetBIOS over TCP (NBT) is not supported for CIFS service over IPv6.

Enabling or disabling CIFS over IPv6

You can enable or disable CIFS over IPv6 by setting the `cifs.ipv6.enable` option to `on` or `off`, respectively.

Before you begin

You must enable IPv6 on the storage system by setting the `ip.v6.enable` option to `on`. For more information about enabling IPv6 on your storage system, see the *Data ONTAP Network Management Guide for 7-Mode*.

About this task

- If you have enabled CIFS over IPv6 and you then disable IPv6 on your storage system by setting the `ip.v6.enable` option to `off`, CIFS is automatically disabled over IPv6.
- You do not need to restart CIFS over IPv6 after restarting the IPv6 global option.
If CIFS over IPv6 is enabled on the storage system, and if you disable and reenable the IPv6 global option, CIFS IPv6 sockets are automatically created to listen for IPv6 addresses.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Enable CIFS over IPv6	<code>options cifs.ipv6.enable on</code>
Disable CIFS over IPv6	<code>options cifs.ipv6.enable off</code>

Note: When CIFS over IPv6 is disabled, no new CIFS sessions are accepted over IPv6, but the existing IPv6 CIFS sessions continue to work over IPv6.

Listing IPv4 or IPv6 CIFS sessions

You can use the `cifs sessions` command to list the IPv4 and IPv6 CIFS sessions running on your storage system. You can also use the `-i` option with the `cifs sessions` command to view the CIFS sessions running over only IPv4 or only IPv6.

Step

1. Perform one of the following actions:

If you	Enter the command...
want to...	

List all
CIFS
sessions

cifs sessions

Example:

```
cifs sessions

Server Registers as 'MACHINE1' in Windows 2000 domain
'IPV6LH1'

Root volume language is not set. Use vol lang.

Selected domain controller \\WIN2K8-204-121 for
authentication

=====

PC IP(PC Name) (user)                #shares    #files

10.73.9.35(machine5-lxp) (ipv6lh1\administrator - root)

                                     1             0

fd20:81be:b255:4204:fd3d:e44:9ab1:6ae5(VISTA204-123)
(ipv6lh1\administrator - root)
```

List the
CIFS
sessions
running
over IPv4

```
cifs sessions -i ipv4
```

Example:

```
cifs sessions -i ipv4

Server Registers as 'MACHINE1' in Windows 2000 domain
'IPV6LH1'

Root volume language is not set. Use vol lang.

Selected domain controller \\WIN2K8-204-121 for
authentication

=====

PC IP(PC Name) (user)                #shares    #files

10.73.9.35(machine5-lxp) (ipv6lh1\administrator - root)

                                1                0
```

If you want to...	Enter the command...						
List the CIFS sessions running over IPv6	<div><pre>cifs sessions -i ipv6</pre><p>Example:</p><pre>cifs sessions -i ipv6</pre><p>Server Registers as 'MACHINE1' in Windows 2000 domain 'IPV6LH1'</p><p>Root volume language is not set. Use vol lang.</p><p>Selected domain controller \\WIN2K8-204-121 for authentication</p><p>=====</p><table><tr><th>PC IP(PC Name) (user)</th><th>#shares</th><th>#files</th></tr><tr><td>fd20:81be:b255:4204:fd3d:e44:9ab1:6ae5(VISTA204-123) (ipv6lh1\administrator - root)</td><td>1</td><td>0</td></tr></table></div>	PC IP(PC Name) (user)	#shares	#files	fd20:81be:b255:4204:fd3d:e44:9ab1:6ae5(VISTA204-123) (ipv6lh1\administrator - root)	1	0
PC IP(PC Name) (user)	#shares	#files					
fd20:81be:b255:4204:fd3d:e44:9ab1:6ae5(VISTA204-123) (ipv6lh1\administrator - root)	1	0					

Listing cumulative IPv4 or IPv6 CIFS sessions

You can use the `-t` option with the `cifs sessions` command to view the cumulative CIFS sessions running over IPv4 or IPv6.

Step

- 1. Enter the following command:

```
cifs sessions -t
```

Example

The following output shows two cumulative CIFS sessions running over IPv4 and one cumulative CIFS session running over IPv6.

```
cifs sessions -t

Using domain authentication. Domain type is Windows 2000.

Root volume language is not set. Use vol lang.

Number of WINS servers: 0

Total CIFS sessions: 2
```

```
CIFS open shares: 2
CIFS open files: 0
CIFS locks: 0
CIFS credentials: 2
CIFS sessions using security signatures: 0
IPv4 CIFS sessions: 1
IPv6 CIFS sessions: 1
Cumulative IPv4 CIFS sessions: 2
Cumulative IPv6 CIFS sessions: 1
```

File sharing between NFS and CIFS

You can optimize Data ONTAP to share files quickly between NFS and CIFS clients without errors.

NFS and CIFS file naming dependencies

File naming conventions depend on both the network clients' operating systems and the file-sharing protocols.

The operating system and the file-sharing protocols determine the following:

- Characters a file name can use
- Case-sensitivity of a file name

Characters a file name can use

If you are sharing a file between clients on different operating systems, you should use characters that are valid in both operating systems.

For example, if you use UNIX to create a file, do not use a colon (:) in the file name because the colon is not allowed in MS-DOS file names. Because restrictions on valid characters vary from one operating system to another, see the documentation for your client operating system for more information about prohibited characters.

Case-sensitivity of a file name

File names are case-sensitive for NFS clients and case-insensitive but case-preserving for CIFS clients.

For example, if a CIFS client creates `Spec.txt`, both CIFS and NFS clients display the file name as `Spec.txt`. However, if a CIFS user later tries to create `spec.txt`, the name is not allowed because, to the CIFS client, that name currently exists. If an NFS user later creates a file named `spec.txt`, NFS and CIFS clients display the file name differently, as follows:

- On NFS clients, you see both file names as they were created, `Spec.txt` and `spec.txt`, because file names are case-sensitive.
- On CIFS clients, you see `Spec.txt` and `Spec~1.txt`.
Data ONTAP creates the `Spec~1.txt` file name to differentiate the two files.

Creating lowercase file names

You can set the `cifs.save_case` option to `off` to force Data ONTAP to ignore the case in which file names are entered and instead force these names to lowercase text.

About this task

Setting this option to `off` provides better compatibility between 16-bit applications and some UNIX tools. However, by default, this option is set to `on`.

Step

1. Enter the following command:

```
options cifs.save_case off
```

How Data ONTAP creates file names

Data ONTAP creates and maintains two file names for files in any directory that has access from a CIFS client: the original long name and a file name in 8.3 format.

For file names that exceed the eight character name or the three character extension limit, Data ONTAP generates an 8.3-format file name as follows:

- It truncates the original file name to six characters, if the file name exceeds six characters.
- It appends a tilde (~) and a number, one through five, to file names that are no longer unique after being truncated.
 If it runs out of numbers because there are more than five similar names, it creates a unique file name that bears no relation to the original file name.
- It truncates the file name extension to three characters.

For example, if an NFS client creates a file named `specifications.html`, the 8.3 format file name created by Data ONTAP is `specif~1.htm`. If this name already exists, Data ONTAP uses a different number at the end of the file name. For example, if an NFS client then creates another file named `specifications_new.html`, the 8.3 format of `specifications_new.html` is `specif~2.htm`.

Controlling the display of dot files from CIFS clients

Dot files typically do not appear on UNIX-based systems; if you want to provide Windows client systems the option of hiding dot files, set the `cifs.show_dotfiles` option to `off`.

About this task

By default, these files are displayed on CIFS client systems, regardless of the Windows Folder Options View setting for showing or hiding hidden files.

Step

1. Enter the following command to set the `cifs.show_dotfiles` option to `off`:

```
options cifs.show_dotfiles off
```

Dot files on this system can be excluded from display when Windows client users select "Do not show hidden files and folders" from the **View** tab on the Folder Options box. (To display the Folder Options box, in Windows Explorer, select **Tools > Folder Options**.)

Enabling file name character translation between UNIX and Windows

If you have legacy file names on both operating systems (Windows and UNIX) that contain characters that are not valid in both operating systems, you can use the `charmap` command to allow CIFS clients to access files with NFS names that would otherwise not be valid for CIFS.

About this task

When files created by NFS clients are accessed by CIFS clients, the storage system looks at the name of the file and if the name is not a valid CIFS file name (for example, if it has an embedded colon “:” character) the storage system returns the 8.3 file name that is maintained for each file. However, this causes problems for applications that encode important information into long file names.

Therefore, if you are sharing a file between clients on different operating systems, you should use characters in the file names that are valid in both operating systems.

However, if you have legacy file names on both operating systems (Windows and UNIX) that contain characters that are not valid in both operating systems, you can define a map that converts the invalid NFS characters into Unicode characters that both CIFS and certain Windows applications can accept.

For more information, see the `na_charmap(1)` man page.

Note: This functionality supports the CATIA MCAD and Mathematica applications as well as other applications that have this requirement.

Step

1. Enter the following command:

```
charmap [volname [mapspec]]
```

volname specifies the name of a volume on a storage system. If you do not specify a value for *volname*, the mapping for all volumes is displayed.

Specify *mapspec* in the following format:

```
hh:hhhh[,hh:hhhh]...
```

Each *hh* represents a hexadecimal value. The first value of each *hh* pair that is separated by a colon is the hexadecimal value of the NFS character you want to translate, and the second value of each *hh* pair is the Unicode value that CIFS will use. If you do not specify a value for *mapspec*, the current mapping, if any, is displayed.

Example

The following command maps characters used by the CATIA application:

```
charmap desvol 3e:ff76,3c:ff77,2a:ff78,3a:ff79,22:ff7a
```

This command maps a set of characters (>, <, *, :, ", ?, \, and |) into Japanese Unicode characters that are not normally used as normal characters in file names. This mapping applies to the volume named “desvol.”

Character restrictions

Make sure that the Unicode characters that are used to represent invalid or illegal characters are those characters that do not normally appear in file names; otherwise, unwanted mappings will occur.

For example, if you try to map a colon (:) to a hyphen (-), but the hyphen (-) was used in the file name correctly, a Windows client trying to access a file named “a-b” would have its request mapped to the NFS name of “a:b” (not the desired outcome).

Refer to the following list of NFS characters when performing your remapping:

- 22 = double quote (")
- 2a = asterisk (*)
- 3a = colon (:)
- 3c = less than (<)
- 3e = greater than (>)
- 3f = question mark (?)
- 5c = backslash (\)
- 7c = pipe (|)
- b1 = (±)

In addition, if you attempt to create or rename a file or directory from a CIFS client with a name that contains the Unicode character 0x0080, an error message appears. The Unicode character 0x0080 is not supported on the storage system.

Clearing a character mapping from a volume

You can use the `charmap` command to clear a character mapping from a volume when you no longer need it. This is useful for example when you remove an application that uses the character mapping from the storage system.

Step

1. To clear a character mapping, enter the following command:

```
charmap volname ""
```

volname specifies the name of a volume on the storage system.

About file locking between protocols

File locking is a method used by client applications to prevent a user from accessing a file previously opened by another user. How Data ONTAP locks files depends on the protocol of the client.

If the client is an NFS client, locks are advisory; if the client is a CIFS client, locks are mandatory.

Because of differences between the NFS and CIFS file locks, an NFS client might fail to access a file previously opened by a CIFS application.

The following occurs when an NFS client attempts to access a file locked by a CIFS application:

- In mixed or NTFS qtrees, file manipulation operations, such as `rm`, `rmdir`, and `mv`, can cause the NFS application to fail.
- NFS read and write operations are denied by CIFS deny-read and deny-write open modes, respectively.
- NFS write operations fail when the written range of the file is locked with an exclusive CIFS `bytelock`.

Breaking locks

You can use the `lock break` command to break file locks. This enables you to correct issues where a file lock prevents client access to files.

Locking prevents multiple processes from modifying the same file at the same time and allows cooperating processes to synchronize access to shared files. Data ONTAP manages locks across the multiprotocol environment.

Sometimes the need arises to manually break locks. An example of a scenario in which you might need to break locks include when you are debugging applications. The `lock break` command manages locks for CIFS, NFSv4, NLM (used with NFSv2 and NFSv3), and FLEXCACHE.

The `lock break` command has five subcommands that serve as filters in determining which locks to break. Three of the subcommands have optional parameters that provide finer granularity in determining which locks to break, allowing you to break only the locks that you choose without breaking other locks that you want to remain intact. You cannot use the `lock break` command without including one of the available subcommands.

To determine which locks you want to break, you can display information about currently held locks using the `lock status` command. The man page for the `lock` command contains detailed information.

Note: If you issue a `lock break` command on files that CIFS clients have currently opened, the CIFS clients must close and reopen the files to be able to further modify them. This is required because breaking locks invalidates the file handles. CIFS clients receive a new valid file handle after reopening the files.

Breaking locks on a specified file

You can use the `lock break` command to break locks on a specified file.

Step

1. To break locks on a specified file, enter the following command:

```
lock break -f file [-o owner -h host] [-p protocol]
```

- The `-f file` option specifies the file on which to break the lock.
You must use the full path name for *file*, for example, `/vol/vol1/file1`.

Note: Although you can use the `lock break` command with the `-f file` option to break NFSv4 locks, it is not recommended because breaking NFSv4 locks might lead to unexpected results on the client. You can prevent the breaking of NFSv4 locks by including the protocol along with the file name.

- The optional `-o owner` option is only pertinent with the CIFS, NLM, and NFSv4 protocols. For NLM and NFSv4 locks, when specifying the `-o owner` option, you must also specify the `-h host` option. The meaning of *owner* varies across protocols:
 - CIFS: `[domain\]user`
 - NLM: Process-ID
 - NFSv4: UID
- The `-h host` option specifies the host on which to break locks.
Only locks held by the CIFS, NFSv4, and NLM protocols can be broken by specifying this option. The meaning of *host* varies across protocols:
 - CIFS: The value can be either a NetBIOS name or a fully qualified domain name (FQDN).
 - NFSv4: The value must be an IP address.
 - NLM: The value can be either a host name (FQDN, host name alias, and so on) or an IP address.

The `lock` command does not resolve the host name to an IP address. Functionally, filtering locks by a host name is not equivalent to filtering locks by the corresponding IP address. If the locks are to be filtered by *host*, then the value of *host* should be obtained from the output of the `lock status -h` command. Such a value of *host* should not be interpreted in any way as it might lead to improper removal of locks.

- You can break all locks for the specified file or use the `-p protocol` option to break the locks for only the specified protocol.

protocol is a case-insensitive string that can take one of the following values:

- `cifs`
- `nfsv4`
- `nlm`
- `flexcache`

Example

The following example breaks the CIFS lock on the `/vol/vol1/Mydocument.doc` file:

```
lock break -f "/vol/vol1/Mydocument.doc" -p cifs
```

Breaking locks on a specified host

You can use the `lock break` command to break locks based on a specified host.

Step

1. To break locks on a specified host, enter the following command:

```
lock break -h host [-o owner] [-f file] [-p protocol]
```

- The `-h host` option specifies the host on which to break locks.
Only locks held by the CIFS, NFSv4, and NLM protocols can be broken by specifying this option. The meaning of *host* varies across protocols:
 - CIFS: The value can be either a NetBIOS name or an IP address.
 - NFSv4: The value must be an IP address.
 - NLM: The value can be either a host name (FQDN, host name alias, and so on) or an IP address.

The `lock` command does not resolve the host name to an IP address. Functionally, filtering locks by a host name is not equivalent to filtering locks by the corresponding IP address. If the locks are to be filtered by host, then the value of *host* should be obtained from the output of the `lock status -h` command. Such a value of *host* should not be interpreted in any way as it might lead to improper removal of locks.

- The optional `-o owner` option is only pertinent with the following protocols:
 - CIFS
 - NLM
 - NFSv4

The meaning of *owner* varies across protocols:

- CIFS: `[domain\]user`
- NLM: `Process-ID`
- NFSv4: `UID`
- You can break all locks for the specified host or use the `-f file` option to break the lock on only the specified file.
You must use the full path name for *file*, for example, `/vol/vol1/file1`.
- You can break all locks for the specified host or use the `-p protocol` option to break the locks for only the specified protocol.

Protocol is a case-insensitive string that can take one of the following values:

- cifs
- nfsv4
- nlm
- flexcache

Example

The following example breaks all CIFS locks on the host1 host:

```
lock break -h host1 -p cifs
```

Breaking locks for a specified lock owner

You can use the `lock break` command to break locks based on a specified lock owner.

Step

1. To break locks for a specified owner, enter the following command:

```
lock break -o owner [-f file] [-p protocol]
```

- Breaking a lock based on owner is only valid for the CIFS protocol.
The value of this option is of the form [domain\]user.
- You can break all locks for the specified owner or use the `-f file` option to break the lock on only the specified file.
You must use the full path name for *file*, for example, `/vol/vol1/file1`.
- You can break all locks for the specified owner or use the `-p protocol` option to break the locks for only the specified protocol.
CIFS is the only supported protocol.

Example

The following example breaks all CIFS locks for the owner MYDOMAIN\user1:

```
lock break -o MYDOMAIN\user1
```

Breaking locks based on a specified protocol

You can use the `lock break` command to break locks based on a specified protocol. You can break locks for the CIFS, NFSv4, NLM, and FLEXCACHE protocols.

Step

1. To break locks for a specified protocol, enter the following command:

```
lock break -p protocol
```

Breaking a lock based on protocol breaks all locks for that protocol. The *protocol* value is a case-insensitive string that can take one of the following values: `cifs`, `nfsv4`, `nlm`, or `flexcache`.

Example

The following example breaks all NFSv4 locks:

```
lock break -p nfsv4
```

Breaking NLM locks on a specified IP network family

You can use the `lock break` command to break NLM locks on a specified IP network family. You can either break locks for the IPv4 network family or the IPv6 network family.

Step

1. To break locks for a specified network family, enter the following command:

```
lock break -net network
```

The `-net network` option breaks all locks for the specified network family. The *network* value is a case-insensitive string that can take one of two possible values: `IPv4` or `IPv6`. This option is only valid for the NLM protocol.

Example

The following example breaks all NLM IPv4 locks:

```
lock break -net IPv4
```


How Data ONTAP treats read-only bits

The read-only bit is a binary digit, which holds a value of 0 or 1, that is set on a file-by-file basis to reflect whether a file is writable (disabled) or read-only (enabled).

CIFS clients that use MS-DOS and Windows can set a per-file read-only bit. NFS clients do not set a per-file read-only bit because NFS clients do not have any protocol operations that use a per-file read-only bit.

Data ONTAP can set a read-only bit on a file when a CIFS client that uses MS-DOS or Windows creates that file. Data ONTAP can also set a read-only bit when a file is shared between NFS clients and CIFS clients. Some software, when used by NFS clients and CIFS clients, requires the read-only bit to be enabled.

For Data ONTAP to keep the appropriate read and write permissions on a file shared between NFS clients and CIFS clients, it treats the read-only bit according to the following rules:

- NFS treats any file with the read-only bit enabled as if it has no write permission bits enabled.
- If an NFS client disables all write permission bits and at least one of those bits had previously been enabled, Data ONTAP enables the read-only bit for that file.
- If an NFS client enables any write permission bit, Data ONTAP disables the read-only bit for that file.
- If the read-only bit for a file is enabled and an NFS client attempts to discover permissions for the file, the permission bits for the file are not sent to the NFS client; instead, Data ONTAP sends the permission bits to the NFS client with the write permission bits masked.
- If the read-only bit for a file is enabled and a CIFS client disables the read-only bit, Data ONTAP enables the owner's write permission bit for the file.
- Files with the read-only bit enabled are writable only by root.

Note: Changes to file permissions take effect immediately on CIFS clients, but might not take effect immediately on NFS clients if the NFS client enables attribute caching.

Deleting files with the read-only bit set

You can set the `cifs.perm_check_ro_del_ok` to allow deletion of files using UNIX delete semantics when the read-only bit is enabled. By default, this behavior is disabled.

About this task

Windows does not allow you to delete a file with the read-only bit enabled. Some multiprotocol source control applications require UNIX delete semantics; files for these applications also cannot be deleted when the read-only bit is enabled.

Step

1. To allow deletion of files using UNIX delete semantics when the read-only bit is enabled, enter the following command:

```
options cifs.perm_check_ro_del_ok on
```

Managing UNIX credentials for CIFS clients

When connecting to your storage system, a user on a CIFS client receives a CIFS credential. The user must also have one or more UNIX credentials to access resources controlled by Data ONTAP.

How CIFS users obtain UNIX credentials

A UNIX credential consists of a UNIX-style user ID (UID) and group IDs (GIDs).

Data ONTAP uses the UNIX credential for the following purposes:

- When a user tries to access files that have UNIX-style security, Data ONTAP uses the UID and the user's GIDs to determine the access rights of the user.
- When you want to use group quotas on a group that contains CIFS users, those CIFS users must have UNIX credentials.

For more information about group quotas, see the *Data ONTAP Storage Management Guide for 7-Mode*.

When a CIFS user tries to connect to the storage system, Data ONTAP tries to determine the UID, the primary GID, and all secondary group GIDs of the CIFS user. If Data ONTAP cannot determine the UID of the CIFS user and a default UNIX user is not defined, the user is denied access.

When retrieving UNIX user and group information, Data ONTAP uses the `/etc/nsswitch.conf` file to determine which name services to use, and in what order to use them. Data ONTAP supports files (passwd), NIS, and LDAP name services for UNIX user and group databases. Data ONTAP obtains UNIX credentials by looking up the UNIX password and group databases, which can be the `/etc/passwd` and `/etc/group` files, NIS maps, or an LDAP directory store to obtain the user and group information for a user. The configured databases contain account information for all users that might access the storage system.

After Data ONTAP determines that the UNIX user name is a valid user name, the UID and GIDs (both primary and secondary) are retrieved from the directory store and are included in the CIFS credential.

Note: If you are using an NIS map, ensure that its entries do not exceed 1,024 characters per line.

For Data ONTAP to obtain the UID and GIDs for a CIFS user, it must first determine the user's UNIX-style name. It does this through user mapping. Data ONTAP does not require that a user's Windows name be identical to the UNIX name. By entering information in the `/etc/usermap.cfg` file, you can specify how each Windows name maps to a UNIX name. If you accept the default mapping, you do not need to enter this information. By default, Data ONTAP uses the Windows name as the UNIX name when it looks up the UID. (The storage system converts uppercase characters in the Windows name to lowercase before the lookup.)

If the user names in the UNIX password database are identical to the Windows names, you need not provide the mapping information in the `/etc/usermap.cfg` file. If the user name is not found in the UNIX password database and the `waf1.default_unix_user` option has been specified, the default login name specified for that option is used. See the `options(1)` man page for more information about setting the `waf1.default_unix_user` option.

Data ONTAP obtains a user's GIDs in the following ways:

- Data ONTAP obtains the user's primary GID from the UNIX password database. Each account in the UNIX password database contains the primary GID for that user.
- Data ONTAP obtains the user's other GIDs from the group database, which can be the NIS group map, the LDAP data store, or the `/etc/group` file.
The group database is where you define membership for various groups.

You can see the UNIX credential of a connected CIFS user when you display CIFS session information.

Ensuring that only intended CIFS users receive UNIX credentials

You can configure which CIFS users receive UNIX credentials by editing the `/etc/usermap.cfg` file, creating UNIX groups and users, and enabling the Windows guest user account.

Steps

1. If some Windows names are different from UNIX names or you want to prevent some CIFS users from accessing the storage system, edit the `/etc/usermap.cfg` file.
2. Create groups in the UNIX group database.
3. For each CIFS user with a mapped UNIX name, enter the user account in the UNIX password database.
4. If you rename the Administrator account, make sure at least one CIFS user maps to the UNIX root account.
5. If you want CIFS users who do not have an entry in the UNIX password database to access the storage system, create a default user account in the UNIX password database and set the `waf1.default_unix_user` option to that user.
6. If you want unauthenticated users to access the storage system, enable the Windows guest user account.

Specifying entries in the `/etc/usermap.cfg` file

Data ONTAP uses the `/etc/usermap.cfg` file to map user names. In its simplest form, each `/etc/usermap.cfg` entry contains a pair of names: the Windows name and the UNIX name. Data ONTAP can translate the Windows name to the UNIX name or vice versa.

About this task

When CIFS is started, if the `/etc/usermap.cfg` file is missing, a default file is created. It contains commented-out sample map entries that are useful for improving security.

When Data ONTAP receives a connection request from a CIFS user, it searches the `/etc/usermap.cfg` file to see whether an entry matches the user's Windows domain name and user name.

If an entry is found, Data ONTAP uses the UNIX name specified in the entry to look up the UID and GID from the UNIX password database. If the UNIX name is a null string, Data ONTAP denies access to the CIFS user.

If an entry is not found, Data ONTAP converts the Windows name to lowercase and considers the UNIX name to be the same as the Windows name. Data ONTAP uses this UNIX name to look up the UID and GID from the UNIX password database.

Data ONTAP scans the file sequentially. It uses the first matching entry for mapping.

For information about character coding of the `/etc/usermap.cfg` file, see the information about the contents of the `/etc` directory in the Storage Management Guide.

Step

1. Specify each entry using the following format:

```
[IP_qualifier:] Windows_name [direction] [IP_qualifier:] UNIX_name
```

After you finish

You can embed comments in the file by beginning the comment lines with `#`. Comments at the end of an entry are also allowed if preceded by `#`. Blank lines are ignored.

About the `IP_qualifier` field

The `IP_qualifier` field is an IP address that qualifies the user name by narrowing the match.

The IP qualifier can be any of the following:

- An IP address in bit notation.
You can specify a subnet by including the number of bits in the subnet mask. For example, `192.4.1.0/24` means the `192.4.1.0` class C subnet.
- A name.

Data ONTAP first considers a name to be a host name. If it cannot find a matching host name in its host name database, it considers the name to be a network name.

- A subnet address.

This includes a network name or IP address and the subnet mask (for example, corpnet/255.255.255.0).

Note: Data ONTAP uses the IP qualifier only for matching. If an IP qualifier is present on the destination side of a map entry, Data ONTAP does not consider the login request to come from that IP qualifier.

About the `Windows_name` field

The `Windows_name` field consists of a Windows domain name, which is optional, and a Windows user name.

On the source side of the map entry, the domain specifies the domain in which the user resides. On the destination side of the map entry, it specifies the domain used for the mapped UNIX entry. If the account name in the entry is a local user account, the Windows domain name is the storage system name.

If you omit the domain name in the `Windows_name` field, it is assumed to be the domain in which the storage system is installed. If the storage system uses local user accounts for authentication, the domain name is the storage system name.

You can use an asterisk (*) as a wildcard character in the following ways:

- You can use it on the source side to indicate that the specified name in any domain maps to the specified UNIX name.
- You can use it on the destination side to indicate that the specified UNIX name maps to a Windows name in any trusted domain.

The trusted domain used for the mapping depends on where Data ONTAP finds the first matching Windows name. Data ONTAP searches only the trusted domains you specify with the `cifs.search_domains` option, in the order in which the trusted domains are specified. If you do not set this option, Data ONTAP searches all trusted domains in an unspecified order.

If the user name contains spaces or a pound sign, enclose the name in double quotation marks, for example, "bob smith" or "eng#lab"\#joe".

Note: Do not enclose the \ in quotation marks.

If the user name is empty or blank (specified as "") on the destination side of the map entry, the matching UNIX name is denied access. Use entries with a blank user name to deny access to some or all UNIX users. If you use these entries in conjunction with `IP_qualifier`, you can exclude all UNIX users except for certain hosts or subnets.

About the `direction` field

The `direction` field indicates the direction of the mapping.

The `direction` field can be one of the values in the following table.

Value of the direction field	Meaning
==	Mapping is bidirectional. The entry maps from Windows to UNIX and from UNIX to Windows. Omitting the direction field has the same meaning as specifying ==.
<=	The entry maps from UNIX to Windows
=>	The entry maps from Windows to UNIX.

About the UNIX_name field

The UNIX_name field is a UNIX name in the UNIX password database.

If the UNIX_name field is empty or blank (specified as " ") on the destination side of the map entry, the specified source name is prevented from logging in. The Windows user cannot log in to the storage system even if the user can see the storage system while browsing the network.

You can use an asterisk (*) in the UNIX name. The asterisk is considered the wildcard character. It means any user. Remember these guidelines when including an asterisk in the Windows name or the UNIX name:

- If the asterisk is on the source side of the mapping, any user maps to the specified name on the destination side.
- If the destination side contains an asterisk but the source side does not, no mapping is done. Data ONTAP does not map an explicitly specified name to a name with an asterisk.
- If both the source and destination sides contain an asterisk, the corresponding name is mapped.

How Data ONTAP interprets domain names in /etc/usermap.cfg

The way in which Data ONTAP interprets a domain name in the `/etc/usermap.cfg` file that contains a dot depends on whether storage system is in a Windows NT domain or a Windows Active Directory domain.

If your storage system is installed in a Windows NT domain, the length of the domain name field affects how the domain name is interpreted.

If your storage system is installed in a Windows Active Directory domain, Data ONTAP interprets the domain names in the same way a Windows server would.

If the storage system is in a Windows NT domain, Data ONTAP follows these rules when interpreting a domain name containing a dot in the `domain\user` format:

- If *domain* is 15 characters or shorter, Data ONTAP recognizes the entire string, including the dot, as the NetBIOS form of the domain name.
For example, `my_company.com` is the NetBIOS form of the domain name in the following name:
`my_company.com\john_smith`
- If *domain* is longer than 15 characters, the dot is treated as a separator, and the string before the first dot is the NetBIOS form of the domain name.

For example, engineering is the NetBIOS form of the domain name in the following name:

engineering.1234567890corporation.com\john_smith

If the storage system is in a Windows Active Directory domain, you can specify a user name in the *domain\user* format. The string before the first dot in *domain* is the NetBIOS form of the domain name, and the entire string in *domain* is the DNS domain name.

For example, engineering is the NetBIOS form of the domain name and engineering.1234567890corporation.com is the DNS domain name in the following name:

engineering.1234567890corporation.com\john_smith

Examples of usermap.cfg entries

There are many different formats for entries in the `usermap.cfg` file using user names, asterisks, IP qualifiers, and more. A list of examples can help you better understand the proper syntax for various entries.

The following table describes some simple `/etc/usermap.cfg` entries:

Entry	Meaning
"Bob Garj" == bobg	The Windows name Bob Garj maps to the UNIX name bobg and vice versa.
mktg\Roy => nobody	The Windows name Roy in the mktg domain maps to the UNIX name nobody. This entry enables Roy to log in with limited access to files with UNIX-style security.
engr\Tom => ""	Disallow login by the user named Tom in the engr domain.

The following table provides some examples with asterisks in the Windows names:

Entry	Meaning
uguest <= *	All UNIX names not yet matched map to Windows user uguest.
*\root => ""	Disallow logins using the Windows name root from all domains.
corporate* == pcuser	Any user in the corporate domain maps to the UNIX name pcuser. No mapping is done for the UNIX name pcuser because an asterisk is used in the Windows user name.

