

IBM System Storage N series



Clustered Data ONTAP 8.2 File Access Management Guide for NFS

Contents

Preface	10
About this guide	10
Supported features	10
Websites	10
Getting information, help, and service	11
Before you call	11
Using the documentation	11
Hardware service and support	12
Firmware updates	12
How to send your comments	12
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	14
LIF configuration requirements for file access management	14
How namespaces and volume junctions affect file access on SVMs with	
FlexVol volumes	15
What namespaces in SVMs with FlexVol volumes are	15
What volume junctions are	15
How volume junctions are used in SMB and NFS namespaces	16
What the typical NAS namespace architectures are	16
How security styles affect data access	19
What the security styles and their effects are	20
Where and when to set security styles	21
How to decide on what security style to use on SVMs with FlexVol	
volumes	21
How security style inheritance works	21
How Data ONTAP preserves UNIX permissions	22
How to manage UNIX permissions using the Windows Security tab	22
NFS and CIFS file naming dependencies	23
Characters a file name can use	23

Case-sensitivity of a file name	23
How Data ONTAP creates file names	23
Use of hard mounts	24
How Data ONTAP supports file access using NFS	25
How Data ONTAP handles NFS client authentication	25
How Data ONTAP grants CIFS file access from NFS clients	25
Supported NFS versions and clients	26
NFSv4.0 functionality supported by Data ONTAP	26
Limitations of Data ONTAP support for NFSv4	26
Data ONTAP support for NFSv4.1	27
Data ONTAP support for parallel NFS	27
Where to find information about NFS support on Infinite Volumes	27
Process for NFS access to UNIX security style data on SVMs with FlexVol volumes	28
Process for NFS access to NTFS security style data on SVMs with FlexVol volumes	28
Setting up file access using NFS	30
Creating and managing data volumes in NAS namespaces	30
Creating data volumes with specified junction points	30
Creating data volumes without specifying junction points	31
Mounting or unmounting existing volumes in the NAS namespace	32
Displaying volume mount and junction point information	33
Configuring security styles	35
Configuring security styles on SVM root volumes	35
Configuring security styles on FlexVol volumes	35
Configuring security styles on qtrees	36
Modifying protocols for SVMs	37
Enabling or disabling NFSv3	38
Enabling or disabling NFSv4.0	38
Enabling or disabling NFSv4.1	38
Enabling or disabling parallel NFS	39
Creating an NFS server	39
Securing NFS access using export policies	40
How export policies control client access to volumes or qtrees	40
Default export policy for SVMs with FlexVol volumes	40
How export rules work	41

How to handle clients with an unlisted security type	42
How security types determine client access levels	45
How to handle superuser access requests	46
Creating an export policy	48
Adding a rule to an export policy	49
Setting an export rule's index number	52
Associating an export policy to a FlexVol volume	53
Assigning an export policy to a qtree	54
Removing an export policy from a qtree	55
Validating qtree IDs for qtree file operations	56
Export policy restrictions and nested junctions for FlexVol volumes	56
Using Kerberos with NFS for strong security	56
Group ID limitation for NFS RPCSEC_GSS	57
Requirements for configuring Kerberos with NFS	57
Specifying the user ID domain for NFSv4	60
Creating a Kerberos realm configuration	61
Creating an NFS Kerberos configuration	62
Configuring SVMs to use LDAP	62
Using LDAP over SSL/TLS to secure communication	63
Creating a new LDAP client schema	65
Creating an LDAP client configuration	66
Enabling LDAP on SVMs	68
How name mappings are used	68
Configuring local UNIX users and groups	75
Creating a local UNIX user	75
Loading local UNIX users from a URI	75
Creating a local UNIX group	76
Loading local UNIX groups from a URI	77
Adding a user to a local UNIX group	77
Loading netgroups into SVMs	78
Creating a NIS domain configuration	79
Support for NFS over IPv6	79
Enabling IPv6 for NFS	79
Where to find information about setting up file access to Infinite Volumes	80
Managing file access using NFS	81
Controlling NFS requests from nonreserved ports	81

Considerations for clients that mount NFS exports using a nonreserved port	82
Commands for managing NFS servers	82
Commands for managing name mappings	82
Commands for managing local UNIX users	83
Commands for managing local UNIX groups	83
Verifying the status of netgroup definitions	83
Commands for managing NIS domain configurations	84
Commands for managing LDAP client configurations	85
Commands for managing LDAP configurations	85
Commands for managing LDAP client schema templates	85
How the access cache works	86
Displaying information about NFS Kerberos configurations	87
Modifying an NFS Kerberos configuration	88
Commands for managing Kerberos realm configurations	88
Commands for managing export policies	89
Commands for managing export rules	89
Managing file locks	90
About file locking between protocols	90
How Data ONTAP treats read-only bits	90
Displaying information about locks	91
Breaking locks	93
Modifying the NFSv4.1 server implementation ID	94
Managing NFSv4 ACLs	95
Benefits of enabling NFSv4 ACLs	95
How NFSv4 ACLs work	95
Enabling or disabling modification of NFSv4 ACLs	96
How Data ONTAP uses NFSv4 ACLs to determine whether it can delete a file	97
Enabling or disabling NFSv4 ACLs	97
Managing NFSv4 file delegations	98
How NFSv4 file delegations work	98
Enabling or disabling NFSv4 read file delegations	99
Enabling or disabling NFSv4 write file delegations	100
Configuring NFSv4 file and record locking	101
About NFSv4 file and record locking	101
Specifying the NFSv4 locking lease period	101

Specifying the NFSv4 locking grace period	102
How NFSv4 referrals work	102
Enabling or disabling NFSv4 referrals	103
Displaying NFS statistics	104
Support for VMware vStorage over NFS	105
Enabling or disabling VMware vStorage over NFS	105
Enabling or disabling rquota support	106
NFSv3 performance improvement by modifying the TCP maximum read and write size	107
Modifying the NFSv3 TCP maximum read and write size	108
Auditing NAS file access events on SVMs with FlexVol volumes	110
How auditing works	110
Basic auditing concepts	110
How the Data ONTAP auditing process works	111
Aggregate space considerations when enabling auditing	112
Auditing requirements and considerations	113
What the supported audit event log formats are	113
Viewing audit event logs	114
SMB file and folder access events that can be audited	114
NFS file and directory access events that can be audited	115
Planning the auditing configuration	116
Creating a file and directory auditing configuration on SVMs	119
Creating the auditing configuration	120
Enabling auditing on the SVM	121
Verifying the auditing configuration	122
Configuring file and folder audit policies	122
Configuring audit policies on NTFS security-style files and directories	122
Configuring auditing for UNIX security style files and directories	126
Displaying information about audit policies applied to files and directories	127
Displaying information about audit policies using the Windows Security tab	127
Displaying information about NTFS audit policies on FlexVol volumes using the CLI	128
Managing auditing configurations	131
Manually rotating the audit event logs	131
Enabling and disabling auditing on SVMs	131

Displaying information about auditing configurations	132
Commands for modifying auditing configurations	134
Deleting an auditing configuration	134
What the process is when reverting	135
Troubleshooting auditing and staging volume space issues	135
How to troubleshoot space issues related to the event log volumes	136
How to troubleshoot space issues related to the staging volumes (cluster administrators only)	136
Using FPolicy for file monitoring and management on SVMs with FlexVol volumes	138
How FPolicy works	138
What the two parts of the FPolicy solution are	138
What synchronous and asynchronous communications are	138
Roles that cluster components play with FPolicy	139
How FPolicy works with external FPolicy servers	140
What the node-to-external FPolicy server communication process is	142
How FPolicy services work across SVM namespaces	143
FPolicy configuration types	144
Requirements, considerations, and best practices for configuring FPolicy	145
Ways to configure FPolicy	145
Requirements for setting up FPolicy	145
Best practices and recommendations when setting up FPolicy	146
Important revert considerations	146
What the steps for setting up an FPolicy configuration are	147
Planning the FPolicy configuration	148
Planning the FPolicy external engine configuration	148
Planning the FPolicy event configuration	154
Planning the FPolicy policy configuration	160
Planning the FPolicy scope configuration	163
Creating the FPolicy configuration	166
Creating the FPolicy external engine	167
Creating the FPolicy policy event	167
Creating the FPolicy policy	168
Creating the FPolicy policy scope	168
Enabling the FPolicy policy	169
Modifying FPolicy configurations	169

Commands for modifying FPolicy configurations	169
Enabling or disabling FPolicy policies	170
Displaying information about FPolicy configurations	170
How the show commands work	171
Commands for displaying information about FPolicy configurations	171
Displaying information about FPolicy policy status	172
Displaying information about enabled FPolicy policies	173
Managing FPolicy server connections	174
Connecting to external FPolicy servers	174
Disconnecting from external FPolicy servers	175
Displaying information about connections to external FPolicy servers	175
Glossary	178
Copyright information	184
Trademark information	185
Index	188

Preface

About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the terms *Cluster-Mode* or *clustered Data ONTAP* are used in this document, they refer to the Data ONTAP features and functionality designed for clusters, which are different from 7-Mode and prior Data ONTAP 7.1, 7.2, and 7.3 release families.

In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in [Websites](#) on page 10).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web

content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 10) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 10).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 10).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

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 and CIFS protocols.

This means clients can access all files on Storage Virtual Machines (SVMs) regardless of what protocol they are connecting with or what type of authentication they require.

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. In addition, the authentication type of the client is one method that can be used to determine whether a client can access data when configuring export policies (optional for CIFS).
- **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 Virtual Machine (SVM).

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
- 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.

LIF configuration requirements for file access management

To properly manage file access control, Data ONTAP must communicate with external services such as NIS, LDAP, and Active Directory servers. The Storage Virtual Machine (SVM) LIFs must be properly configured to allow these communications.

The communication with external services usually happens over the data LIF of the SVM. Therefore, you must ensure that the SVM has a data LIF properly configured to reach all required external services on each node.

In addition, in some situations, communication over the data LIF might fail or must be made on a node that does not host data LIFs for the SVM. In this case, the storage system attempts to use node-management and cluster-management LIFs instead. If your environment allows this, you should also ensure that the node-management and cluster-management LIFs in the cluster can reach these external services as well.

For more information about LIF configuration, see the *Clustered Data ONTAP Network Management Guide*.

How namespaces and volume junctions affect file access on SVMs with FlexVol volumes

You must understand what namespaces and volume junctions are and how they work to correctly configure file access on Storage Virtual Machines (SVMs) in your storage environment.

What namespaces in SVMs with FlexVol volumes are

A namespace is a logical grouping of volumes that are joined together at junction points to create a single, logical file system that derives from the Storage Virtual Machine (SVM) root volume. Each SVM has a namespace.

CIFS and NFS servers on a data SVM can store and access data across the namespace. Each client can access the entire namespace by mounting an export or accessing a single SMB share at the top of the namespace.

Alternatively, SVM administrators can create exports at each volume junction so that clients can create mount points at intermediate locations in the namespace, or they can create SMB shares that point to any directory path in the namespace.

Volumes can be added at any time by mounting them to any location in the namespace. Clients can immediately access the newly added volume, provided that the volume junction is under the point at which they are accessing the namespace and provided that they have sufficient permissions.

What volume junctions are

Volume junctions are a way to join individual volumes together into a single, logical namespace. Volume junctions are transparent to CIFS and NFS clients. When NAS clients access data by traversing a junction, the junction appears to be an ordinary directory.

A junction is formed when a volume is mounted to a mount point below the root and is used to create a file-system tree. The top of a file-system tree is always the root volume, which is represented by a slash (/). A junction points from a directory in one volume to the root directory of another volume.

A volume must be mounted at a junction point in the namespace to allow NAS client access to contained data:

- Although specifying a junction point is optional when a volume is created, data in the volume cannot be exported and a share cannot be created until the volume is mounted to a junction point in the namespace.
- A volume that was not mounted during volume creation can be mounted post-creation.
- New volumes can be added to the namespace at any time by mounting them to a junction point.
- Mounted volumes can be unmounted; however, unmounting a volume disrupts NAS client access to all data in the volume and to all volumes mounted at child junction points beneath the unmounted volume.

- Junction points can be created directly below a parent volume junction, or they can be created on a directory within a volume.

For example, a path to a volume junction for a volume named “vol3” might be `/vol1/vol2/vol3`, or it might be `/vol1/dir2/vol3`, or even `/dir1/dir2/vol3`.

For more information, see the *Clustered Data ONTAP File Access Management Guide for CIFS* or the *Clustered Data ONTAP File Access Management Guide for NFS*.

How volume junctions are used in SMB and NFS namespaces

You can mount volumes at junction points anywhere within the namespace to create a single, logical namespace. If you specify a junction point when the volume is created, the volume is automatically mounted at the time the volume is created and is available for NAS access. You can create SMB shares and NFS exports on the mounted volume.

If you do not specify a junction point, the volume is online but is not mounted for NAS file access. You must mount a volume to a junction point before it can be used for NAS file access.

What the typical NAS namespace architectures are


All Storage Virtual Machine (SVM) name spaces derive from the root volume; however, there are several typical NAS namespace architectures that you can use as you create your SVM name space. You can choose the namespace architecture that matches your business and workflow needs.

The top of the namespace is always the root volume, which is represented by a slash (/). The namespace architecture under the root falls into three basic categories:

- A single branched tree, with only a single junction to the root of the namespace
- Multiple branched trees, with multiple junction points to the root of the namespace
- Multiple stand-alone volumes, each with a separate junction point to the root of the name space

Namespace with single branched tree

An architecture with a single branched tree has a single insertion point to the root of the SVM namespace. The single insertion point can be either a junctioned volume or a directory beneath the root. All other volumes are mounted at junction points beneath the single insertion point (which can be a volume or a directory).




For example, a typical volume junction configuration with the above namespace architecture might look like the following configuration, where all volumes are junctioned below the single insertion point, which is a directory named “data”:

Vserver	Volume	Junction Active	Junction Path	Junction Path Source
vs1	corp1	true	/data/dir1/corp1	RW_volume
vs1	corp2	true	/data/dir1/corp2	RW_volume
vs1	data1	true	/data/data1	RW_volume
vs1	eng1	true	/data/data1/eng1	RW_volume
vs1	eng2	true	/data/data1/eng2	RW_volume
vs1	sales	true	/data/data1/sales	RW_volume
vs1	vol1	true	/data/vol1	RW_volume
vs1	vol2	true	/data/vol2	RW_volume
vs1	vol3	true	/data/vol3	RW_volume
vs1	vs1_root	-	/	-

Namespace with multiple branched trees

An architecture with multiple branched trees has multiple insertion points to the root of the SVM namespace. The insertion points can be either junctioned volumes or directories beneath the root. All other volumes are mounted at junction points beneath the insertion points (which can be volumes or directories).




For example, a typical volume junction configuration with the above namespace architecture might look like the following configuration, where there are three insertion points to the root volume of the SVM. Two insertion points are directories named “data” and “projects”. One insertion point is a junctioned volume named “audit”:

Vserver	Volume	Junction Active	Junction Path	Junction Path Source
vs1	audit	true	/audit	RW_volume
vs1	audit_logs1	true	/audit/logs1	RW_volume
vs1	audit_logs2	true	/audit/logs2	RW_volume
vs1	audit_logs3	true	/audit/logs3	RW_volume
vs1	eng	true	/data/eng	RW_volume
vs1	mktg1	true	/data/mktg1	RW_volume
vs1	mktg2	true	/data/mktg2	RW_volume
vs1	project1	true	/projects/project1	RW_volume
vs1	project2	true	/projects/project2	RW_volume
vs1	vs1_root	-	/	-

Namespace with multiple stand-alone volumes

In an architecture with stand-alone volumes, every volume has an insertion point to the root of the SVM namespace; however, the volume is not junctioned below another volume. Each volume has a unique path, and is either junctioned directly below the root or is junctioned under a directory below the root.



For example, a typical volume junction configuration with the above namespace architecture might look like the following configuration, where there are five insertion points to the root volume of the SVM, with each insertion point representing a path to one volume.

Vserver	Volume	Junction Active	Junction Path	Junction Path Source
vs1	eng	true	/eng	RW_volume
vs1	mktg	true	/vol/mktg	RW_volume
vs1	project1	true	/project1	RW_volume
vs1	project2	true	/project2	RW_volume
vs1	sales	true	/sales	RW_volume
vs1	vs1_root	-	/	-

How security styles affect data access

Each volume and qtree on the storage system has a security style. The security style determines what type of permissions are used for data on volumes when authorizing users. You must understand what the different security styles are, when and where they are set, how they impact permissions, how they differ between volume types, and more.

For more information about security styles, see the *Clustered Data ONTAP File Access Management Guide for CIFS* or *Clustered Data ONTAP File Access Management Guide for NFS*.

What the security styles and their effects are

There are four different security styles: UNIX, NTFS, mixed, and unified. Each security style has a different effect on how permissions are handled for data. You must understand the different effects to ensure that you select the appropriate security style for your purposes.

It is important to understand that security styles do not determine what client types can or cannot access data. Security styles only determine the type of permissions Data ONTAP uses to control data access and what client type can modify these permissions.

For example, if a volume uses UNIX security style, SMB clients can still access data (provided that they properly authenticate and authorize) due to the multiprotocol nature of Data ONTAP. However, Data ONTAP uses UNIX permissions that only UNIX clients can modify using native tools.

Security style	Clients that can modify permissions	Permissions that clients can use	Resulting effective security style	Clients that can access files
UNIX	NFS	NFSv3 mode bits	UNIX	NFS and SMB
		NFSv4.x ACLs	UNIX	
NTFS	SMB	NTFS ACLs	NTFS	
Mixed	NFS or SMB	NFSv3 mode bits	UNIX	
		NFSv4.x ACLs	UNIX	
		NTFS ACLs	NTFS	
Unified (only for Infinite Volumes)	NFS or SMB	NFSv3 mode bits	UNIX	
		NFSv4.1 ACLs	UNIX	
		NTFS ACLs	NTFS	

For more information about the unified security style, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

When the security style is mixed or unified, the effective permissions depend on the client type that last modified the permissions because users set the security style on an individual basis. If the last client that modified permissions was an NFSv3 client, the permissions are UNIX NFSv3 mode bits. If the last client was an NFSv4 client, the permissions are NFSv4 ACLs. If the last client was an SMB client, the permissions are Windows NTFS ACLs.

Note: Data ONTAP initially sets some default file permissions. By default, the effective security style on all data in UNIX, mixed, and unified security style volumes is UNIX and the effective permissions type is UNIX mode bits (0755 unless specified otherwise) until configured by a client as allowed by the default security style. By default, the effective security style on all data in NTFS security style volumes is NTFS and has an ACL allowing full control to everyone.

Where and when to set security styles

Security styles can be set on FlexVol volumes (both root or data volumes) and qtrees. Security styles can be set manually at the time of creation, inherited automatically, or changed at a later time.

Note: Infinite Volumes always use the unified security style. You cannot configure or change the security style of an Infinite Volume.

How to decide on what security style to use on SVMs with FlexVol volumes

To help you decide what security style to use on a volume, you should consider two factors. The primary factor is the type of administrator that manages the file system. The secondary factor is the type of user or service that accesses the data on the volume.

When you configure the security style on a volume, you should consider the needs of your environment to ensure that you select the best security style and avoid issues with managing permissions. The following considerations can help you decide:

Security style	Choose if...
UNIX	<ul style="list-style-type: none"> The file system is managed by a UNIX administrator. The majority of users are NFS clients. An application accessing the data uses a UNIX user as the service account.
NTFS	<ul style="list-style-type: none"> The file system is managed by a Windows administrator. The majority of users are SMB clients. An application accessing the data uses a Windows user as the service account.
Mixed	The file system is managed by both UNIX and Windows administrators and users consist of both NFS and SMB clients.

How security style inheritance works

If you do not specify the security style when creating a new FlexVol volume or qtree, it inherits its security style.

Security styles are inherited in the following manner:

- A FlexVol volume inherits the security style of the root volume of its containing Storage Virtual Machine (SVM).
- A qtree inherits the security style of its containing FlexVol volume.
- A file or directory inherits the security style of its containing FlexVol volume or qtree.

Infinite Volumes cannot inherit security styles. All files and directories in Infinite Volumes always use the unified security style. The security style of an Infinite Volume and the files and directories it contains cannot be changed.

How Data ONTAP preserves UNIX permissions

When files in a FlexVol volume that currently have UNIX permissions are edited and saved by Windows applications, Data ONTAP can preserve the UNIX permissions.

When applications on Windows clients edit and save files, they read the security properties of the file, create a new temporary file, apply those properties to the temporary file, and then give the temporary file the original file name.

When Windows clients perform a query for the security properties, they receive a constructed ACL that exactly represents the UNIX permissions. The sole purpose of this constructed ACL is to preserve the file's UNIX permissions as files are updated by Windows applications to ensure that the resulting files have the same UNIX permissions. Data ONTAP does not set any NTFS ACLs using the constructed ACL.

How to manage UNIX permissions using the Windows Security tab

If you want to manipulate UNIX permissions of files or folders in UNIX or mixed security-style qtrees or volumes on Storage Virtual Machines (SVMs) with FlexVol volumes, you can use the Security tab on Windows clients. Alternatively, you can use applications that can query and set Windows ACLs.

- **Modifying UNIX permissions**

You can use the Windows Security tab to view and change UNIX permissions for a UNIX security-style volume or qtree. This is also true for a mixed security-style volume or qtree where the files and folders have a UNIX effective security style.

If mode permissions are used, you can directly change the mode permissions for the listed UID, GID, and others (everyone else with an account on the computer). For example, if the displayed UID has r-x permissions, you can change the UID permissions to rwx.

- **Changing UNIX permissions to NTFS permissions**

You can use the Windows Security tab to replace UNIX security objects with Windows security objects on a mixed security-style volume or qtree where the files and folders have a UNIX effective security style.

You must first remove the listed entries and then replace them with the desired Windows User and Group objects. You can then configure NTFS-based ACLs on the Windows User and Group objects. By removing UNIX security objects and adding Windows Users and Groups to a file or folder in a mixed security-style volume or qtree, you change the effective security style on the file or folder from UNIX to NTFS.

When changing permissions on a folder, the default Windows behavior is to propagate these changes to all subfolders and files. Therefore, you must change the propagation choice to the desired setting if you do not want to propagate a change in security style to all child folders, subfolders, and files.

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.

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`.

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.

How Data ONTAP supports file access using NFS

You can export and unexport volumes or qtrees on your storage system, making them available or unavailable, respectively, for mounting by NFS clients.

How Data ONTAP handles NFS client authentication

NFS clients must be properly authenticated before they can access data on the Storage Virtual Machine (SVM). Data ONTAP authenticates the clients by checking their UNIX credentials against name services you configure.

When an NFS client connects to the SVM, Data ONTAP obtains the UNIX credentials for the user by checking different name services, depending on the name services configuration of the SVM. Data ONTAP can check credentials for local UNIX accounts, NIS domains, and LDAP domains. At least one of them must be configured so that Data ONTAP can successfully authenticate the user. You can specify multiple name services and the order in which Data ONTAP searches them.

In a pure NFS environment with UNIX volume security styles, this configuration is sufficient to authenticate and provide the proper file access for a user connecting from an NFS client.

If you are using mixed, NTFS, or unified volume security styles, Data ONTAP must obtain a CIFS user name for the UNIX user for authentication with a Windows domain controller. This can happen either by mapping individual users using local UNIX accounts or LDAP domains, or by using a default CIFS user instead. You can specify which name services Data ONTAP searches in which order, or specify a default CIFS user.

How Data ONTAP grants CIFS file access 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 with NTFS permissions.

Data ONTAP does this by converting the user's UNIX User ID (UID) into a CIFS credential, and 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.

Supported NFS versions and clients

Before you can use NFS in your network, you need to know which NFS versions and clients Data ONTAP supports.

For the latest information about which NFS versions and clients Data ONTAP supports, see the N series Interoperability Matrices website (accessed and navigated as described in [Websites](#) on page 10).

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.

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 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.
- Named attributes are not supported.

- 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.

Data ONTAP support for NFSv4.1

Data ONTAP supports the NFSv4.1 protocol to allow access for NFSv4.1 clients.