Entry	Meaning
<code>Engineer == *</code>	Any UNIX name maps to the Windows name <code>Engineer</code> in the storage system's domain. No mapping is done for the Windows name <code>Engineer</code> because an asterisk is used in the UNIX user name.
<p>Either of the following entries:</p> <ul style="list-style-type: none"> • <code>homeusers* *</code> • <code>homeusers* == *</code> 	<p>All UNIX users map to the corresponding names in the <code>homeusers</code> domain. For example, a UNIX user named <code>bob</code> maps to <code>homeusers\bob</code>.</p> <p>All Windows users from the <code>homeusers</code> domain map to their corresponding UNIX names. For example, a Windows user named <code>john</code> in the <code>homeusers</code> domain maps to the UNIX name <code>john</code>.</p>

The following table provides some examples with IP qualifiers.

Entry	Meaning
<code>Engineering* <= sunbox2:*</code>	UNIX names from the host named <code>sunbox2</code> map to the same names in the <code>Engineering</code> domain.
<code>Engineering* <= 192.9.200.70:*</code>	UNIX names from the IP address <code>192.9.200.70</code> map to the same names in the <code>Engineering</code> domain.
<code>"" <= 192.9.200.0/24:*</code>	All NFS requests from the <code>192.9.200.0</code> subnet are denied because UNIX names from this subnet map to a null string.
<code>192.9.200.0/24:test-dom* => ""</code>	All users in the <code>test-dom</code> domain are denied access from the <code>192.9.200.0</code> subnet.

Entry	Meaning
<code>** == corpnet/255.255.0.0:*</code>	<p>All user names from all domains map to the corresponding UNIX names. If user names are not unique across domains, this entry might cause different Windows names to map to the same UNIX name.</p> <p>Because IP qualifiers are only for matching, specifying <code>corpnet/255.255.0.0:</code> does not affect the result of Windows to UNIX mapping.</p> <p>Because the mapping is bidirectional, all UNIX names from the <code>corpnet/255.255.0.0</code> network map to the same names in one of the storage system's trusted domains.</p>

Guidelines for mapping user names

You should follow some guidelines to keep entries simple and easy to understand.

Keep the following guidelines in mind when performing mapping:

- Keep Windows user names and UNIX user names the same whenever possible.
If the names are identical, you do not need to create map entries in the `/etc/usermap.cfg` file.
- Avoid creating confusing map entries such as these:


```
"tome s" => tjs
bill <= tjs
```
- Avoid using IP qualifiers to map users differently.
For example, it is confusing if you map UNIX user `tjs` from `UHOST1` to Windows user "Tom S" but UNIX user `tjs` from `UHOST2` to Windows user Smith. Use IP qualifiers only to restrict access.

Recommended entries for increased security

You should add several entries to the `/etc/usermap.cfg` file to prevent unauthorized users from accessing the storage system.

Remember that the order of entries is important when you copy these recommended entries to your file, because Data ONTAP uses the first matching entry to determine the mapping.

Map entry	Meaning
<code>*\root => nobody</code>	Any Windows users named <code>root</code> can log in, but they do not have UNIX permissions. For any instances of a Windows user named <code>root</code> that should map differently, you explicitly add a map entry earlier in the <code>/etc/usermap.cfg</code> file.

Map entry	Meaning
<code>guest <= administrator</code> <code>guest <= root</code>	The first entry prevents spoofing the Windows Administrator account from UNIX (if the Administrator account has not been renamed). The second entry maps the UNIX user root to the Windows guest account. Type the second entry near the end of the <code>/etc/usermap.cfg</code> file after any explicit map entries for root-privileged UNIX hosts or subnets.
<code>** => ""</code> <code>"" <= *</code>	These entries, placed at the end of the file, prevent any other mapping from occurring. They defeat the default behavior that if an entry is not matched, the same name is tried.

Restricting NFS access to mapped users or clients

In a multiprotocol environment, you can restrict NFS access to allow only users or clients that have been mapped in the `/etc/usermap.cfg` file. This optional security restriction enables you to restrict NFS access to only specified users or clients.

About this task

Typically you want Data ONTAP to automatically map users and only use the `/etc/usermap.cfg` file to specify what you do not want to map automatically. If you enable this option, all users must be mapped in the `/etc/usermap.cfg` file for successful NFS access. If a user cannot be mapped, access fails. By default, this option is disabled.

Step

- 1. To restrict NFS access to mapped users or clients, enter the following command:
`options nfs.require_valid_mapped_uid on`

Mapping a Windows account to root

If you have only CIFS clients in your environment and your storage system was set up as a multiprotocol storage system, you must have at least one Windows account that has root privilege for accessing files on the storage system; otherwise, you cannot manage the storage system because you do not have access to files with UNIX-style security, which might include some configuration files in the `/etc` directory.

About this task

If your storage system was set up as NTFS-only, however, the `/etc` directory has a file-level ACL that enables the Administrators group to access the Data ONTAP configuration files.

Step

1. Perform one of the following actions.

If you want to map...	Then...
Administrator accounts to root	<p>Verify that the <code>wabl.nt_admin_priv_map_to_root</code> option is set to on.</p> <p>All accounts in the Administrators group are considered root, even if you do not have an <code>/etc/usermap.cfg</code> entry mapping the accounts to root. If you create a file using an account that belongs to the Administrators group, the file is owned by root when you view the file from a UNIX client.</p>
Selected accounts to root	<p>For each account that maps to root, add an <code>/etc/usermap.cfg</code> entry.</p> <p>Note: It is important to have at least one Windows account that maps to root on a multiprotocol storage system. Otherwise, no accounts can access the configuration files in the <code>/etc</code> directory.</p> <p>Then disable the <code>wabl.nt_admin_priv_map_to_root</code> option by entering the following command:</p> <pre>options wabl.nt_admin_priv_map_to_root off</pre> <p>Accounts in the Administrators group no longer map to root. You can use only those accounts that you map to root in the <code>/etc/usermap.cfg</code> file to access files with UNIX-style security. Each account in the Administrators group has a separate UNIX ID.</p>

Mapping UNIX names to UIDs and GIDs

For a CIFS user to have a UID and GIDs, you must create a UNIX account in the UNIX password database that corresponds to the user’s UNIX name.

About this task

For each UNIX name, Data ONTAP obtains the UID and the primary GID from the UNIX password database. Data ONTAP obtains secondary GIDs for the UNIX name from the UNIX group database. A CIFS user whose UNIX name does not exist in the password database can still obtain a UID if you enable the default UNIX user account.

If your storage system is an NIS client before you run `cifs setup`, Data ONTAP does not automatically create the `/etc/passwd` file. If NIS is not enabled when you run `cifs setup`, Data ONTAP automatically creates the `/etc/passwd` file.

If the NIS server fails and the storage system does not have the `/etc/passwd` file, CIFS users cannot connect to the storage system. You can create the `/etc/passwd` file to ensure that the storage system can obtain UNIX credentials for CIFS users even when NIS is unavailable.

The default `/etc/passwd` file contains entries for these UNIX names:

- root

- pcuser
- nobody

For information about the format of the `/etc/group` and `/etc/passwd` files, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Step

1. Perform one of the following actions:

If you use...	Then...
NIS but not the <code>/etc/passwd</code> file	Add the UNIX name of each CIFS user to the NIS password map.
The <code>/etc/passwd</code> file but not NIS	Add an entry in the <code>/etc/passwd</code> file for the UNIX name of each user. Because Data ONTAP does not support a command for creating a password entry, use a UNIX host that supports the <code>passwd</code> command to create the <code>/etc/passwd</code> file on the host. Then copy the file from the host to the storage system.

Enabling or disabling the default UNIX user account

You should create a default UNIX user account if there are users who need to connect to the storage system occasionally but do not need to have individual entries in the UNIX password database. These users can use the default user account to connect to the storage system.

About this task

The default UNIX name of the default user is `pcuser`. You can specify another name through the `wapl.default_unix_user` option. If this option is set to a null string, no one can access the storage system as a UNIX default user. That is, each user must have an account in the password database before they can access the storage system.

For a user to connect to the storage system using the default user account, the user must meet the following prerequisites:

- The user is authenticated.
- The user is in a trusted domain.
- The user name does not map to a null string in the `/etc/usermap.cfg` file.

If quotas are enabled, the default user account is subject to quota restrictions in the same way as other users. For example, if the default user name is `pcuser` and a default user quota applies to the `/vol/vol0` volume, `pcuser` is restricted by this default user quota. For more information about quotas for the default user, see the section about how disk space owned by default users is counted in the chapter about disk space management using quotas in the *Data ONTAP Storage Management Guide*.

Step

1. Perform one of the following actions.

If you want to...	Then...
Disable the default UNIX user account	<p>Enter the following command:</p> <pre>options wafl.default_unix_user ""</pre> <p>Only users with accounts in the password database can access the storage system.</p>
Enable the default UNIX user account	<p>Create an entry either in the NIS password database or the <code>/etc/passwd</code> file for the pcuser account.</p>
Change the name of the default UNIX user account from pcuser to another name	<p>Set the <code>wafl.default_unix_user</code> option to the new name for the default UNIX user account</p> <p>For example, enter the following command to change the default user name to someuser:</p> <pre>options wafl.default_unix_user someuser</pre>

Enabling or disabling the Windows guest user account

The effect of enabling the Windows guest user account depends on how your storage system authenticates users.

About this task

Here are the possibilities:

- If the storage system uses the domain controller or local user accounts to authenticate users, enabling the Windows guest user account means that users who log in from untrusted domains can connect to the storage system.
These users use the UNIX UID that you create specifically for the Guest account. A user logged in as Guest does not have a home directory.
- If the storage system uses the UNIX password database to authenticate users, enabling the Windows guest user account has the same effect as enabling the default UNIX account, except that the user logged in as Guest does not have a home directory.

Step

1. Perform one of the following actions.

If you want to...	Then...
Disable the Windows guest user account	<p>Enter the following command:</p> <pre>options cifs.guest_account ""</pre>

If you want to...	Then...
Enable the Windows guest user account	<p>Create a user account in the NIS password database or the <code>/etc/passwd</code> file to be used by the guest account.</p> <p>Then enter the following command to specify the guest user account name used in the UNIX password database:</p> <pre>options cifs.guest_account <i>unix_name</i></pre> <p><i>unix_name</i> is the name of the user account in the UNIX password database.</p>

Managing the SID-to-name map cache

CIFS frequently is required to map security identifiers (SIDs) to user and group names and vice versa for user authentication, quota management, console command processing, and various RPC responses. The SID-to-name map cache contains entries that map SIDs to pre-Windows 2000 user and group names.

About this task

The storage system obtains the SID-to-name mapping information by querying the domain controller. To minimize multiple lookups of the same names, SID-to-name information received from the domain controller is saved in the SID-to-name map cache on the storage system.

The SID-to-name map cache is enabled on the storage system by default. You can manually control the cache by changing the lifetime of the entries, clearing entries, or turning SID-to-name map caching off or on. The cache persists if CIFS is terminated or restarted, but it does not persist across a reboot or a takeover and giveback.

When the storage system requires SID-to-name mapping information, it first looks for a matching entry in the SID-to-name map cache. If a matching entry is not found or if an expired matching entry is found, the storage system queries the appropriate domain controller for current mapping information. If the domain controller is not available, an expired mapping entry might be used by the storage system.

Here are the main benefits of using the SID-to-name map cache for name lookup:

- Increased performance for authorization
- Faster user response for console commands that perform mapping operations

Enabling or disabling the SID-to-name map cache

You can enable or disable the SID-to-name map cache by setting the `cifs.sidcache.enable` option to `on` or `off`, respectively.

Step

1. Perform one of the following actions.

If you want the SID-to-name map cache...	Then...
Enabled	Enter the following command: options cifs.sidcache.enable on
Disabled	Enter the following command: options cifs.sidcache.enable off

Changing the lifetime of SID-to-name mapping entries

You can change the lifetime of SID-to-name mapping entries by setting the `cifs.sidcache.lifetime` option.

Step

1. Enter the following command:

```
options cifs.sidcache.lifetime time
```

time is the number of minutes that new mapping entries are used before they expire.

Clearing all or part of the SID-to-name map cache

Periodically, expired entries that are more than one week old are automatically cleared from the SID-to-name map cache; however, you might want to manually clear entries in the SID-to-name map cache when users change their accounts or user names. Alternatively, you might want to manually clear all SID-to-name map cache entries to prevent the storage system from using an expired entry when the domain controller is not available.

Step

1. Perform one of the following actions:

If you want to clear the SID-to-name map cache entries for...	Enter the command...
All Windows domains, users, groups, and SIDS	cifs sidcache clear all
A specific Windows domain	cifs sidcache clear [domain] <i>domain</i> is the Windows domain of the cache entries you want to clear. If you do not specify the domain, entries for the storage system's home domain are cleared from the cache.

If you want to clear the SID-to-name map cache entries for...	Enter the command...
A specific user or group	<pre>cifs sidcache clear user <i>username</i></pre> <p><i>username</i> is the specific Windows user or group entry you want to clear from the cache. The user name can be specified in the following ways:</p> <ul style="list-style-type: none">• <i>domain\username</i>• <i>username</i> <p>When the user name is specified without a domain, the storage system’s home domain is used for the domain.</p>
A specific SID	<pre>cifs sidcache clear sid <i>textualSid</i></pre> <p><i>textualSid</i> is the textual form of the SID you want to clear from the cache. Specify the SID using standard “S-1-5...” syntax.</p> <p>Example:</p> <pre>cifs sidcache clear sid S-1-5-21-4503-17821-16848-500</pre>

Using LDAP services

An LDAP server enables you to centrally maintain user information. If you store your user database on an LDAP server in your environment, you can configure your storage system to look up user information in your existing LDAP database.

About this task

As a result, you do not have to maintain separate configuration files for each storage system that is on your network. If you have several storage systems on your network, maintaining user information centrally saves you from updating these files on each storage system every time you add or delete a user or a group.

Data ONTAP supports LDAP for user authentication, file access authorization, user lookup and mapping services between NFS and CIFS, and LDAP over the Secure Sockets Layer (SSL).

Data ONTAP supports connections to LDAP servers that require signing. LDAP signing support is enabled by default.

Configuring LDAP services

This section provides information to help you configure Data ONTAP to connect to your LDAP database.

Specifying the general search base and scope

The LDAP base is the distinguished name of the LDAP tree in which user information is stored. All lookup requests sent to the LDAP server will be limited to the search base and scope specified by the `ldap.base` option value, unless further restricted by a more specific base and scope lookup value, such as `ldap.base.passwd`, `ldap.base.group`, or `ldap.base.netgroup`.

Step

1. Enter the following command:

```
options ldap.base name
```

name specifies the base distinguished name. Use quotes around names with embedded spaces.

Example

```
options ldap.base "o=examplecompany,c=us"
```

Overriding general base and scope values for user password, group, and netgroup lookups

Although it is not required, you can specify base and scope values for user password, group, and netgroup lookups, to limit such lookup queries to a specific branch of your LDAP database. Limiting the search base and scope of these queries can significantly improve performance.

About this task

After you specify the search base and scope values for the `ldap.base.passwd`, `ldap.base.group`, and `ldap.base.netgroup` options, these values take precedence over the search base and scope set for `ldap.base`, for user password, group, and netgroup lookups, respectively.

Steps

1. Set the base and scope search values for user password lookups by entering the following command:

```
options ldap.base.passwd "distinguished_name"
```

distinguished_name is the specific branch for user password lookups as defined in your LDAP database.

2. Set the base and scope search values for group lookups by entering the following command:

```
options ldap.base.group "distinguished_name"
```

distinguished_name is the specific branch for group lookups as defined in your LDAP database.

3. Set the base and scope search values for netgroup lookups by entering the following command:

```
options ldap.base.netgroup "distinguished_name"
```

distinguished_name is the specific branch for netgroup lookups as defined in your LDAP database.

Specifying LDAP servers

You can specify the LDAP servers to be used for LDAP queries by setting the `ldap.servers` option.

Step

1. Enter the following command:

```
options ldap.servers "name[ name...]"
```

name is the name of an LDAP server. You can enter multiple server names using a space-separated list enclosed in quotes. Data ONTAP attempts to establish connections in the order in which you specify these servers.

Note: A Windows LDAP server uses simple authentication instead of SASL unless the following conditions are met: you specify the Windows LDAP server as a name, not an IP address, and you specify the IP address and name of the Windows LDAP server in the `/etc/hosts` file. For information about editing the `/etc/hosts` file, see the *Data ONTAP System Administration Guide for 7-Mode*.

Example

```
options ldap.servers "server1 server2"
```

Specifying preferred LDAP servers

You can set the `ldap.servers.preferred` to specify preferred LDAP servers. This allows you to improve performance by directing to specific LDAP servers that are on faster links.

Step

1. Enter the following command:

```
options ldap.servers.preferred "name [ name...]"
```

name specifies the name of a preferred LDAP server. You can enter multiple server names using a space-separated list enclosed in quotes.

Example

```
options ldap.servers.preferred "server1 server2"
```

Enabling or disabling LDAP

You can enable or disable LDAP by setting the `ldap.enable` option to `on` or `off`, respectively.

Step

1. Perform one of the following actions.

If you want to...	Then...
Enable LDAP	Enter the following command: <code>options ldap.enable on</code>
Disable LDAP	Enter the following command: <code>options ldap.enable off</code>

Enabling or disabling SSL for LDAP traffic

You can enable or disable secure sockets layer (SSL) encrypting of LDAP traffic by setting the `ldap.ssl.enable` option to `on` or `off`, respectively.

Before you begin

You must have a root authority-signed certificate installed on your storage system before enabling SSL for LDAP.

Note: The same certificate-signing authority must issue both the certificate on the storage system and the certificate on the server.

Step

1. Perform one of the following actions:

If you want SSL for LDAP...	Then...
Enabled	Enter the following command: <code>options ldap.ssl.enable on</code>
Disabled	Enter the following command: <code>options ldap.ssl.enable off</code>

After you finish

If you installed the root certificate on the storage system after enabling this option, you must disable and then reenable it so the storage system reads the certificate.

Installing a root certificate for SSL for LDAP traffic

You can install a root certificate for use for Secure Sockets Layer (SSL) encrypting of LDAP traffic on your storage system by using the `keymgr` command.

Steps

1. Download a certificate from your preferred trusted signing authority to the storage system. Remember the certificate's location on the storage system.

2. Enter the following command:

```
keymgr install root certificate_filename
```

certificate_filename is the complete file name for the certificate. After the `keymgr` command installs the certificate, you can remove the copy you placed on the storage system.

Example

```
keymgr install root /etc/my_cert
```

Note: The same certificate-signing authority must issue both the certificate on the storage system and the certificate on the server.

3. Set the LDAP port to port 636.

Adding the ldap entry to the `/etc/nsswitch.conf` file

You can add the `ldap` entry to the `/etc/nsswitch.conf` file to enable LDAP for UNIX client authentication.

Steps

1. Open the `/etc/nsswitch.conf` file on the storage system for editing.
2. Enter the following at the password, group, and netgroup lines:

```
ldap
```

You can optionally add `files` and/or `nis` to the password line, but they must be entered after `ldap` if you want to use LDAP as the primary mechanism to retrieve user information.

Example

```
passwd: ldap files nis
```

3. Save the file.

Specifying the administrative user name

If anonymous authentication does not work in your environment, you need to specify an administrative user name to be used for administrative queries for looking up UIDs and GIDs.

Step

1. Enter the following command:

```
options ldap.name name
```

name is the LDAP distinguished name to be used for administrative queries. You should use the name of a user with read-only access to the LDAP database. Use quotes around names with embedded spaces.

Example

```
options ldap.name "cn=root,o=examplecompany,c=us"
```

Specifying the administrative password

You can set the administrative password by setting the `ldap.passwd` option.

Step

1. Enter the following command:

```
options ldap.passwd password
```

password is the password for the administrative user.

The password is displayed as a series of asterisks.

Specifying the LDAP port

You can set the `ldap.port` option to specify the port for LDAP queries. This is useful if the LDAP server has been set up to use a port other than the default for LDAP, port 389.

Step

1. Enter the following command:

```
options ldap.port N
```

N specifies the LDAP port number.

LDAP server option precedence

Data ONTAP chooses an LDAP server based on your LDAP server option settings.

Server designation option	Server selection order
<code>ldap.preferred.servers</code>	When specified, servers listed in this option value will be tried first, according to list order.
<code>ldap.servers</code>	When no <code>ldap.preferred.servers</code> are specified, or specified servers are not available, servers designated in this option value will be tried, according to list order.
<code>ldap.ADdomain</code>	When no <code>ldap.preferred.servers</code> and no <code>ldap.servers</code> are specified or available, servers designated in this option value will be tried using domain controller selection methodology.

Managing client authentication and authorization

You can enable LDAP authentication of UNIX and Windows clients; in addition, you can enable LDAP authorization of Windows client access to UNIX files and UNIX client access to NTFS or mixed files.

Enabling LDAP-based UNIX client authentication

You can enable LDAP-based UNIX client authentication by making sure `ldap` is entered on the password line of the `/etc/nsswitch.conf` file.

Enabling LDAP-based Windows client authentication

You can authenticate Windows clients through an LDAP server by performing steps in addition to adding `ldap` to the `passwd` line of the `/etc/nsswitch.conf` file.

Steps

1. Run `cifs setup` on the storage system to be accessed, and specify NIS/LDAP as the authentication method to be used for CIFS clients on that storage system.
2. Configure the local security settings of each Windows client to use clear text (unencrypted) password authentication rather than Kerberos or other encrypted authentication methods.
3. Verify that your Windows clients have their `userpassword` attribute configured in the LDAP user database.

Enabling LDAP authorization for NFS file access from Windows clients

You can enable authorization of Windows client access to UNIX files on a storage system that uses LDAP authentication by performing two tasks.

Steps

1. On the storage system to be accessed, verify that every CIFS user who needs to access UNIX files is mapped to an associated UNIX user name in the `usermap.cfg` file.
2. Verify that every associated UNIX user name has an entry in the LDAP database.

Enabling LDAP authorization for NTFS or mixed file system access from UNIX clients

You can enable authorization of UNIX client access to an NTFS or mixed file system on a storage system that uses LDAP authentication by performing several tasks.

Steps

1. Verify that every UNIX user that needs to access an NTFS or mixed file system has an entry in the LDAP database.
2. On the storage system to be accessed, verify that every UNIX user that needs to access an NTFS or mixed file system is mapped to an associated CIFS user name in the `usermap.cfg` file.

Managing LDAP user-mapping services

You can use LDAP services to map between UNIX and Windows user accounts, instead of using NIS data or to adding entries to the `/etc/usermap.cfg` file. By default, Data ONTAP uses the same (one-to-one) user account resolution process in both directions: UNIX-to-Windows mapping and Windows-to-UNIX mapping.

About this task

By default, LDAP-based user-mapping is disabled. (Data ONTAP retrieves user-mapping information from the `/etc/usermap.cfg` file.)

When converting to LDAP from file-based user-mapping, you must remove mapping entries (except for null session entries) from the `/etc/usermap.cfg` file. If mapping entries are present in that file, they will be used for user-mapping instead of LDAP records.

If you've configured Data ONTAP for null sessions, make sure you leave the null session client entry in the `/etc/usermap.cfg` file.

To allow Data ONTAP access to LDAP lookup services, if your UNIX user account information is stored in a non-Active Directory LDAP server, that LDAP server must be configured to allow either simple authentication or anonymous user searches.

Steps

1. From the Data ONTAP command line, specify a value for the option `ldap.usermap.attribute.windowsaccount`:

```
options ldap.usermap.attribute.windowsaccount account_name
```

account_name is the user object attribute Data ONTAP will use for Windows account lookups.

2. Extend your LDAP schema to include the user object attribute you entered in Step 1.
3. From the Data ONTAP command line, specify a value for the `ldap.usermap.attribute.unixaccount` option:

```
options ldap.usermap.attribute.unixaccount account_name
```

account_name is the user object attribute Data ONTAP will use for UNIX account lookups.

4. Extend your LDAP schema to include the values you entered in Step 2 and Step 3.
5. Enter the following command:

```
options ldap.usermap.enable on
```

If you have a significant load on your LDAP server, you might want to improve performance by setting a separate search base or search base and scope for user-mapping.

Specifying base and scope values for user-mapping

LDAP options allow you to set search base and scope, to limit attribute searches to the appropriate areas of your LDAP database. Setting these options will improve the speed of LDAP lookups.

Step

1. Use the following syntax when specifying search base and scope.

Base and scope values must correspond to the structure of your LDAP data:

```
options ldap.usermap.base "base[:scope][;base2[:scope2]]"
```

Examples

Entering this command sets the search base for user-mapping lookups to `ou=People,dc=domain0` and the (unspecified) search scope defaults to SUBTREE:

```
options ldap.usermap.base ou=People,dc=domain0"
```

The use of parentheses applies the specified search scope (BASE) to `ou=People,dc=domain0`. The unspecified search scope for the `o ("org")` object defaults to SUBTREE.

```
options ldap.usermap.base "(ou=People,dc=domain0):BASE;o=org"
```


After you finish

For more information about setting search base and scope values, see your LDAP documentation.

Managing Active Directory LDAP servers

Data ONTAP provides the ability to connect to Active Directory for LDAP lookup services.

Using Active Directory LDAP servers

You can use Active Directory for LDAP services by entering the fully qualified Active Directory domain into the Data ONTAP `ldap.ADdomain` option.

About this task

As Windows-to-UNIX mapping is performed using Active Directory, Data ONTAP does the following:

- Verifies that the user account exists within the Active Directory domain specified for that account
- Performs a query to the Active Directory domain specified in the `ldap.ADdomain` option
- Returns the UNIX user account information and verifies that the user account exists

Requirements for Active Directory LDAP servers

You need several things to use Active Directory as your LDAP server.

You need these things to use Active Directory as your LDAP server:

- A valid CIFS license
- Your storage system joined to an Active Directory domain
- A two-way trust relationship established between your storage system's domain and your LDAP server's domain, if they are different

Enabling Active Directory LDAP lookup services

You can enable Active Directory for LDAP lookup services by performing several tasks.

Steps

1. If your UNIX user account information is not in Active Directory, or if it is not in an LDAP server that is configured to allow anonymous user searches, enter the user name and password to be used for LDAP lookups into the `ldap.name` and `ldap.passwd` options, respectively.


```
options ldap.name user_name
options ldap.passwd password
```
2. In the `/etc/nsswitch.conf` file, specify `ldap` for the `passwd` entry, the `group` entry, or both, to designate LDAP as the lookup service to use.
3. If you have a custom schema, enter values for NSSMAP options.
4. From the Data ONTAP command line, enter the following command:

```
options ldap.ADDomain fully_qualified_domain_name
```

Example

```
options ldap.ADDomain group.company.com
```

Note: The domain you enter must either be the local domain or a domain that shares a trust relationship with the local domain.

Monitoring Active Directory LDAP server connections

To monitor Active Directory LDAP server connection, you can display Active Directory LDAP server information and connection status for all LDAP server types.

Step

- 1. Perform one of the following actions.

If you want to...	Enter the command...
Display Active Directory LDAP server information	cifs domaininfo Following the list of domain controller connections and domain controller selection preferences, a list of Active Directory LDAP server connections is displayed, followed by the list of LDAP server selection preferences.
Display connection status for all LDAP server types	netstat Both Active Directory and non-Active Directory LDAP server connection state information is shown on port 389 (or the non-default value assigned using the <code>ldap.port</code> option).

Troubleshooting Active Directory LDAP server connections

You can instruct Data ONTAP to log all domain controller address discovery and connection activities by setting the `cifs.trace_dc_connection` option to `on`.

Step

- 1. Enter the following command:

```
options cifs.trace_dc_connection on
```

Data ONTAP logs all domain controller address discovery and connection activities to the system log.

About Active Directory LDAP server connection pooling and selection

Data ONTAP performs several operations to improve LDAP performance.

These operations include the following operations:

- Active Directory LDAP server connections are pooled on a per-domain basis.

- When no response is received from the current LDAP server, subsequent connections are made to the next best available LDAP server.
- Once every minute, Data ONTAP performs a check to see whether a better LDAP server has become available.
- Every four hours, Data ONTAP discovers the available Active Directory LDAP servers and reorders the list, sorting servers in the following order:
 - Preferred servers, left in the order specified by the `prefdc` command.
 - Favored servers, sorted by fastest response time
 - Other Active Directory LDAP servers, sorted by fastest response time

LDAP server authentication limitations

Data ONTAP connects to servers specified by `ldap.servers` and `ldap.preferred.servers` options and attempts to authenticate using a simple bind. Because simple binds do not provide sufficient authentication to establish a connection with Active Directory servers, do not specify Active Directory servers within these two option values.

Managing LDAP schema

By default, Data ONTAP supports LDAP servers that comply with RFC 2307, which specifies a Network Information Service (NIS)-style schema. You can replace the default values of LDAP options with your custom attribute names to configure Data ONTAP to query your custom (not RFC 2307-compliant) schema.

About this task

Your RFC 2307-compliant schema must be extended on the LDAP servers that you want to use for LDAP queries.

For more information refer to RFC 2307 or to documentation by third-party directory integration vendors.

About the default schema

By default, the Data ONTAP's schema variables are set to the appropriate RFC 2307 values.

Option	Default value (per RFC 2307)
<code>ldap.nssmap.objectClass.posixAccount</code>	<code>posixAccount</code>
<code>ldap.nssmap.objectClass.posixGroup</code>	<code>posixGroup</code>
<code>ldap.nssmap.attribute.groupname</code>	<code>cn</code>
<code>ldap.nssmap.attribute.netgroupname</code>	<code>cn</code>
<code>ldap.nssmap.attribute.nisNetGroupTriple</code>	<code>nisNetGroupTriple</code>

Option	Default value (per RFC 2307)
<code>ldap.nssmap.attribute.memberUid</code>	<code>memberUid</code>
<code>ldap.nssmap.attribute.uid</code>	<code>uid</code>
<code>ldap.nssmap.attribute.uidNumber</code>	<code>uidNumber</code>
<code>ldap.nssmap.attribute.gidNumber</code>	<code>gidNumber</code>
<code>ldap.nssmap.attribute.userPassword</code>	<code>userPassword</code>
<code>ldap.nssmap.attribute.homeDirectory</code>	<code>homeDirectory</code>
<code>ldap.nssmap.attribute.loginShell</code>	<code>loginShell</code>
<code>ldap.nssmap.attribute.gecos</code>	<code>gecos</code>

Modifying the custom schema options to match your LDAP schema

You can change Data ONTAP's schema to match your LDAP schema by changing the appropriate `ldap.nssmap.*` options.

Step

1. Enter the following command:

```
options ldap.nssmap.attribute.gidNumber object
```

object specifies the object that contains Group ID (GID) numbers. The default is `gidNumber`.

Examples

For a custom LDAP schema in which the object containing GID numbers is “groupid,” you would enter the following command:

```
options ldap.nssmap.attribute.gidNumber groupid
```

Enabling Storage-Level Access Guard using the `fsecurity` command

Beginning in Data ONTAP 7.2.2, storage administrators can set security (permissions and auditing) on volumes and qtrees using the `fsecurity` command. This feature is called *Storage-Level Access Guard*.