By default NFSv4.1 is disabled. You can enable it by specifying the `-v4.1` option and setting it to enabled when creating an NFS server on the Storage Virtual Machine (SVM).

Data ONTAP does not support NFSv4.1 directory and file level delegations.

Data ONTAP support for parallel NFS

Data ONTAP supports parallel NFS (pNFS). The pNFS protocol offers performance improvements by giving clients direct access to the data of a set of files distributed across multiple nodes of a cluster. It helps clients locate the optimal path to a volume.

Where to find information about NFS support on Infinite Volumes

For information about the NFS versions and functionality that Infinite Volumes support, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

Process for NFS access to UNIX security style data on SVMs with FlexVol volumes

Understanding the process used for NFS access to UNIX security style data is helpful when designing a file access configuration that provides appropriate security settings.

When an NFS client connects to a storage system to access data with UNIX security style, Data ONTAP goes through the following steps:

1. Obtain the UNIX credentials for the user.
Data ONTAP checks local UNIX accounts, NIS servers, and LDAP servers, depending on the Storage Virtual Machine (SVM) configuration.
2. Authorize the user.
Data ONTAP checks the UNIX credentials for the user against the UNIX permissions of the data to determine what type of data access the user is allowed, if any.

In this scenario, name mapping is not performed because CIFS credentials are not required.

Process for NFS access to NTFS security style data on SVMs with FlexVol volumes

Understanding the process used for NFS access to NTFS security style data is helpful when designing a file access configuration that provides appropriate security settings.

When an NFS client connects to a storage system to access data with NTFS security style, Data ONTAP goes through the following steps:

1. Obtain the UNIX credentials for the user.
Data ONTAP checks local UNIX accounts, NIS servers, and LDAP servers, depending on the Storage Virtual Machine (SVM) configuration.
2. Map the UNIX user to a CIFS name.
Data ONTAP checks local name mapping rules, LDAP mapping rules, and the default CIFS user, depending on the SVM configuration.
3. Establish a connection to a Windows domain controller.
Data ONTAP uses a cached connection, queries DNS servers, or uses a specified preferred domain controller.
4. Authenticate the user.
Data ONTAP connects to the domain controller and performs pass-through authentication.
5. Authorize the user.

Data ONTAP checks the CIFS credentials for the user against the NTFS permissions of the data to determine what type of data access the user is allowed, if any.

Setting up file access using NFS

You must complete a number of steps to allow clients access to files on Storage Virtual Machines (SVMs) using NFS. There are some additional steps that are optional depending on the current configuration of your environment.

For clients to be able to access files on SVMs using NFS, you must complete the following tasks:

1. Enable the desired NFS protocol versions on the SVM.
You can specify the versions of NFS that clients can use to access files on the SVM.
2. Create an NFS server on the SVM.
An NFS server is a logical entity on the SVM that enables the SVM to serve files over NFS.
3. Configure the NFS server with the appropriate security, name mapping, and other settings depending on the network and storage environment.

Note: The SVM must exist before you can set up file access using NFS. For more information about SVMs, see the *Clustered Data ONTAP System Administration Guide for SVM Administrators*

Creating and managing data volumes in NAS namespaces

To manage file access in a NAS environment, you must manage data volumes and junction points on your Storage Virtual Machine (SVM) with FlexVol volumes. This includes planning your namespace architecture, creating volumes with or without junction points, mounting or unmounting volumes, and displaying information about data volumes and NFS server or CIFS server namespaces.

Creating data volumes with specified junction points

You can specify the junction point when you create a data volume. The resultant volume is automatically mounted at the junction point and is immediately available to configure for NAS access.

Before you begin

The aggregate in which you want to create the volume must already exist.

Steps

1. Create the volume with a junction point by using the following command:

```
volume create -vserver vservers_name -volume volume_name -aggregate
aggregate_name -size {integer[KB|MB|GB|TB|PB]} -security-style {ntfs|
unix|mixed} -junction-path junction_path
```

The junction path must start with the root (/) and can contain both directories and junctioned volumes. The junction path does not need to contain the name of the volume. Junction paths are independent of the volume name.

Specifying a volume security style is optional. If you do not specify a security style, Data ONTAP creates the volume with the same security style that is applied to the root volume of the Storage Virtual Machine (SVM). However, the root volume's security style might not be the security style you want applied to the data volume you create. The recommendation is to specify the security style when you create the volume to minimize difficult-to-troubleshoot file-access issues.

There are many optional parameters that you can use to customize a data volume. To learn more about them, see the man pages for the `volume create` command.

2. Verify that the volume was created with the desired junction point:

```
volume show -vserver vs1 -volume home4 -junction
```

Example

The following example creates a volume named “home4” located on SVM vs1 that has a junction path `/eng/home`:

```
cluster1::> volume create -vserver vs1 -volume home4 -aggregate aggr1 -size
1g -junction-path /eng/home
[Job 1642] Job succeeded: Successful
```

```
cluster1::> volume show -vserver vs1 -volume home4 -junction
```

Vserver	Volume	Active	Junction Path	Junction Path Source
vs1	home4	true	/eng/home	RW_volume

Creating data volumes without specifying junction points

You can create a data volume without specifying a junction point. The resultant volume is not automatically mounted, and is not available to configure for NAS access. You must mount the volume before you can configure SMB shares or NFS exports for that volume.

Before you begin

The aggregate in which you want to create the volume must already exist.

Steps

1. Create the volume without a junction point by using the following command:

```
volume create -vserver vs1 -volume home4 -aggregate aggr1 -size {integer[KB|MB|GB|TB|PB]} -security-style {ntfs|unix|mixed}
```

Specifying a volume security style is optional. If you do not specify a security style, Data ONTAP creates the volume with the same security style that is applied to the root volume of the Storage Virtual Machine (SVM). However, the root volume's security style might not be the security style you want applied to the data volume. The recommendation is to specify the security style when you create the volume to minimize difficult-to-troubleshoot file-access issues.

There are many optional parameters that you can use to customize a data volume. To learn more about them, see the man pages for the `volume create` command.

- 2. Verify that the volume was created without a junction point:

```
volume show -vserver vs1 -volume volume_name -junction
```

Example

The following example creates a volume named “sales” located on SVM vs1 that is not mounted at a junction point:

```
cluster1::> volume create -vserver vs1 -volume sales -aggregate aggr3 -size 20GB
[Job 3406] Job succeeded: Successful

cluster1::> volume show -vserver vs1 -junction
```

Vserver	Volume	Junction Active	Junction Path	Junction Path Source
vs1	data	true	/data	RW_volume
vs1	home4	true	/eng/home	RW_volume
vs1	vs1_root	-	/	-
vs1	sales	-	-	-

Mounting or unmounting existing volumes in the NAS namespace

A volume must be mounted on the NAS namespace before you can configure NAS client access to data contained in the Storage Virtual Machine (SVM) volumes. You can mount a volume to a junction point if it is not currently mounted. You can also unmount volumes.

About this task

If you unmount a volume, all data within the junction point, including data in volumes with junction points contained within the unmounted volume's namespace, are inaccessible to NAS clients. When you unmount a volume, data within the volume is not lost. Additionally, existing volume export policies and SMB shares created on the volume or on directories and junction points within the unmounted volume are retained. If you remount the unmounted volume, NAS clients can access the data contained within the volume using existing export policies and SMB shares.

Steps

- 1. Perform the desired action:

If you want to...	Enter the command...
Mount a volume	<code>volume mount -vserver vs1 -volume volume_name -junction-path junction_path</code>
Unmount a volume	<code>volume unmount -vserver vs1 -volume volume_name</code>

2. Verify that the volume is in the desired mount state:

```
volume show -vserver vs1 -volume volume_name -junction
```

Examples

The following example mounts a volume named “sales” located on SVM vs1 to the junction point /sales:

```
cluster1::> volume mount -vserver vs1 -volume sales -junction-path /sales

cluster1::> volume show -vserver vs1 -junction
```

Vserver	Volume	Active	Junction Path	Junction Path Source
vs1	data	true	/data	RW_volume
vs1	home4	true	/eng/home	RW_volume
vs1	vs1_root	-	/	-
vs1	sales	true	/sales	RW_volume

The following example unmounts a volume named “data” located on SVM vs1:

```
cluster1::> volume unmount -vserver vs1 -volume data

cluster1::> volume show -vserver vs1 -junction
```

Vserver	Volume	Active	Junction Path	Junction Path Source
vs1	data	-	-	-
vs1	home4	true	/eng/home	RW_volume
vs1	vs1_root	-	/	-
vs1	sales	true	/sales	RW_volume

Displaying volume mount and junction point information

You can display information about mounted volumes for Storage Virtual Machines (SVMs) and the junction points to which the volumes are mounted. You can also determine which volumes are not mounted to a junction point. You can use this information to understand and manage your SVM namespace.

Step

1. Perform the desired action:

If you want to display...	Enter the command...
Summary information about mounted and unmounted volumes on the SVM	<code>volume show -vserver vs1 -junction</code>
Detailed information about mounted and unmounted volumes on the SVM	<code>volume show -vserver vs1 -volume vs1 -instance</code>
Specific information about mounted and unmounted volumes on the SVM	<p>a. If necessary, you can display valid fields for the <code>-fields</code> parameter by using the following command:</p> <pre>volume show -fields ?</pre> <p>b. Display the desired information by using the <code>-fields</code> parameter:</p> <pre>volume show -vserver vs1 -fields fieldname,...</pre>

Examples

The following example displays a summary of mounted and unmounted volumes on SVM vs1:

```
cluster1::> volume show -vserver vs1 -junction
```

Vserver	Volume	Active	Junction Path	Junction Path Source
vs1	data	true	/data	RW_volume
vs1	home4	true	/eng/home	RW_volume
vs1	vs1_root	-	/	-
vs1	sales	true	/sales	RW_volume

The following example displays information about specified fields for volumes located on SVM vs2:

```
cluster1::> volume show -vserver vs2 -fields vserver,volume,aggregate,size,state,type,security-style,junction-path,junction-parent,node
```

vserver	volume	aggregate	size	state	type	security-style	junction-path	junction-parent	node
vs2	data1	aggr3	2GB	online	RW	unix	-	-	node3
vs2	data2	aggr3	1GB	online	RW	ntfs	/data2	vs2_root	node3
vs2	data2_1	aggr3	8GB	online	RW	ntfs	/data2/d2_1	data2	node3
vs2	data2_2	aggr3	8GB	online	RW	ntfs	/data2/d2_2	data2	node3
vs2	pubs	aggr1	1GB	online	RW	unix	/publications	vs2_root	node1
vs2	images	aggr3	2TB	online	RW	ntfs	/images	vs2_root	node3
vs2	logs	aggr1	1GB	online	RW	unix	/logs	vs2_root	node1
vs2	vs2_root	aggr3	1GB	online	RW	ntfs	/	-	node3

Configuring security styles

You configure security styles on FlexVol volumes and qtrees to determine the type of permissions Data ONTAP uses to control access and what client type can modify these permissions.

For information about the security style of Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

Configuring security styles on SVM root volumes

You configure the Storage Virtual Machine (SVM) root volume security style to determine the type of permissions used for data on the root volume of the SVM.

Steps

1. Perform one of the following actions:

Create the SVM with the...	Specify the security style by...
<code>vserver setup</code> command	Entering the desired root volume security style when prompted by the CLI wizard.
<code>vserver create</code> command	Including the <code>-rootvolume-security-style</code> parameter with the desired security style.

The possible options for the root volume security style are `unix`, `ntfs`, or `mixed`. You cannot use `unified` security style because it only applies to Infinite Volumes.

For more information about the `vserver setup` or `vserver create` commands, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. To display the configuration, including the security style of the SVM you created, enter the following command:

```
vserver show -vserver vserver_name
```

Configuring security styles on FlexVol volumes

You configure the FlexVol volume security style to determine the type of permissions used for data on FlexVol volumes of the Storage Virtual Machine (SVM).

Steps

1. Perform one of the following actions:

If the FlexVol volume...	Use the command...
Does not yet exist	<code>volume create</code> and include the <code>-security-style</code> parameter to specify the security style.
Already exists	<code>volume modify</code> and include the <code>-security-style</code> parameter to specify the security style.