About this task

With the Storage-Level Access Guard security in place, any storage object can contain up to three types of security layers:

- NTFS/UNIX/NFSv4 security.
Exists on the directory or file that represents the storage object. This security is also the same security you can set from a client.
- Storage-Level Access Guard file security.
Applies to every file within the storage object. Applying this security will not affect access to, or auditing of, directories.
- Storage-Level Access Guard directory security.
Applies to every directory within the storage object. Applying this security will not affect access to, or auditing of, files.

Note: At this time, only NTFS access permissions are supported for Storage-Level Access Guard. For a UNIX user to perform a security check on qtrees or volumes where Storage-Level Access Guard has been applied, the UNIX user must be mapped to a Windows user.

Storage-Level Access Guard security applies to files and directories but is not inherited by them. If you view the security settings on a file or directory, you will not see the Storage-Level Access Guard security.

However, access to a file or directory in Data ONTAP is determined by the combined effect of both the native permissions applied to files and/or directories and the Storage-Level Access Guard permissions set on qtrees and/or volumes. Both levels of security are evaluated to determine what the effective permissions a file or directory has.

About the `fsecurity` command

Using the `fsecurity` command, storage administrators can apply security over large directories without experiencing significant performance degradation, because security settings are being managed locally on the storage system, not from remote clients. In addition, storage administrators can set security on many files and directories at once by using the same command.

Note: For a list of all `fsecurity` commands, enter `fsecurity help` at the storage system command line or refer to the `fsecurity(1)` man page.

Generating and editing the job definition file

You can generate a job definition file to apply Storage-Level Access Guard security to a qtree or volume, or to set bulk permissions on files and directories.

About this task

The job definition file is a Unicode text file that contains information such as security descriptors and paths that define discretionary access control lists (DACLS) and system access control lists (SACLs).

This information is encoded using the Security Descriptor Definition Language (SDDL).

After you create or edit the file and copy it to the storage system, you can use the `fsecurity apply` command to validate and apply the file's security definitions. Running the command on the file creates a job that runs in the background on the storage system. After the job is complete, you can view the results from the storage system console.

There are no requirements for the name and storage system location of the job definition file. In these examples, the following name and location are used:

```
/vol/vol0/templates/security-base.sec
```

The job definition file format must be ASCII or Unicode (UCS-2).

There are two ways to create and update the job definition file:

- Using the `secedit` utility (available from the NOW site)
- Using a text editor

Managing the job definition file with the `secedit` utility

You can create and edit the job definition file using the `secedit` utility.

Steps

1. Download the `secedit.exe` executable file from the NetApp Support Site.
2. Follow the instructions in the `Secedit_Readme.txt` file to create the job definition file.

Managing the job definition file with a text editor

You can generate, update, and then validate the job definition file using a text editor.

Steps

1. Create a text file (for example, `security-base.sec`) or edit an existing job definition file.
2. Copy the new or updated file to a directory on your storage system (for example, `/vol/vol0/templates/`).
3. Check the validity of the file before you apply the definitions to jobs by running the `fsecurity` `apply` command with the `-c` option.

Note: If any line in the definition file is invalid, the security job will not be created when the `fsecurity apply` command is run.

Specifying job definition file elements

When you are defining your security settings in the job definition file, you can apply bulk security settings (permissions and auditing) by specifying a propagation mode.

About this task

Specifying a propagation mode allows you to quickly and effectively configure these settings without the performance degradation caused by applying them over a network.

The propagation modes are as follows:

- 0 = Propagate inheritable permissions to all subfolders and files (Propagate).

- 1 = Do not allow permissions on this file or folders to be replaced (Ignore); this mode is not currently available.
- 2 = Replace existing permissions on all subfolders and files with inheritable permissions (Replace).

The following is an example of an fsecurity job description file.

```
cb56f6f4
1,0,"/vol/vol0/qt1",0,"D:(A;CIOI;0x1f01ff;;DOMAIN\Administrator)"
1,1,"/vol/vol0/qt2",0,"D:(D;CIOI;0x000002;;Everyone)"
```

The first line, the string `cb56f6f4`, is mandatory, and is always the same. The following table describes how the elements in the second line of the example apply security settings to a qtree called `/vol/vol0/qt1`.

Sample element	Description
1	NTFS security type
0	Standard security; Storage-Level Access Guard security not set
"/vol/vol0/qt1"	Path of the target storage object (double quotes are required for this field)
0	Propagation mode (0 stands for “propagate” in this example)
"D:(A;CIOI;0x1f01ff;;DOMAIN\Administrator)"	SDDL representation of a DACL that gives the domain administrator Full Control (double quotes are required for this field)

For more information about the format and syntax of the job definition file, see the `fsecurity(5)` man page.

Creating a security job and applying it to the storage object

The `fsecurity apply` command is used to create a security job based on the job definition file. This command is also used to apply Storage-Level Access Guard to a qtree or volume, or bulk security settings to files and directories. Using this command also allows you to set SACs for auditing at the qtree and volume level.

About this task

You can apply the following options when creating a security job:

- The `-c` option lets you check the validity of the job without actually applying the contents.
- The `-i` option lets you ignore errors and continue to process the job.
- The `-v` lets you view each task within the job as it is generated.

For a complete description of the `fsecurity apply` command and its options, refer to the `fsecurity_apply(1)` man page.

Security jobs can be run simultaneously by different administrators, and can conflict with one another.

Step

1. Enter the following command:

```
fsecurity apply job_definition_file_path
```

Example

```
fsecurity apply /vol/vol10/templates/security-base.sec
```

```
Added security job 94089
```

The job ID is used to monitor the status of, or cancel, the job.

Checking the status of or canceling a security job

The `fsecurity status` command can be used to view the status of jobs that are currently running and the completion status of the previous 15 jobs.

About this task

The `fsecurity cancel` command can be used to stop all of the currently running jobs. If a job ID is specified, only that job will stop.

Note: Completed jobs cannot be canceled.

For a complete description of these commands, refer to the `fsecurity_status(1)` and `fsecurity_cancel(1)` man pages.

Step

1. Perform one of the following actions.

If you want to...	Enter the command...
View job status	<code>fsecurity status [job_id]</code>
Cancel a job	<code>fsecurity cancel [job_id] all</code>
The job ID is used to cancel a specific job, and the option <code>all</code> is to cancel all jobs.	

Displaying the security settings on files and directories

The `fsecurity show` command can be used to view the security settings on files and directories.

About this task

The output of this command contains the security style of the qtree or volume that the file or directory resides in. The current security style varies in mixed qtree environments and depends on which security style is currently active on the storage object.

When specifying a file or directory path, wildcards can be used to list the security for the contents of a directory.

For a complete description of this command, refer to the `fsecurity_show(1)` man page.

Step

1. Enter the following command:

```
fsecurity show file_directory_qtree_path [option]
```

You can also specify the inode number of the file or directory (instead of the file or directory path), as shown in the following example.

```
fsecurity show -v volume_name -i inode_number [option]
```

For a complete listing of options and description of command output, see the `fsecurity_show(1)` man page.

Removing the Storage-Level Access Guard

The `fsecurity remove-guard` command can be used to remove the Storage-Level Access Guard from a qtree or volume. A qtree cannot be deleted if Storage-Level Access Guard is applied to it. For more information, refer to the `fsecurity remove-guard(1)` man page.

Step

1. Enter the following command:

```
fsecurity remove-guard volume_qtree_path
```

Note: Removing the Storage-Level Access Guard does not remove the standard file-level security (such as NTFS security) that is present on the files and directories within a qtree or volume.

Auditing system access events

Data ONTAP audits logon, logoff, and file access events similarly to Windows. There are some differences, however, in how you enable auditing and how you manage the files that log audit event information.

About auditing

When you configure Data ONTAP for auditing, the event log file and the settings for all options persist across a reboot or if CIFS is terminated or restarted.

Data ONTAP auditing can be performed in two ways:

- CIFS auditing refers to auditing access events from Windows clients that access data on the storage system using the CIFS protocol.
- NFS auditing refers to auditing access events from UNIX clients that access data on the storage system using the NFS protocol.

Both CIFS and NFS auditing can be configured on a storage system. Each type has different configuration requirements and audit capabilities.

Auditing is not currently supported for other file access protocols.

Events that Data ONTAP can audit

You can enable auditing for several categories of events.

The following categories can be audited:

- Logon and logoff events (available only with CIFS auditing enabled)
- Local user and group account management (available only with CIFS auditing enabled)
- File access events at the file and directory level

Note: You must activate access auditing for individual files and directories.

- File access events at the qtree or volume level

Note: Auditing of events at the qtree or volume level is available only by applying Storage-Level Access Guard security.

Event ID	Event	Description	Category
516	AdtEvtntDiscard	Audit events were lost	Audit Log
517	AdtLogClear	Audit log was cleared	Audit Log
528	AdtSuccessfulLogon	Local logon	Logon/Logoff
529	AdtUnknownUser	Unknown user name or bad password	Logon/Logoff

Event ID	Event	Description	Category
530	AdtCantLogonNow	Account logon time restriction	Logon/Logoff
531	AdtAccountDisabled	Account currently disabled	Logon/Logoff
532	AdtUserAccountExpired	User account has expired	Logon/Logoff
533	AdtCantLogonHere	User cannot log on to this computer	Logon/Logoff
534	AdtLogonTypeRestricted	User not granted logon type here	Logon/Logoff
535	AdtPasswordExpired	User password has expired	Logon/Logoff
536	AdtNetLogonInactive	NetLogon component is not active	Logon/Logoff
537	AdtUnsuccessfulLogon	Logon failed for other reasons	Logon/Logoff
538	AdtUserLogoff	Local or network user logoff	Logon/Logoff
539	AdtLockedOut	Account locked out	Logon/Logoff
540	AdtSuccessfulNetLogon	Network (CIFS) logon	Logon/Logoff
560	AdtObjOpen	Object (file or directory) open	File Access
562	AdtHandleClosed	Handle that resulted in AdtObjOpen is closed	File Access
563	AdtObjOpenForDelete	Object (file or directory) open for deletion	Logon/Logoff
567	AdtObjAccessAttempt	Object access (read, write, etc.)	File Access
612	AdtPolicyChange	Audit policy changed	Policy Change
624	AdtUserCreated	User created	Account Management
630	AdtUserDeleted	User deleted	Account Management
635	AdtGroupCreated	Group created	Account Management
636	AdtLclGrpMemberAdded	Security enabled local group member added	Account Management

Event ID	Event	Description	Category
637	AdtLclGrpMemberRemoved	Security enabled local group member removed	Account Management
638	AdtGroupDeleted	Group deleted	Account Management

Configuring system event auditing

You must perform several tasks to configure system event auditing.

Steps

1. Determine what events you want to audit.
For example, if you want to audit all the events on a volume or qtree, apply the Storage-Level Access Guard security using the `fsecurity` command.
2. If you want to audit file and directory access events, set your system access control lists (SACLs).
3. Enable CIFS auditing and NFS auditing, as applicable.
4. If you want to use Live View to manage auditing, enable Live View.
Otherwise, familiarize yourself with audit log management.
5. Use Event Viewer to display audit events.

Setting SACLs

System access control lists (SACLs) can be used to enable auditing access on files and directories.

About this task

There are three ways to set SACLs for auditing access:

- If you want to audit access events on all files and directories within a volume or qtree, you should set SACLs by applying Storage-Level Access Guard security.
- If you want to audit access events on individual files and directories, you can set SACLs in two ways:
 - Using your Windows Explorer GUI.
 - Using the `fsecurity` command

Note: Make sure that you select only the events you need to audit, as selecting too many audit options might impact system performance.

To enable auditing access on individual files and directories, complete the following steps on the Windows administration host.

Steps

1. Select the file or directory for which you want to enable auditing access.

2. Right-click on the file or directory, and select **Properties**.
3. Select the **Security** tab.
4. Click **Advanced**.
5. Select the **Auditing** tab.
6. Add, edit, or remove the auditing options you want.

For more information on how to set these options, see your Windows documentation.

Configuring Data ONTAP for CIFS auditing

When you enable or disable CIFS auditing, you enable auditing of policy change events. There is not a separate CIFS option to enable policy change events at this time.

Before you begin

Following are the prerequisites for CIFS auditing:

- CIFS must be licensed and enabled on the storage system before enabling auditing.
- The file or directory to be audited must be in a mixed or NTFS volume or qtree. You cannot audit CIFS events for a file or directory in a UNIX volume or qtree unless Storage-Level Access Guard is enabled.
- You must specify access events to record.
- Event auditing is turned off by default.

To identify events for auditing, you must enable individual options and enable auditing.

Step

1. Perform one of the following actions.

If you want to turn auditing on or off for...	Enter the command...
File access events	<code>options cifs.audit.file_access_events.enable {on off}</code>
Logon and logoff events	<code>options cifs.audit.logon_events.enable {on off}</code>
Local account management events	<code>options cifs.audit.account_mgmt_events.enable {on off}</code> Note: You use MMC Event Viewer to view changes to the account management events.

If you want to turn auditing on or off for...	Enter the command...
All events	<pre>cifs audit {start stop}</pre> <p>Alternatively, you can start and stop CIFS auditing using the <code>cifs.audit.enable</code> option. For example, entering the following command is the equivalent of using the <code>cifs audit start</code> command:</p> <pre>options cifs.audit.enable {on off}</pre> <p>Use <code>on</code> to start CIFS auditing or <code>off</code> to stop auditing.</p> <p>Note: CIFS auditing is disabled by default.</p>

Configuring Data ONTAP for NFS auditing

NFS auditing can record access events for files and directories, but it cannot record logon, logoff, and other events supported by CIFS auditing. The file or directory to be audited can be in a volume or qtree of any security style (NTFS, UNIX, or mixed).

Before you begin

- CIFS must be licensed and enabled on the storage system.
- CIFS auditing must be enabled on the storage system.
Auditing is disabled by default.
- Events to record must be identified.

Specifying NFS audit events

To specify events for NFS auditing in an NTFS or mixed security style qtree or volume, you must set system access control lists (SACLs) on files and directories.

Steps

1. Create the log filter file (usually called `/etc/log/nfs-audit`) on the storage system.

This file is used to identify which file events get included in the audit log by default. The filter file has no content.

Note: You must create the NFS log filter file in an NTFS or mixed style volume or qtree. If you do not, you will not be able to set a SACL on the filter file, which is required for auditing.

2. Set the `cifs.audit.nfs.filter.filename` option to identify the filter file.

For more information about the `cifs.audit.nfs.filter.filename` option, see the `options(1)` man page.

3. Set the filter file's system access control list (SACL).

After you finish

You can create an NFS filter file for auditing events in NTFS or mixed security style qtrees, but SACLs set on individual files and directories take precedence over the SACL set on the filter file.

How the filter file controls NFS audit events

The log filter file controls file audit events by means of the SACL you set on it. Setting a SACL on the filter file has the same effect as setting the same SACL on every file and directory on the storage system.

Note: Because the log filter file SACL can potentially generate audit events from every file and directory on the storage system, enabling NFS auditing with the log filter file can affect system performance.

The effect of the filter file depends on the security setting of the qtree in which the files are located.

When an operation is performed on files in a UNIX security style, the event is logged depending on the SACL on the filter file.

When an operation is performed on files in an NTFS or mixed security-style qtree that has no SACL set, the event is logged depending on the SACL on the filter file.

However, if SACLs are set on individual files or directories, these SACLs take precedence over the SACL set on the filter file.

Enabling NFS auditing

You can enable NFS auditing by performing several tasks.

About this task

For more information about the options described in these steps, see the options(1) man page.

Steps

1. In the `/etc/log` directory on the storage system, create a file called `nfs-audit`.

Note: Steps 1 and 2 are mandatory for auditing in a UNIX security style qtree but optional for auditing in NTFS or mixed security style qtrees.

2. To identify the NFS log filter file, enter the following command:

```
options cifs.audit.nfs.filter.filename /etc/log/nfs-audit
```

3. To enable auditing of file access events, enter the following command:

```
options cifs.audit.file_access_events.enable on
```

Note: Auditing of file access and logon events is turned off by default.

4. To enable NFS auditing, enter the following command:

```
options cifs.audit.nfs.enable on
```

5. Configure audit log management.
6. On the Windows administration host, set the filter file's system access control list (SACL).

Configuring Live View

When Live View is enabled, an Access Logging Facility (ALF) daemon runs once a minute, flushing audit events from memory to the internal log file `/etc/log/cifsaudit.alf` on disk.

About this task

The ALF daemon also attempts to save and convert ALF records to EVT records that can be viewed by Event Viewer. It does so either once every minute, or when the `.alf` file becomes 75 percent full.

EVT records are stored in three files in the `/etc/log` directory:

- `fixedsection`
- `varsectiona`
- `varsectionb`

The ALF daemon uses these files to service Eventlog RPC requests from Windows clients running Event Viewer. When Live View is enabled, Event Viewer displays the most recent audit events up to 5,000 records.

Each time new records are saved from the internal log file, they are written to the Live View files and they are also backed up into EVT files. The backup files are saved in the `/etc/log` directory with a timestamp as part of their name.

Audit events can be viewed in real-time and backup EVT files can be viewed as static files using Event Viewer.

Note: Beginning in Data ONTAP 7.2.2, Live View can be enabled together with `cifs.audit.autosave` options, which control the size of the internal audit file and how it is saved.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Enable or disable Live View	<pre>options cifs.audit.liveview.enable {on off}</pre> <p>Use <code>on</code> to enable Live View or <code>off</code> to disable it.</p> <p>Note: Before enabling Live View, you must enable auditing on the storage system. Live View is disabled by default.</p>

If you want to...	Enter the command...
Clear the current ALF and EVT files	cifs audit clear The internal <code>cifsaudit.alf</code> log file and the current EVT log files in the <code>/etc/log</code> directory are cleared. However, any backup EVT files with timestamps are not affected by this command.

Saving and clearing audit events

You can specify when automatic saves occur, the maximum number of automatically-saved files, and the maximum size of the `cifsaudit.alf` file. You can also clear the `cifsaudit.alf` file.

Where Data ONTAP logs audit event information

Audit event information is stored in an internal log file, `/etc/log/cifsaudit.alf`. If you do not use Live View, you should periodically save the contents of this file to an external EVT event log file either manually or by setting up automatic saving of this file.

By default, the external event log is the `/etc/log/adtlog.evt` file. You can specify another file as the event log. If the specified file does not already exist, Data ONTAP creates the file when it saves information to the file. The directory containing the file, however, must exist; otherwise, an error message appears when you specify the file.

Size and format of the internal and external log files

You can specify the maximum size of the internal `cifsaudit.alf` log file between 524,288 bytes (512K) and 68,719,476,736 bytes (64 GB). The default size is 524,288 bytes.

The external event log (`.evt` file) that is generated from the `cifsaudit.alf` file will be larger, because the compressed contents of the `cifsaudit.alf` file are expanded and reformatted in the external event log file. The external event log is in Windows format. You can view it with Event Viewer. The `cifsaudit.alf` log file is internally formatted and cannot be viewed with Event Viewer.

Data ONTAP event log updates

Data ONTAP updates the event log under certain conditions to ensure that audit event information is saved.

To save audit event information to the external event log, you can issue the `cifs audit save` or `cifs audit clear` command, or enable automatic saving of the event information. Data ONTAP does not update the event log when the log is being viewed by a client. However, the file access information gathered when the event log is open is not lost.

It is important to issue the `cifs audit save` command frequently or enable frequent automatic saves to prevent loss of event information. If your event generation rate is very high, the `cifsaudit.alf` file fills quickly and might overwrite older events before they are saved to the event log.

Specifying the external event log location

If you prefer to store event logs in a different location, you can use the `cifs.audit.saveas` option to specify the location.

Step

1. To specify where Data ONTAP logs audit event information, enter the following command:

```
options cifs.audit.saveas filename
```

filename is the complete path name of the file to which Data ONTAP logs audit event information. You must use `.evt` as the file extension. You must use quotes around path names that contain a space.

Examples

```
options cifs.audit.saveas /etc/log/mylog.evt
```

```
options cifs.audit.saveas "/home/my event log/audit.evt"
```

Saving audit events to the event log manually

You can use the `cifs audit save` command to update the event log manually.

About this task

You do not have to manually save audit events after executing the `cifs audit clear` command; in this case, Data ONTAP saves audit events automatically.

Step

1. Enter the following command to update the event log:

```
cifs audit save [-f]
```

The `-f` option allows you to overwrite the existing event log. If the event log does not exist, you can omit the `-f` option.

Result

Data ONTAP writes to the event log the event information gathered since the last event log update.

Automatic saving of audit events

You can specify that the storage system saves audit events automatically to the event log based on a time interval or the size of the internal log file—that is, how full the `cifsaudit.alf` file is.

If you specify both a size threshold and a time interval, audit events are saved to the event log whenever the size threshold or the time interval is reached. The default for the size threshold is 75 percent. The default for the time interval is one day.

Each time the internal log file is automatically saved to the external event file, an extension is added to the base name of the event file. You can select one of the following types of extensions to be added:

- counter
- timestamp

If one of these extensions is not specified, a timestamp is used as the file extension; however, the value `timestamp` is not displayed.

The storage system saves the event files for up to six weeks. You can specify a limit to the number of event files that can be saved.

Enabling automatic saves based on internal log file size

If you have enabled automatic saves based on the size of the internal log file, you can specify the size threshold.

About this task

The default size threshold for the internal log file is 75 percent, so that whenever the internal log file is 75 percent full, the contents are automatically saved to the external event file. You can specify the threshold as a percentage of the size of the internal log file or as an absolute size.

The following table shows the units of measure and values you can use to specify the size threshold of the internal log file for automatic saves.

Units of measure	Values
% (percentage of the <code>cifsaudit.alf</code> file)	1 to 100
k (kilobytes)	1 to 67108864
m (megabytes)	1 to 65526
g (gigabytes)	1 to 64

Step

1. Perform one of the following actions:

If you want to...	Then...
Specify the size threshold at which the internal log file is automatically saved	<div>Enter the following command:</div> <div>options cifs.audit.autosave.onsize.threshold <i>Nsuffix</i></div> <div><i>N</i> is the value of the size threshold.</div> <div><i>suffix</i> is the unit of measure.</div> <div>Example:</div> <div>options cifs.audit.autosave.onsize.threshold 90%</div>
Enable or disable automatic saves based on the size of the internal log file	<div>Enter the following command:</div> <div>options cifs.audit.autosave.onsize.enable {on off}</div>

Enabling automatic saves based on a time interval

If you have enabled automatic saves based on a time interval, you can specify the time interval.

About this task

The following table shows the units of measure and values you can use to specify the time interval for automatic saves.

Units of measure	Values
s (seconds)	1 to 60
m (minutes)	1 to 60
h (hours)	1 to 24
d (days)	1 to 7

Step

1. Perform one of the following actions:

If you want to...	Then...
Specify a different time interval for automatically saving the internal log file to the external event file	<p>Enter the following command:</p> <pre>options cifs.audit.autosave.ontime.interval <i>Nsuffix</i></pre> <p><i>N</i> is the value of the time interval.</p> <p><i>suffix</i> is the unit of measure.</p> <p>Example:</p> <pre>options cifs.audit.autosave.ontime.interval 1d</pre>
Enable automatic saves based on a time interval	<p>Enter the following command:</p> <pre>options cifs.audit.autosave.ontime.enable {on off}</pre>

Specifying counter extensions

If you select “counter” for automatic file naming, the extension is a number value.

About this task

When an automatic save occurs, the old event files are renamed using sequentially numbered extensions. The newest event file does not have a number value added to it.

For example, if the base file name is `eventlog`, when an automatic save occurs, the newest event file is named `eventlog.evt`, the previous `eventlog.evt` file is copied to `eventlog1.evt`, the `eventlog1.evt` file is copied to `eventlog2.evt`, and so on.

Step

1. Enter the following command:

```
options cifs.audit.autosave.file.extension counter
```

Specifying timestamp extensions

If you select `timestamp` for automatic file naming, the file name is in a timestamp format.

About this task

The timestamp file name is constructed using the following variables:

`base_name_of_event_file.YYYYMMDDHHMMSS.evt`

Variable	Description
YYYY	The 4-digit year

Variable	Description
<i>MM</i>	The 2-digit month
<i>DD</i>	The 2-digit day
<i>HH</i>	The 2-digit hour
<i>MM</i>	The 2-digit minute
<i>SS</i>	The 2-digit second

Step

1. Enter the following command:

```
options cifs.audit.autosave.file.extension timestamp
```

Specifying the maximum number of automatically saved files

You can use the `cifs.audit.autosave.file.limit` option to specify the maximum number of event files that can be saved automatically.

Step

1. Enter the following command:

```
options cifs.audit.autosave.file.limit value
```

value is a number from 0 to 999. If you set this value to 0, there is no limit to how many event files the storage system saves automatically. If you set this value to anything other than 0, the oldest event file is always overwritten after the storage system autosave file limit is reached.

Note: If you set this value to 0, you should regularly monitor the `/etc/log` directory and clear out unneeded log files. Too many log files in this directory can cause system performance degradation.

Specifying the maximum size of the `cifsaudit.alf` file

You can use the `cifs.audit.logsize` option to specify the maximum size of the `cifsaudit.alf` file.

Step

1. Enter the following command:

```
options cifs.audit.logsize size
```

size is the number of bytes. If you enter an invalid number, a message displays the range of acceptable values.

Note: Data ONTAP overwrites the oldest data after the `cifsaudit.alf` file reaches the maximum size. To prevent loss of event data, you should save the `cifsaudit.alf` file before

it is filled. By default, when the file is 75 percent full, a warning message is issued. Additional warning messages are sent when the file is nearly full and data is about to be overwritten, and when data has already been overwritten.

SNMP traps for auditing events

Data ONTAP includes SNMP traps to provide a trigger for certain actions (such as notification) based on information about certain auditing events.

If you want CIFS clients to receive SNMP traps for auditing events, you must register the clients using the SNMP feature of Data ONTAP. Registered clients must have SNMP software that listens for SNMP traps.

An SNMP trap is issued whenever any of the following occurs:

- The specified time interval is reached and the `cifsaudit.alf` file is saved.
- The specified size threshold is reached and the `cifsaudit.alf` file is saved.
- The default size threshold, 75 percent full, is reached and the `cifsaudit.alf` file is in danger of wrapping and overwriting event data, but the file is not saved because the `cifs.audit.autosave.onsize.enable` and `cifs.audit.autosave.ontime.enable` options are turned off.
- The `cifsaudit.alf` file has wrapped and event data has been overwritten, because none of the automatic save options are turned on.

Clearing the `cifsaudit.alf` file

If you want to remove existing information and start over with an empty log file, you can use the `cifs audit clear` command to clear the internal `cifsaudit.alf` file.

Step

1. Enter the following command:

```
cifs audit clear
```

Result

If the audit has started, the internal `cifsaudit.alf` log file is cleared. If the audit has stopped, the `cifsaudit.alf` file is deleted. After you execute this command, Data ONTAP automatically saves the event log.

Viewing and understanding event detail displays

You can view real-time audit events captured with Live View, the external event log (`.evt` file) that you saved, or a backup log file created by Live View.

About this task

The following event detail displays are available:

- Network logon
- Unsuccessful network logon
- Network logoff
- Windows file access
- UNIX file access
- Unsuccessful file access
- Lost record event
- Clear audit log event

Ways to view and display audit events

You can view audit events with Microsoft Event Viewer from a Windows client, either from Administrative Tools in the Control Panel or from the Microsoft Management Console (MMC).

There are two ways to view audit events:

- In a real-time display.
When the Live View feature is enabled, the EVT event log file is automatically refreshed every minute. This provides a continuous up-to-date view in Event Viewer of the 5,000 most recent audit events.
Note: To use the Live View feature, your Windows client must be using Windows 2000 or later.
- In a static display.
You can manage the EVT event log yourself, either manually or by setting up automatic saving. In this case Event Viewer displays the most recently saved version of the log file contents, depending on how you manage the file.

Viewing real-time audit events with Live View

You can use the Windows Event Viewer to view real-time audit events captured with LiveView.

Before you begin

Before viewing real-time audit events, you must configure Live View.

Steps

1. From a Windows client, start Event Viewer from Administrative Tools in the Control Panel or from the Microsoft Management Console.
2. Select **Action > Connect to Another Computer**.
3. In the dialog box, enter the name of the storage system you want to audit and click **OK**.
4. Select the **Security** entry on the left side of the application.

The right side of the application is populated with the latest audit events captured on the storage system (up to 5,000 events).

Viewing static event log files

You can use a Windows client to view the external event log (.evt file) that you saved, or to view a backup log file created by Live View.

Steps

1. From a Windows client, start Event Viewer from Administrative Tools in the Control Panel or from the Microsoft Management Console.

2. Select **Log > Open**.

Note: Do not try to open the event log by selecting **Log > Select Computer** and double-clicking the storage system name. If you do, the Event Viewer displays the error message “The RPC server is unavailable,” because Data ONTAP does not communicate with the Event Viewer with RPC calls unless Live View is enabled.

3. Choose the event log on the storage system.

Windows file access detail displays

Windows file access detail displays show many types of information.

The following table describes the fields of Windows file access detail displays.

Field	Description
Object Server	The name of the subsystem server process calling the audit check function. This is always SECURITY because this is a security log.
Object Type	The type of object being accessed.
Object Name	The name (such as a file name) of the object being accessed.
New Handle ID	The new handle identifier of the open object.
Operation ID	A unique identifier associating multiple events resulting from a single operation.
Process ID	The identifier of the client process accessing the object.
Primary User Name	The user name of the user requesting the object access. When impersonation is taking place, this is the user name with which the server process is logged on.

Field	Description
Primary Domain	The name of the computer, or SYSTEM if the user identified by Primary User Name is SYSTEM. If the computer is a member of a Windows NT Server domain, this can also be the name of the domain containing the primary user's account.
Primary Logon ID	A unique identifier assigned when the primary user logged on.
Client User Name	Your login name.
Client Domain	The name of your computer or the domain containing the client user's account.
Client Logon ID	A unique identifier assigned when the client user logged on.
Accesses	The types of accesses to the object that were attempted.
Privileges	Your privileges.

UNIX file access detail displays

UNIX file access detail displays show the same kind of information as the Windows file access detail displays, but NFS access appears instead of an object name, because the file is accessed through NFS.

In addition, UNIX file access detail displays show the following information about the file that you are auditing:

- The ID of the volume in which the file is located
- The ID of the latest Snapshot copy in which the file is located
- The inode of the file

This information enables you to find the file using the `find -inum` command from an NFS client.

Unsuccessful file access and lost record event detail displays

Unsuccessful file access detail displays show failed attempts to access a file. Furthermore, if Data ONTAP cannot create an audit record, the lost record event detail displays give a reason

For example, an unsuccessful file access occurs when a user tries to access a file but does not have permission to access it. The display shows the ID of the user who tried to access the file and an indication that the access attempt was unsuccessful.

If Data ONTAP cannot create an audit record, the lost record event detail displays give a reason, such as the following:

```
Internal resources allocated for the queueing of audit messages have
been exhausted, leading to the loss of some audits.
Number of audit records discarded:          1
```

Controlling CIFS access to symbolic links

A symbolic link is a special file created by NFS clients that points to another file or directory. A symbolic link is, in some respects, similar to a “shortcut” in the Windows environment.

About this task

There are two kinds of symbolic links:

- Absolute symbolic links begin with a slash (/) and are treated as a path derived from the root of the file system.
- Relative symbolic links begin with a character other than a slash (/) and are treated as a path relative to the parent directory of the symbolic link.

CIFS clients cannot create symbolic links, but they can follow the symbolic links created by NFS clients.

There are special requirements to enable CIFS access to the following types of symbolic links:

- Absolute symbolic links.
Since the destination of an absolute symbolic link depends on the type of UNIX mount, CIFS clients need additional information to interpret absolute symbolic links.
- Relative symbolic links to destinations on the same storage system outside the share in which the relative symbolic link is located.
By default, Data ONTAP does not allow a CIFS client to follow a symbolic link outside the share to which the CIFS client is authenticated.

Enabling CIFS clients to follow symbolic links

You can use the `cifs.symlinks.enable` option to enable CIFS access to symbolic links after they have been disabled.

About this task

The `cifs.symlinks.enable` option is enabled by default.

Step

1. Enter the following command to enable CIFS access to symbolic links:

```
options cifs.symlinks.enable on
```

Result

CIFS clients will directly follow relative symbolic links to destinations in the same share

Specifying how CIFS clients interact with symbolic links

You can specify how CIFS clients interact with symbolic links by creating Map entries in the `/etc/symlink.translations` file (absolute symbolic links only), creating Widelink entries in the `/etc/symlink.translations` file (absolute symbolic links only), or disabling NT share boundary checking for symbolic links (relative and absolute symbolic links).

About this task

Use the following table to help determine which options you want to implement. The table shows for each option the types of destinations that symbolic links will be able to point to.

Symbolic link destination can be...	Map entries	Widelink entries	No share boundary check
The same share on the same storage system	X	X	X
Another share on the same storage system		X	X
A non-shared area of the same storage system			X
A share on another storage system		X	
A share on another CIFS server or a desktop PC		X	

Why you should avoid symbolic links to files

You should prevent CIFS clients from following symbolic links that point to files because Data ONTAP can update the wrong files.

The wrong files might be updated because many CIFS client applications perform operations such as writing to a temporary file, renaming the original file to a backup name, then renaming the temporary file to the original name.

When client applications perform these operations, if the original file was targeted directly by a symbolic link, that file would be stored in the directory where the symbolic link was, and the renamed symbolic link would point to the original file rather than to the updated file.

Note: CIFS clients following symbolic links to directories, rather than to individual files, do not experience this problem.

About Map entries

Map entries are used to redirect absolute symbolic links on the storage system. You create Map entries in the `/etc/symlink.translations` file. Map entries allow CIFS clients to follow absolute symbolic links to target destinations within the same share.

Note: CIFS client users who follow symlinks to resources outside the link's share do not work, unless the `cifs share -nosymlink_strict_security` option has been specified for the source share.

Map entries have the following requirements:

- To resolve an absolute symbolic link, there must be a Map entry in the `/etc/symlink.translations` file that determines the destination of the link.
- The symbolic link destination must be in the same share as the link itself, or the link must be in a share for which the `-nosymlink_strict_security` option has been specified.

When you use Map entries to redirect absolute symbolic links, Windows share security is preserved for both the symbolic link and the destination, because they are in the same share. If you have both Map entries and Widelink entries in the `symlink.translations` file, the storage system uses the first matching entry it finds.

About widelink entries

Widelink entries are another way to redirect absolute symbolic links on your storage system. You create widelink entries in the `/etc/symlink.translations` file.

Widelink entries allow CIFS clients to follow absolute symbolic links to target destinations either on the same storage system or outside the storage system. They are enabled on a per-share basis.

Widelink entries have the following requirements:

- The share in which the absolute symbolic links are located must be enabled for widelinks.
- In order to resolve an absolute symbolic link, there must be a widelink entry in the `/etc/symlink.translations` file that determines the destination of the link.
- The destination of the widelink entry must be one of the following:
 - The same share as the symbolic link
 - Another share on the same storage system
 - A share on another storage system
 - A share on another CIFS server or desktop PC
- The CIFS client must have client-side support for Microsoft Distributed File System (DFS). Windows NT and later clients support DFS by default.

To follow widelink entries, the CIFS client automatically requests and receives a DFS referral from the storage system to establish an authenticated connection with the target share. This preserves NT share security for both the symbolic link and the destination. After the connection is established, the CIFS client can make new requests directly to the target share or server, thereby increasing performance.

If you have both map entries and widelink entries in the `/etc/symlink.translations` file, the storage system uses the first matching entry it finds.

Widelink entries have the following limitations:

- Even if the destination of the widelink is a file, it appears as a directory in directory listings. The system API for opening the file will correctly follow the widelink, but this might confuse certain applications. To avoid this problem, you should create a widelink that resolves to a directory, rather than a file.
- Windows 95, Windows 98, and Windows ME clients cannot follow a widelink to another widelink.
- Windows NT clients cannot display or modify ACLs in a share enabled for widelinks. This restriction does not apply to Windows 2000 and later clients.
- Widelinks cannot direct a client to a non-shared area on the destination machine.

About disabling share boundary checking for symbolic links

When you disable share boundary checking for symbolic links, CIFS clients can follow symbolic links anywhere on the storage system. This behavior is set on a per-share basis and affects both relative and absolute symbolic links.

Disabling share boundary checking for symbolic links has the following requirements:

- The share in which the symbolic links are located must be set to `nosymlink_strict_security`.
- In order to resolve an absolute symbolic link, there must be a Map entry in the `/etc/symlink.translations` file that determines the destination of the link.
- The destinations for relative symbolic links and for mapped absolute symbolic links might be in any shared or non-shared area of the storage system.

Disabling share boundary checking for symbolic links has the following limitations:

- Relative symbolic links cannot be used to span volumes; you must use absolute symbolic links.
- Symbolic links cannot be followed off the storage system to other systems.
- NT share security is preserved for the symbolic link itself because the CIFS client has to authenticate to connect to the share in which the symbolic link is located.
- NT share security is preserved for the destination of the symbolic link only if the destination is in the same share.
- NT share security is not preserved for the destination of the symbolic link if the destination is outside the share, because the CIFS client does not have to authenticate to the destination (which might or might not be a CIFS share).

Note: If you disable share boundary checking for symbolic links, be sure to secure any areas of the storage system that you do not want users to access. This is necessary because a user can create a symbolic link to any path on the storage system.

Redirecting absolute symbolic links

You can redirect absolute symbolic links on the storage system by creating Map entries in the `/etc/symlink.translations` file or creating Widelink entries in the `/etc/symlink.translations` file.

About this task

NFS clients interpret the file system location represented by an absolute symbolic link based on how the file systems are mounted on the client. CIFS clients do not have access to NFS clients' mount information.

To allow CIFS clients to follow absolute symbolic links on the storage system, you must redirect the absolute symbolic link so that CIFS clients can interpret the file system location represented by the absolute symbolic link. You can redirect absolute symbolic links by creating entries in the `/etc/symlink.translations` file. The `/etc/symlink.translations` file performs the same role on the storage system as automounter tables on UNIX servers

Creating Map entries

You can create Map entries by editing the `/etc/symlink.translations` file.

Steps

1. Open the `/etc/symlink.translations` file for editing.
2. Enter one or more lines in the file using the following format:

Map *template result*

template is used to match absolute symbolic links.

result specifies a storage system path that is substituted for the matching absolute symbolic link.

Note: To specify a space or pound (#) character in a file path, you must prepend a backslash (\) escape character.

Examples

```
Map /u/users/charlie/* /home/charlie/*
```

```
Map /templ/* /vol/vol2/util/t/*
```

```
Map /u/users/bob\ smith/* /home/bob\ smith/*
```

Creating Widelink entries

You can create Widelink entries by editing the `/etc/symlink.translations` file.

Steps

1. Open the `/etc/symlink.translations` file for editing.
2. Enter one or more lines in the file using the following format:

Widelink *template* [*qtree*] *result*

template specifies the UNIX path name.

result specifies the CIFS UNC path name.

qtree allows multiple entries in different qtrees to have the same template value.

Note: Unlike in a Map entry, you can specify a space and pound (#) character in a file path without prepending a backslash (\) escape character. In a Widelink entry, a backslash character is a standard file path character in accordance with the Universal Naming Convention.

Examples

In the following examples, the result uses CIFS path name syntax, with backslashes as separators, and allows an embedded space. The wildcard character (*) in the template path name represents zero or more characters, including the backslash character (\). In the result path name, the wildcard character represents text from the corresponding match in the template path name:

```
Widelink /eng/proj/* @/vol/vol2 \\filer\hw\proj\*
```

```
Widelink /eng/proj/* \\filer\sw\proj\*
```

How the storage system uses Map and Widelink entries

To allow CIFS clients to follow absolute symbolic links, the storage system searches the entries in the `/etc/symlink.translations` file in sequential order until a matching entry is found or the lookup fails.

The storage system uses the first matching entry it finds to generate a path to the destination. Therefore, it is important to put the most restrictive entries first to prevent premature mapping errors.

This example shows how to list Map entries. `/u/home/*` is more specific than `/u/*`:

```
Map /u/home/* /vol/vol2/home/*
```

```
Map /u/* /vol/vol0/*
```

This example shows how to list Widelink entries:

```
Widelink /u/docs/* \\filer\engr\tech pubs\*
```



```
Widelink /u/* \\filer\engr\*
```

Optimization of NFS directory access for CIFS clients

You can optimize CIFS client access to an NFS directory by configuring Data ONTAP to convert non-Unicode directories to Unicode format when either CIFS clients or NFS clients access directories. You can also configure Data ONTAP to create only Unicode-formatted directories which eliminates the need for Unicode conversion.

By default, Data ONTAP configures directories created by NFS clients in non-Unicode format and directories created by CIFS clients in Unicode format. As a result, CIFS directories are directly accessible to NFS clients, but NFS directories are not directly accessible to CIFS clients.

Before a CIFS client can access an NFS directory, Data ONTAP must convert the NFS directory to CIFS format, which includes assigning 8.3-format names to each entry and converting the file name to Unicode representation. This conversion is triggered the first time a CIFS client attempts to access an NFS directory that was created by an NFS client and previously has never been accessed from a CIFS client. The conversion is done automatically (“on the fly”), as the storage system receives the access request.

This automatic Unicode conversion is slow and CPU-intensive and can have a significant performance impact on the storage system and clients accessing the storage system. It runs at a high priority and might make the storage system appear to be unresponsive. Running the `sysstat` command displays high CPU usage. The conversion process also delays processing of WAFL's consistency points, until the conversion is complete. If a consistency point is delayed more than 10 minutes, Data ONTAP displays a message on the system console. Depending on the number of subdirectories and files the directory contains, Unicode conversion can take hours or even days.

If you intend to share files between CIFS and NFS clients, you should instead configure Data ONTAP to create directories in Unicode format immediately after installing Data ONTAP. This step ensures that all new directories are created in Unicode format, and it eliminates the need for conversion later upon first access by a CIFS client.

Creating new directories in Unicode format

You can use the `vol options` command to configure Data ONTAP to create all new directories in Unicode format. This prevents the need for Unicode conversion later when a CIFS client first accesses the directory.

Step

1. Enter the following command:

```
vol options volume_name create_unicode on
```

This operation might fail if a nondisruptive volume movement is being performed on the target volume.

Converting existing directories to Unicode format

When a CIFS client first accesses a directory, it triggers the Unicode conversion. If your storage environment already contains existing NFS directories with large numbers of subdirectories and files, the automatic on-first-access conversion can have a performance impact on the storage system. There are methods to prevent this performance impact.

About this task

You can proactively convert existing NFS directories to Unicode format manually under controlled conditions. You can also configure Data ONTAP to trigger Unicode conversion upon access from both NFS and CIFS clients to distribute the load.

Step

1. Perform one or more of the following actions:

If you want to...	Then...
Trigger Unicode conversion manually	<p>From a CIFS client, use Windows Search or equivalent function to scan all CIFS accessible volumes on the storage system.</p> <p>This triggers Unicode conversion on NFS directories that previously have not been accessed by CIFS clients.</p> <p>Note: Perform this step only during non-business hours or a maintenance window. The action triggers Unicode conversion on multiple directories, leading to a significant performance impact on the storage system while conversion is taking place. Depending on the number of subdirectories and files the directories contain, Unicode conversion can take hours or even days, so plan accordingly.</p>
Reduce the time required for Unicode conversion of a directory	<p>a. If you have a directory that contains more than 50,000 files, create a new CIFS directory from a Windows client on the same volume.</p> <p>b. In the same qtree as the directory you want to convert, use the NFS <code>mv</code> command to move the files into the directory you just created.</p> <p>c. Optionally, remove the old directory and assign its name to the new directory.</p>
Trigger Unicode conversion when directories are accessed by either CIFS or NFS clients	<p>Enter the following command:</p> <pre>vol options volume_name convert_unicode on</pre> <p>Note: Do not enable the <code>convert_unicode</code> option when you have directories that contain more than 50,000 files. This operation might fail if a nondisruptive volume movement is being performed on the target volume.</p>

If you want to...	Then...
Move non-Unicode volume contents to a Unicode volume	<p>a. Verify that on the source volume the <code>convert_ucose</code> option is set to <code>off</code>.</p> <p>b. Create a new volume.</p> <p>c. Configure Unicode on the new volume by entering the following commands:</p> <pre>vol options volume_name create_ucose on vol options volume_name convert_ucose on</pre> <p>d. Use <code>ndmcopy</code> or <code>qtree SnapMirror</code> to move the data from the source volume to the destination volume.</p> <p>For more information about <code>ndmcopy</code>, see the <i>Data ONTAP Data Protection Tape Backup and Recovery Guide for 7-Mode</i>.</p> <p>For more information about <code>qtree SnapMirror</code>, see the <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>

Preventing CIFS clients from creating uppercase file names

You can set the `cifs.save_case` option to `off` to prevent CIFS clients from creating uppercase filenames.

About this task

Older, 16-bit CIFS clients that open and save files change the file name by changing the lowercase or mixed-case characters to all uppercase characters. You can prevent these uppercase file names by forcing Data ONTAP to store CIFS file names using lowercase characters.

Step

1. Enter the following command:

```
options cifs.save_case off
```

Accessing CIFS files from NFS clients

Data ONTAP uses Windows NT File System (NTFS) security semantics to determine whether a UNIX user, on an NFS client, has access to a file in a mixed or NTFS `qtree`.

About this task

Data ONTAP does this by converting the user's UNIX User ID (UID) into a CIFS credential, then using the CIFS credential to verify that the user has access rights to the file. A CIFS credential consists of a primary Security Identifier (SID), usually the user's Windows user name, and one or more group SIDs that correspond to Windows groups of which the user is a member.

The time Data ONTAP takes converting the UNIX UID into a CIFS credential can be from tens of milliseconds to hundreds of milliseconds because the process involves contacting a domain controller. Data ONTAP maps the UID to the CIFS credential and enters the mapping in a credential cache to reduce the verification time caused by the conversion.

Adding mapping entries to the WAFL credential cache

You can add mapping entries to the WAFL credential cache at any time. Normally, this is not necessary because entries are created automatically as the storage system is accessed.

Before you begin

You must have the names and IP addresses of the entries you want to add to the WAFL credential cache.

About this task

The best way to add entries is in a script that loads the WAFL credential cache with entries at boot time. This immediately puts the entries in the WAFL credential cache rather than waiting for Data ONTAP to create the entries in the course of accessing the files.

Note: The cache is limited to 10,000 entries. If you exceed this limit, the older entries are deleted.

Step

1. Enter the following command:

```
wcc -a -u uname -i ipaddress
```

uname specifies the UNIX name of a user.

ipaddress specifies the IP address of the host that the user is on.

Deleting mapping entries from the WAFL credential cache

You can delete entries from the WAFL credential cache at any time. You might want to delete entries after making security changes, to ensure they take effect immediately.

Before you begin

You must have the name for the entry you want to delete from the WAFL credential cache. To further narrow down the selection, you can optionally specify an IP address.

About this task

Security changes might not take effect immediately when you change a user's rights. For example if you remove a user from a group and a mapping for that user already exists in the WAFL credential cache, the user will continue to have that group's access to files until the entry in the WAFL credential cache times out automatically. The default credential cache timeout period is 20 minutes.

Step

1. Enter the following command:

```
wcc -x name
```

name is one of the following specifications: `-s` followed by the Windows user name or group name found in the CIFS credential or `-u` followed by the UNIX name found in the CIFS credential.

Note: If *name* is the name of a group, this procedure deletes all members of that group from the WAFL credential cache.

You can further narrow the specification of a user by adding `-i`, followed by the IP address of the host that the user is on. If you do not specify *name*, all entries are deleted.

Example

```
wcc -x -u jdoe -i 10.100.4.41
```

Setting how long mapping entries are valid

Increasing the time that the CIFS credential remains in the WAFL credential cache after Data ONTAP updates it improves performance. Performance is improved because Data ONTAP doesn't have to take the time to create a CIFS credential to verify access to a file.

About this task

The disadvantage of increasing the time that CIFS credentials remain in the WAFL credential cache is that if you change a user's access rights, the change does not take effect until Data ONTAP updates the WAFL credential cache. In this case, the user might temporarily retain rights to a file to which you have just denied access.

If you do not expect problems of this type, you can increase the time that the credential entry is valid. If you need to see access right updates as they occur and slower performance is not an issue, you can use a smaller value than the default.

Step

1. Enter the following command:

```
options wafl.wcc_minutes_valid n
```

n is the number of minutes you want each entry to be valid. It can range from 1 through 20,160. The default value is 20.

Monitoring WAFL credential cache statistics

By monitoring WAFL credential cache statistics, you can view what entries are currently cached and the UNIX UID-to-CIFS credential mapping. This information is useful when you need to know what entries are in the WAFL credential cache or what the access rights are for users listed in the entries.

Step

1. Enter the following command:

```
wcc -d uname
```

uname is the UNIX name of a user. Omit *uname* to list all credential entries in the WAFL credential cache. You can get more detailed information by appending *-v* to the command line. You can have up to three instances of the *-v* option (*-vvv*) per command; each instance represents an increasing level of detail.

Example

The following sample shows the output of statistics with the *-d* option:

```
wcc -d
```

```
tday (UID 10350) from 10.121.4.41 => NT-DOMAIN\tday*
Total WCC entries: 3; oldest is 127 sec.
Total Administrator-privileged entries: 1
* indicates members of "BUILTIN\Administrators" group
```

The following sample shows the output of statistics with the *-v* option used twice:

```
wcc -dvv
```

```
jdoe (UID 1321) from 10.121.4.41 => NT-DOMAIN\jdoe
*****
UNIX uid = 1321
NT membership
    NT-DOMAIN\jdoe
    NT-DOMAIN\Domain Users
NT-DOMAIN\SU Users
NT-DOMAIN\Installers
NT-DOMAIN\tglob
    NT-DOMAIN\Engineering
BUILTIN\Users
User is also a member of Everyone, Network Users,
    Authenticated Users
*****
tday (UID 10350) from 10.121.4.41 => NT-DOMAIN\tday*
*****
    UNIX uid = 10350
NT membership
```

```

NT-DOMAIN\tday
NT-DOMAIN\Domain Users
NT-DOMAIN\Domain Admins
NT-DOMAIN\SU Users
NT-DOMAIN\Installers
BUILTIN\Users
                BUILTIN\Administrators
User is also a member of Everyone, Network Users,
Authenticated Users
*****
bday (UID 1219) from 10.121.4.41 => NT-DOMAIN\bday
                *****
                UNIX uid = 1219
NT membership
NT-DOMAIN\bday
NT-DOMAIN\Domain Users
NT-DOMAIN\Installers
NT-DOMAIN\SU Users
BUILTIN\Users
User is also a member of Everyone, Network Users,
Authenticated Users
*****
Total WCC entries: 3; oldest is 156 sec.
Total Administrator-privileged entries: 1
* indicates members of "BUILTIN\Administrators" group

```

Managing mapping inconsistencies

You can manage mapping inconsistencies by performing several tasks.

About this task

If a user cannot access a file that should be accessible, there are several possible reasons:

- You granted access recently and the WAFL credential cache does not have the new mapping entry.
You can determine mapping inconsistencies between recently granted rights and the WAFL credential cache by comparing CIFS credential mappings. You can display mapping results for the user's UNIX name or user's Windows name.
- The NFS client could not obtain CIFS credentials.
You can determine whether an NFS client can perform a CIFS login to the storage system by tracing CIFS logins.
- Depending on the NFS client, it might be necessary to wait for the NFS attribute cache to time out before changes to the CIFS credential take effect.

Steps

1. Display the current CIFS credential mapping of a UNIX name by entering the following command:

```
wcc -s uname
```

uname is the Windows user name. You can further narrow the specification of the user by adding *-i*, followed by the IP address of the host that the user is on. You can get more detailed information by appending *-v* to the command line. You can have up to three instances of the *-v* option (*-vvv*) per command; each instance represents an increasing level of detail.

2. Note the CIFS credential information.
3. To display information about all connected users, enter the following command:

```
cifs sessions -s
```

4. Locate the user's information in the output.
5. Compare the two CIFS credential mappings.
6. If the CIFS credential mappings are different, disconnect the client by entering the following command:

```
cifs terminate workstation
```

Result

When the client reconnects, the CIFS credential mappings will be correct.

Tracing CIFS logins

You can trace CIFS logins by monitoring any attempt by an NFS client to obtain a CIFS credential.

About this task

Use CIFS login tracing carefully because it reports every CIFS login. Persistent use can result in excessive console and log messages, which can affect system performance. By default, the `cifs.trace_login` option is disabled. The option should only be enabled temporarily for diagnostic purposes. Keep it disabled all other times.

Step

1. Enter the following command:

```
options cifs.trace_login {on | off}
```

Use `on` to enable or `off` to disable CIFS login tracing.

Tracing domain controller connections

You can configure Data ONTAP to send messages to the console when it tries to improve the domain controller connection every few minutes.

About this task

Because tracing functions send frequent messages to the console and system log, do not persistently enable this option. By default, this feature is disabled.

Step

1. Enter the following command:

```
options cifs.trace_dc_connection {on | off}
```

Allowing CIFS clients without UNIX "execute" permissions to run .dll and .exe files

You can set the `cifs.grant_implicit_exe_perm` option to `on` to allow CIFS clients to run `.dll` and `.exe` files even when the UNIX executable bit is not set.

Step

1. Enter the following command:

```
options cifs.grant_implicit_exe_perm on
```

Result

Executables with only "read" UNIX permissions are implicitly granted execute permissions when run from a CIFS client.

How file editing with Windows applications affects UNIX permissions

Many Windows applications incorrectly interpret the ACLs when reading files that have UNIX security. When the application saves the file, the original UNIX permissions are lost. Using the `cifs.preserve_unix_security` option avoids this problem.

You should set the `cifs.preserve_unix_security` option to `on` if you serve files under the following conditions:

- The files have UNIX permissions (that is, mode bits are set using the `chmod` or `umask` commands).
- NFS v4 Access Control Lists (ACLs) are not applied to the files.
- The files are in a `qtree` with UNIX or mixed security.
- The files are edited using Windows applications.

Note: When this option is enabled, a UNIX-style `qtree` appears as an NTFS volume instead of a FAT volume when viewed from a Windows client.

When the `cifs.preserve_unix_security` option is set, you can view and edit UNIX permissions using the Security tab in the Windows Properties dialog box. However, you cannot modify permissions from a Windows client if the operation is not permitted by the UNIX system. For example, you cannot change the ownership of a file you do not own, because the UNIX system

does not permit this operation. This restriction prevents Windows clients from bypassing UNIX permissions set on the storage system.

For more information about the `cifs.preserve_unix_security` option, see the `options(1)` man page.

File access using FTP

You can enable and configure the File Transfer Protocol (FTP) server to let users of Windows and UNIX FTP clients access the files on your storage system.

Managing FTP

You can manage FTP by enabling or disabling it, configuring it, and viewing statistics related to it.

Enabling or disabling the FTP server

You can enable or disable the FTP server by modifying the `ftpd.enable` option. This allows clients to access files using FTP. By default, this option is `off`.

Step

1. Perform one of the following actions:

If you want the FTP server to be...	Enter the command...
Enabled	<code>options ftpd.enable on</code> The FTP server begins listening for FTP requests on standard FTP port 21.
Disabled	<code>options ftpd.enable off</code>

Enabling or disabling the TFTP server

You can enable or disable the TFTP server by modifying the `tftpd.enable` option. This allows clients to access files using TFTP. By default, this option is `off`.

Step

1. Perform one of the following actions:

If you want the TFTP server to be...	Enter the command...
Enabled	<code>options tftpd.enable on</code> The TFTP server begins listening for TFTP requests on standard FTP port 69.
Disabled	<code>options tftpd.enable off</code>

Enabling or disabling FTP file locking

To prevent users from modifying files while the FTP server is transferring them, you can enable FTP file locking. Otherwise, you can disable FTP file locking. By default, FTP file locking is disabled.

Step

1. Perform one of the following actions.

If you want FTP file locking to be...	Enter the command...
Enabled for deleting and renaming	<code>options ftpd.locking delete</code>
Enabled for deleting, renaming, and writing	<code>options ftpd.locking write</code>
Disabled	<code>options ftpd.locking none</code>

Specifying the FTP authentication style

To configure the FTP server to use UNIX, Windows, or both authentication styles, you can set the `ftpd.auth_style` option to `unix`, `ntlm`, or `mixed`, respectively. By default, this option is `mixed`.

About this task

When you specify the UNIX authentication style, the FTP server authenticates users using the `/etc/passwd` file, NIS, or LDAP.

When you specify the NTLM authentication style, the FTP server authenticates users using the Windows domain controller. The NTLM authentication style is more secure than the UNIX authentication style because it uses encrypted user names and passwords.

When you specify the mixed authentication style, the FTP server uses the NTLM authentication style for users with names containing a backslash (`\`) or `"@"` character; it uses the UNIX authentication style for all other users.

Steps

1. Enter the following command:

```
options ftpd.auth_style style
```

style is `unix`, `ntlm`, or `mixed`.

2. Perform one of the following actions:

If you specified...	Then...
ntlm	<p>Specify the CIFS home directory in the <code>/etc/cifs_homedir.cfg</code> file and then enter the following command:</p> <pre>cifs homedir load</pre> <p>The home directory of a user is a combination of the path you specify in <code>/etc/cifs_homedir.cfg</code> and the user ID of the user. The path you specify in <code>/etc/cifs_homedir.cfg</code> is case-sensitive; however, the user ID is not case-sensitive. For example, if the path is <code>\home</code> and the user name is <code>JOHN</code>, the home directory for the user is <code>\home\john</code>.</p>
unix and previously enabled NIS by setting the <code>nis.enable</code> option to on	<p>Add an appropriate <code>passwd</code> entry to the <code>/etc/nsswitch</code> file.</p> <ul style="list-style-type: none"> To authenticate users using the <code>/etc/passwd</code> file only, add the following entry: <pre>passwd: files</pre> To authenticate users using NIS only, add the following entry: <pre>passwd: nis</pre> To authenticate users using both the <code>/etc/passwd</code> file and NIS, add the following entry: <pre>passwd: files nis</pre>

Limitations of the NTLM authentication style

The NTLM authentication style has some limitations.

These limitations include the following:

- NTLMv2 relies on domain controller-based services that do not exist on the storage system. For this reason, only NTLMv1 and earlier can be used to connect to storage systems operating in workgroup mode.
- Workgroup storage system Windows clients that use NTLM authentication should have “LAN Manager authentication level” set to a level other than “NTLMv2 Only.”
Setting this option changes the registry value for “LMCompatibilityLevel” to 0, 1, or 2. These are the only NTLM settings supported by the storage system for workgroup environments.
- Although domain-based clients in an Active Directory environment can perform authentication using NTLMv2 (because requests are passed along from the storage system to the domain controller), no connection information for local storage system accounts is available to the domain controller.
For this reason, local storage system accounts would fail authentication during attempts to connect to a storage system in such an environment.

Enabling or disabling the bypassing of FTP traverse checking

You can enable or disable the bypassing of FTP traverse checking by setting the `ftpd.bypass_traverse_checking` option to `on` or `off`, respectively. By default, this option is set to `off`.

About this task

If the `ftpd.bypass_traverse_checking` option is set to `off`, when a user attempts to access a file using FTP, Data ONTAP checks the traverse (execute) permission for all directories in the path to the file. If any of the intermediate directories does not have the "X" (traverse permission), Data ONTAP denies access to the file. If the `ftpd.bypass_traverse_checking` option is set to `on`, when a user attempts to access a file, Data ONTAP does not check the traverse permission for the intermediate directories when determining whether to grant or deny access to the file.

Step

1. Perform one of the following actions.

If you want the bypassing of FTP traverse checking to be...	Then...
Enabled	Enter the following command: options ftpd.bypass_traverse_checking on
Disabled	Enter the following command: options ftpd.bypass_traverse_checking off

Restricting FTP access

You can restrict FTP access by blocking FTP users and restricting FTP users to a specific directory (either their home directories or a default directory).

Blocking specific FTP users

To prevent specific FTP users from accessing the storage system, you can add them to the `/etc/ftpusers` file.

Steps

1. Access the `/etc` directory on the storage system's default volume (`/vol1/vol0` by default) from an NFS or CIFS client.
2. Open the `/etc/ftpusers` file in a text editor. (If the file does not exist, create it.)
3. Add the user names of the users (one name per line) to whom you want to deny access.
For NTLM authentication, you must specify user names using one of the following formats:
 - `Domain\username`

- `Username@domain`

Note: In the preceding formats, you must specify the exact name of the domain; otherwise, the FTP server will not deny access to the user. For example, if the name of a domain includes a ".com" suffix, you must include that suffix.

4. Save the `/etc/ftputers` file.

Restricting FTP users to a specific directory

To restrict FTP users to a specific directory, you can set the `ftpd.dir.restriction` option to `on`; otherwise, to let FTP users access the entire storage system, you can set the `ftpd.dir.restriction` option to `off`. By default, this option is `on`.

Step

1. Perform one of the following actions.

If you want to...	Then...
Restrict FTP users to their home directories or a default directory	Enter the following command: options ftpd.dir.restriction on
Let FTP users access the entire storage system	Enter the following command: options ftpd.dir.restriction off

After you finish

If you set the `ftpd.dir.restriction` option to `on`, you can use the `ftpd.dir.override` option to specify whether FTP users can access their home directories or a default directory.

Restricting FTP users to their home directories or a default directory

To restrict FTP users to a default directory, you can set the `ftpd.dir.override` option. Otherwise, to restrict FTP users to their home directories, you can clear the `ftpd.dir.override` option. By default, this option is cleared.