The possible options for the FlexVol volume security style are `unix`, `ntfs`, or `mixed`. You cannot use `unified` security style because it only applies to Infinite Volumes.

If you do not specify a security style when creating a FlexVol volume, the volume inherits the security style of the root volume.

For more information about the `volume create` or `volume modify` commands, see the *Clustered Data ONTAP Logical Storage Management Guide*.

- To display the configuration, including the security style of the FlexVol volume you created, enter the following command:

```
volume show -volume volume_name -instance
```

Configuring security styles on qtrees

You configure the qtree volume security style to determine the type of permissions used for data on qtrees.

Steps

- Perform one of the following actions:

If the qtree...	Use the command...
Does not exist yet	<code>volume qtree create</code> and include the <code>-security-style</code> parameter to specify the security style.
Already exists	<code>volume qtree modify</code> and include the <code>-security-style</code> parameter to specify the security style.

The possible options for the qtree security style are `unix`, `ntfs`, or `mixed`. You cannot use `unified` security style because it only applies to Infinite Volumes.

If you do not specify a security style when creating a qtree, the default security style is `mixed`.

For more information about the `volume qtree create` or `volume qtree modify` commands, see the *Clustered Data ONTAP Logical Storage Management Guide*.

- To display the configuration, including the security style of the qtree you created, enter the following command:

```
volume qtree show -qtree qtree_name -instance
```

Modifying protocols for SVMs

Before you can configure and use NFS or SMB on Storage Virtual Machines (SVMs), you must enable the protocol. This is typically done during SVM setup, but if you did not enable the protocol during setup, you can enable it later by using the `vserver modify` command.

Steps

1. Check which protocols are currently enabled for the SVM by entering the following command:

```
vserver show -vserver vserver_name -fields allowed-protocols
```

2. Modify the list of enabled protocols for the SVM by entering the following command:

```
vserver modify vserver vserver_name -allowed-protocols  
protocol_name[,protocol_name,...]
```

You must enter the complete list of protocols you want to be enabled on the SVM, including the protocols that are already enabled. Any protocol not specified with the command is automatically disabled and moved to the disallowed protocol list.

You can also use the SVM setup wizard to modify protocols for the SVM by using the `vserver setup` command.

See the man page for each command for more information.

3. Confirm that the allowed protocol list was updated correctly by entering the following command:

```
vserver show -vserver vserver_name -fields allowed-protocols
```

Examples

The following command displays which protocols are currently enabled on the SVM named `vs1`.

```
vs1::> vserver show -vserver vs1 -fields allowed-protocols
vserver allowed-protocols
-----
vs1      nfs
```

The following command allows access over SMB by adding `cifs` to the list of enabled protocols on the SVM named `vs1`.

```
vs1::> vserver modify -vserver vs1 -allowed-protocols nfs,cifs
```

Enabling or disabling NFSv3

You can enable or disable NFSv3 by modifying the `-v3` 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>vserver nfs modify -vserver vserver_name -v3 enabled</code>
Disable NFSv3	<code>vserver nfs modify -vserver vserver_name -v3 disabled</code>

Enabling or disabling NFSv4.0

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

About this task

NFSv4.0 is not supported for Storage Virtual Machines (SVMs) with Infinite Volume.

Step

1. Perform one of the following actions:

If you want to...	Enter the following command...
Enable NFSv4.0	<code>vserver nfs modify -vserver vserver_name -v4.0 enabled</code>
Disable NFSv4.0	<code>vserver nfs modify -vserver vserver_name -v4.0 disabled</code>

Enabling or disabling NFSv4.1

You can enable or disable NFSv4.1 by modifying the `-v4.1` option. This allows file access for clients using the NFSv4.1 protocol. By default, NFSv4.1 is disabled.

Step

1. Perform one of the following actions:

If you want to...	Enter the following command...
Enable NFSv4.1	<code>vserver nfs modify -vserver vserver_name -v4.1 enabled</code>
Disable NFSv4.1	<code>vserver nfs modify -vserver vserver_name -v4.1 disabled</code>

Enabling or disabling parallel NFS

To enable or disable parallel NFS (pNFS), you can modify the `-v4.1-pnfs` option. By default pNFS is enabled.

Before you begin

NFSv4.1 support is required to be able to use pNFS.

If you want to enable parallel NFS, you must first disable NFS referrals. They cannot be both enabled at the same time.

Step

1. Perform one of the following actions:

If you want to...	Enter the command...
Enable pNFS	<code>vserver nfs modify -vserver vserver_name -v4.1-pnfs enabled</code>
Disable pNFS	<code>vserver nfs modify -vserver vserver_name -v4.1-pnfs disabled</code>

Creating an NFS server

The NFS server is necessary to provide NFS clients with access to the Storage Virtual Machine (SVM). You can use the `vserver nfs create` command to create an NFS server.

Before you begin

The cluster administrator might also want to configure an NTP server for the SVM. See the *Clustered Data ONTAP System Administration Guide for SVM Administrators* for more information.

Before creating an NFS server, you might want to configure an NIS domain for the SVM. If not, the NFS server uses local-users and local-groups.

Step

1. Use the `vserver nfs create` command to create an NFS server.

Example

The following command creates an NFS server on the SVM named vs1 with NFSv3 disabled, NFSv4.0 enabled, and NFSv4.0 ACLs enabled.

```
vs1::> vserver nfs create -vserver vs1 -v3 disabled -v4.0 enabled -
v4.0-acl enabled
```

Securing NFS access using export policies

You can use export policies to restrict NFS access to volumes or qtrees to clients that match specific parameters.

For information about how export policies affect Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

How export policies control client access to volumes or qtrees

Export policies contain one or more *export rules* that process each client access request. The result of the process determines whether the client is denied or granted access and what level of access. An export policy with export rules must exist on the SVM for clients to access data.

You associate exactly one export policy with each volume or qtree to configure client access to the volume or qtree. The SVM can contain multiple export policies. This enables you to do the following for SVMs with multiple volumes or qtrees:

- Assign different export policies to each volume or qtree of the SVM for individual client access control to each volume or qtree in the SVM.
- Assign the same export policy to multiple volumes or qtrees of the SVM for identical client access control without having to create a new export policy for each volume or qtree.

If a client makes an access request that is not permitted by the applicable export policy, the request fails with a permission-denied message. If a client does not match any rule in the export policy, then access is denied. If an export policy is empty, then all accesses are implicitly denied.

You can modify an export policy dynamically on a system running Data ONTAP.

Default export policy for SVMs with FlexVol volumes

Each Storage Virtual Machine (SVM) with FlexVol volumes has a default export policy that contains no rules. An export policy with rules must exist before clients can access data on the SVM, and each FlexVol volume contained in the SVM must be associated with an export policy.

When you create your SVM with FlexVol volumes, the storage system automatically creates a default export policy called `default` for the root volume of the SVM. You must create one or more rules for the default export policy before clients can access data on the SVM. Alternatively, you can

create a custom export policy with rules. You can modify and rename the default export policy, but you cannot delete the default export policy.

When you create a FlexVol volume in its containing SVM with FlexVol volume, the storage system creates the volume and associates the volume with the default export policy for the root volume of the SVM. By default, each volume created in the SVM is associated with the default export policy for the root volume. You can use the default export policy for all volumes contained in the SVM, or you can create a unique export policy for each volume. You can associate multiple volumes with the same export policy.

How export rules work

Export rules are the functional elements of an export policy. Export rules match client access requests to a volume or qtree against specific parameters you configure to determine how to handle the client access requests.

An export policy must contain at least one export rule to allow access to clients. If an export policy contains more than one rule, the rules are processed in the order in which they appear in the export policy. The rule order is dictated by the rule index number. If a rule matches a client, the permissions of that rule are used and no further rules are processed. If no rules match, the client is denied access.

You can configure export rules to determine client access permissions using the following criteria:

- The file access protocol used by the client sending the request, for example, NFSv4 or SMB.
- A client identifier, for example, host name or IP address.
- The security type used by the client to authenticate, for example, Kerberos v5, NTLM, or AUTH_SYS.

If a rule specifies multiple criteria, and the client does not match one or more of them, the rule does not apply.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule any`

The client access request is sent using the NFSv3 protocol and the client has the IP address 10.1.17.37.

Even though the client access protocol matches, the IP address of the client is in a different subnet from the one specified in the export rule. Therefore, client matching fails and this rule does not apply to this client.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule any`

The client access request is sent using the NFSv4 protocol and the client has the IP address 10.1.16.54.

The client access protocol matches and the IP address of the client is in the specified subnet. Therefore, client matching is successful and this rule applies to this client. The client gets read-write access regardless of its security type.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule krb5,ntlm`

Client #1 has the IP address 10.1.16.207, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

The client access protocol and IP address matches for both clients. The read-only parameter allows read-only access to all clients regardless of the security type they authenticated with. Therefore both clients get read-only access. However, only client #1 gets read-write access because it used the approved security type Kerberos v5 to authenticate. Client #2 does not get read-write access.

How to handle clients with an unlisted security type

When a client presents itself with a security type that is not listed in an access parameter of an export rule, you have the choice of either denying access to the client or mapping it to the anonymous user ID instead by using the option `none` in the access parameter.

A client might present itself with a security type that is not listed in an access parameter because it was authenticated with a different security type or was not authenticated at all (security type `AUTH_NONE`). By default, the client is automatically denied access to that level. However, you can add the option `none` to the access parameter. As a result, clients with an unlisted security style are

mapped to the anonymous user ID instead. The `-anon` parameter determines what user ID is assigned to those clients. The user ID specified for the `-anon` parameter must be a valid user that is configured with permissions you deem appropriate for the anonymous user.

Valid values for the `-anon` parameter range from 0 to 65535.

User ID assigned to <code>-anon</code>	Resulting handling of client access requests
0 - 65533	The client access request is mapped to the anonymous user ID and gets access depending on the permissions configured for this user.
65534	The client access request is mapped to the user nobody and gets access depending on the permissions configured for this user. This is the default.
65535	The access request from any client is denied when mapped to this ID and the client presents itself with security type AUTH_NONE. The access request from clients with user ID 0 is denied when mapped to this ID and the client presents itself with any other security type.

When using the option `none`, it is important to remember that the read-only parameter is processed first. Consider the following guidelines when configuring export rules for clients with unlisted security types:

Read-only includes <code>none</code>	Read-write includes <code>none</code>	Resulting access for clients with unlisted security types
No	No	Denied
No	Yes	Denied because read-only is processed first
Yes	No	Read-only as anonymous
Yes	Yes	Read-write as anonymous

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule sys,none`
- `-rwrule any`
- `-anon 70`

Client #1 has the IP address 10.1.16.207, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

Client #3 has the IP address 10.1.16.234, sends an access request using the NFSv3 protocol, and did not authenticate (meaning security type AUTH_NONE).

The client access protocol and IP address matches for all three clients. The read-only parameter allows read-only access to clients with their own user ID that authenticated with AUTH_SYS. The read-only parameter allows read-only access as the anonymous user with user ID 70 to clients that authenticated using any other security type. The read-write parameter allows read-write access to any security type, but in this case only applies to clients already filtered by the read-only rule.

Therefore, clients #1 and #3 get read-write access only as the anonymous user with user ID 70. Client #2 gets read-write access with its own user ID.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule sys,none`
- `-rwrule none`
- `-anon 70`

Client #1 has the IP address 10.1.16.207, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

Client #3 has the IP address 10.1.16.234, sends an access request using the NFSv3 protocol, and did not authenticate (meaning security type AUTH_NONE).

The client access protocol and IP address matches for all three clients. The read-only parameter allows read-only access to clients with their own user ID that authenticated with AUTH_SYS. The read-only parameter allows read-only access as the anonymous user with user ID 70 to clients that authenticated using any other security type. The read-write parameter allows read-write access only as the anonymous user.

Therefore, client #1 and client #3 get read-write access only as the anonymous user with user ID 70. Client #2 gets read-only access with its own user ID but is denied read-write access.

How security types determine client access levels

The security type that the client authenticated with plays a special role in export rules. You must understand how the security type determines the levels of access the client gets to a volume or qtree.