Step

1. Perform one of the following actions:

If you want to restrict FTP users to...	Then...
Their home directories	Enter the following command: options ftpd.dir.override ""

If you want to restrict FTP users to...	Then...
A default directory	Enter the following command: options ftpd.dir.override <i>directory</i> <i>directory</i> is the name of the default directory to which you want to restrict FTP users.

After you finish

Make sure the FTP users have read access to the directory you created in Step 1. For more information, see the *Data ONTAP Storage Management Guide*.

Managing FTP log files

You can manage FTP log files by viewing FTP log files, specifying the maximum number of FTP log files, and specifying the maximum size of the current FTP log files.

How the FTP server manages its log files

Data ONTAP logs all FTP server requests in the `/etc/log/ftp.cmd` file and all file transfers in the `/etc/log/ftp.xfer` file.

Data ONTAP writes data to an FTP log file (either the `/etc/log/ftp.cmd` or `/etc/log/ftp.xfer` file) until the FTP log file reaches its maximum size. Then Data ONTAP performs the following tasks:

1. If the total number of FTP log files is equal to the maximum number of FTP log files, it deletes the oldest FTP log file.
2. It increments the suffixes of the old FTP log files.
3. It adds the `.1` suffix to the current FTP log file, thereby making it an old FTP log file.
4. It creates a new FTP log file.

Example

Assuming that the maximum number of log files is 6, when the `/etc/log/ftp.xfer` log file reaches its maximum size, Data ONTAP performs the following tasks:

1. It deletes the `/etc/log/ftp.xfer.5` file, if the file exists.
2. It renames `/etc/log/ftp.xfer.4` to `/etc/log/ftp.xfer.5`, `/etc/log/ftp.xfer.3` to `/etc/log/ftp.xfer.4`, and so on.
3. It renames `/etc/log/ftp.xfer` to `/etc/log/ftp.xfer.1`.
4. It creates a new `/etc/log/ftp.xfer` log file.

The `/etc/log/ftp.xfer` log file format

The `/etc/log/ftp.xfer` file contains information on all files that the FTP server transfers.

The following table describes the fields in the `/etc/log/ftp.xfer` file.

Field	Description
timestamp	Timestamp of the log record
xferTime	Duration, in seconds, of the file transfer
clientIP	IP address of the FTP client
xferCount	Byte count of transferred file
filename	File name of the transferred file
xferType	Can be “a” (ascii), “e” (ebcdic), or “b” (binary)
xferDirection	Can be “o” (outbound) or “i” (inbound)
accessType	Can be “a” (anonymous), “r” (real user), or “g” (guest)

The `/etc/log/ftp.cmd` log file format

The `/etc/log/ftp.cmd` file contains information on all commands that the FTP server receives.

The following table describes the fields in the `/etc/log/ftp.cmd` file.

Field	Description
timestamp	Timestamp of the log record
serialNo	Serial number of the FTP connection
command	FTP command

Viewing an FTP log file

To view an FTP log file, you can open it in a text editor or viewer.

About this task

The FTP server maintains two log files:

- The `/etc/log/ftp.cmd` file contains information on all commands that the FTP server receives.
- The `/etc/log/ftp.xfer` file contains information on all files that the FTP server transfers.

Steps

1. Access the `/etc/log` directory on the storage system's default volume (`/vol/vol0` by default) from an NFS or CIFS client.
2. Open the log file in a text editor or viewer.

Specifying the maximum number of FTP log files

You can set the `ftpd.log.nfiles` option to specify the maximum number of FTP log files. By default, the maximum number of FTP log files is 6.

Step

1. Enter the following command:

```
options ftpd.log.nfiles n
```

n is the maximum number of log files. For more information, see the `na_options(1)` man page.

Specifying the maximum size of the current FTP log files

You can set the `ftpd.log.filesize` option to specify the maximum size of the current FTP log files (the `/etc/log/ftp.cmd` and `/etc/log/ftp.xfer` log files). By default, the maximum size of the current FTP log files is 512 KB.

Step

1. Enter the following command:

```
options ftpd.log.filesize filesize
```

filesize is an integer followed by `K` or `k` (for KB) or `G` or `g` (for GB). For more information, see the `na_options(1)` man page.

Example

The following command sets the maximum size of the current FTP log files to 1 GB:

```
options ftpd.log.filesize 1G
```

Viewing SNMP traps that the FTP server generates

To view SNMP traps that the FTP server generates, you can start and configure SNMP on the storage system and view the SNMP traps on a UNIX client.

SNMP traps that the FTP server generates

The FTP server generates several SNMP traps.

The FTP server generates SNMP traps when the following events occur:

- Concurrent connections reach the `ftpd.max_connections_threshold` value.
- Concurrent connections reach the `ftpd.max_connections` value.
- The FTP daemon process stops due to an error.

For more information about SNMP, see the *Data ONTAP Network Management Guide for 7-Mode*.

Starting and configuring SNMP on the storage system

To start SNMP on the storage system, you can use the `snmp` command.

Steps

1. Enter the following command:

```
snmp init 1
```

2. Enter the following command:

```
snmp traphost add hostname
```

hostname is the host name of the UNIX client that will receive SNMP traps that the FTP server generates.

After you finish

You must enable SNMP traps on the UNIX client that you specified in Step 2.

Viewing SNMP traps on a UNIX client

To view SNMP traps on a UNIX client, you can enter the `snmptrapd -P` command.

Before you begin

Before you can view SNMP traps on a UNIX client, you must start and configure SNMP on the storage system.

Step

1. Enter the following command:

```
snmptrapd -P
```

Viewing FTP statistics

To view FTP statistics, you can enter the `ftp stat` command.

Step

1. Enter the following command:

```
ftp stat -p native
```

You can also use this command to view SFTP-only, explicit-FTPS-only, implicit-FTPS-only, IPv4-only, or IPv6-only statistics. For more information, see the `na_ftp(1)` man page.

Result

The `ftp stat` command displays the following statistics:

- Current number of FTP connections
- Highest number of simultaneous FTP connections
- Total number of FTP connections since FTP statistics were reset

Resetting FTP statistics

To reset FTP statistics, you can use the `ftp stat -z` command.

Step

1. Enter the following command:

```
ftp stat -z
```

Specifying the maximum number of FTP connections

To specify the maximum number of FTP connections that the FTP server allows, you can use the `ftpd.max_connections` option. By default, the maximum number of FTP connections is 500.

Step

1. Enter the following command:

```
options ftpd.max_connections n
```

`n` is the maximum number of FTP connections that the FTP server allows. If you set the `ftpd.max_connections` option to a value that is less than the current number of FTP connections, the FTP server refuses new connections until the number falls below the new maximum. The FTP server does not interrupt existing FTP connections.

In a HA configuration, the maximum number of FTP connections doubles automatically when the storage system is in takeover mode.

Specifying the maximum number of TFTP connections

To specify the maximum number of TFTP connections that the TFTP server allows, you can use the `tftpd.max_connections` option. The default number of TFTP connections is 8. The maximum number of connections supported is 32.

Step

1. Enter the following command:

```
options tftpd.max_connections n
```

n is the maximum number of TFTP connections that the TFTP server allows. If you set the `tftpd.max_connections` option to a value that is less than the current number of TFTP connections, the TFTP server refuses new connections until the number falls below the new maximum. The TFTP server does not interrupt existing TFTP connections.

In an HA configuration, the maximum number of TFTP connections doubles automatically when the storage system is in takeover mode.

Setting the FTP connection threshold

To specify how close the number of FTP connections must come to the maximum number of FTP connections before the FTP server adds an entry to the system log and (optionally) triggers an SNMP trap, you can set the `ftpd.max_connections_threshold` option. By default, this option is 0 (off).

Step

1. Enter the following command:

```
options ftpd.max_connections_threshold n
```

n is the percentage (0 through 99) of the value of `ftpd.max_connections`.

Specifying the TCP window size for FTP operations

To specify the TCP window size for FTP operations, you can use the `ftpd.tcp_window_size` option. By default, the TCP window size for FTP operations is 28,960.

Before you begin

Change the TCP window size for FTP operations only when your network configuration requires it. A change can strongly impact FTP performance.

Step

1. Enter the following command:

```
options ftpd.tcp_window_size n
```

n is the new TCP window size (the number of bytes the FTP server is willing to take from the FTP client at one time) for FTP operations.

Specifying the FTP idle timeout

You can set the `ftpd.idle_timeout` option to specify the FTP idle timeout value. This is the amount of time an FTP connection can be idle before the FTP server terminates it. By default, the FTP idle timeout value is 900 seconds.

Step

1. Enter the following command:

`options ftpd.idle_timeout n s | m | h`

n is the new timeout value. Append the letter *s*, *m*, or *h* to specify whether *n* represents seconds, minutes, or hours, respectively.

Managing anonymous FTP access

You can manage anonymous FTP access by enabling or disabling anonymous FTP access and specifying the home directory and user name for anonymous users.

Enabling or disabling anonymous FTP access

To enable or disable anonymous FTP access, you can set the `ftpd.anonymous.enable` option to `on` or `off`, respectively. This allows clients to access files using FTP without individual authentication. By default, this option is `off`.

Step

1. Perform one of the following actions.

If you want anonymous FTP access to be...	Enter the command...
Enabled	<code>options ftpd.anonymous.enable on</code>
Disabled	<code>options ftpd.anonymous.enable off</code>

After you finish

If you enable anonymous FTP access, you must perform the following tasks:

- Specify the user name for anonymous FTP users.
- Specify the home directory for anonymous FTP users.

Specifying the user name for anonymous FTP users

To specify the user name for anonymous FTP users, you can set the `ftpd.anonymous.name` option. By default, the user name for anonymous FTP users is "anonymous."

About this task

If the FTP authentication style is `unix`, the user name that you specify with this option overrides the user name that you specified for the FTP user in the `/etc/passwd` file.

Step

1. Enter the following command:

`options ftpd.anonymous.name username`

username is the name of the anonymous user.

Specifying the home directory for anonymous FTP users

To specify the home directory for anonymous FTP users (that is, the only directory to which anonymous FTP users have access), you can use the `ftpd.anonymous.home_dir` option.

About this task

If the FTP authentication style is `unix`, the home directory that you specify with this option overrides the home directory that you specified for the `ftp` user in the `/etc/passwd` file or NIS.

When the FTP server authenticates an anonymous FTP user with the NTLM authentication style, the FTP user has the same access privileges as the `null` user.

Steps

1. Create the home directory for anonymous FTP users.
2. Enter the following command:

```
options ftpd.anonymous.home_dir homedir
```

homedir is the name of the home directory for anonymous FTP users.

After you finish

Make sure that anonymous FTP users have read access to the directory you created in Step 1. For more information, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Managing the Secure File Transfer Protocol (SFTP)

You can manage the Secure File Transfer Protocol (SFTP) by enabling or disabling it and setting various options for it.

About SFTP

The Secure File Transfer Protocol (SFTP) is a secure replacement for the File Transfer Protocol (FTP). SFTP is based on the Secure Shell protocol.

Similar to FTP, SFTP is an interactive file transfer program that performs all operations over an encrypted SSH transport. Unlike FTP, SFTP encrypts both commands and data, providing effective protection against common network security risks. The SSH client and server provide both command-line SFTP tools and a graphical user interface for Windows users. SFTP encrypts the session, preventing the casual detection of your user name, password, or anything you have transmitted. This protocol assumes that it runs over a secure channel, that the server has already authenticated the user at the client end, and that the identity of the client user is externally available to the server implementation. SFTP runs from the SSH Connection Protocol as a subsystem.

Data ONTAP implements SFTP in accordance with version 03 of the Internet-Draft of the SSH File Transfer Protocol, which is available at tools.ietf.org/html/draft-ietf-secsh-filexfer-03.

Limitations of Data ONTAP support for SFTP

You should be aware of several limitations of Data ONTAP support for SFTP.

Since WAFL can write a maximum of 64 KB of data in one I/O operation, SFTP packet size is limited to 64 KB as well. If Data ONTAP receives a packet larger than 64 KB over SFTP, it closes the session and resets the connection. If Data ONTAP receives a read request for more than 64 KB of data, it reads only 64 KB of data.

Enabling or disabling SFTP

To enable or disable SFTP, you can set the `sftp.enable` option to `on` or `off`, respectively. This allows clients to access files using SFTP. By default, this option is `off`.

Before you begin

SSH must already be set up and running. For more information, see the `na_secureadmin(1)` man page or the *Data ONTAP System Administration Guide for 7-Mode*.

Step

1. Perform one of the following actions:

If you want SFTP to be...	Enter the command...
Enabled	<code>options sftp.enable on</code>
Disabled	<code>options sftp.enable off</code>

Enabling or disabling SFTP file locking

To prevent users from modifying files while the SFTP server is transferring them, you can enable SFTP file locking. By default, SFTP file locking is disabled.

Step

1. Perform one of the following actions:

If you want SFTP file locking to be...	Enter the command...
Disabled	<code>options sftp.locking none</code>
Enabled so that files being retrieved cannot be deleted or renamed	<code>options sftp.locking delete</code>
Enabled so that files being retrieved cannot be opened for writing or deletion	<code>options sftp.locking write</code>

SFTP authentication styles

When users connect using SFTP and authenticate, you can use three different authentication styles: mixed, NTLM, or UNIX. This determines the authentication style used and acceptable user name format.

In mixed mode, the authentication style used is selected by the format of the user name. User names with "\" are authenticated using NTLM and those without are authenticated using UNIX. Setting the authentication style explicitly to NTLM or UNIX forces the respective authentication type to be used regardless of the format of the user name.

To authenticate using NTLM with a domain user account, you must specify the user name in the format *domain\user_name*.

To authenticate using NTLM with a local user account, you must specify the user name in the format *storagesystem_hostname\user_name*.

By default, clients use user authentication methods in the following order: public-key, keyboard-interactive, and password authentication (if available). Public-key and certificate authentication are combined into the public-key authentication method. Data ONTAP allows public-key and password authentication by default.

Specifying the SFTP authentication style

To configure SFTP to use UNIX, Windows, or both authentication styles, you can set the `sftp.auth_style` option to `unix`, `ntlm`, or `mixed`, respectively. By default, this option is `mixed`.

Step

1. Perform one of the following actions:

If you want the authentication style to be...	Then...
Mixed	Enter the following command: options sftp.auth_style mixed
NTLM	Enter the following command: options sftp.auth_style ntlm
UNIX	Enter the following command: options sftp.auth_style unix

If you set the authentication style to mixed mode, Data ONTAP decides which authentication style to use based on the username format. If a username contains "\" or "@", it authenticates using NTLM. If it does not, it authenticates using UNIX.

If you set the authentication style to NTLM or UNIX, Data ONTAP uses only the specified authentication style regardless of the format of the username.

Enabling or disabling SFTP bypass traverse checking

You can enable or disable SFTP traverse checking by setting the `sftp.bypass_traverse_checking` option to `on` or `off`, respectively. By default, this option is set to `off`.

About this task

If the `sftp.bypass_traverse_checking` option is set to `off`, when a user attempts to access a file using SFTP, Data ONTAP checks the traverse (execute) permission for all directories in the path to the file. If any of the intermediate directories does not have the "X" (traverse permission), Data ONTAP denies access to the file. If the `sftp.bypass_traverse_checking` option is set to `on`, when a user attempts to access a file, Data ONTAP does not check the traverse permission for the intermediate directories when determining whether to grant or deny access to the file.

Step

1. Perform one of the following options:

If you want the SFTP bypass traverse checking to be...	Then...
Enabled	Enter the following command: <code>options sftp.bypass_traverse_checking on</code>
Disabled	Enter the following command: <code>options sftp.bypass_traverse_checking off</code>

Enabling or disabling SFTP user home directory restrictions

You can enable or disable SFTP user home directory restrictions by setting the `sftp.dir_restriction` option to `on` or `off`, respectively. By default, this option is `off`.

About this task

If this option is set to `on`, regular users are restricted to their home directories or, if you specify a directory using the `sftp.dir_override` option, they are restricted to the override directory. If this option is `off`, regular users are not restricted to a particular directory.

Note: This option has no effect on the default user account.

Step

1. Perform one of the following actions:

If you want SFTP user home directory restrictions to be...	Then...
Enabled	Enter the following command: options sftp.dir_restriction on
Disabled	Enter the following command: options sftp.dir_restriction off

Specifying the SFTP override path for user home directories

You can specify the SFTP override path for user home directories by setting the `sftp.dir_override` option. By default, this option is " " (null).

About this task

If you set this option to null (the default value), regular users are placed in their home directories when they log in. Otherwise, regular users are placed in the directory you specify.

Note: This option has no effect on the default user account.

Step

1. Enter the following command:

```
options sftp.dir_override path
```

path is the SFTP override path.

Enabling or disabling the overriding of UNIX permissions

To enable or disable the overriding of UNIX permissions specified by SFTP clients, you can use the `sftp.override_client_permissions` option. By default, this option is `off`.

About this task

If this option is enabled, Data ONTAP sets the UNIX permissions on newly created files and directories to 0755 regardless of the permissions specified by the SFTP client.

Step

1. Perform one of the following actions:

If you want overriding of UNIX permissions to be...	Then...
Enabled	Enter the following command: options sftp.override_client_permissions on

If you want overriding of UNIX permissions to be...	Then...
Disabled	Enter the following command: <code>options sftp.override_client_permissions off</code>

Managing SFTP log files

You can manage SFTP log files by enabling or disabling them, specifying their maximum size, and specifying their maximum number.

Enabling or disabling SFTP log files

You can enable or disable SFTP log files by setting the `sftp.log_enable` option to `on` or `off`, respectively. This enables you to enable or disable SFTP event logging. By default, this option is `on`.

About this task

When this option is enabled, Data ONTAP logs SFTP commands and data transfer operations to the `/etc/log/sftp.cmd.*` log files.

Step

1. Perform one of the following actions:

If you want SFTP log files to be...	Enter the command...
Enabled	<code>options sftp.log_enable on</code>
Disabled	<code>options sftp.log_enable off</code>

Specifying the maximum number of SFTP log files

To specify the maximum number of SFTP log files, you can set the `sftp.log_nfiles` option. By default, the maximum number of SFTP log files is 6.

Step

1. Enter the following command:

```
options sftp.log_nfiles n
```

`n` is the maximum number of log files from 1 to 100. For more information, see the `na_options(1)` man page.

Specifying the maximum size of the current SFTP log files

To specify the maximum size of the current `/etc/log/sftp.cmd.*` SFTP log files, you can set the `sftp.log_filesize` option. By default, the maximum size of the current SFTP log files is 512 KB.

Step

1. Enter the following command:

```
options sftp.log_filesize filesize
```

filesize is the maximum size of the current SFTP files expressed as a value from 1K to 4 GB. You can specify the value in gigabytes (G), megabytes (M), kilobytes (K), or bytes (blank). For more information, see the `na_options(1)` man page.

Example

The following example sets the maximum size of the current SFTP log files to 1 GB:

```
options sftpd.log_filesize 1G
```

Viewing SFTP statistics

To view SFTP statistics, you can enter the `sftp stat` command.

Step

1. Enter the following command:

```
sftp stat
```

Resetting SFTP statistics

To reset SFTP statistics, you can enter the `sftp stat -z` command.

Step

1. Enter the following command:

```
sftp stat -z
```

Specifying the maximum number of SFTP connections

To specify the maximum number of SFTP connections, you can use the `sftp.max_connections` option. By default, the maximum number of FTP connections is 15.

About this task

The maximum number of connections cannot exceed 15. Furthermore, the maximum number of SSH connections is reduced by the maximum number of SFTP connections.

Step

1. Enter the following command:

```
options sftp.max_connections n
```

n is the maximum number of SFTP connections from 0 to 15.

If you set the `sftp.max_connections` option to a value that is less than the current number of SFTP connections, Data ONTAP refuses new connections until the number falls below the new maximum. Data ONTAP does not interrupt existing SFTP connections.

Specifying the SFTP idle timeout value

To specify the SFTP idle timeout value (that is, the amount of time an SFTP connection can be idle before it becomes a candidate for termination), you can set the `sftp.idle_timeout` option. By default, the SFTP idle timeout value is 900 seconds.

Step

1. Enter the following command:

```
options sftp.idle_timeout timeout
```

timeout is an SFTP timeout value that is between 300s and 48h. You can specify the SFTP timeout value in seconds (s), minutes (m), or hours (h).

Managing FTP over SSL (FTPS)

You can manage FTP over SSL (FTPS) by enabling or disabling implicit FTPS, enabling or disabling explicit FTPS, and specifying whether explicit FTPS connections can be opened in secure mode.

Differences between implicit and explicit FTPS

FTP over SSL (FTPS) allows FTP software to perform secure file transfers.

Typically, data sent over an FTP connection, whether over a control connection or a data connection, is sent in clear text and without any freshness or integrity guarantees. FTPS provides an extension to the FTP protocol that allows FTP software to perform secure file transfers over an implicit FTPS connection or an explicit FTPS connection.

Implicit FTPS

Data ONTAP provides an industry-standard implementation of implicit FTPS; there is no corresponding RFC. With implicit FTPS, security is achieved by encrypting and decrypting data in the transport layer by SSL. In particular, implicit FTPS works as follows:

- Data ONTAP listens on port 990.
- The FTPS client connects to port 990.

A SSL handshake is initiated on connection. If the handshake fails, no further communication is allowed.

- After the completion of a successful SSL handshake, all further FTP communication goes through SSL and is secure.
- The command channel cannot be restored back to clear text.
New packets that some clients might require are `PBSZ` and `PROT`. The only argument that Data ONTAP supports for the `PROT` command is `P`, meaning private encrypted communications.
- The default port for the data channel is 989.

Explicit FTPS

Data ONTAP implements explicit FTPS in accordance with RFC 2228 and RFC 4217. In particular, explicit FTPS works as follows:

- Data ONTAP listens on port 21 (the standard FTP port).
- The FTP client connects to port 21 over a normal TCP connection.
Any communication over the connection is clear text to begin with. The connection can be made secure by issuing the `AUTH` command.
- After receiving the `AUTH` command, Data ONTAP initiates an SSL handshake.
You can use the `CCC` command to restore the command channel back to clear text.
- Before starting a data connection the client must issue `PBSZ` and `PROT` commands.
Without these commands the data connection would be clear text. The only arguments Data ONTAP supports for the `PROT` command are `C` and `P`, meaning clear text or private data channels.
- As specified in the RFC, the `PBSZ` command should be preceded by a successful authentication data exchange, and the `PROT` command should be preceded by a successful `PBSZ` command.
- The default port for the data channel is 20.

Enabling or disabling explicit FTPS

You can enable or disable explicit FTPS by setting the `ftpd.explicit.enable` option to `on` or `off`, respectively. By default, this option is `off`.

Before you begin

Before you can enable FTPS, you must set up and start SSL using the `secureadmin` command. For more information, see the `na_secureadmin(1)` man page or the *Data ONTAP System Administration Guide for 7-Mode*.

Note: SSL certificates for vFiler units are shared with the default vFiler unit. Therefore, you must set up and start SSL on the default vFiler unit.

Before you can enable explicit FTPS, you must set the `ftpd.enable` option to `on`.

About this task

When the `ftpd.explicit.enable` option is enabled, Data ONTAP allows explicit FTPS connections on port 21. When this option is disabled, Data ONTAP does not allow explicit FTPS connections on port 21.

Step

1. Perform one of the following actions:

If you want explicit FTPS to be...	Enter the command...
Enabled	<code>options ftpd.explicit.enable on</code>
Disabled	<code>options ftpd.explicit.enable off</code>

Allowing or preventing the opening of explicit FTPS data connections in secure mode

You can allow or prevent the opening of explicit FTPS data connections in secure mode by setting the `ftpd.explicit.allow_secure_data_conn` option to `on` or `off`, respectively. By default, this option is `on`.

Before you begin

Explicit FTPS must be enabled. Otherwise, this option has no effect.

About this task

When this option is enabled, Data ONTAP allows explicit FTPS connections to open data connections in secure mode (that is, by sending the `PROT P` command). When this option is disabled, Data ONTAP prevents explicit FTPS connections from opening data connections in secure mode.

Step

1. Perform one of the following actions:

If you want the opening of explicit FTPS data connections in secure mode to be...	Enter the command...
Allowed	<code>options ftpd.explicit.allow_secure_data_conn on</code>
Prevented	<code>options ftpd.explicit.allow_secure_data_conn off</code>

Enabling or disabling implicit FTPS

You can enable or disable implicit FTPS by setting the `ftpd.implicit.enable` option to `on` or `off`, respectively. By default, this option is `off`.

Before you begin

FTP must be enabled.

About this task

When this option is enabled, Data ONTAP allows implicit FTPS connections on port 990. When this option is disabled, Data ONTAP does not allow implicit FTPS connections on port 990.

Step

1. Perform one of the following actions:

If you want implicit FTPS to be...	Enter the command...
Enabled	<code>options ftpd.implicit.enable on</code>
Disabled	<code>options ftpd.implicit.enable off</code>

Managing FTP over IPv6

Starting with Data ONTAP 7.3.1, you can allow FTP clients to access the files on your storage system over IPv6.

Enabling or disabling FTP over IPv6

You can enable or disable FTP over IPv6 by setting the `ftpd.ipv6.enable` option to `on` or `off`, respectively.

Before you begin

You must enable IPv6 on the storage system by setting the `ip.v6.enable` option to `on`. For more information about enabling IPv6 on your storage system, see the *Data ONTAP Network Management Guide*.

About this task

- If you enable FTP over IPv6 and you then disable IPv6 on your storage system by setting the `ip.v6.enable` option to `off`, clients can no longer connect to the FTP server using IPv6.
- You do not need to restart FTP over IPv6 after restarting the IPv6 global option.
If FTP over IPv6 is enabled on the storage system, and if you disable and reenable the IPv6 global option, FTP IPv6 sockets are automatically created to listen for IPv6 addresses.

Step

1. Enable or disable FTP over IPv6.

If you want FTP over IPv6 to be...	Enter the command...
Enabled	<code>options ftpd.ipv6.enable on</code>

If you want FTP over IPv6 to be...	Enter the command...
Disabled	<code>options ftpd.ipv6.enable off</code>

Viewing FTP over IPv6 statistics

You can use the `ftp stat` command to view FTP over IPv6 statistics.

Step

1. Enter the following command:

```
ftp stat -i ipv6
```

For more information, see the `na_ftp(1)` man page.

File access using HTTP

To let HTTP clients (web browsers) access the files on your storage system, you can enable and configure Data ONTAP's built-in HyperText Transfer Protocol (HTTP) server. Alternatively, you can purchase and connect a third-party HTTP server to your storage system.

Managing the Data ONTAP HTTP server

Managing the HTTP server that is built into Data ONTAP involves several tasks.

Enabling or disabling the Data ONTAP HTTP server

You can use the `httpd.enable` option to enable or disable the HTTP server that is built into Data ONTAP. By default, this option is `off`. When this option is enabled, web browsers can access all of the files in the HTTP server's root directory.

Step

- 1. Perform one of the following actions:

If you want HTTP to be...	Enter the command...
Enabled	<code>options httpd.enable on</code>
Disabled	<code>options httpd.enable off</code>

Enabling or disabling the bypassing of HTTP traverse checking

You can enable or disable the bypassing of HTTP traverse checking by setting the `httpd.bypass_traverse_checking` option to `on` or `off`, respectively. By default, this option is set to `off`.

About this task

If the `httpd.bypass_traverse_checking` option is set to `off`, when a user attempts to access a file using the HTTP protocol, Data ONTAP checks the traverse (execute) permission for all directories in the path to the file. If any of the intermediate directories does not have the "X" (traverse permission), Data ONTAP denies access to the file. If the `http.bypass_traverse_checking` option is set to `on`, when a user attempts to access a file, Data ONTAP does not check the traverse permission for the intermediate directories when determining whether to grant or deny access to the file.

Step

- 1. Perform one of the following actions.

If you want the bypassing of HTTP traverse checking to be...	Then...
Enabled	Enter the following command: options httpd.bypass_traverse_checking on
Disabled	Enter the following command: options httpd.bypass_traverse_checking off

Specifying the root directory for the Data ONTAP HTTP server

You can set the `httpd.rootdir` option to specify the root directory for the HTTP server that is built into Data ONTAP. This is the directory that contains all of the files that an HTTP client can access.

Step

1. Enter the following command:
options httpd.rootdir *directory*
directory is the full path to the HTTP server's root directory.

Example

The following command sets the HTTP server's root directory to `/vol0/home/users/pages`:

```
options httpd.rootdir /vol0/home/users/pages
```

Specifying the maximum size of the log file for the Data ONTAP HTTP server

You can set the `ftpd.log.filesize` option to specify the maximum size of the log file for the HTTP server that is built into Data ONTAP. This option specifies the maximum log file size of the HTTP and FTP log files in the `/etc/log` directory, including the `ftp.cmd`, `ftp.xfer`, and `httpd.log` files. By default, this option is set to 512 kilobytes.

Step

1. Enter the following command:
options ftpd.log.filesize *bytes*
bytes is the new maximum size of the HTTP server's log file.

Testing the Data ONTAP HTTP server

To confirm that the HTTP server that is built into Data ONTAP is working, you can copy an HTML file into the HTTP server's root directory and then access the file from a web browser. You can also

access the HTTP server's root directory (or a subdirectory of the HTTP server's root directory) directly from a web browser.

Steps

1. Copy an HTML file into the HTTP server's root directory.
2. From a web browser running on a separate system, access the file you copied into the HTTP server's root directory.

The URL is `http://www.hostname.com/myfile.html`, where *hostname* is the host name of the storage system and *myfile.html* is the name of the file you copied into the HTTP server's root directory. You should see the contents of the file.

3. Optionally access the HTTP server's root directory (or a subdirectory of the HTTP server's root directory) directly from a web browser running on a separate client.

The URL is `http://www.hostname.com`, where *hostname* is the host name of the storage system.

The HTTP server looks for the following files in the following order in the directory that you specify:

- a. `index.html`
- b. `default.htm`
- c. `index.htm`
- d. `default.html`

If none of these files exists, the storage system automatically generates an HTML version of the directory listing for that directory (if the `httpd.autoindex.enable` option is `on`) or responds with the “403” (forbidden) error code (if the `httpd.autoindex.enable` option is `off`). For more information about the `httpd.autoindex.enable` option, see the `na_options(1)` man page.

Specifying how the Data ONTAP HTTP server maps MIME content types to file name extensions

You can create or edit the `/etc/httpd.mimetypes` file to specify how the HTTP server that is built into Data ONTAP maps Multipurpose Internet Mail Extensions (MIME) content types to file name extensions. If the `/etc/httpd.mimetypes` file does not exist, the HTTP server uses the mappings in the `/etc/httpd.mimetypes.sample` file. For more information, see the `na_httpd.mimetypes(5)` man page.

About this task

Web browsers interpret files according to their MIME content type. For example, if the MIME content type of a file is an image type, web browsers render the file as an image using a graphics program.

Note: For more information about MIME, see RFC 1521.