The three possible access levels are as follows:

1. Read-only
2. Read-write
3. Superuser (for clients with user ID 0)

Because the access level by security type is evaluated in this order, you must observe the following rules when constructing access level parameters in export rules:

For a client to get access level...	These access parameters must match the client's security type...
Normal user read-only	Read-only (<code>-rorule</code>)
Normal user read-write	Read-only (<code>-rorule</code>) and read-write (<code>-rwrule</code>)
Superuser read-only	Read-only (<code>-rorule</code>) and <code>-superuser</code>
Superuser read-write	Read-only (<code>-rorule</code>) and read-write (<code>-rwrule</code>) and <code>-superuser</code>

The following are valid security types for each of these three access parameters:

- `any`
- `none`
- `never`

This security type is not valid for use with the `-superuser` parameter.

- `krb5`
- `ntlm`
- `sys`

When matching a client's security type against each of the three access parameters, there are three possible outcomes:

If the client's security type...	Then the client...
Matches one specified in the access parameter.	Gets access for that level with its own user ID.
Does not match one specified, but the access parameter includes the option <code>none</code> .	Gets access for that level but as the anonymous user with the user ID specified by the <code>-anon</code> parameter.

If the client's security type...	Then the client...
Does not match one specified and the access parameter does not include the option none.	Does not get any access for that level. This does not apply to the <code>-superuser</code> parameter because it always includes none even when not specified.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule sys,krb5`
- `-superuser krb5`

Client #1 has the IP address 10.1.16.207, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

Client #3 has the IP address 10.1.16.234, has user ID 0, sends an access request using the NFSv3 protocol, and did not authenticate (AUTH_NONE).

The client access protocol and IP address matches for all three clients. The read-only parameter allows read-only access to all clients regardless of security type. The read-write parameter allows read-write access to clients with their own user ID that authenticated with AUTH_SYS or Kerberos v5. The superuser parameter allows superuser access to clients with user ID 0 that authenticated with Kerberos v5.

Therefore, client #1 gets superuser read-write access because it matches all three access parameters. Client #2 gets read-write access but not superuser access. Client #3 gets read-only access but not superuser access.

How to handle superuser access requests

When you configure export policies, you need to consider what you want to happen if the storage system receives a client access request with user ID 0, meaning as a superuser, and set up your export rules accordingly.

In the UNIX world, a user with the user ID 0 is known as the superuser, typically called root, who has unlimited access rights on a system. Using superuser privileges can be dangerous for several reasons, including breach of system and data security.

By default, Data ONTAP maps clients presenting with user ID 0 to the anonymous user. However, you can specify the `-superuser` parameter in export rules to determine how to handle clients

presenting with user ID 0 depending on their security type. The following are valid options for the `-superuser` parameter:

- `any`
- `none`
- `krb5`
- `ntlm`
- `sys`

This is the default setting if you do not specify the `-superuser` parameter.

There are two different ways how clients presenting with user ID 0 are handled, depending on the `-superuser` parameter configuration:

If the <code>-superuser</code> parameter and the client's security type...	Then the client...
Match	Gets superuser access with user ID 0.
Do not match	Gets access as the anonymous user with the user ID specified by the <code>-anon</code> parameter and its assigned permissions. This is regardless of whether the read-only or read-write parameter specifies the option <code>none</code> .

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule krb5,ntlm`
- `-anon 127`

Client #1 has the IP address 10.1.16.207, has user ID 746, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

The client access protocol and IP address matches for both clients. The read-only parameter allows read-only access to all clients regardless of the security type they authenticated with. However, only client #1 gets read-write access because it used the approved security type Kerberos v5 to authenticate.

Client #2 does not get superuser access. Instead, it gets mapped to anonymous because the `-superuser` parameter is not specified. This means it defaults to `none` and automatically maps

user ID 0 to anonymous. Client #2 also only gets read-only access because its security type did not match the read-write parameter.

Example

The export policy contains an export rule with the following parameters:

- `-protocol nfs3`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule krb5,ntlm`
- `-superuser krb5`
- `-anon 0`

Client #1 has the IP address 10.1.16.207, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with Kerberos v5.

Client #2 has the IP address 10.1.16.211, has user ID 0, sends an access request using the NFSv3 protocol, and authenticated with AUTH_SYS.

The client access protocol and IP address matches for both clients. The read-only parameter allows read-only access to all clients regardless of the security type they authenticated with. However, only client #1 gets read-write access because it used the approved security type Kerberos v5 to authenticate. Client #2 does not get read-write access.

The export rule allows superuser access for clients with user ID 0. Client #1 gets superuser access because it matches the user ID and security type for the read-only and `-superuser` parameters. Client #2 does not get read-write or superuser access because its security type does not match the read-write parameter or the `-superuser` parameter. Instead, client #2 is mapped to the anonymous user, which in this case has the user ID 0.

Creating an export policy

Before creating export rules, you must create an export policy to hold them. You can use the `vserver export-policy create` command to create an export policy.

Steps

1. To create an export policy, enter the following command:

```
vserver export-policy create -vserver vserver_name -policyname policy_name
```

`-vserver vserver_name` specifies the Storage Virtual Machine (SVM) name.

`-policyname policy_name` specifies the name of the new export policy. The policy name can be up to 256 characters long.
2. Assign the export policy to a volume or qtree by performing one of the following actions:

To assign the export policy to a...	Enter the command...
Volume	<code>volume modify -vserver <i>vserver_name</i> -volume <i>volume_name</i> -policy <i>export_policy_name</i></code>
Qtree	<code>volume qtree modify -vserver <i>vserver_name</i> -volume <i>volume_name</i> -qtree <i>qtree_name</i> -export-policy <i>export_policy_name</i></code>

Example

The following command creates an export policy named `rs1` on the SVM named `vs1`:

```
vs1::> vserver export-policy create -vserver vs1 -policyname rs1
```

Adding a rule to an export policy

You can use the `vserver export-policy rule create` command to create an export rule for an export policy. This enables you to define client access to data.

Before you begin

The export policy you want to add the export rules to must already exist.

Step

1. To create an export rule, enter the following command:

```
vserver export-policy rule create -vserver vserver_name -policyname policy_name -ruleindex integer -protocol {any|nfs3|nfs|cifs|nfs4|flexcache},... -clientmatch text -rorule {any|none|never|krb5|ntlm|sys},... -rwrule {any|none|never|krb5|ntlm|sys},... -anon user_ID -superuser {any|none|krb5|ntlm|sys},... -allow-suid {true|false} -allow-dev {true|false}
```

`-vserver vserver_name` specifies the Storage Virtual Machine (SVM) name.

`-policyname policy_name` specifies the name of the existing export policy to add the rule to.

`-ruleindex integer` specifies the index number for the rule. Rules are evaluated according to their order in the list of index numbers; rules with lower index numbers are evaluated first. For example, the rule with index number 1 is evaluated before the rule with index number 2.

`-protocol {any|nfs3|nfs|cifs|nfs4|flexcache}` specifies the access protocols that the export rule applies to. You can specify a comma-separated list of multiple access protocols for an export rule. If you specify the protocol as `any`, do not specify any other protocols in the list. If you do not specify an access protocol, the default value of `any` is used.

`-clientmatch text` specifies the client to which the rule applies. You can specify the match in any of the following formats:

Client match format	Example
Domain name preceded by the "." character	<code>.example.com</code>
Host name	<code>host1</code>
IPv4 address	<code>10.1.12.24</code>
IPv4 address with a subnet mask expressed as a number of bits	<code>10.1.12.10/4</code>
IPv4 address with a network mask	<code>10.1.16.0/255.255.255.0</code>
IPv6 address in dotted format	<code>::1.2.3.4</code>
IPv6 address with a subnet mask expressed as a number of bits	<code>ff::00/32</code>
A single netgroup with the netgroup name preceded by the @ character	<code>@netgroup1</code>

Note: Entering an IP address range, such as 10.1.12.10-10.1.12.70, is not allowed. Entries in this format are interpreted as a text string and treated as a host name. Entering an IPv6 address with a network mask, such as ff::12/ff::00, is not allowed.

When specifying individual IP addresses in export rules for granular management of client access, do not specify IP addresses that are dynamically (for example, DHCP) or temporarily (for example, IPv6) assigned. Otherwise, the client loses access when its IP address changes.

When specifying host or domain names, ensure that they can be properly resolved through DNS, both forward and reverse.

`-rorule {any|none|never|krb5|ntlm|sys|}` provides read-only access to clients that authenticate with the specified security types.

`-rwrule {any|none|never|krb5|ntlm|sys|}` provides read-write access to clients that authenticate with the specified security types.

Note: A client can only get read-write access for a specific security type if the export rule allows read-only access for that security type as well. If the read-only parameter is more restrictive for a security type than the read-write parameter, the client cannot get read-write access.

You can specify a comma-separated list of multiple security types for a rule. If you specify the security type as `any` or `never`, do not specify any other security types. Choose from the following valid security types:

- `any`

A matching client can access the data regardless of security type.

- none

If listed alone, clients with any security type are granted access as anonymous. If listed with other security types, clients with a specified security type are granted access and clients with any other security type are granted access as anonymous.

- never

A matching client cannot access the data regardless of security type.

- krb5

A matching client can access the data if it is authenticated by Kerberos 5.

- ntlm

A matching client can access the data if it is authenticated by CIFS NTLM.

- sys

A matching client can access the data if it is authenticated by NFS AUTH_SYS.

-anon *user_ID* specifies a UNIX user ID or user name that is mapped to client requests that arrive with a user ID of 0 (zero), which is typically associated with the user name root. The default value is 65534, which is typically associated with the user name nobody.

-superuser {any|none|krb5|ntlm|sys|} provides superuser access to clients that authenticate with the specified security types.

Note: A client can only get superuser access for a specific security type if the export rule allows read-only access for that security type as well. If the read-only parameter is more restrictive for a security type than the superuser parameter, the client cannot get superuser access.

-allow-suid {true|false} specifies whether to allow access to set user ID (SUID) and set group ID (SGID). If this parameter is set to true, clients can modify the SUID or SGID of files, directories, and volumes. If this parameter is set to false, clients can modify the SUID or SGID of directories and volumes but not files. The default is true.

-allow-dev {true|false} specifies whether to allow creation of devices. The default is true.

Examples

The following command creates an export rule on the SVM named vs1 in an export policy named rs1. The rule has the index number 1. The rule matches all clients. The rule enables all NFS access. It enables read-only access by all clients and requires Kerberos authentication for read-write access. Clients with the UNIX user ID 0 (zero) are mapped to user ID 65534 (which typically maps to the user name nobody). The rule enables SUID and SGID modification but does not enable the creation of devices.

```
vs1::> vserver export-policy rule create -vserver vs1
-policyname rs1 -ruleindex 1 -protocol nfs -clientmatch 0.0.0.0/0
-rorule any -rwrule krb5 -anon 65534 -allow-suid true -allow-dev
false
```

The following command creates an export rule on the SVM named vs2 in an export policy named expol2. The rule has the index number 21. The rule matches clients to members of the netgroup dev_netgroup_main. The rule enables all NFS access. It enables read-only access for clients that authenticated with AUTH_SYS and requires Kerberos authentication for read-write access. Clients with the UNIX user ID 0 (zero) are mapped to user ID 65534 (which typically maps to the user name nobody). The rule enables SUID and SGID modification but does not enable the creation of devices.

```
vs1::> vserver export-policy rule create -vserver vs2
-policyname expol2 -ruleindex 21 -protocol nfs -clientmatch
@dev_netgroup_main -rorule sys -rwrule krb5 -anon 65534
-allow-suid true -allow-dev false
```

Setting an export rule's index number

You can use the `vserver export-policy rule setindex` command to manually set an existing export rule's index number. This enables you to rearrange the order in which Data ONTAP processes export rules.

About this task

If the new index number is already in use, the command inserts the rule at the specified spot and reorders the list accordingly.

Step

1. To modify the index number of a specified export rule, enter the following command:

```
vserver export-policy rule setindex -vserver virtual_server_name -  
policyname policy_name -ruleindex integer -newruleindex integer
```

`-vserver virtual_server_name` specifies the Storage Virtual Machine (SVM) name.

`-policyname policy_name` specifies the policy name.

`-ruleindex integer` specifies the current index number of the export rule. If the specified index number is already in use by an existing export policy rule, the index numbers of that and all successive export policy rules are incremented by one.

`-newruleindex integer` specifies the new index number of the export rule.

Example

The following command changes the index number of an export rule at index number 3 to index number 2 in an export policy named rs1 on the SVM named vs1.

```
vs1::> vserver export-policy rule setindex -vserver vs1
-policyname rs1 -ruleindex 3 -newruleindex 2
```

Associating an export policy to a FlexVol volume

Each FlexVol volume contained in the Storage Virtual Machine (SVM) must be associated with an export policy that contains export rules for clients to access data in the volume.

About this task

When you create the SVM, Data ONTAP creates a default export policy called *default* for the SVM. Data ONTAP assigns the *default* export policy to the SVM volumes. You can associate another custom export policy that you create instead of the default policy to a volume. Before you associate a custom export policy to a volume, you must create one or more export rules that allow the desired access to data in the volume and assign those export rules to the export policy that you want to associate with the volume.

You can associate an export policy to a volume when you create the volume or at any time after you create the volume. You can associate one export policy to the volume.

Steps

1. Use either the `volume create` or `volume modify` command with the `-policy` option to associate an export policy to the volume.

Example

```
cluster::> volume create -vserver vs1 -volume vol1
-aggregate aggr2 -state online -policy cifs_policy -security-style
ntfs
-junction-path /dept/marketing -size 250g -space-guarantee volume
```

For more information about the `volume create` and `volume modify` commands, see the man pages.

2. Verify that the export policy associated with the volume has the desired access configuration using the `vserver export-policy rule show` command with the `-policyname` option.

Example

```
cluster::> vserver export-policy rule show -policyname cifs_policy
```

Virtual Server	Policy Name	Rule Index	Access Protocol	Client Match	RO Rule
vs1	cifs_policy	1	cifs	0.0.0.0/0	any

For more information about the `vserver export-policy rule show` command, see the man pages.

The command displays a summary for that export policy, including a list rules applied to that policy. Data ONTAP assigns each rule a rule index number. After you know the rule index number, you can use it to display detailed information about the specified export rule.

3. Verify that the rules applied to the export policy are configured correctly by using the `vserver export-policy rule show` command and specify the `-policyname`, `-vserver`, and `-ruleindex` options.

Example

```
cluster::> vserver export-policy rule show -policyname cifs_policy -
vserver vs1 -ruleindex 1
                                Virtual Server: vs1
                                Policy Name: cifs_policy
                                Rule Index: 1
                                Access Protocol: cifs
Client Match Hostname, IP Address, Netgroup, or Domain: 0.0.0.0/0
                                RO Access Rule: any
                                RW Access Rule: any
User ID To Which Anonymous Users Are Mapped: 0
                                Superuser Security Flavors: any
                                Honor SetUID Bits In SETATTR: true
                                Allow Creation of Devices: true
```

Assigning an export policy to a qtree

Instead of exporting an entire volume, you can also export a specific qtree on a volume to make it directly accessible to clients. You can export a qtree by assigning an export policy to it. You can assign the export policy either when you create a new qtree or by modifying an existing qtree.

Before you begin

The export policy must exist.

About this task

By default, qtrees use the parent export policy of the containing volume if not otherwise specified at the time of creation.

Steps

1. Perform one of the following actions:

To assign an export policy to...	Enter the command...
A new qtree	<code>volume qtree create -vserver vserver_name -qtree-path /vol/volume_name/qtree_name -export-policy export_policy_name</code>
An existing qtree	<code>volume qtree modify -vserver vserver_name -qtree-path /vol/volume_name/qtree_name -export-policy export_policy_name</code>

2. Verify that the export policy was assigned correctly by entering the following command:

```
volume qtree show -qtree qtree_name -fields export-policy
```

Examples

The following command creates a new qtree named dev1 on volume vol_eng on the SVM named vs1 and assigns it the export policy dev1_export.

```
vs1::> volume qtree create -vserver vs1 -qtree-path /vol/vol_eng/dev1 -export-policy dev1_export
```

The following command modifies an existing qtree named dev2 on volume vol_eng on the SVM named vs1 and assigns it the export policy dev2_export.

```
vs1::> volume qtree modify -vserver vs1 -qtree-path /vol/vol_eng/dev2 -export-policy dev2_export
```

Removing an export policy from a qtree

If you decide you do not want a specific export policy assigned to a qtree any longer, you can remove the export policy by modifying the qtree to inherit the export policy of the containing volume instead. You can do this by using the `volume qtree modify` command with the `-export-policy` parameter and an empty name string ("").

About this task

You must perform this task to remove all export policies assigned to qtrees before downgrading to a Data ONTAP release earlier than 8.2.1 that does not support qtree exports.

Steps

1. To remove an export policy from a qtree, enter the following command:

```
volume qtree modify -vserver vserver_name -qtree-path /vol/volume_name/qtree_name -export-policy ""
```

- 2. Verify that the qtree was modified accordingly by entering the following command:
`volume qtree show -qtree qtree_name -fields export-policy`

Validating qtree IDs for qtree file operations

Data ONTAP can perform an optional additional validation of qtree IDs. This validation ensures that client file operation requests use a valid qtree ID and that clients can only move files within the same qtree. You can enable or disable this validation by modifying the `-validate-qtrees-export` parameter. This parameter is enabled by default.

About this task

This parameter is only effective when you have assigned an export policy directly to one or more qtrees on the Storage Virtual Machine (SVM).

Steps

- 1. Set the privilege level to advanced:
`set -privilege advanced`
- 2. Perform one of the following actions:

If you want qtree ID validation Enter the following command...
to be...

Enabled	<code>vserver nfs modify -vserver vserver_name - validate-qtrees-export enabled</code>
---------	--

Disabled	<code>vserver nfs modify -vserver vserver_name - validate-qtrees-export disabled</code>
----------	---

- 3. Return to the admin privilege level:
`set -privilege admin`

Export policy restrictions and nested junctions for FlexVol volumes

If you configured export policies to set a less restrictive policy on a nested junction but a more restrictive policy on a higher level junction, access to the lower level junction might fail.

You should ensure that higher level junctions have less restrictive export policies than lower level junctions.

Using Kerberos with NFS for strong security

You can use Kerberos to provide strong authentication between Storage Virtual Machines (SVMs) and NFS clients to ensure secure NFS communication.

Beginning with version 3, NFS supports generic security services for RPC (RPCSEC_GSS), which enables the use of Kerberos 5. Kerberos provides strong secure authentication for client/server

applications. Authentication provides verification of a user's or process's identity to a server. In the Data ONTAP environment, Kerberos provides authentication between SVMs and NFS clients.

Group ID limitation for NFS RPCSEC_GSS

Although Data ONTAP supports generic security services for RPC (RPCSEC_GSS) for NFSv3 and later, you should be aware of a limitation related to the number of group IDs for UNIX users.

Data ONTAP supports up to 32 group IDs when handling NFS user credentials using RPCSEC_GSS authentication. If a user has more than 32 group IDs in their credentials, the remaining group IDs are truncated.

To avoid authentication issues due to this limitation, ensure that users do not belong to more than 32 groups.

Requirements for configuring Kerberos with NFS

Before you configure Kerberos with NFS on the storage system, you must verify that certain items in your network and storage environment are properly configured.

Note: The steps to configure your environment depend on what version and type of client operating system, domain controller, Kerberos, DNS, etc., that you are using. Documenting all these variables is beyond the scope of this document. For more information, see the respective documentation for each component.

The following items should be configured first:

Network environment requirements

- Kerberos
You must have a working Kerberos setup with a key distribution center (KDC), such as Windows Active Directory based Kerberos or MIT Kerberos.
NFS servers must use “nfs” as the primary component of their machine principal.
- DES encryption
The domain controller must have DES encryption enabled if you are running Windows Active Directory.
- NTP
You must have a working time server running NTP. This is necessary to prevent Kerberos authentication failure due to time skew.
- Domain name resolution (DNS)
Each UNIX client and each Storage Virtual Machine (SVM) LIF must have a proper service record (SRV) registered with the KDC under forward and reverse lookup zones. All participants must be properly resolvable via DNS.
- User accounts
Each client must have a user account in the Kerberos realm. NFS servers must use “nfs” as the primary component of their machine principal.

NFS client requirements

- NFS

Each client must be properly configured to communicate over the network using NFSv3 or NFSv4.

Clients must support RFC1964 and RFC2203. For more information, see the N series Interoperability Matrices website (accessed and navigated as described in [Websites](#) on page 10).
- Kerberos

Each client must be properly configured to use Kerberos authentication, including the following details:

 - Encryption for TGS communication is enabled.
DES for Windows Active Directory based Kerberos or MIT Kerberos, or DES3 for MIT Kerberos only.
 - The most secure encryption type for TGT communication is enabled.
 - The Kerberos realm and domain are configured correctly.
 - GSS is enabled.
- DNS

Each client must be properly configured to use DNS for correct name resolution.
- NTP

Each client must be synchronizing with the NTP server.
- Host and domain information

Each client's `/etc/hosts` and `/etc/resolv.conf` files must contain the correct host name and DNS information, respectively.
- Keytab files

Each client must have a keytab file from the KDC. The realm must be in uppercase letters. The encryption type must be DES-CBC-MD5 (for Windows Active Directory based Kerberos or MIT Kerberos) or DES3-CBC-SHA1 (for MIT Kerberos only). For Windows, you must allow weak cryptography to use DES.
- Optional: For best performance, clients benefit from having at least two network interfaces: one for communicating with the local area network and one for communicating with the storage network.

Storage system requirements

- IPv4

You must have IPv4 network connectivity configured. Kerberos is not supported over IPv6.
- NFS license

The storage system must have a valid NFS license installed. For more information about managing feature licenses, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.
- CIFS license

The CIFS license is optional. It is only required for checking Windows credentials when using multiprotocol name mapping. It is not required in a strict UNIX-only environment. For more

information about managing feature licenses, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

- SVM

You must have at least one SVM configured on the storage system. For more information about configuring SVMs, see the *Clustered Data ONTAP Software Setup Guide*.

- DNS on the SVM

You must have configured DNS on each SVM. For more information about configuring DNS on SVMs, see the *Clustered Data ONTAP Software Setup Guide*.

- NFS server

You must have configured NFS on the SVM.

- CIFS server

If you are running a multiprotocol environment, you must have configured CIFS on the SVM. The CIFS server is required for multiprotocol name mapping.

- Volumes

You must have a root volume and at least one data volume configured for use by the SVM. For more information about configuring volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

- Root volume

The root volume of the SVM must have the following configuration:

Set the...	To...
Security style	UNIX
UID	root or ID 0
GID	root or ID 0
UNIX permissions	777

In contrast to the root volume, data volumes can have either security style.

- UNIX groups

The SVM must have the following UNIX groups configured:

Group name	Group ID
daemon	1
root	0
pcuser	65534 (created automatically by Data ONTAP when you create the SVM)

- UNIX users

The SVM must have the following UNIX users configured:

User name	User ID	Primary group ID	Comment
nfs	500	0	Required for GSS INIT phase The first component of the server SPN is used as the user.
pcuser	65534	65534	Required for NFS and CIFS multiprotocol use Created and added to the pcuser group automatically by Data ONTAP when you create the SVM.
root	0	0	Required for mounting

The nfs user is not required if a Kerberos-UNIX name mapping exists for the SPN that is bound to the data LIF.

- Export policies and rules
You must have configured export policies with the necessary export rules for the root and data volumes and qtrees. If all volumes of the SVM are accessed over Kerberos, you can set the export rule for the root volume to `anon=0, -rorule, -rwrule, -superuser, and -krb`.
- Kerberos-UNIX name mapping
The SVM must have the Kerberos principal `nfs/fqdn@REALM` mapped to the UNIX user root.

Specifying the user ID domain for NFSv4

To specify the user ID domain, you can set the `-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:

```
vserver nfs modify -vserver vserver_name -v4-id-domain NIS_domain_name
```

Creating a Kerberos realm configuration

If you want to use Kerberos authentication for client access, you must first configure the Storage Virtual Machine (SVM) to use an existing Kerberos realm. You can use the `vserver services kerberos-realm create` command to configure the SVM to use a Kerberos realm.