Step

1. Create entries in the `/etc/httpd.mimetypes` file with the desired mappings.

Entries use the following format:

An optional comment.

suffixContent-Type

The text after the pound character (#) is a comment. The file name suffix is not case-sensitive.

suffix is the file name extension to which you want to map a MIME content type.

Content-Type is the MIME Content-Type type. The first field of the MIME Content-Type describes the general type of data contained in the file; the second field is the data subtype, which shows the specific format in which the data is stored.

Example

The following entries to the `/etc/httpd.mimetypes` file map the `/image/pict` MIME content type to files with `.pct` and `.pict` file name extensions:

```
# My clients' browsers can now use
# PICT graphics files.
pct image/pict
pict image/pict
```

Now, if it is configured properly, web browsers will start a graphics program as a helper application, allowing users to view `.pct` and `.pict` files as graphics files.

Specifying how the Data ONTAP HTTP server translates HTTP requests

You can add `map`, `redirect`, `pass`, or `fail` translation rules to the `/etc/httpd.translations` configuration file. This allows you to specify how the HTTP server that is built into Data ONTAP responds to HTTP requests.

How the Data ONTAP HTTP server translations file works

The HTTP server that is built into Data ONTAP processes the rules in the `/etc/httpd.translations` file in the order they are listed, applying a rule if the URL matches the template. After the first successful match, the HTTP server stops processing the remaining rules.

You can use an asterisk (*) as a wildcard character in the *template* and *result* fields of `map`, `redirect`, and `pass` rules that you add to the `/etc/httpd.translations` file.

In the *template* field, the wildcard character matches zero or more characters, including the slash (/) character.

In the *result* field, the wildcard character represents the text expanded from the match in the template field. You should include the wildcard character in the result field only if you use a wildcard in the template field.

If you use multiple wildcard characters, the first one in the result field corresponds to the first one in the template field, the second one in the result field corresponds to the second one in the template field, and so on.

Adding a map rule

You can add a map rule to the `/etc/httpd.translations` file to specify that the HTTP server should map a URL to another location.

Steps

1. Open the `/etc/httpd.translations` file in a text editor.
2. Add the following rule:

```
map template result
```

template is the component of a URL that you want to map to another location (for example, `/image-bin/graphics/`).

result specifies the new location.

3. Save the file.

Example

The following map rule in the `/etc/httpd.translations` file maps a URL containing an `/image-bin` component to the `/usr/local/http/images` directory:

```
map /image-bin/* /usr/local/http/images/*
```

Adding a redirect rule

You can add a redirect rule to the `/etc/httpd.translations` file to specify that the HTTP server should redirect a URL containing a specific component to a new location.

Steps

1. Open the `/etc/httpd.translations` file in a text editor.
2. Add the following entry:

```
redirect template result
```

template is a component of a URL to redirect.

result specifies the new location.

Note: You must specify the result field for the redirect rule as a complete URL beginning with `http://` and the host name.

3. Save the file.

Example

The following entry in the `/etc/httpd.translations` file redirects Common Gateway Interface (CGI) requests to an HTTP server named `cgi-host`:

```
redirect /cgi-bin/* http://cgi-host/*
```

Adding a pass rule

You can add a pass entry to the `/etc/httpd.translations` file to specify that the HTTP server should process a rule as is, disregarding other rules.

Steps

1. Open the `/etc/httpd.translations` file in a text editor.
2. Add the following entry:

```
pass template [result]
```

template is a component of a URL

result is an optional location to which the HTTP server redirects the URL.

3. Save the file.

Example

The following entry in the `/etc/httpd.translations` file processes a request for a URL containing `/image-bin` as is:

```
pass /image-bin/*
```

Adding a fail rule

You can add a fail rule to the `/etc/httpd.translations` file to specify that the HTTP server should deny access to a URL containing a specific component.

Steps

1. Open the `/etc/httpd.translations` file in a text editor.
2. Add the following entry:

```
fail template
```


template is the URL component to which the HTTP server should deny access.

3. Save the file.

Example

The following entry in the `/etc/httpd.translations` file denies access to the `/usr/` forbidden directory:

```
fail /usr/forbidden/*
```

Configuring MIME Content-Type values

You can configure the storage system to send the appropriate MIME (Multipurpose Internet Mail Extensions) Content-Type value in each response to a get request from a client by mapping the file name suffix, for example, `.gif`, `.html`, or `.mpg`, according to information in the `/etc/httpd.mimetypes` file.

About this task

The MIME Content-Type value of a file tells a browser on a client how to interpret the file. For example, if the MIME Content-Type value shows that a file is an image file, and the client is configured properly, the browser can render the image by using a graphics program.

For more information about MIME, see RFC 1521.

Step

1. Edit the entries in the `/etc/httpd.mimetypes` file.

Entries are in the following format:

```
# An optional comment.
```

```
suffixContent-Type
```

Lines preceded by the `#` sign are comments. The file name suffix is not case-sensitive.

Example

The following are sample entries:

```
# My clients' browsers can now use
# PICT graphics files.
pct      image/pict
pict     image/pict
```

In the sample entries, files whose names end with `.pct` or `.pict` are mapped to the MIME Content-Type value of `image/pict`. The first field in the Content-Type value describes the general type of data contained in the file; the second field is the data subtype, which shows the

specific format in which the data is stored. If the browser on the client is configured to start a graphics program as a helper application, the user can view a file named `file.pict` as a graphics file on the client.

Maintaining security for the Data ONTAP HTTP server

You can maintain security for the HTTP server that is built into Data ONTAP by using the HTTP options to restrict access, using an HTTP virtual firewall, protecting web pages with user authentication, disabling support for the HTTP TRACE method, and specifying how long Data ONTAP keeps idle HTTP connections open.

Using HTTP options to restrict access

The HTTP options restrict access to HTTP services from specified hosts and from specified interfaces.

Step

1. Perform one of the following actions:

If you want to...	Configure the option...
Restrict access to HTTP services	<code>httpd.access</code>
Enable or disable support for the HTTP TRACE method	<code>httpd.method.trace.enable</code> By default, this option is <code>off</code> . The HTTP TRACE method allows an HTTP client to see what is being received at the other end of the request chain, for debugging purposes. (For more information, see RFC 2616.) However, attackers can leverage the HTTP TRACE method in conjunction with cross-domain browser vulnerabilities to read sensitive header information from third-party domains. For more information, search for Vulnerability Note 867593 in the United States Computer Emergency Readiness Team Vulnerability Notes Database at www.cert.org .

Examples

In the following example, only host Host1 is allowed access through interface e3 to the HTTPD services on storage system Filer1:

```
Filer1> options httpd.access host=Host1 AND if=e3
```

Using an HTTP virtual firewall

An HTTP virtual firewall provides security on your storage system by restricting HTTP access through the subnet interface over which the HTTP requests arrive.

About this task

You restrict HTTP access by marking the subnet interface as untrusted. An untrusted subnet interface provides only read-only HTTP access to the storage system. By default, a subnet interface is trusted.

Mark a subnet interface as untrusted if it meets all the following conditions:

- You know you are going to service HTTP requests over that interface.
- You do not want to allow requests through protocols other than HTTP.
- You want to restrict access to the storage system through that interface to read-only access.

Step

1. Enter the following command:

```
ifconfig interface_name [trusted | untrusted]
```

interface_name is the specific interface to set as trusted or untrusted.

Use *trusted* to allow full HTTP access or *untrusted* to restrict HTTP access.

Example

The following command marks the f0 interface as untrusted:

```
ifconfig f0 untrusted
```

Protecting web pages

You can restrict HTTP access, and thereby protect web pages, by preventing unauthorized users from accessing web pages. In this way, only specified users or groups can access directories containing the web pages.

About this task

Data ONTAP provides the following two methods of authentication for HTTP access:

- Basic
- NTLM

You specify the method of authentication to use in the `/etc/httpd.access` file. Both authentication methods can coexist on a storage system, but you can specify only one authentication method per directory in the HTTP subtree.

Basic authentication

You use the following three configuration files to set up authentication for the HTTP service: `/etc/httpd.access`, `/etc/httpd.passwd`, and `/etc/httpd.group`.

The `/etc/httpd.access` file contains the method of authentication, the directories for which you want to restrict access, and the list of users and groups authorized to access these directories.

The `/etc/httpd.passwd` file contains the encrypted form of the password that a user, specified in the `/etc/httpd.access` file, uses to gain access to the directories specified in the `/etc/httpd.access` file. The `/etc/httpd.passwd` file uses the same format that the `/etc/passwd` file uses.

The `/etc/httpd.group` file contains group and user IDs of the members of each group who are authorized to access the directories specified in the `/etc/httpd.access` file. The `/etc/httpd.group` file uses the same format that the `/etc/group` file uses.

NTLM authentication for directory security

You can use Windows Domain Authentication instead of basic authentication for a directory. Data ONTAP uses the Domain Controller (DC) to authenticate users accessing the directories containing the web pages.

You must specify the directories in the `/etc/httpd.access` file for which you want the domain controller to authenticate users.

A user accessing a directory for which NTLM authentication has been set up must specify a domain with the user name. If a domain is not specified, the domain of the storage system is assumed as a default. The users can specify the domain in either of the following formats:

- `user_name@domain_name`
- `domain_name\user_name`

Note: You must have CIFS running on your storage system to use the NTLM authentication for HTTP access.

You do not need to maintain information in the `/etc/httpd.passwd` and `/etc/httpd.group` files, thus centralizing user administration. And, if you use Internet Explorer (IE) as your browser, NTLM authentication is a more secure method of authenticating users because user name and password are not transmitted in plain text.

Note: Netscape browsers send user names and passwords in plain text, providing no security advantage for NTLM.

Editing the `/etc/httpd.access` file

The `/etc/httpd.access` file contains options that govern the access to and appearance of each directory.

About this task

The storage system supports the following options:

- `Directory`
Specifies the directory you want to protect. The directory option encloses all other options.
- `AuthName`
Specifies an alias for the directory that appears instead of the directory name in the browser password dialog box when a user tries to access the directory.
- `require user`
Specifies the users who can access the directory.
- `require group`
Specifies the groups that can access the directory.

Note: The options `require user` and `require group` are only required for basic authentication.

Option information for each directory in the `/etc/httpd.access` file is given in the following format:

```
<Directory directory>
```

```
option ...
```

```
</Directory>
```

directory is the specific directory tree name for which you want to enable authorized access.

Steps

1. Open the `/etc/httpd.access` file for editing.
2. Specify the directory tree you want to protect in the following line:

```
<Directory directory>
```

directory specifies the directory tree name you want protected.

3. If you are configuring basic authentication using `/etc/httpd.passwd` and `/etc/httpd.group` files, specify the alias for the directory in the following line:

```
AuthName title_phrase
```

title_phrase is any string you specify that appears instead of the directory name in the browser password dialog box when a user tries to access the directory. This name can contain spaces. For example:

AuthName Secured Area

4. Otherwise, if you are configuring NTLM authentication, specify the following, exactly as shown:

AuthName Windows(tm) Authentication

5. Specify the users who can access the directory in the following line:

```
require user user_id[, user_id, ...]
```

user_id specifies the user ID for each user who should have access to the directory.

6. Specify the groups that can access the directory in the following line:

```
require group group_id[, group_id, ...]
```

group_id specifies the group ID for each group that should have access to the directory.

7. End the option or list of options for the specified directory using the following line:

```
</Directory>
```

8. Save the file.

Example

The following example shows the use of multiple Directory options in a `/etc/httpd.access` file to specify either Basic or NTLM authentication on a storage system:

```
<Directory /vol/vol0/web1>
AuthName Windows(tm) Authentication
</Directory>
<Directory /vol/vol0/web2>
AuthName Web2 directory
require user test1
require group testg1
</Directory>
<Directory /vol/vol0/web3>
AuthName Windows(tm) Authentication
</Directory>
<Directory /vol/vol0/web4>
AuthName Web4 directory
require user test2
</Directory>
```

In this example, web1 and web3 use NTLM authentication and web2 and web4 use basic authentication. Access to web2 is limited to user test1 and members of group testg1, and access to web4 is limited to user test2.

Creating and editing the `httpd.passwd` file

The `/etc/httpd.passwd` file contains encrypted passwords of users listed in the `/etc/httpd.access` file. This file is only required if you are using basic authentication to authenticate users.

About this task

If you have an HTTP server that uses a user name and password method to authenticate users, you can copy user IDs and encrypted passwords from it. You must edit the `/etc/httpd.passwd` file to remove users that you do not want to have access.

If an HTTP server is not available, you can copy an existing `/etc/passwd` file from a UNIX server and save it on the storage system as the `/etc/httpd.passwd` file.

Steps

1. Open the `/etc/httpd.passwd` file.
2. Remove the user IDs and encrypted passwords of users that you do not want to have access to the directory you specified in the `/etc/httpd.access` file.
3. Save the edits.

Creating and editing the `httpd.group` file

The `/etc/httpd.group` file contains the group names and the users belonging to those groups. This file is only required if you are using basic authentication to authenticate users.

About this task

If you have an HTTP server that authenticates groups of users, you can copy the group names and user IDs from it. You must edit the `/etc/httpd.group` file to remove groups that you do not want to have access.

If an HTTP server is not available, you can copy an existing `/etc/group` file from a UNIX server and save it on the storage system as the `/etc/httpd.group` file.

Steps

1. In the `/etc/httpd.group` file, edit the following line:

```
group_id:user_id [, user_id, ...]
```

The lists are copied in from a server that has a similar list.

2. Add or remove groups and users. Group and user information is listed in the following format:

```
group_id: user_id[user_id ...]
```

group_id is the group name.

user_id is the name of each user who belongs to the group.

3. Save the file.

Configuring HTTP virtual hosting

In Data ONTAP 7.3 and later releases, you can configure HTTP virtual hosting by adding alias IP addresses to a physical interface. Data ONTAP no longer uses *vh* interfaces for this purpose.

Steps

1. Enable HTTP by entering the following command:

```
options httpd.enable on
```

2. Add one or more alias IP addresses to the physical interface that you will be using for HTTP virtual hosting by entering the following command:

```
ifconfig physical_interface_name [IP_address_family] alias IP_address
```

Example

The following command adds the 192.225.37.102 alias IP address to the *e0a* physical interface:

```
ifconfig e0a alias 192.225.37.102
```

For more information, see the *na_ifconfig(1)* man page.

3. Add entries to the */etc/httpd.hostprefixes* file that map the alias IP addresses you specified in Step 2 to one or more subdirectories of the HTTP root directory.

The *httpd.rootdir* option determines the HTTP root directory.

Example

The following entry to the */etc/httpd.hostprefixes* file maps the 192.225.37.102 alias IP address to the */httpdir1* subdirectory:

```
/httpdir1 192.225.37.102
```

4. Test your HTTP virtual hosting configuration by using an HTTP client to connect to the alias IP addresses you created and mapped in Steps 2 and 3, respectively.

Displaying HTTP server statistics

You can use the *httpstat* command to display five types of statistics about operations of the HTTP server that is built into Data ONTAP.

About this task

The five statistics types include the following:

- Request
- Detailed

- Error
- Service
- Timeout

Step

1. Enter the following command:

```
httpstat [-dersta]
```

-d displays detailed statistics.

-e displays error statistics.

-r displays request statistics.

-s displays service statistics.

-t displays timeout statistics.

-a displays all HTTP statistics.

If you use no arguments, `httpstat` displays HTTP request statistics.

For detailed information about the `httpstat` command, see the `httpstat(1)` man page.

Request statistics

If you specify request statistics, Data ONTAP displays the following statistics.

Label of statistic	Description
Accept	Number of new connections accepted by the storage system
Reuse	Number of new requests received on existing connections
Response	Number of responses sent
InBytes	Number of bytes received for all incoming requests
OutBytes	Number of bytes sent, including all HTTP headers, but not including data generated by servlets

Detailed statistics

If you specify detailed statistics, Data ONTAP displays the following statistics.

Label of statistic	Description
Get	Number of requests for files received
Head	Number of requests for file information received
Redirect	Number of requests redirected to another file

Label of statistic	Description
NotMod	Number of times clients (browsers) are told that requested files are not modified
Post	Number of POST requests received
Put	Number of PUT requests received
Servlet	Number of servlet requests received

Error statistics

If you specify error statistics, Data ONTAP displays the following statistics

Label of statistic	Description
Errors	Number of HTTP protocol error responses returned
BadReq	Number of unrecognized requests received
LogDiscard	Number of log entries discarded because the log was full
UnAuth	Number of requests denied because they lacked authorization
RcvErr	Number of requests aborted because of errors on the input socket

Service statistics

If you specify service statistics, Data ONTAP displays the following statistics.

Label of statistic	Description
Open	Number of currently open connections
Peak	Maximum number of connections ever achieved
Waits	Current number of connections accepted, but waiting for a connection structure

Timeout statistics

If you specify timeout statistics, Data ONTAP displays the following statistics.

Label of statistic	Description
Pending	Number of connection structures reclaimed after the network connection was started, but before any data was sent to the storage system
Active	Number of connection structures reclaimed after the network connection was started and a partial request was sent, but before the complete request arrived

Label of statistic	Description
Idle	Number of connections that were reclaimed after a complete request, but before the open connection could receive another request

Resetting statistics for the Data ONTAP HTTP server

You can use the `httpstat -z` command to reset statistics for the HTTP server that is built into Data ONTAP.

Step

1. Enter the following command:

```
httpstat -z[derta]
```

`-zd` displays detailed statistics.

`-ze` displays error statistics.

`-zr` displays request statistics.

`-zt` displays timeout statistics.

`-za` displays all HTTP statistics except the service statistics.

Note: You cannot reset the service statistics.

For detailed information about the `httpstat` command, see the `httpstat(1)` man page.

Viewing HTTP server connection information

You can view many types of information in the `/etc/log/httpd.log` file for each connection established by the HTTP server that is built into Data ONTAP.

Steps

1. Access the `/etc/log` directory on the storage system default volume (`/vol/vol0` by default) from an NFS or CIFS client.
2. Use a text viewer or text editor to open and view the `httpd.log` file.
3. Close the log file when you are finished viewing it.

Result

Data ONTAP displays the following types of information:

- IP address of HTTP client
- Names of authorized users making requests.

If the page is protected, Data ONTAP lists authorized names it gets from the `/etc/httpd.passwd` file. If the page is not protected, dashes appear instead of a name.

- Time of connection
Greenwich Mean Time (GMT), in dd/mm/yy:hh:mm:ss format
- Request line from connecting host, for example, `get /my_company.html`
- Status code returned by the server, as defined in the HTTP 1.0 specifications
- Total bytes sent in response by the storage system, not including the MIME header

Example

```
192.9.77.2 - - [26/Aug/2003:16:45:50] "GET /top.html" 200 1189
192.9.77.2 - - [26/Aug/2003:16:45:50] "GET /header.html" 200 531
192.7.15.6 - - [26/Aug/2003:16:45:51] "GET /logo.gif" 200 1763
198.9.200.2 - - [26/Aug/2003:16:45:57] "GET /task/top.html" 200 334
192.9.20.5 authuser [26/Aug/2003:16:45:57] "GET /task/head.html"
200 519
```

Changing the `/etc/log/httpd.log` file format

The default format of the `/etc/log/httpd.log` file shows the IP address of the HTTP clients and the HTTP path accessed, but not which virtual host is accessed. You can change the format of the `/etc/log/httpd.log` file so that it distinguishes HTTP messages by virtual hosts by setting the `httpd.log.format` option.

Step

1. Enter the following command:

```
options httpd.log.format alt1
```

To revert the setting for log format, change this option from `alt1` to the default value, `common`.

Purchasing and connecting a third-party HTTP server to your storage system

You can work around limitations of the HTTP server that is built into Data ONTAP by purchasing and connecting a third-party HTTP server to your storage system.

About this task

The Data ONTAP HTTP server has the following limitations:

- No support for Secure HTTP (HTTPS)
- No support for more than one HTTP root directory
- No support for scripts (that is, the HTTP supports file serving only)
- Scalability and performance problems if there are a large number of file operations on a large number of small files

Steps

1. Purchase a third-party HTTP server.
2. Connect the third-party HTTP server to your storage system using the NFS protocol.

For more information, see the documentation that comes with your third-party HTTP server.

Support for HTTP and HTTPS over IPv6

Starting with Data ONTAP 7.3.1, you can allow HTTP clients to access the files on your storage system over IPv6. You can also enable HTTPS for secure administrative access to the storage system.

You can use HTTP over IPv6 only for accessing files from your storage system.

Enabling or disabling HTTP and HTTPS over IPv6

You can enable or disable HTTP and HTTPS over IPv6 by setting the `httpd.ipv6.enable` option to `on` or `off`, respectively.

Before you begin

IPv6 must already be enabled. For more information about enabling IPv6 on your storage system, see the *Data ONTAP Network Management Guide for 7-Mode*.

About this task

- If you have enabled HTTP and HTTPS over IPv6 and you then disable IPv6 on your storage system by setting the `ip.v6.enable` option to `off`, HTTP and HTTPS are automatically disabled over IPv6.
- You do not need to restart HTTP and HTTPS over IPv6 after restarting the IPv6 global option. If HTTP and HTTPS over IPv6 are enabled on the storage system, and if you disable and reenable the IPv6 global option, HTTP and HTTPS IPv6 sockets are automatically created to listen for IPv6 addresses.

Step

1. Perform one of the following actions:

If you want HTTP and HTTPs over IPv6 to be...	Enter the command...
Enabled	<code>options httpd.ipv6.enable on</code>
Disabled	<code>options httpd.ipv6.enable off</code>

Listing HTTP connections over IPv4 or IPv6

You can use the `httpstat` command to view all HTTP connections over IPv4 and IPv6. You can also use the `-i` option with the `httpstat` command to view the HTTP connections over only IPv4 or only IPv6.

Step

- 1. Perform one of the following actions:

If you want to...	Enter the command...
List all HTTP connections	<div>httpstat</div> <div>Example:</div> <div><pre>httpstat Requests Accept Reuse Response InBytes OutBytes Total Stats: 73 0 39 27620 7247 IPv4 Stats: 43 0 10 9460 9 IPv6 Stats: 30 0 29 18160 7238</pre></div>
List HTTP connections over IPv4	<div>httpstat -i ipv4</div> <div>Example:</div> <div><pre>httpstat -i ipv4 Requests Accept Reuse Response InBytes OutBytes 43 0 10 9460 9</pre></div>

If you want to...	Enter the command...															
List HTTP connections over IPv6	httpstat -i ipv6 Example: <div><pre>httpstat -i ipv6</pre><table><tr><th colspan="5">Requests</th></tr><tr><th>Accept</th><th>Reuse</th><th>Response</th><th>InBytes</th><th>OutBytes</th></tr><tr><td>30</td><td>0</td><td>29</td><td>18160</td><td>7238</td></tr></table></div>	Requests					Accept	Reuse	Response	InBytes	OutBytes	30	0	29	18160	7238
Requests																
Accept	Reuse	Response	InBytes	OutBytes												
30	0	29	18160	7238												

File access using WebDAV

To let users use WebDAV interoperable, collaborative applications, you can add the WebDAV Web-based Distributed Authoring and Versioning) protocol to your existing HTTP service. Alternatively, you can purchase and connect a third-party WebDAV server to your storage system.

Understanding WebDAV

The WebDAV protocol defines the HTTP extensions that enable distributed web authoring tools to be broadly interoperable, while supporting user needs. WebDAV enables you to create HTTP directories.

The WebDAV protocol provides support for remote software development teams through a wide-range of collaborative applications. WebDAV leverages the success of HTTP and acts as a standard access layer for a wide range of storage repositories. HTTP gives read access, WebDAV gives write access.

Major features of this protocol include the following:

- **Locking**
Long-duration exclusive and shared write locks prevent two or more collaborators from writing to the same resource without first merging changes. To achieve robust Internet-scale collaboration, where network connections might be disconnected arbitrarily, and for scalability, because each open connection consumes server resources, the duration of DAV locks is independent of any individual network connection.
- **Properties**
XML properties provide storage for arbitrary metadata, such as a list of authors on web resources. These properties can be set, deleted, and retrieved efficiently using the DAV protocol. The DASL (DAV Searching and Locating) protocol provides searches of web resources based on the values in XML properties.
- **Namespace manipulation**
Because resources sometimes need to be copied or moved as the web evolves, DAV supports copy and move operations. Collections, similar to file system directories, can be created and listed.
- **HTTP feature support**
The Data ONTAP WebDAV implementation supports your HTTP configuration settings, such as redirect rules, authentication, and access restrictions. To use WebDAV, you need to have HTTP service enabled and configured.
- **CIFS feature support**
The Data ONTAP WebDAV implementation supports CIFS home directories when the storage system has a valid CIFS license installed.

Managing the Data ONTAP WebDAV server

Managing the WebDAV server that is built into Data ONTAP includes tasks of enabling or disabling the WebDAV protocol and pointing a WebDAV client to a home directory.

Enabling or disabling the Data ONTAP WebDAV server

You can set the `webdav.enable` option to `on` or `off`, respectively, to enable or disable the WebDAV server that is built into Data ONTAP. By default, this option is `on`.

Before you begin

Before you can enable the Data ONTAP WebDAV server, you must enable the Data ONTAP HTTP server. The WebDAV server supports your HTTP configuration settings, such as redirect rules, authentication, and access restrictions.

Furthermore, the WebDAV server supports CIFS home directories when you have valid CIFS licenses and you have enabled and configured CIFS home directories.

Step

1. Perform one of the following actions:

If you want WebDAV to be...	Enter the command...
Enabled	<code>options webdav.enable on</code>
Disabled	<code>options webdav.enable off</code>

Pointing a WebDAV client to a home directory

You can point a WebDAV client to a home directory by appending a tilde (~) character to the URL that you enter in the WebDAV client's navigation field.

Step

1. In the navigation (or default directory) field of your WebDAV client, enter a URL with the following syntax:

`http://host[:port]/~`

host is the host name or IP address for the storage system

port is the port through which you want to access the storage system. The tilde (~) character specifies the user's home directory.

Examples

The following are examples of valid WebDAV home directory URLs

```
http://eng_filer.lab.company.com/~  
http://10.120.83.104:80/~
```

Purchasing and connecting a third-party WebDAV server to your storage system

You can work around limitations of the WebDAV server that is built into Data ONTAP by purchasing and connecting a third-party WebDAV server to your storage system.

About this task

The Data ONTAP WebDAV server has the following limitations:

- Supports values that contain two-byte Unicode characters only; Data ONTAP will not properly record larger Unicode characters.
- Supports the core WebDAV protocols only; Data ONTAP does not support the secondary WebDAV protocols.
- Does not support home directory features for virtual IP addresses. URLs that specify a virtual IP address as the host will not be resolved.

Steps

1. Purchase a third-party WebDAV server.
2. Connect the third-party WebDAV server to your storage system via the NFS protocol.

For more information, see the documentation that comes with your third-party WebDAV server.

Event log and audit policy mapping

Some Event Log and Audit group policies are applied differently by Data ONTAP than by Windows systems.

If Group Policy Object (GPO) support is enabled on your storage system, Data ONTAP processes and applies all relevant GPOs. Most of the relevant group policy settings are applied uniformly on storage systems running Data ONTAP and Windows systems.

However, two types of policy—Event Log and Audit (Local Policies)—are applied differently on storage systems because the underlying logging and auditing technologies are different. Event Log and Audit GPOs are applied to storage systems by mapping and setting corresponding Data ONTAP options. The effect of mapping these options is similar but not identical to Event Log and Audit policy settings.

The following tables show the Data ONTAP options that are set when the corresponding GPOs are applied. For more information about the options, see the options(1) man page.

Event Log mapping values

For each row in the following table, the right column shows the Data ONTAP options that are set when the Event Log policies (and settings and examples, if appropriate) in the left column are applied.

Policy name	Setting	Data ONTAP options
Maximum security log size	n/a	<code>cifs.audit.logsize</code>
Retention method for security log	Overwrite events by days (Example: 7 days)	<code>cifs.audit.autosave.file.extension</code> <code>timestamp</code> <code>cifs.audit.autosave.file.limit 0</code> <code>cifs.audit.autosave.onsize.threshold 100</code> <code>cifs.audit.autosave.onsize.enable on</code> <code>cifs.audit.autosave.ontime.interval 7d</code> <code>cifs.audit.autosave.ontime.enable on</code> <code>cifs.audit.saveas /etc/log/adtlog.evt</code> <code>cifs.audit.enable on</code>

Policy name	Setting	Data ONTAP options
Retention method for security log	Overwrite events as needed	<pre>cifs.audit.autosave.file.extension timestamp cifs.audit.autosave.file.limit 1 cifs.audit.autosave.onsize.threshold 100 cifs.audit.autosave.onsize.enable on cifs.audit.autosave.ontime.enable off cifs.audit.saveas /etc/log/adtlog.evt cifs.audit.enable on</pre>
Retention method for security log	Do not overwrite events (clear log manually)	<pre>cifs.audit.autosave.file.extension timestamp cifs.audit.autosave.file.limit 0 cifs.audit.autosave.onsize.threshold 100 cifs.audit.autosave.onsize.enable on cifs.audit.autosave.ontime.enable off cifs.audit.saveas /etc/log/adtlog.evt cifs.audit.enable on</pre>

Audit mapping values

For each row in the following table, the right column shows the Data ONTAP options that are set when the Audit policies (and settings and examples, if appropriate) in the left column are applied.

Policy name	Setting	Data ONTAP options
Audit account logon events Audit logon events	Both policies are defined but neither sets audit attempts.	<pre>cifs.audit.logon_events .enable off</pre>
Audit account logon events Audit logon events	Both policies are defined, and audit attempts are set to Success, Failure, or Success & Failure	<pre>cifs.audit.enable on cifs.audit.logon_events .enable on</pre>
Audit directory service access Audit object access	Both policies are defined but neither sets audit attempts.	<pre>cifs.audit.file_access_ events.enable off</pre>
Audit directory service access Audit object access	Both policies are defined, and audit attempts are set to Success, Failure, or Success & Failure	<pre>cifs.audit.enable on cifs.audit.file_access_ events.enable on</pre>

Policy name	Setting	Data ONTAP options
Other Audit policies and settings.	No mapping action is performed.	None

Glossary

To understand the file access and protocols management concepts in this document, you might need to know how certain terms are used.

A

ACL	Access control list.
active/active configuration	<ul style="list-style-type: none"> In the Data ONTAP 7.2 and 7.3 release families, a pair of storage systems or V-Series systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. Also sometimes referred to as <i>active/active pairs</i>. In the Data ONTAP 8.x release family, this functionality is referred to as a <i>high-availability (HA) configuration</i> or an <i>HA pair</i>. In the Data ONTAP 7.1 release family and earlier releases, this functionality is referred to as a <i>cluster</i>.
adapter	A SCSI card, network card, hot-swap adapter, serial adapter, or VGA adapter that plugs into an expansion slot. Sometimes called <i>expansion card</i> .
address resolution	The procedure for determining an address corresponding to the address of a LAN or WAN destination.
administration host	The client you specify during system setup for managing your storage system. The setup program automatically configures the storage system to accept telnet and rsh connections from this client, to give permission to this client for mounting the / and /home directories, and to use this client as the mail host for sending AutoSupport email messages. At any time after you run the setup program, you can configure the storage system to work with other clients in the same way as it does with the administration host.
agent	A process that gathers status and diagnostic information and forwards it to network management stations, for example, <i>SNMP agent</i> .
appliance	A device that performs a single, well-defined function and is simple to install and operate.
ATM	Asynchronous Transfer Mode. A network technology that combines the features of cell-switching and multiplexing to offer reliable and efficient network services. ATM provides an interface between devices such as workstations and routers, and the network.
authentication	A security step performed by a domain controller for the storage system's domain, or by the storage system itself, using its <code>/etc/passwd</code> file.