Before you begin

The cluster administrator should have configured NTP on the storage system to avoid authentication issues. Time differences between a client and server (clock skew) are a common cause of authentication failures.

Step

1. Use the `vserver services kerberos-realm create` command to configure a Kerberos realm.

Examples

The following command creates a Kerberos realm configuration that uses a Microsoft Active Directory server as the KDC server. The Kerberos configuration is named AUTH. The Kerberos realm is AUTH.EXAMPLE.COM. The Active Directory server is named ad-1 and its IP address is 10.10.8.14. The permitted clock skew is 300 seconds (the default). The IP address of the KDC server is 10.10.8.14, and its port number is 88 (the default). The encryption type is Data Encryption Standard (DES). The comment is "Microsoft Kerberos config".

```
vs1::> vserver services kerberos-realm create -configname AUTH
-realm AUTH.EXAMPLE.COM -adserver-name ad-1 -adserver-ip 10.10.8.14
-clock-skew 300 -kdc-ip 10.10.8.14 -kdc-port 88 -kdc-vendor
Microsoft -comment "Microsoft Kerberos config"
```

The following command creates a Kerberos realm configuration that uses an MIT KDC. The Kerberos configuration is named SECURE. The Kerberos realm is SECURITY.EXAMPLE.COM. The permitted clock skew is 300 seconds and the Active Directory server and IP address are SUSAN and 10.10.3.1, respectively. The IP address of the KDC server is 10.10.9.1, and its port number is 88. The KDC vendor is Other to indicate a UNIX vendor. The IP address of the administrative server is 10.10.9.1, and its port number is 749 (the default). The IP address of the password server is 10.10.9.1, and its port number is 464 (the default). The encryption type is DES. The comment is "UNIX Kerberos config".

```
vs1::> vserver services kerberos-realm create -configname SECURE
-realm SECURITY.EXAMPLE.COM. -clock-skew 300 -adserver-name SUSAN -
adserver-ip 10.10.3.1
-kdc-ip 10.10.9.1 -kdc-port 88 -kdc-vendor Other -adminserver-ip
10.10.9.1 -adminserver-port 749
```

```
-passwordserver-ip 10.10.9.1 -passwordserver-port 464 -comment
"UNIX Kerberos config"
```

Creating an NFS Kerberos configuration

You can use the `vserver nfs kerberos-config modify` command to enable Kerberos and create a Kerberos configuration for Storage Virtual Machines (SVMs). This enables the SVM to use Kerberos security services for NFS.

About this task

For best performance, the SVM should have at least two data LIFs. One for use by the SPN for UNIX and CIFS-related Kerberos traffic, and another one for non-Kerberos traffic.

If you are using Microsoft Active Directory Kerberos, the first 15 characters of any service principal names (SPNs) used in the domain must be unique. Microsoft Active Directory (AD) has a limitation for SPNs of 15 characters maximum and does not allow duplicate SPNs.

Step

1. Enter the following command:

```
vserver nfs kerberos-config modify -vserver vserver_name -lif
logical_interface -kerberos {enabled|disabled} -spn
service_principal_name
```

If you need to create the SPN in a different OU of the Kerberos realm, you can specify the optional `-ou` parameter.

See the man page for the command for more information.

Example

The following example creates an NFS Kerberos configuration for the SVM named `vs1` on the logical interface `ves03-d1` with the SPN `nfs/ves03-d1.lab.example.com@TEST.LAB.EXAMPLE.COM` in the OU `lab2ou`.

```
vs1::> vserver nfs kerberos-config modify -lif ves03-d1 -vserver
vs2 -kerberos enabled -spn nfs/ves03-
d1.lab.example.com@TEST.LAB.EXAMPLE.COM -ou "ou=lab2ou"
```

Configuring SVMs to use LDAP

You can configure Storage Virtual Machines (SVMs) to use LDAP servers in your environment for obtaining network information. You modify the `-ns-switch` parameter to configure the network

information sources and the `-nm-switch` parameter to configure the name mapping sources for SVMs.

Step

1. Perform one of the following actions:

If you want to configure SVMs to use LDAP as a...	Enter the command...
Network information source	<code>vserver modify -vserver vserver_name -ns-switch file,ldap</code>
Network information source and for name mapping.	<code>vserver modify -vserver vserver_name -ns-switch file,ldap -nm-switch file,ldap</code>

`-ns-switch` specifies the sources to search for network information and in what order.

`-nm-switch` specifies the sources to search for name mapping information and in what order. In a UNIX-only environment, this switch is not necessary. Name mapping is only required in a mixed environment using both UNIX and Windows.

See the man page for the command for more information.

Using LDAP over SSL/TLS to secure communication

You can use LDAP over SSL/TLS to secure communication between the Storage Virtual Machine (SVM) LDAP client and the LDAP server. This allows LDAP to encrypt all traffic to and from the LDAP server.

LDAP over SSL/TLS concepts

You must understand certain terms and concepts about how Data ONTAP uses SSL/TLS to secure LDAP communication. Data ONTAP can use LDAP over SSL/TLS for setting up authenticated sessions between Active Directory-integrated LDAP servers or UNIX-based LDAP servers.

Terminology

There are certain terms that you should understand about how Data ONTAP uses LDAP over SSL to secure LDAP communication.

LDAP (Lightweight Directory Access Protocol) A set of protocols for accessing and managing information directories. LDAP is used as information directory for storing objects such as users, groups, and netgroups. LDAP also provides directory services that manage these objects and fulfill LDAP requests from LDAP clients.

SSL	(Secure Sockets Layer) A secure protocol developed for sending information securely over the Internet. SSL is used to provide either server or mutual (server and client) authentication. SSL provides encryption only. If a method to ensure data integrity is needed, it must be provided by the application using SSL.
TLS	(Transport Layer Security) An IETF standards track protocol that is based on the earlier SSL specifications. It is the successor to SSL.
LDAP over SSL/TLS	(Also known as <i>LDAPS</i>) A protocol that uses SSL or TLS to secure communication between LDAP clients and LDAP servers. The terms <i>SSL</i> and <i>TLS</i> are often used interchangeably unless referring to a specific version of the protocol.
Start TLS	(Also known as <i>start_tls</i> , <i>STARTTLS</i> , and <i>StartTLS</i>) A mechanism to provide secure communication by using the TLS/SSL protocols.

How Data ONTAP uses LDAP over SSL/TLS

By default, LDAP communications between client and server applications are not encrypted. This means that it is possible to use a network monitoring device or software and view the communications between LDAP client and server computers. This is especially problematic when an LDAP simple bind is used because the credentials (user name and password) used to bind the LDAP client to the LDAP server are passed over the network unencrypted.

The SSL and TLS protocols run above TCP/IP and below higher-level protocols, such as LDAP. They use TCP/IP on behalf of the higher-level protocols, and in the process, permit an SSL-enabled server to authenticate itself to an SSL-enabled client and permit both machines to establish an encrypted connection. These capabilities address fundamental security concerns about communication over the Internet and other TCP/IP networks. Data ONTAP uses the START TLS method to set up the secured connection.

Data ONTAP supports SSL server authentication, which enables the Storage Virtual Machine (SVM) LDAP client to confirm the LDAP server's identity during the bind operation. SSL/TLS-enabled LDAP clients can use standard techniques of public-key cryptography to check that a server's certificate and public ID are valid and have been issued by a certificate authority (CA) listed in the client's list of trusted CAs.

This version of Data ONTAP supports the following:

- LDAP over SSL/TLS for SMB-related traffic between the Active Directory-integrated LDAP servers and the SVM
- LDAP over SSL/TLS for LDAP traffic for name mapping
Either Active Directory-integrated LDAP servers or UNIX-based LDAP servers can be used to store information for LDAP name mapping.
- Self-signed root CA certificates
When using an Active-Directory integrated LDAP, the self-signed root certificate is generated when the Windows Server Certificate Service is installed in the domain. When using an UNIX-based LDAP server for LDAP name mapping, the self-signed root certificate is generated and saved by using means appropriate to that LDAP application.

Data ONTAP does not support signing (integrity protection) and sealing (encryption) of the data. The default is not to enable LDAP over SSL/TLS.

Installing the self-signed root CA certificate on the SVM

Before you can use secure LDAP authentication when binding to LDAP servers, you must install the self-signed root CA certificate on the Storage Virtual Machine (SVM).

Steps

1. Install the self-signed root CA certificate:

- a) Enter the following command to begin the certificate installation:

```
security certificate install -vserver vservice_name -type server-ca
```

The console output displays the following message:

Please enter Certificate: Press <Enter> when done

- b) Open the certificate .pem file with a text editor, copy the certificate, including the lines beginning with -----BEGIN CERTIFICATE----- and ending with -----END CERTIFICATE-----, and paste the certificate on the console.
 - c) Verify that the certificate is displayed after the console prompt.
 - d) To complete the installation, press **Enter**.
2. Verify that the certificate is installed:

```
security certificate show -vserver vservice_name
```

Creating a new LDAP client schema

Data ONTAP provides three LDAP schemas: one for Active Directory Services for UNIX compatibility, one for Active Directory Identity Management for UNIX compatibility, and one for RFC-2307 LDAP compatibility. If the LDAP schema in your environment differs from these, you must create a new LDAP client schema for Data ONTAP.

About this task

The default LDAP schemas provided by Data ONTAP cannot be modified. To create a new schema, you create a copy and then modify the copy accordingly.

Steps

1. Display the existing LDAP client schema templates to identify the one you want to copy:

```
vserver services ldap client schema show
```

2. Set the privilege level to advanced:

```
set -privilege advanced
```

3. Make a copy of an existing LDAP client schema:

```
vserver services ldap client schema copy -vserver vserver_name -schema
existing_schema_name -new-schema-name new_schema_name
```

4. Use the `vserver services ldap client schema modify` command to modify the new schema and customize it for your environment.

See the man page for the command for more information.

5. Return to the admin privilege level:

```
set -privilege admin
```

Creating an LDAP client configuration

You can use the `vserver services ldap client create` command to create an LDAP client configuration. You must set up an LDAP client first to be able to use LDAP services.

Steps

1. To create an LDAP client configuration, enter the following command:

```
vserver services ldap client create -vserver vserver_name -client-config
client_config_name {-servers LDAP_server_list | -ad-domain ad_domain -
preferred-ad-servers preferred_ad_server_list -bind-as-cifs-server
{true|false}} -schema schema -port port -query-timeout integer -min-
bind-level {anonymous|simple|sasl} -bind-dn LDAP_DN -bind-password
password -base-dn LDAP_DN -base-scope {base|onelevel|subtree}
```

`-vserver vserver_name` specifies the name of the SVM for which you want to create an LDAP client configuration.

`-client-config client_config_name` specifies the name of the new LDAP client configuration.

`-servers LDAP_server_list` specifies one or more LDAP servers by IP address in a comma-delimited list.

`-ad-domain ad_domain` specifies the AD domain.

`-preferred-ad-servers preferred_ad_server_list` specifies one or more preferred Active Directory servers by IP address in a comma-delimited list.

`-bind-as-cifs-server {true|false}` specifies whether to bind using the CIFS credentials of the SVM. The default is `false`. If you want to set this parameter to `true`, you must first install a CIFS license and create a CIFS server.

`-schema schema` specifies the schema template to use. You can either use one of the provided default schemas or create your own schema by copying a default schema (they are read-only) and modifying the copy.

`-port port` specifies the LDAP server port. The default is 389.

`-query-timeout integer` specifies the query timeout in seconds. The allowed range is 0-10 seconds. The default is 3 seconds.

`-min-bind-level {anonymous|simple|sasl}` specifies the minimum bind authentication level. The default is `anonymous`.

`-bind-dn LDAP_DN` specifies the Bind user. For Active Directory servers, specify the user in the account (`DOMAIN\user`) or principal (`user@domain.com`) form. Otherwise, specify the user in distinguished name (`CN=user,DC=domain,DC=com`) form.

`-bind-password password` specifies the Bind password.

`-base-dn LDAP_DN` specifies the base DN. The default is `" "` (root).

`-base-scope {base|onelevel|subtree}` specifies the base search scope. The default is `subtree`.

`-use-start-tls {true|false}` specifies whether to use Start TLS for the LDAP connection. If set to `true` and the LDAP server supports it, the LDAP client uses an encrypted TLS/SSL connection to the server. The default is `false`. You must install a self-signed root CA certificate of the LDAP server to use this option.