AutoSupport	A storage system daemon that triggers email messages from the customer site to technical support or another specified email recipient when there is a potential storage system problem.
B	
big-endian	A binary data format for storage and transmission in which the most significant byte comes first.
C	
CIFS	See <i>Common Internet File System (CIFS)</i> .
client	A workstation or PC in a client-server architecture; that is, a computer system or process that requests services from and accepts the responses of another computer system or process.
cluster interconnect	The cables and adapters with which two nodes (storage systems) in an HA pair are connected, and over which heartbeat and WAFL log information are transmitted when both nodes are running.
cluster monitor	The software that administers the relationship of nodes in a cluster.
community	A logical relationship between an SNMP agent and one or more SNMP managers. A community is identified by name, and all members of the community have the same access privileges.
console	The physical or virtual terminal that is used to monitor and control a storage system.
Copy-On-Write (COW)	The technique for creating Snapshot copies without consuming excess disk space.
D	
degraded mode	The operating mode of a storage system when a disk in the RAID group fails or the batteries on the NVRAM card are low.
disk ID number	The number assigned by the storage system to each disk when it probes the disks at startup.
disk shelf	A shelf that contains disk drives and is attached to a storage system.
E	
emulated storage system	A software copy of a failed storage system that is hosted by its takeover storage system. The emulated storage system appears to users and administrators to be a functional version of the failed storage system. For example, it has the same name as the failed storage system.
Ethernet adapter	An Ethernet interface card.

expansion card A SCSI card, NVRAM card, network card, hot-swap card, or console card that plugs into a storage system expansion slot. Sometimes called an *adapter*.

expansion slot The slots on the storage system board into which you insert expansion cards.

F

failed storage system A physical storage system that has ceased operating. In a high-availability configuration, it remains the failed storage system until a giveback succeeds.

FDDI adapter A Fiber Distributed Data Interface (FDDI) interface card.

FDDI-fiber An FDDI adapter that supports a fiber-optic cable.

FDDI-TP An FDDI adapter that supports a twisted-pair cable.

FPolicy The Data ONTAP proprietary file policy feature that provides the ability to control access permissions based on file properties, such as file type.

G

GID See *Group ID (GID)*.

giveback The return of identity from an emulated storage system to the failed system, resulting in a return to normal operation. The reverse of *takeover*.

group In Data ONTAP operating in 7-Mode, a group of users defined in the storage system's `/etc/group` file.

Group ID (GID) The number used by UNIX systems to identify groups.

H

HA (high availability)

- In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

heartbeat A repeating signal transmitted from one storage system to the other that indicates that the storage system is in operation. Heartbeat information is also stored on disk.

hot spare disk A disk installed in the storage system that can be used to substitute for a failed disk. Before the disk failure, the hot spare disk is not part of the RAID disk array.

hot swap The process of adding, removing, or replacing a disk while the storage system is running.

hot swap adapter An expansion card that makes it possible to add or remove a hard disk with minimal interruption to file system activity.

I

inode A data structure containing information about files on a storage system and in a UNIX file system.

interrupt switch A switch on some storage system front panels used for debugging purposes.

L

LAN Emulation (LANE) The architecture, protocols, and services that create an Emulated LAN using ATM as an underlying network topology. LANE enables ATM-connected end systems to communicate with other LAN-based systems.

local storage system The system you are logged in to.

M

magic directory A directory that can be accessed by name but does not show up in a directory listing. The .snapshot directories, except for the one at the mount point or at the root of the share, are magic directories.

mailbox disk One of a set of disks owned by each storage system that is used to store the HA configuration state information of a storage system. If that system stops operating, the takeover system uses the information in the mailbox disks in constructing a virtual storage system. Mailbox disks are also used as file system disks.

mail host The client host responsible for sending automatic email to technical support when certain storage system events occur.

Maintenance mode An option when booting a storage system from a system boot disk. Maintenance mode provides special commands for troubleshooting hardware and configuration.

MIB Management Information Base. ASCII files that describe the information that the SNMP agent sends to network management stations.

MIME Multipurpose Internet Mail Extensions. A specification that defines the mechanisms for specifying and describing the format of Internet message bodies. An HTTP response containing the MIME Content-Type header allows the HTTP client to invoke the application that is appropriate for the data received.

MultiStore In Data ONTAP operating in 7-Mode, an optional software product that enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network.

N

NDMP	Network Data Management Protocol. A protocol that allows storage systems to communicate with backup applications and provides capabilities for controlling the robotics of multiple tape backup devices.
network adapter	An Ethernet, FDDI, or ATM card.
network management station	See <i>NMS</i> .
NMS	Network Management Station. A host on a network that uses third-party network management application (SNMP manager) to process status and diagnostic information about a storage system.
normal mode	The state of a storage system when there is no takeover in the high-availability configuration.
null user	The Windows NT machine account used by applications to access remote data.
NVRAM cache	Nonvolatile RAM in a storage system, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a storage system or power failure.
NVRAM card	An adapter that contains the storage system's NVRAM cache.
NVRAM mirror	A synchronously updated copy of the contents of the storage system NVRAM (nonvolatile random access memory) contents kept on the partner storage system.
P	
panic	A serious error condition causing the storage system or V-Series system to halt. Similar to a software crash in the Windows system environment.
parity disk	The disk on which parity information is stored for a RAID4 disk drive array. In RAID groups using RAID-DP protection, two parity disks store the parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.
partner mode	The method you use to communicate through the command-line interface with a virtual storage system during a takeover.
partner node	From the point of view of the local node (storage system), the other node in a high-availability configuration.
PCI	Peripheral Component Interconnect. The bus architecture used in newer storage system models.
pcnfsd	A storage system daemon that permits PCs to mount storage system file systems. The corresponding PC client software is called <i>PC-NFS</i> .
POST	Power-on self-tests. The tests run by a storage system after the power is turned on.

PVC	Permanent Virtual Circuit. A link with a static route defined in advance, usually by manual setup.
Q	
qtree	A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.
R	
RAID	Redundant Array of Independent Disks. A technique that protects against disk failure by computing parity information based on the contents of all the disks in an array. Storage systems use either RAID4, which stores all parity information on a single disk, or RAID-DP, which stores all parity information on two disks.
RAID disk scrubbing	The process in which a system reads each disk in the RAID group and tries to fix media errors by rewriting the data to another disk area.
S	
SCSI adapter	An expansion card that supports SCSI disk drives and tape drives.
SCSI address	The full address of a disk, consisting of the disk's SCSI adapter number and the disk's SCSI ID, such as 9a.1.
SCSI ID	The number of a disk drive on a SCSI chain (0 to 6).
serial adapter	An expansion card for attaching a terminal as the console on some storage system models.
serial console	An ASCII or ANSI terminal attached to a storage system's serial port. Used to monitor and manage storage system operations.
share	A directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known as a <i>CIFS share</i> .
SID	Security identifier used by the Windows operating system.
Snapshot copy	An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copies enable users to restore files and to back up the storage system to tape while the storage system is in use.
SVC	Switched Virtual Circuit. A connection established through signaling. The user defines the endpoints when the call is initiated.
system board	A printed circuit board that contains a storage system's CPU, expansion bus slots, and system memory.

T

takeover	The emulation of the failed node identity by the takeover node in a high-availability configuration; the opposite of <i>giveback</i> .
takeover mode	The method you use to interact with a node (storage system) when it has taken over its partner. The console prompt indicates when the node is in takeover mode.
takeover node	A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.
trap	An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the storage system.
tree quota	A type of disk quota that restricts the disk usage of a directory created by the quota qtree command. Different from user and group quotas that restrict disk usage by files with a given UID or GID.
U	
UID	user identification number.
Unicode	A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.
V	
VCI	Virtual Channel Identifier. A unique numerical tag defined by a 16-bit field in the ATM cell header that identifies a virtual channel over which the cell is to travel.
vFiler unit	In Data ONTAP operating in 7-Mode, a virtual storage system that you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage systems on the network.
VGA adapter	An expansion card for attaching a VGA terminal as the console.
volume	<ul style="list-style-type: none"> For Data ONTAP, a logical entity that holds user data that is accessible through one or more of the supported access protocols, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Fibre Channel (FC), and Internet SCSI (iSCSI). V-Series treats an IBM volume as a disk. For IBM, the area on the storage array that is available for a V-Series system or non V-Series host to read data from or write data to. The V-Series documentation uses the term <i>array LUN</i> to describe this area.
VPI	Virtual Path Identifier. An eight-bit field in the ATM cell header that indicates the virtual path over which the cell should be routed.

W**WAFL**

Write Anywhere File Layout. A file system designed for the storage system to optimize write performance.

WINS

Windows Internet Name Service.

workgroup

A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

Copyright information

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- /etc/ad directory [128](#)
- /etc/cifs_nbalias.cfg
 - creating NetBIOS aliases in [143](#)
- /etc/exports
 - editing [16](#)
- /etc/httpd.access
 - editing [341](#)
- /etc/httpd.group [343](#)
- /etc/httpd.passwd
 - creating and editing [343](#)
- /etc/httpd.translation
 - adding fail rules [336](#)
- /etc/httpd.translations
 - adding map rules [335](#)
 - adding pass rules [336](#)
 - adding redirect rules [335](#)
 - defined [334](#)
- /etc/log/ftp.cmd
 - log file format [313](#)
- /etc/log/ftp.xfer
 - log file format [313](#)
- /etc/log/httpd.log
 - changing format [348](#)
- /etc/nsswitch.conf
 - adding LDAP entry [260](#)
 - enabling LDAP client authentication [262](#)
- /etc/symlink.translations
 - creating map entries [295](#)
 - creating widelink entries [296](#)
 - map entries [293](#)
 - redirecting absolute symbolic links [295](#)
 - using map and widelink entries [296](#)
 - widelink entries [293](#)
- /etc/usermap.cfg
 - direction [245](#)
 - increasing security [249](#)
 - interpreting domain names [246](#)
 - IP_qualifier [244](#)
 - mapping user names [249](#)
 - mapping Windows accounts to root [250](#)
 - restricting NFS access using [250](#)
 - sample entries [247](#)
 - specifying entries [244](#)
 - UNIX name [246](#)
 - Windows name [245](#)

- 8.3-format file names
 - creating [233](#)

A

- access
 - FTP, restricting [310](#)
 - restricting using HTTP options [338](#)
 - why allowed or denied [156](#)
- access cache
 - adding entries [22](#)
 - explained [21](#)
 - optimizing performance [24](#)
 - removing entries [23](#)
 - setting timeout values [24](#)
 - viewing statistics [23](#)
- access control
 - troubleshooting [154](#)
- access control lists
 - See* ACLs
- access control lists (ACLs)
 - compatibility between NFSv4 and NTFS [44](#)
 - file-level, displaying and changing [92](#)
 - managing [85](#)
 - NFSv4, benefits of enabling [44](#)
 - NFSv4, managing [42](#)
 - share-level, adding users or groups from MMC [86](#)
 - share-level, changing from CLI [90](#)
 - share-level, displaying and changing [86](#)
 - share-level, displaying and changing from MMC [88](#)
 - share-level, removing users or groups using CLI [91](#)
 - share-level, removing users or groups using MMC [89](#)
 - share-level, specifying group IDs [94](#)
- access control lists)
 - See* ACLs
- access-based enumeration
 - defined [83](#)
 - enabling or disabling [83](#)
 - executing commands from Windows clients [84](#)
- accounts
 - adding, displaying, and removing local user [117](#)
 - local users, limitations of [117](#)
 - machine, preventing data access [142](#)

- machine, using for access in Kerberos environments [141](#)
- ACEs
 - relationship to ACLs and inheritance [85](#)
- ACL permissions
 - NFSv3/v4 clients, displaying [91](#)
- ACLs
 - enabling or disabling NFSv4 [44](#)
 - how inheritance works [85](#)
 - NFSv4, how they work [42](#)
- ACLs (access control lists)
 - compatibility between NFSv4 and NTFS [44](#)
 - file-level, displaying and changing [92](#)
 - managing [85](#)
 - NFSv4, benefits of enabling [44](#)
 - NFSv4, managing [42](#)
 - share-level, adding users or groups from MMC [86](#)
 - share-level, changing from CLI [90](#)
 - share-level, displaying and changing [86](#)
 - share-level, displaying and changing from MMC [88](#)
 - share-level, removing users or groups using CLI [91](#)
 - share-level, removing users or groups using MMC [89](#)
 - share-level, specifying group IDs [94](#)
- ACLs)
 - how Data ONTAP uses share-level [85](#)
- Active Directory
 - LDAP lookup services, enabling [265](#)
 - LDAP servers, connection pooling and selection [266](#)
 - LDAP servers, managing [265](#)
 - LDAP servers, monitoring connections [266](#)
 - LDAP servers, requirements [265](#)
 - LDAP servers, troubleshooting connections [266](#)
 - LDAP servers, using [265](#)
 - simple binds [267](#)
- adding
 - HTTP fail rules [336](#)
 - HTTP map rules [335](#)
 - HTTP pass rules [336](#)
 - HTTP redirect rules [335](#)
 - mapping entries to WAFL credential cache [300](#)
 - users to local groups from MMC [119](#)
- aliases
 - NetBIOS, creating [142](#)
 - NetBIOS, creating from CLI [142](#)
 - NetBIOS, creating in /etc/cifs_nbalias.cfg [143](#)
 - NetBIOS, displaying [143](#)
- APIs
 - supported VMware vStorage, for NFS [38](#)
- audit events
 - saving [283](#)
- audit policies
 - mapping [355](#)
 - mapping values [356](#)
- auditing
 - CIFS, configuring [277](#)
 - clearing events [281](#)
 - clearing internal audit log file [287](#)
 - configuring automatic saves by log file size [283](#)
 - defined [274](#)
 - displaying events [287, 288](#)
 - enabling automatic saves by time interval [284](#)
 - event log location [281](#)
 - failed access attempts [290](#)
 - lost record events [290](#)
 - NFS, controlling events with filter file [279](#)
 - NFS, enabling [279](#)
 - NFS, specifying events [278](#)
 - prerequisites for configuring Data ONTAP for NFS auditing [278](#)
 - saving events [281](#)
 - saving events manually [282](#)
 - SNMP traps for events [287](#)
 - specifying log counter extensions [285](#)
 - specifying log timestamp extensions [285](#)
 - specifying maximum auto save files [286](#)
 - specifying maximum size of internal audit log file [286](#)
 - UNIX file access details [290](#)
 - updating event logs [281](#)
 - viewing events [288](#)
 - viewing static event logs [289](#)
 - Windows file access details [289](#)
- authentication
 - FTP, specifying [308](#)
 - Kerberos [134](#)
 - managing for clients [262](#)
 - method, displaying [117](#)
 - methods of HTTP access [339](#)
 - NTLM, limitations of [309](#)
 - UNIX [133](#)
 - UNIX client, enabling for LDAP [262](#)
 - Windows client, enabling for LDAP [262](#)
 - Windows workgroup [133](#)
- authentication style
 - SFTP, specifying [321](#)
- authentication styles
 - SFTP [321](#)
- authentication-based
 - restrictions [13](#)

authorization
managing for clients [262](#)

B

basic concepts
BranchCache [105](#)

boundary checking
enabling or disabling for symbolic links from shares [79](#)

BranchCache
about using to cache CIFS shares at branch offices [103](#)
basic concepts [105](#)
changing server secret key [114](#)
configuring on the storage system [108](#)
considerations when enabling on CIFS shares [110](#)
creating enabled CIFS shares [111](#)
Data ONTAP version requirements [108](#)
disabling on existing CIFS shares [116](#)
displaying hash statistics by file size [113](#)
displaying number of hashes flushed [114](#)
Distributed Cache mode [104](#)
enabling on a CIFS share [111](#)
enabling on existing CIFS shares [111](#)
hash statistics [113](#)
Hosted Cache mode [104](#)
hosted cache server [104](#)
how clients retrieve content with [106](#)
how Data ONTAP implements [106](#)
how Data ONTAP manages hashes [112](#)
introduction to enabling on CIFS share [110](#)
modes [104](#)
network protocol support requirements [108](#)
reasons for invalidating hashes [112](#)
terms [105](#)
what happens when you disable [115](#)
where to get information about configuring clients and cache server [109](#)
Windows hosts version requirements [108](#)

breaking locks
based on file [237](#)
based on network family [240](#)
based on protocol [240](#)
based on specified host [238](#)
based on specified owner [239](#)
reasons for [236](#)

browsing
enabling or disabling [81](#)

bypass traverse checking

SFTP, enabling or disabling [322](#)

C

cache server
where to get BranchCache configuration information [109](#)

caches
flushing from name server database [50](#)
how name server database works [50](#)

caching
enabling or disabling [82](#)
setting client-side properties [83](#)

case-sensitivity
of file names [232](#)

character mapping
clearing from volumes [235](#)

character restrictions
for file names [235](#)

character translation
enabling for file names [234](#)

CIFS
about using BranchCache for caching at branch offices [103](#)
auditing, configuring [277](#)
client events [175](#)
clients, optimizing NFS directory access [297](#)
configuring [59](#)
configuring shutdown messages [149](#)
consideration when enabling BranchCache on shares [110](#)
considerations when reconfiguring [64](#)
controlling access to symbolic links [291](#)
disabling [148](#)
disabling BranchCache on existing shares [116](#)
enabling BranchCache shares [111](#)
enabling clients to follow symbolic links [291](#)
file locking [236](#)
file naming dependencies [232](#)
file sharing with NFS [232](#)
files, accessing from NFS clients [299](#)
giving clients permission to run .dll and .exe files [305](#)
how Data ONTAP uses share-level ACLs [85](#)
introduction to enabling BranchCache on shares [110](#)
Kerberos authentication [134](#)
license, configuring [57](#)
modifying existing share for BranchCache [111](#)
monitored events [166](#)
monitoring activity [144](#)

- preventing clients from creating uppercase file names [299](#)
 - read-only bits explained [241](#)
 - reconfiguring [65](#)
 - recovering from issues when creating share comments [74](#)
 - resource limitations [147](#)
 - restarting service [150](#)
 - server name requirements [59](#)
 - share naming conventions [74](#)
 - specifying how clients interact with symbolic links [292](#)
 - stopping server screening for disconnected requests [173](#)
 - supported clients [58](#)
 - supported domain controllers [58](#)
 - tracing logins [304](#)
 - users, obtaining UNIX credentials for [242](#)
 - what the cifs setup command does [58](#)
 - what you need to know when creating shares [73](#)
 - CIFS sessions
 - cumulative [230](#)
 - over IPv4 [228](#)
 - over IPv6 [228](#)
 - cifs setup
 - server name requirements [59](#)
 - cifs setup command
 - tasks you can perform with [58](#)
 - CLI
 - fpolicy ext[ension] inc[lude] add *PolicyName* ext-list [203–206](#)
 - clients
 - CIFS, supported [58](#)
 - disconnecting using MMC [147](#)
 - managing authentication and authorization [262](#)
 - where to get BranchCache configuration information [109](#)
 - commands
 - tasks with cifs setup [58](#)
 - computer account passwords
 - configuring automatic changing of on storage system [152](#)
 - manually changing on storage system [151](#)
 - configuration files
 - creating and copying Kerberos [30](#)
 - configurations
 - considerations when enabling BranchCache on CIFS shares [110](#)
 - configuring
 - BranchCache on the storage system [108](#)
 - CIFS [59](#)
 - HTTP MIME types [333](#)
 - HTTP requests [334](#)
 - licenses [57](#)
 - SFTP connection limit [325](#)
 - connections
 - SFTP, limit of [325](#)
 - creating
 - BranchCache enabled CIFS shares [111](#)
 - CIFS shares from CLI [75](#)
 - directories in a home directory path [100](#)
 - file names [233](#)
 - map entries [295](#)
 - Widelink entries [296](#)
 - credential cache
 - WAFL, adding mapping entries to [300](#)
 - WAFL, deleting mapping entries from [300](#)
 - WAFL, monitoring statistics [302](#)
 - credentials
 - UNIX, managing for CIFS clients [242](#)
 - UNIX, obtaining for CIFS users [242](#)
 - UNIX, specifying CIFS users [243](#)
- ## D
- databases
 - flushing cache from name server [50](#)
 - deleting
 - files with the read-only bit set [241](#)
 - servers from the prefcd list [138](#)
 - shares from the CLI [84](#)
 - descriptions
 - displaying and changing [151](#)
 - directories
 - converting to Unicode format [298](#)
 - creating Unicode-formatted [297](#)
 - displaying security settings for [273](#)
 - FTP, restricting users [311](#)
 - matching with a user [96](#)
 - specifying permissions for newly created [80](#)
 - directory access
 - NFS, optimizing for CIFS clients [297](#)
 - directory create operations
 - configuring FPolicy to monitor [191](#)
 - directory create request monitoring
 - defined [190](#)
 - directory create requests
 - registering FPolicy to monitor [191](#)
 - directory delete operations
 - configuring FPolicy to monitor [188, 189](#)

- directory delete request monitoring
 - defined [188](#)
 - directory delete requests
 - registering FPolicy to monitor [189](#)
 - directory events [176](#)
 - directory operations [176](#)
 - directory rename operations
 - configuring FPolicy to monitor [190](#)
 - directory rename request monitoring
 - defined [189](#)
 - directory rename requests
 - registering FPolicy to monitor [190](#)
 - disable
 - BranchCache, what happens when you [115](#)
 - disabling
 - CIFS [148](#)
 - fencing [19](#)
 - FPolicy [168](#)
 - FTP server [307](#)
 - FTPS, explicit [327](#)
 - HTTP [331](#)
 - implicit FTPS [328](#)
 - lease oplocks [130](#)
 - NFS [15](#)
 - NFSv2 [37](#)
 - NFSv3 [37](#)
 - NFSv4 [42](#)
 - oplocks [130](#)
 - SMB 2.x [67](#)
 - SSL for LDAP [259](#)
 - TFTP server [307](#)
 - the default UNIX user account [252](#)
 - the Windows guest user account [253](#)
 - vStorage over NFS [38](#)
 - WebDAV [353](#)
 - disconnecting
 - clients using the MMC [147](#)
 - users from CLI [148](#)
 - displaying
 - BranchCache hash statistics by file size [113](#)
 - durable handle information [68](#)
 - export options for file system paths [21](#)
 - file system paths [20](#)
 - Group Policy Objects (GPOs) [126](#)
 - HTTP server connection information [347](#)
 - HTTP server statistics [344](#)
 - information about file policies [171](#)
 - lease oplock status [131](#)
 - NFS statistics [36](#)
 - permission tracing filters [155](#)
 - preferred domain controllers and LDAP servers [139](#)
 - UNIX file access details [290](#)
 - Distributed Cache mode
 - benefits of [104](#)
 - BranchCache [104](#)
 - how clients retrieve content in [106](#)
 - domain controllers
 - CIFS, supported [58](#)
 - ensuring successful authentication with [137](#)
 - tracing connections [304](#)
 - domains
 - changing storage system [60](#)
 - displaying preferred controllers [139](#)
 - reestablishing connection with [139](#)
 - selecting controllers [135](#)
 - specifying preferred controllers [137](#)
 - understanding the controller discovery process [135](#)
 - dot files
 - displaying on CIFS clients [233](#)
 - durable handles
 - definition of [68](#)
 - monitoring status of [68](#)
- ## E
- enabling
 - fencing [19](#)
 - file name character translation [234](#)
 - FPolicy [168](#)
 - FTP server [307](#)
 - FTPS, explicit [327](#)
 - HTTP [331](#)
 - implicit FTPS [328](#)
 - Kerberos for NFS [32](#)
 - lease oplocks [130](#)
 - NFS [15](#)
 - NFSv2 [37](#)
 - NFSv3 [37](#)
 - NFSv4 [42](#)
 - oplocks [130](#)
 - SMB 2.x [67](#)
 - SSL for LDAP [259](#)
 - TFTP server [307](#)
 - the default UNIX user account [252](#)
 - the Windows guest user account [253](#)
 - vStorage over NFS [38](#)
 - WebDAV [353](#)
 - enforcing
 - SMB 2.0 signing [72](#)
 - enumeration

- access-based, defined [83](#)
- access-based, enabling or disabling [83](#)
- access-based, executing commands from Windows clients [84](#)
- error messages
 - FPolicy [218–226](#)
 - fpolicy.server.fqdn.unavail severity="ERR" [222](#)
- event log
 - external, specifying location [282](#)
- event logs
 - mapping [355](#)
 - mapping values [355](#)
 - updating [281](#)
- events
 - audit, saving and clearing [281](#)
 - auditable [274](#)
 - monitored through CIFS [166](#)
 - monitored through NFS [167](#)
 - saving manually to audit event log [282](#)
 - screened for NFS and CIFS clients [175](#)
 - system access, auditing [274](#)
 - system, configuring auditing [276](#)
- explicit
 - FTPS, allowing or preventing [328](#)
- explicit FTPS
 - differences between implicit [326](#)
- export options
 - displaying for file system paths [21](#)
- exporting
 - file system paths [17, 18](#)
- extensions
 - adding to exclude list [204](#)
 - adding to include list [203](#)
 - adding to include or exclude list [203](#)
 - displaying [203](#)
 - displaying list of, in exclude list [203](#)
 - displaying list of, in include list [203](#)
 - removing from exclude list [205](#)
 - removing from include list [204](#)
 - removing from include or exclude list [204](#)
 - resetting include list [206](#)
 - resetting include or exclude list [206](#)
 - resetting the exclude list [206](#)
 - screening with wildcards [202](#)
 - setting exclude list [206](#)
 - setting include list [205](#)
 - setting or replacing include or exclude list [205](#)

F

- features
 - unsupported Windows [57](#)
- fencing
 - enabling or disabling [19](#)
- file access
 - how Data ONTAP controls [13](#)
- file and record locking
 - NFSv4, described [49](#)
- file blocking
 - configuring native [165](#)
- file close operations
 - configuring FPolicy to monitor [180](#)
- file close request monitoring
 - defined [179](#)
- file close requests
 - registering FPolicy for monitoring [180](#)
- file create operations
 - configuring FPolicy to monitor [178, 179](#)
- file create requests
 - registering FPolicy to monitor [179](#)
- file delegations
 - enabling or disabling NFSv4 read [46](#)
 - enabling or disabling NFSv4 write [46](#)
 - how they work for NFSv4 [45](#)
 - NFSv4, managing [45](#)
- file delete operations
 - configuring FPolicy to monitor [182, 183](#)
- file delete request monitoring
 - defined [182](#)
- file delete requests
 - registering FPolicy to monitor [183](#)
- file events [176](#)
- file link operations
 - configuring FPolicy to monitor [186](#)
- file link requests
 - registering FPolicy to monitor [186](#)
- file locking
 - explained [236](#)
 - SFTP, enabling or disabling [320](#)
- file locks
 - breaking [236](#)
 - breaking based on file [237](#)
 - breaking based on specified host [238](#)
 - breaking based on specified owner [239](#)
- file lookup operations
 - configuring FPolicy to monitor [192](#)
- file lookup request monitoring
 - defined [192](#)

- file lookup requests
 - registering FPolicy to monitor [192](#)
- file management
 - using Windows administrative tools [153](#)
- file names
 - case-sensitivity [232](#)
 - character restrictions [235](#)
 - creating [233](#)
 - creating in lowercase [233](#)
 - dependencies for NFS and CIFS [232](#)
 - enabling character translation [234](#)
 - preventing CIFS clients from creating uppercase [299](#)
 - valid characters for [232](#)
- file open operations
 - configuring FPolicy to monitor [177](#)
- file open request monitoring
 - defined [176](#)
- file open requests
 - registering FPolicy to monitor [178](#)
- file operations [176](#)
- file policies
 - creating [169](#)
 - destroying [172](#)
 - disabling [172](#)
 - displaying information [171](#)
 - enabling [170](#)
- file read operations
 - configuring FPolicy to monitor [185](#)
 - registering FPolicy to monitor [185](#)
- file read request monitoring
 - defined [184](#)
- file rename operations
 - configuring FPolicy to monitor [181](#)
- file rename request monitoring
 - defined [180](#)
- file rename requests
 - registering FPolicy to monitor [182](#)
- file request monitoring
 - defined [178](#)
- file screening
 - disabling server connection [207](#)
 - displaying server information [207](#)
 - specifying [170](#)
- file screening server
 - managing [207](#)
- file sharing
 - between NFS and CIFS [232](#)
- file symlink operations
 - configuring FPolicy to monitor [187](#)
- file symlink requests
 - registering FPolicy to monitor [188](#)
- File System ID (FSID) [39](#)
- file system paths
 - displaying [20](#)
 - displaying export options [21](#)
 - enabling or disabling fencing [19](#)
 - exporting [17, 18](#)
 - exporting or unexporting [16](#)
 - synchronizing [19](#)
 - unexporting [18, 19](#)
- file write operations
 - configuring FPolicy to monitor [184](#)
- file write request monitoring
 - defined [183](#)
- file write requests
 - registering FPolicy to monitor [184](#)
- file-based
 - restrictions [13](#)
- files
 - accessing over FTP [307](#)
 - accessing over HTTP [331](#)
 - accessing over WebDAV [352](#)
 - audit, access details [289, 290](#)
 - avoiding symbolic links [292](#)
 - displaying security settings for [273](#)
 - enabling or disabling SFTP log [324](#)
 - failed access attempts [290](#)
 - lost record events [290](#)
 - specifying maximum number of SFTP log [324](#)
 - specifying permissions for newly created [80](#)
- filter file
 - controlling NFS audit events [279](#)
- firewall
 - virtual HTTP, using [339](#)
- forcegroup option
 - about using the [75](#)
- FPolicy
 - about [157](#)
 - adding operations to monitor [209](#)
 - assigning secondary servers [208](#)
 - CLI commands [211](#)
 - communicating with the storage environment [161](#)
 - defined [157](#)
 - enabling or disabling [168](#)
 - error messages
 - fpolicy.server.fqdn.unavail severity="ERR" [222](#)
 - Frequently asked questions [213](#)
 - Frequently asked questions, access [215](#)
 - Frequently asked questions, file screening [216](#)
 - Frequently asked questions, general [213](#)

Frequently asked questions, performance [215](#)

Frequently asked questions, server [217](#)

introduction [157](#)

limitations [162](#)

monitoring operations [209, 210](#)

registering to monitor file create requests [179](#)

removing operations to monitor [210](#)

removing secondary servers [209](#)

secondary servers, defined [208](#)

setting up [168](#)

using [159, 168](#)

warning messages [222](#)

work flowchart [159](#)

fpolicy.fscreen.cfg.pCreateErr severity="ERR" [221](#)

fpolicy.fscreen.cfg.pCreateInfo severity="WARNING" [225](#)

fpolicy.fscreen.request.pathError severity="ERR" [221](#)

fpolicy.fscreen.server.closeError severity="ERR" [218](#)

fpolicy.fscreen.server.completionInconsistent
severity="WARNING" [223](#)

fpolicy.fscreen.server.completionRequestLost
severity="WARNING" [223](#)

fpolicy.fscreen.server.completionUnexpectedState
severity="ERR" [220](#)

fpolicy.fscreen.server.connectedNone
severity="WARNING" [222](#)

fpolicy.fscreen.server.connectError severity="ERR" [218](#)

fpolicy.fscreen.server.connecting.badOperationList
severity="WARNING" [224](#)

fpolicy.fscreen.server.connecting.badParameter
severity="WARNING" [224](#)

fpolicy.fscreen.server.connecting.internalError
severity="ERR" [220](#)

fpolicy.fscreen.server.connecting.privError
severity="ERR" [221](#)

fpolicy.fscreen.server.droppedConn
severity="WARNING" [225](#)

fpolicy.fscreen.server.pingRejected severity="ERR" [219](#)

fpolicy.fscreen.server.requestError severity="ERR" [219](#)

fpolicy.fscreen.server.requestRejected severity="ERR" [219](#)

fpolicy.fscreen.server.requestStatusError severity="ERR" [220](#)

fpolicy.fscreen.server.requestTO severity="ERR" [222](#)

fpolicy.fscreen.server.unexpectedFileDataResponse
severity="WARNING" [226](#)

fpolicy.fscreen.vol.i2p.off severity="WARNING" [225](#)

fpolicy.srv.conn.badOptParam severity="WARNING" [224](#)

Frequently asked questions

FPolicy [213](#)

FPolicy server [217](#)

FPolicy, access [215](#)

FPolicy, file screening [216](#)

FPolicy, general [213](#)

FPolicy, performance [215](#)

fsecurity

creating and applying security jobs [271](#)

defined [269](#)

enabling Storage-Level Access Guard [268](#)

FSID (File System ID) [39](#)

FTP

accessing files [307](#)

anonymous access, enabling or disabling [318](#)

authentication style, specifying [308](#)

blocking users [310](#)

bypassing traverse checking, enabling or disabling [310](#)

file locking, enabling or disabling [308](#)

managing [307](#)

managing anonymous access [318](#)

managing log files [312](#)

resetting statistics [316](#)

restricting access [310](#)

restricting users [311](#)

server, enabling or disabling [307](#)

setting connection threshold [317](#)

SNMP traps generated by [314](#)

specifying anonymous user name [318](#)

specifying home directory for anonymous users [319](#)

specifying idle timeout [317](#)

specifying maximum number of connections [316](#)

specifying maximum number of log files [314](#)

specifying maximum size of log files [314](#)

specifying TCP window size [317](#)

viewing log files [313](#)

viewing SNMP traps [314](#)

viewing statistics [315](#)

FTP connections

over IPv6 [330](#)

FTPS

differences between implicit and explicit [326](#)

explicit, allowing or preventing [328](#)

explicit, enabling or disabling [327](#)

implicit, enabling or disabling [328](#)

G

get attributes operations

configuring FPolicy to monitor [193](#)

- get attributes request monitoring
 - defined [193](#)
- get attributes requests
 - registering FPolicy to monitor [194](#)
- GIDs
 - mapping UNIX user names [251](#)
- glossary [358](#)
- GPOs
 - applying [120](#)
- GPOs (Group Policy Objects)
 - creating file system security [122](#)
 - displaying [126](#)
 - enabling or disabling support for [122](#)
 - managing [122](#)
 - requirements for using [121](#)
 - troubleshooting update problems [127](#)
 - updating settings [126](#)
- group ID formats
 - differences in NFSv3 and v4 [51](#)
- group IDs
 - allowing NFSv4 as numeric strings [51](#)
 - disallowing NFSv4 as numeric strings [51](#)
- Group Policy Objects
 - See* GPOs
- Group Policy Objects (GPOs)
 - creating file system security [122](#)
 - displaying [126](#)
 - enabling or disabling support for [122](#)
 - event logs and audit policies mapping [355](#)
 - managing [122](#)
 - requirements for using [121](#)
 - troubleshooting update problems [127](#)
 - updating settings [126](#)
- groups
 - adding, displaying, and removing local from CLI [118](#)
 - local, adding from MMC [118](#)
 - local, adding users to from MMC [119](#)
 - local, managing [118](#)
 - local, removing using MMC [120](#)
 - local, working with SnapMirror [120](#)

H

- hard mounts [36](#)
- hash statistics
 - BranchCache [113](#)
- hash timeout value
 - determining when to increase [114](#)
- hashes

- displaying number of flushed BranchCache [114](#)
- how Data ONTAP manages BranchCache [112](#)
- reasons for invalidating [112](#)
- home directories
 - CIFS, accessing using share aliases [102](#)
 - creating directories in [99](#), [100](#)
 - creating subdirectories when using extensions [101](#)
 - defined [95](#)
 - disabling [103](#)
 - displaying paths [98](#)
 - enabling access from other users [102](#)
 - managing [95](#)
 - specifying naming style [99](#)
 - specifying paths [97](#)
 - specifying the SFTP override path for [323](#)
 - specifying using UNC, syntax [101](#)
 - WebDAV, accessing [353](#)
- home directory restrictions
 - SFTP, enabling or disabling [322](#)
- Hosted Cache mode
 - benefits [104](#)
 - BranchCache [104](#)
 - how clients retrieve content in [106](#)
- HTTP
 - adding fail rules [336](#)
 - adding map rules [335](#)
 - adding pass rules [336](#)
 - adding redirect rules [335](#)
 - basic authentication [340](#)
 - built-in server, enabling or disabling [331](#)
 - built-in server, managing [331](#)
 - bypassing of traverse checking, enabling or disabling [331](#)
 - changing /etc/log/httpd.log format [348](#)
 - configuring MIME types [333](#)
 - configuring requests [334](#)
 - configuring virtual hosting [344](#)
 - creating and editing /etc/httpd.group [343](#)
 - creating and editing /etc/httpd.passwd [343](#)
 - detailed statistics [345](#)
 - editing /etc/httpd.access [341](#)
 - error statistics [346](#)
 - file access [331](#)
 - methods of access authentication [339](#)
 - NTLM authentication for directory security [340](#)
 - request statistics [345](#)
 - resetting statistics [347](#)
 - restricting access using [338](#)
 - service statistics [346](#)
 - specifying maximum log file size [332](#)

- specifying root directory [332](#)
- testing server [332](#)
- timeout statistics [346](#)
- translations file [334](#)
- using a virtual firewall [339](#)
- using third-party server [348](#)
- viewing connection information [347](#)

HTTP connections

- over IPv4 [350](#)
- over IPv6 [350](#)

HTTP server

- displaying statistics [344](#)

I

ID formats

- differences in NFSv3 and v4 user and group [51](#)

idle

- sessions, timing out [145](#)

idle timeout

- FTP, specifying [317](#)

implicit

- FTPS, enabling or disabling [328](#)

implicit FTPS

- differences between explicit [326](#)

include extensions lists

- setting [205](#)

inheritance

- and ACEs and ACLs [85](#)

IP_qualifier

- explained [244](#)

IPv6

- disable over FTP [329](#)
- disabling NFS over [56](#)
- disabling over CIFS [228](#)
- disabling over HTTP [349](#)
- enable over FTP [329](#)
- enabling NFS over [56](#)
- enabling over CIFS [228](#)
- enabling over HTTP [349](#)
- NFS support for [56](#)
- support for CIFS [227](#)
- support for FTP [329](#)
- support for HTTP [349](#)
- textual representation of addresses [56](#)

J

job definition file

- for Storage-Level Access Guard, generating and editing [269](#)
- managing with secedit utility [270](#)
- managing with text editor [270](#)
- specifying elements [270](#)

K

Kerberos

- authentication [134](#)
- configuring for NFS [26](#), [28](#)
- creating and copying configuration files [30](#)
- enabling [25](#)
- enabling for NFS [32](#)
- NFS clients that support v5 security services [26](#)
- preventing passive replay attacks [135](#)
- rules for configuring for NFS [29](#)
- supported principal formats [34](#)
- using machine accounts for access [141](#)

keytab files

- configuring Kerberos [30](#)

keywords

- screening operations [226](#)

krb5.conf files

- configuring Kerberos [30](#)

L

LDAP

- Active Directory lookup services, enabling [265](#)
- Active Directory servers, connection pooling and selection [266](#)
- Active Directory servers, managing [265](#)
- Active Directory servers, monitoring connections [266](#)
- Active Directory servers, requirements [265](#)
- Active Directory servers, troubleshooting connections [266](#)
- Active directory servers, using [265](#)
- adding entry to /etc/nsswitch.conf [260](#)
- configuring [256](#)
- default schema [267](#)
- displaying preferred servers [139](#)
- enabling authorization for NFS file access from Windows clients [263](#)
- enabling authorization for NTFS or mixed file system access from UNIX clients [263](#)
- enabling or disabling [259](#)
- enabling or disabling SSL [259](#)
- enabling UNIX client authentication [262](#)

- enabling Windows client authentication [262](#)
- installing SSL root certificate [260](#)
- managing schema [267](#)
- mapping users [263](#)
- modifying schema options [268](#)
- selecting servers [135](#)
- server selection order [262](#)
- setting administrative password [261](#)
- simple binds [267](#)
- specifying administrative user names [261](#)
- specifying base and scope values [257](#)
- specifying ports [261](#)
- specifying preferred servers [137](#), [258](#)
- specifying search base and scope [257](#)
- specifying servers [258](#)
- user mapping, specifying base and scope values [264](#) using [256](#)
- lease oplocks
 - enabling or disabling [130](#)
 - enabling or disabling on qtrees [130](#)
 - how they provide enhanced locking [129](#)
 - improving client performance with [129](#)
 - monitoring status of [131](#)
- licenses
 - configuring [57](#)
 - enabling or disabling NFS [15](#)
- limitations
 - CIFS resources [147](#)
 - of Data ONTAP support for NFSv4 [40](#)
 - of local user accounts [117](#)
 - SFTP [320](#)
- link operations
 - configuring FPolicy to monitor [186](#)
- link requests
 - monitoring [186](#)
- lists
 - setting include extensions [205](#)
- Live View
 - configuring [280](#)
 - displaying audit events [288](#)
 - displaying events [287](#)
 - viewing events [288](#)
- local groups
 - adding, displaying, and removing from CLI [118](#)
- local user accounts
 - adding, displaying, and removing [117](#)
 - when to create [116](#)
- locking grace period
 - NFSv4, specifying [50](#)
- locking lease period

- specifying NFSv4 [49](#)
- locks
 - breaking based on file [237](#)
 - breaking based on network family [240](#)
 - breaking based on protocol [240](#)
 - breaking based on specified host [238](#)
 - breaking based on specified owner [239](#)
 - reasons for breaking [236](#)
- log
 - external event, specifying location [282](#)
- log files
 - enabling or disabling SFTP [324](#)
 - FTP, managing [312](#)
 - FTP, specifying maximum number of [314](#)
 - FTP, specifying maximum size of [314](#)
 - FTP, viewing [313](#)
 - HTTP, specifying maximum size of [332](#)
 - size and format [281](#)
 - specifying maximum number of SFTP [324](#)
 - specifying maximum size for SFTP [325](#)
- logins
 - CIFS, tracing [304](#)

M

- machine accounts
 - preventing data access [142](#)
 - using for access in Kerberos environments [141](#)
- managing
 - HTTP [331](#)
 - WebDAV [353](#)
- map cache
 - SID-to-name, clearing [255](#)
 - SID-to-name, enabling or disabling [254](#)
 - SID-to-name, managing [254](#)
- map entries
 - creating [295](#)
 - defined [293](#)
 - using [296](#)
- mapping
 - managing inconsistencies [303](#)
 - UNIX names to UIDs and GIDs [251](#)
 - users with LDAP [263](#)
 - Windows accounts to root [250](#)
- mapping entries
 - adding to WAFL credential cache [300](#)
 - configuring valid time [301](#)
 - deleting from WAFL credential cache [300](#)
 - SID-to-name, changing lifetime of [255](#)
- messages

- sending to users [150](#)
- MIME types
 - configuring [337](#)
 - HTTP, configuring [333](#)
- MMC
 - adding users or groups to share-level ACLs [86](#)
 - adding users to local group [119](#)
 - connecting to storage system [58](#)
 - deleting shares [84](#)
 - disconnecting clients [147](#)
 - displaying and changing share properties [76](#)
 - displaying and changing share-level ACLs [88](#)
 - removing local groups [120](#)
 - removing users or groups from share-level ACLs [89](#)
 - running the Share a Folder wizard [74](#)
- monitor list
 - adding operations [209](#)
 - removing operations [210](#)
- monitoring
 - durable handle status [68](#)
 - lease oplock status [131](#)
 - WAFL credential cache statistics [302](#)
- mount requests
 - controlling NFS, from nonreserved ports [36](#)
- mount service statistics
 - displaying [35](#)
- mountd requests
 - tracing [35](#)
- mounting problems
 - debugging [35](#)
- mountpoints
 - affected by NFSv4 pseudo-fs [41](#)
- mounts [36](#)
- multiple server configuration
 - defined [161](#)
- multiprotocol
 - changing, effects of [62](#)

N

- name server database cache
 - how it works [50](#)
- name server databases
 - flushing cache from [50](#)
- names
 - valid characters for file [232](#)
- naming styles
 - domain [99](#)
 - non-domain [100](#)
 - specifying for home directories [99](#)

- native file blocking
 - configuring [165](#)
 - defined [163](#)
 - using [164](#)
- NetBIOS
 - creating aliases [142](#)
 - creating aliases from CLI [142](#)
 - creating aliases in /etc/cifs_nbalias.cfg [143](#)
 - displaying aliases [143](#)
 - over TCP, disabling [143](#)
- NFS
 - auditing, enabling [279](#)
 - benefits of enabling v4 ACLs [44](#)
 - client events [175](#)
 - clients that support Kerberos v5 security services [26](#)
 - clients, accessing CIFS files [299](#)
 - compatibility between v4 and NTFS ACLs [44](#)
 - configuring Kerberos [26](#), [28](#)
 - controlling audit events [279](#)
 - controlling requests from nonreserved ports [36](#)
 - differences in v3 and v4 user and group ID formats [51](#)
 - displaying open delegation statistics, v4 [47](#), [48](#)
 - displaying statistics [36](#)
 - enabling Kerberos [25](#), [32](#)
 - enabling or disabling [15](#)
 - enabling or disabling v2 [37](#)
 - enabling or disabling v3 [37](#)
 - enabling or disabling v4 [42](#)
 - enabling or disabling v4 ACLs [44](#)
 - enabling or disabling v4 read file delegations [46](#)
 - enabling or disabling v4 write file delegations [46](#)
 - file access from Windows clients, enabling LDAP authorization [263](#)
 - file locking [236](#)
 - file naming dependencies [232](#)
 - file sharing with CIFS [232](#)
 - how v4 ACLs work [42](#)
 - managing v4 ACLs [42](#)
 - monitored events [167](#)
 - optimizing directory access for CIFS clients [297](#)
 - prerequisites for configuring Data ONTAP for auditing [278](#)
 - pseudo-fs affecting mountpoints [41](#)
 - read-only bits explained [241](#)
 - restricting user access [250](#)
 - rules for configuring Kerberos for [29](#)
 - specifying audit events [278](#)
 - specifying user ID domain for v4 [42](#)
 - support for, over IPv6 [56](#)

- supporting v4 clients [40](#)
- v3/v4 clients, displaying Windows ACL permissions [91](#)
- v4 file and record locking described [49](#)
- v4 support [40](#)
- v4, determining file deletion [44](#)
- v4, displaying open delegation statistics [47, 48](#)
- v4, how file delegations work [45](#)
- v4, limitations of Data ONTAP support [40](#)
- v4, managing file delegations [45](#)
- v4, specifying locking grace period [50](#)
- v4, specifying locking lease period [49](#)
- NFS clients
 - enabling or disabling fencing [19](#)
- NFSv4
 - allowing user and group IDs as numeric strings [51](#)
 - disallowing user and group IDs as numeric strings [51](#)
- NLM locks
 - breaking based on network family [240](#)
- NTFS
 - ACLs, compatibility with NFSv4 [44](#)
- NTLM
 - authentication, limitations of [309](#)
- NTLM authentication
 - HTTP directory security [340](#)
- null sessions
 - providing access [140](#)
 - using for access in non-Kerberos environments [140](#)
- null users
 - granting access to shares [141](#)

O

- ONTAPI [157](#)
 - open delegation statistics
 - displaying, NFSv4 [47, 48](#)
 - NFSv4, displaying [47, 48](#)
 - operations
 - adding to monitor list [209](#)
 - removing from monitor list [210](#)
 - setting or replacing list of monitored [210](#)
 - oplocks
 - changing delay time for sending breaks [132](#)
 - enabling or disabling [130](#)
 - enabling or disabling on qtrees [130](#)
 - improving client performance with [129](#)
 - write cache data-loss considerations [129](#)
 - optimizing
 - access cache performance [24](#)
 - optionsusing HTTP to restrict access [338](#)
 - organizational units (OUs)
 - associating with [121](#)
 - OUs (organizational units)
 - associating with [121](#)
- ## P
- passwords
 - administrative, setting for LDAP [261](#)
 - changing storage system computer account manually [151](#)
 - configuring automatic computer account changes [152](#)
 - path names
 - configuration requirements [128](#)
 - PC-NFS
 - creating user entries [53](#)
 - defining umask for files and directories [54](#)
 - supporting clients [52](#)
 - pcnfsd daemon
 - enabling or disabling [52](#)
 - explained [52](#)
 - performance
 - optimizing for access cache [24](#)
 - using lease oplocks to improve client [129](#)
 - using oplocks to improve client [129](#)
 - permissions
 - adding tracing filters [154](#)
 - displaying tracing filters [155](#)
 - removing tracing filters [155](#)
 - specifying for newly created files and directories [80](#)
 - ports
 - controlling NFS mount requests from nonreserved [36](#)
 - specifying for LDAP [261](#)
 - principals
 - supported client and user formats [34](#)
 - protecting
 - web pages, methods of [339](#)
 - protocol modes
 - changing [61](#)
 - protocols
 - disabling SMB 2.x [67](#)
 - enabling SMB 2.x [67](#)
 - support requirements for BranchCache [108](#)
 - supported [13](#)
 - supported SMB 2.1 enhancements [66](#)
 - pseudo-fs
 - NFSv4, affecting mountpoints [41](#)

Q

qtrees

- enabling or disabling lease oplocks on [130](#)
- enabling or disabling oplocks on [130](#)

queries

- statistics, saving and reusing [147](#)

R

read file delegations

- enabling or disabling NFSv4 [46](#)

read-only bits

- deleting files [241](#)
- explained [241](#)

read-only domain controllers

- ensuring successful authentication with [137](#)

Remote Procedure Calls (RPC) [157](#)

removing

- entries from access cache [23](#)
- Storage-Level Access Guards [273](#)

reports

- generating BranchCache hash statistics by file size [113](#)

requirements

- BranchCache network protocol support [108](#)
- CIFS server name [59](#)
- Data ONTAP version requirements for BranchCache [108](#)
- for using GPOs [121](#)
- Windows host version requirements for BranchCache [108](#)

resetting

- HTTP statistics [347](#)

resources

- CIFS, limitations [147](#)

restricting

- FTP users to directories [311](#)

restrictions

- authentication-based [13](#)
- file-based [13](#)

retrieving

- content with BranchCache [106](#)
- content with Distributed Cache mode [106](#)
- content with Hosted Cache mode [106](#)

root directory

- enabling for WebNFS [55](#)
- HTTP, specifying [332](#)
- setting for WebNFS [55](#)
- specifying name for WebNFS [55](#)

RPC (Remote Procedure Calls) [157](#)**S**

SACLs (System access control lists)

- setting [276](#)

screening

- by extension, defined [201](#)
- by volume, defined [195](#)
- using wildcards [196, 202](#)

screening operations

- keywords [226](#)

screening timeout

- setting [174](#)

scripts

- startup, shutdown [128](#)

secedit utility

- managing job definition file with [270](#)

secondary servers

- assigning [208](#)
- defined [208](#)
- removing [209](#)

security

- increasing for user access [249](#)

security jobs

- applying to storage objects [271](#)
- canceling [272](#)
- checking status of [272](#)
- creating from job definition file [271](#)

security level

- setting the minimum [134](#)

security settings

- displaying for files and directories [273](#)

server screening

- stopping for disconnected CIFS requests [173](#)

server secret key

- changing BranchCache [114](#)

server timeout

- setting [173](#)

servers

- connection pooling and selection [266](#)
- deleting from prefcd list [138](#)
- enabling or disabling NFS [15](#)
- LDAP, selection order [262](#)

sessions

- displaying information [145](#)
- idle, timing out [145](#)

set attributes operations

- configuring FPolicy to monitor [194, 195](#)

set attributes request monitoring

- defined [194](#)
- set attributes requests
 - registering FPolicy to monitor [195](#)
- SFTP
 - authentication style, specifying [321](#)
 - authentication styles [321](#)
 - bypass traverse checking, enabling or disabling [322](#)
 - connection limit [325](#)
 - enabling or disabling [320](#)
 - enabling or disabling log files [324](#)
 - file locking, enabling or disabling [320](#)
 - home directory restrictions, enabling or disabling [322](#)
 - idle timeout, specifying [326](#)
 - limitations [320](#)
 - specifying maximum number of log files [324](#)
 - specifying maximum size of log files [325](#)
 - specifying the override path for user home directories [323](#)
- share boundary checking
 - disabling for symbolic links [294](#)
- share comments
 - recovering from issues when creating CIFS [74](#)
- share properties
 - enabling or disabling virus scanning [81](#)
- shares
 - changing properties from CLI [78](#)
 - considerations when enabling BranchCache on CIFS [110](#)
 - creating from CLI [75](#)
 - creating from MMC on Windows clients [74](#)
 - deleting [84](#)
 - deleting from CLI [84](#)
 - deleting from MMC [84](#)
 - disabling BranchCache on existing CIFS [116](#)
 - displaying and changing properties [76](#)
 - displaying and changing properties from MMC [76](#)
 - displaying properties from CLI [78](#)
 - enabling BranchCache on CIFS [111](#)
 - enabling BranchCache when creating CIFS [111](#)
 - introduction to enabling BranchCache on CIFS [110](#)
 - naming conventions [74](#)
 - what you need to know when creating CIFS [73](#)
- shutdown messages
 - configuring for CIFS [149](#)
- SID-to-name map cache
 - clearing [255](#)
 - enabling or disabling [254](#)
 - managing [254](#)
- SID-to-name mapping entries
 - changing lifetime of [255](#)
- simple binds
 - Active Directory, LDAP [267](#)
- SMB
 - 2.0, enforcing signing [72](#)
 - enabling or disabling 2.x [67](#)
 - enabling signing [71](#)
 - how Data ONTAP uses share-level ACLs [85](#)
 - signing policies affecting communications [69](#)
 - signing, performance impact of [70](#)
 - signing, support for [69](#)
 - support for v1.0 [66](#)
 - support for v2.0 [66](#)
 - supported 2.1 enhancements [66](#)
 - when to enable v2.0 [67](#)
 - when to enable v2.1 [67](#)
- SMB 2.1
 - how lease oplocks enhance locking [129](#)
- SMB named pipe
 - enabling or disabling multiple open instances [175](#)
- SMBv2.x durable handles
 - definition of [68](#)
- SnapMirror
 - working with local groups [120](#)
- SNMP
 - configuring [315](#)
 - FTP, viewing traps [314](#)
 - starting [315](#)
 - traps for auditing events [287](#)
 - traps generated by FTP [314](#)
 - viewing traps on UNIX clients [315](#)
- soft mounts [36](#)
- specifying
 - FTP authentication style [308](#)
 - FTP idle timeout [317](#)
 - home directory for anonymous FTP users [319](#)
 - home directory paths [97](#)
 - HTTP server root directory [332](#)
 - LDAP administrative password [261](#)
 - LDAP administrative user names [261](#)
 - LDAP search base and scope [257](#)
 - LDAP server port [261](#)
 - maximum number of auto save files [286](#)
 - maximum number of FTP log files [314](#)
 - maximum size of cifsaudit.alf file [286](#)
 - maximum size of FTP log files [314](#)
 - maximum size of HTTP log files [332](#)
 - preferred domain controllers and LDAP servers [137](#)
 - preferred LDAP servers [258](#)
 - SFTP idle timeout [326](#)

SSL

- enabling or disabling for LDAP [259](#)
- installing root certificate for LDAP [260](#)

statistics

- displaying for VMware vStorage over NFS [39](#)
- displaying NFS [36](#)
- FTP, resetting [316](#)
- FTP, viewing [315](#)
- HTTP server, displaying [344](#)
- monitoring WAFL credential cache [302](#)
- saving and reusing queries [147](#)
- tracking [145](#)
- viewing [146](#)
- viewing for access cache [23](#)

storage systems

- changing computer account password manually [151](#)
- configuring automatic computer account password changes [152](#)
- configuring BranchCache on [108](#)

Storage-Level Access Guard

- enabling [268](#)
- generating and editing job definition file [269](#)

Storage-Level Access Guards

- removing [273](#)

strings

- allowing NFSv4 user and group IDs as numeric [51](#)
- disallowing NFSv4 user and group IDs as numeric [51](#)

supported protocols [13](#)

symbolic links

- absolute, redirecting [295](#)
- avoiding [292](#)
- controlling CIFS access to [291](#)
- disabling share boundary checking [294](#)
- enabling CIFS clients to follow [291](#)
- enabling or disabling boundary checking for [79](#)
- specifying how CIFS clients interact with [292](#)
- wide, enabling or disabling [79](#), [103](#)
- with home directories [97](#)

symlink operations

- configuring FPolicy to monitor [187](#)

symlink request monitoring

- defined [187](#)

symlink requests

- registering FPolicy to monitor [188](#)

synchronizing

- file system paths [19](#)

System access control lists (SACLs)

- setting [276](#)

system access events

- auditing [274](#)

system event auditing

- configuring [276](#)

T

TCP window size

- FTP, specifying [317](#)

testing

- HTTP server [332](#)

TFTP

- server, enabling or disabling [307](#)
- specifying maximum number of connections [316](#)

timeout

- SFTP, specifying [326](#)

timeout values

- setting for access cache [24](#)

timing out

- idle sessions [145](#)

tracing filters

- adding for permissions [154](#)
- displaying for permissions [155](#)
- removing for permissions [155](#)

traps

- SNMP, for auditing events [287](#)

U

UIDs

- mapping UNIX user names [251](#)

umask

- defining for files and directories of PC-NFS users [54](#)
- explained [53](#)

unexporting

- file system paths [18](#), [19](#)

Unicode

- converting directories [298](#)
- formatted, creating directories [297](#)

UNIX

- authentication [133](#)
- credentials, managing for CIFS clients [242](#)
- credentials, obtaining for CIFS users [242](#)
- credentials, specifying for CIFS users [243](#)
- enabling LDAP authorization for NTFS or mixed file system access [263](#)
- enabling or disabling default user accounts [252](#)
- file access details [290](#)
- LDAP-based client authentication, enabling [262](#)
- mapping user names to UIDs and GIDs [251](#)
- viewing SNMP traps [315](#)

- UNIX permissions
 - overriding, enabling or disabling [323](#)
- unsupported features
 - Windows [57](#)
- user account names
 - Windows, specifying [63](#)
- user accounts
 - adding, displaying, and removing local [117](#)
 - enabling or disabling Windows guest [253](#)
 - UNIX, enabling or disabling [252](#)
- user ID domains
 - specifying for NFSv4 [42](#)
- user ID formats
 - differences in NFSv3 and v4 [51](#)
- user IDs
 - allowing NFSv4 as numeric strings [51](#)
 - disallowing NFSv4 as numeric strings [51](#)
- user mapping
 - LDAP, specifying base and scope values [264](#)
- user names
 - administrative, specifying for LDAP [261](#)
 - mapping [249](#)
 - translating between Windows and UNIX [244](#)
 - UNIX, mapping to UIDs and GIDs [251](#)
- users
 - disconnecting from CLI [148](#)
 - FTP, blocking [310](#)
 - FTP, restricting [311](#)
 - limitations of local accounts [117](#)
 - local, managing [116](#)
 - null, granting access to shares [141](#)
 - restricting [311](#)
 - sending messages to [150](#)
 - specifying [144](#)

V

- virtual hosting
 - configuring [344](#)
- virus scanning
 - enabling or disabling on the share [81](#)
- VMware
 - displaying statistics for vStorage over NFS [39](#)
 - enabling or disabling vStorage over NFS [38](#)
 - supported vStorage APIs for NFS [38](#)
- volumes
 - adding to exclude list [198](#)
 - adding to include list [198](#)
 - adding to include or exclude list [198](#)
 - clearing character mapping [235](#)

- displaying [197](#)
- removing from exclude list [199](#)
- removing from include list [199](#)
- removing from include or exclude list [199](#)
- resetting exclude list [201](#)
- resetting include list [200](#)
- screening with wildcards [196](#)
- setting exclude list [200](#)
- setting include list [200](#)
- specifying or replacing in include or exclude list [200](#)

- vStorage
 - displaying statistics for VMware, over NFS [39](#)
 - enabling or disabling over NFS [38](#)
 - supported APIs for NFS [38](#)

W

- WAFL
 - credential cache, adding mapping entries to [300](#)
 - credential cache, deleting mapping entries from [300](#)
 - credential cache, monitoring statistics [302](#)
- warning messages
 - FPolicy [222](#)
- web pages
 - methods of restricting HTTP access [339](#)
- WebDAV
 - accessing home directories [353](#)
 - enabling or disabling [353](#)
 - explained [352](#)
 - file access [352](#)
 - managing [353](#)
 - using third-party server [354](#)
- WebNFS
 - enabling or disabling [54](#)
 - enabling root directory [55](#)
 - setting root directory [55](#)
 - specifying name of root directory [55](#)
 - supporting clients [54](#)
- widelink entries
 - creating [296](#)
 - defined [293](#)
 - using [296](#)
- widelinks
 - enabling or disabling [79](#), [103](#)
- wildcards
 - using for screening [196](#), [202](#)
- Windows
 - enabling LDAP authorization for NFS file access [263](#)
 - enabling or disabling guest user account [253](#)

- file access details [289](#)
- LDAP-based client authentication, enabling [262](#)
- specifying user account names [63](#)
- unsupported features [57](#)
- workgroup authentication [133](#)
- Windows applications, preserving UNIX permissions [305](#)
- Windows Server 2008 domain controllers
 - ensuring successful authentication with [137](#)
- WINS servers
 - specifying [59](#)
- write cache
 - data loss considerations when using oplocks [129](#)
- write file delegations
 - enabling or disabling NFSv4 [46](#)