2. Verify that the LDAP client configuration was created successfully:

```
vserver services ldap client show -client-config client_config_name
```

Examples

The following command creates a new LDAP client configuration named “ldap1” to work with an Active Directory server for LDAP:

```
cluster1::> vserver services ldap client create -vserver vs1 -
client-config ldapclient1 -ad-domain addomain.example.com -bind-as-
cifs-server true -schema AD-SFU -port 389 -query-timeout 3 -min-
bind-level simple -base-dn DC=addomain,DC=example,DC=com -base-
scope subtree -preferred-ad-servers 172.17.32.100
```

The following command modifies the LDAP client configuration named “ldap1” to work with an Active Directory server for LDAP by specifying an Active Directory domain and to bind to the server using the CIFS credentials of the SVM:

```
cluster1::> vserver services ldap client modify -vserver vs1 -
client-config ldap1 -bind-as-cifs-server true -ad-domain
addomain.example.com
```

The following command modifies the LDAP client configuration named “ldap1” by specifying the base DN:

```
cluster1::> vserver services ldap client modify -vserver vs1 -
client-config ldap1 -base-dn CN=Users,DC=addomain,DC=example,DC=com
```

Enabling LDAP on SVMs

You can use the `vserver services ldap create` command to enable LDAP on Storage Virtual Machines (SVMs) and configure it to use an LDAP client you created. This enables SVMs to use LDAP for name mapping.

Before you begin

An LDAP client configuration must exist on the SVM.

Step

1. To enable LDAP on SVMs, enter the following command:

```
vserver services ldap create -vserver vserver_name -client-config  
client_config_name -client-enabled true
```

`-vserver vserver_name` specifies the SVM name.

`-client-config client_config_name` specifies the client configuration name.

`-client-enabled` specifies whether the LDAP client is enabled. The default is `true`.

Example

The following command enables LDAP on the SVM named “vs1” and configures it to use the LDAP client configuration named “ldap1”.

```
cluster1::> vserver services ldap create -vserver vs1 -client-  
config ldap1 -client-enabled true
```

How name mappings are used

Data ONTAP uses name mapping to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, and UNIX identities to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client or a CIFS client.

Name mapping is usually required due to the multiprotocol nature of Data ONTAP, which supports CIFS and NFS client access to the same data. Data stored on Storage Virtual Machines (SVMs) with FlexVol volumes uses either UNIX- or NTFS-style permissions. To authorize a client, the credentials must match the security style. Consider the following scenarios:

If a CIFS client wants to access data with UNIX-style permissions, Data ONTAP cannot directly authorize the client because it cannot use CIFS credentials with UNIX-style permissions. To properly authorize the client request, Data ONTAP must first map the CIFS credentials to the appropriate UNIX credentials so that it can then use the UNIX credentials to compare them to the UNIX-style permissions.

If an NFS client wants to access data with NTFS-style permissions, Data ONTAP cannot directly authorize the client because it cannot use UNIX credentials with NTFS-style permissions. To properly authorize the client request, Data ONTAP must first map the UNIX credentials to the appropriate NTFS credentials so that it can then use the NTFS credentials to compare them to the NTFS-style permissions.

There are two exceptions where you do not have to use name mapping:

- You configure a pure UNIX environment and do not plan to use CIFS access or NTFS security style on volumes.
In this scenario, name mapping is not required because Data ONTAP can use the UNIX credentials of the NFS clients to directly compare them to the UNIX-style permissions.
- You configure the default user to be used instead.
In this scenario, name mapping is not required because instead of mapping every individual client credential all client credentials are mapped to the same default user.

Note that you can use name mapping only for users, not for groups. It is not possible to map CIFS users to a group ID (GID), or UNIX users to a group in the Active Directory (AD). Similarly, it is not possible to map a GID to a group or a user in AD, or an AD group to a UNIX UID or GID.

However, you can map a group of individual users to a specific user. For example, you can map all AD users that start or end with the word SALES to a specific UNIX user and to the user's UID. As a result, you can rename certain users in AD and use regular expressions to effectively emulate group actions. This type of mapping also works in reverse.

How name mapping works

Data ONTAP goes through a number of steps when attempting to map user names. They include checking the local name mapping database and LDAP, trying the user name, and using the default user if configured.

When Data ONTAP has to map credentials for a user, it first checks the local name mapping database and LDAP server for an existing mapping. Whether it checks one or both and in which order is determined by the `-nm-switch` parameter of the Storage Virtual Machine (SVM) configuration.

- For Windows to UNIX mapping
If no mapping is found, Data ONTAP checks whether the lowercase Windows user name is a valid user name in the UNIX domain. If this does not work, it uses the default UNIX user provided it is configured. If the default UNIX user is not configured and it cannot obtain a mapping this way either, mapping fails and an error is returned.
- For UNIX to Windows mapping
If no mapping is found, Data ONTAP tries to find a Windows account that matches the UNIX name in the CIFS domain. If this does not work, it uses the default CIFS user, provided it is configured. If the default CIFS user is not configured and it cannot obtain a mapping this way either, mapping fails and an error is returned.

Multidomain searches for UNIX user to Windows user name mappings

Data ONTAP supports multidomain searches when mapping UNIX users to Windows users. All discovered trusted domains are searched for matches to the replacement pattern until a matching result is returned. Alternatively, you can configure a list of preferred trusted domains, which is used instead of the discovered trusted domain list and is searched in order until a matching result is returned.

How domain trusts affect UNIX user to Windows user name mapping searches

To understand how multidomain user name mapping works, you must understand how domain trusts work with Data ONTAP. Active Directory trust relationships with the CIFS server's home domain can be a bidirectional trust or can be one of two types of unidirectional trusts, either an inbound trust or an outbound trust. The home domain is the domain to which the CIFS server on the Storage Virtual Machine (SVM) belongs.

- Bidirectional trust***
 With bidirectional trusts, both domains trust each other. If the CIFS server's home domain has a bidirectional trust with another domain, the home domain can authenticate and authorize a user belonging to the trusted domain and vice versa.
 UNIX user to Windows user name mapping searches can be performed only on domains with bidirectional trusts between the home domain and the other domain.
- Outbound trust***
 With an outbound trust, the home domain trusts the other domain. In this case, the home domain can authenticate and authorize a user belonging to the outbound trusted domain.
 A domain with an outbound trust with the home domain is *not* searched when performing UNIX user to Windows user name mapping searches.
- Inbound trust***
 With an inbound trust, the other domain trusts the CIFS server's home domain. In this case, the home domain cannot authenticate or authorize a user belonging to the inbound trusted domain.
 A domain with an inbound trust with the home domain is *not* searched when performing UNIX user to Windows user name mapping searches.

How wildcards (*) are used to configure multidomain searches for name mapping

Multidomain name mapping searches are facilitated by the use of wildcards in the domain section of the Windows user name. The following table illustrates how to use wildcards in the domain part of a name mapping entry to enable multidomain searches:

Pattern	Replacement	Result
root	*\\administrator	The UNIX user “root” is mapped to the user named “administrator”. All trusted domains are searched in order until the first matching user named “administrator” is found.

Pattern	Replacement	Result
*	**	Valid UNIX users are mapped to the corresponding Windows users. All trusted domains are searched in order until the first matching user with that name is found.

How multidomain name searches are performed

You can choose one of two methods for determining the list of trusted domains used for multidomain name searches:

- Use the automatically discovered bidirectional trust list compiled by Data ONTAP
 - The advantage to this method is that there is no management overhead and that the list is made of trusted domains that Data ONTAP has determined are valid.
 - The disadvantage is that you cannot choose the order that the trusted domains are searched.
- Use the preferred trusted domain list that you compile
 - The advantage to this method is that you can configure the list of trusted domains in the order that you want them searched.
 - The disadvantage is that there is more management overhead and that the list might become outdated, with some listed domains not being valid, bidirectionally trusted domains.

If a UNIX user is mapped to a Windows user with a wildcard used for the domain section of the user name, the Windows user is looked up in all the trusted domains as follows:

- If a preferred trusted-domain list is configured, the mapped Windows user is looked up in this search list only, in order.
 The search ends as soon as the Windows user is found. If the same Windows user name exists in two different trusted domains, then the user belonging to the domain listed first in the preferred trusted-domain list is returned. If the Windows user is not found in any domains in the preferred list, an error is returned.
 If you want the home domain to be included in the search, it must be included in the preferred trusted domain list.
- If a preferred list of trusted domains is not configured, then the Windows user is looked up in all the bidirectional trusted domains of the home domain.
 The search ends as soon as the Windows user is found. If the same Windows user name exists in two different trusted domains, the user belonging to the domain listed first in the automatically discovered trusted-domain list is returned. You cannot control the order of the trusted domains in the automatically discovered list. If the Windows user is not found in any of the discovered trusted domains, the user is then looked up in the home domain.
- If there are no bidirectionally trusted domains for the home domain, the user is looked up in the home domain.

For more information about managing the list of bidirectional trusted domains, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

Name mapping conversion rules

A Data ONTAP system keeps a set of conversion rules for each Storage Virtual Machine (SVM). Each rule consists of two pieces: a *pattern* and a *replacement*. Conversions start at the beginning of the appropriate list and perform a substitution based on the first matching rule. The pattern is a UNIX-style regular expression. The replacement is a string containing escape sequences representing subexpressions from the pattern, as in the UNIX `sed` program.

It is possible to allow NFS access to volumes with NTFS security style for users in a different domain from the one that the storage system belongs to, provided that the proper name mapping rule exists.

If a user matches a rule to map to a user in a different domain, the domain must be trusted. To ensure successful mapping to users in other domains for both SMB and NFS access, there must be a bidirectional trust relationship between the domains.

If a user matches a rule but the user cannot authenticate in the other domain because it is untrusted, the mapping fails.

The SVM automatically discovers all bidirectional trusted domains, which are used for multi-domain user mapping searches. Alternatively, you can configure a list of preferred trusted domains that are used for name mapping searches instead of the list of automatically discovered trusted domains.

Regular expressions are not case-sensitive when mapping from Windows to UNIX. However, they are case-sensitive for Kerberos-to-UNIX and UNIX-to-Windows mappings.

As an example, the following rule converts the Windows user named “jones” in the domain named “ENG” into the UNIX user named “jones”.

Pattern	Replacement
ENG\\jones	jones

Note that the backslash is a special character in regular expressions and must be escaped with another backslash.

The caret (^), underscore (_), and ampersand (&) characters can be used as prefixes for digits in replacement patterns. These characters specify uppercase, lowercase, and initial-case transformations, respectively. For instance:

- If the initial pattern is `(.+)` and the replacement pattern is `\1`, then the string `jOe` is mapped to `jOe` (no change).
- If the initial pattern is `(.+)` and the replacement pattern is `_1`, then the string `jOe` is mapped to `joe`.
- If the initial pattern is `(.+)` and the replacement pattern is `\^1`, then the string `jOe` is mapped to `JOE`.
- If the initial pattern is `(.+)` and the replacement pattern is `\&1`, then the string `jOe` is mapped to `Joe`.

If the character following a backslash-underscore (_), backslash-caret (\^), or backslash-ampersand (\&) sequence is not a digit, then the character following the backslash is used verbatim.

The following example converts any Windows user in the domain named “ENG” into a UNIX user with the same name in NIS.

Pattern	Replacement
ENG\\(.+)	\1

The double backslash (\\) matches a single backslash. The parentheses denote a subexpression but do not match any characters themselves. The period matches any single character. The asterisk matches zero or more of the previous expression. In this example, you are matching ENG\ followed by one or more of any character. In the replacement, \1 refers to whatever the first subexpression matched. Assuming the Windows user ENG\jones, the replacement evaluates to jones; that is, the portion of the name following ENG\.

Note: If you are using the CLI, you must delimit all regular expressions with double quotation marks ("). For instance, to enter the regular expression (.+) in the CLI, type "(.+) " at the command prompt. Quotation marks are not required in the Web UI.

For further information about regular expressions, see your UNIX system administration documentation, the online UNIX documentation for `sed` or `regex`, or *Mastering Regular Expressions*, published by O'Reilly and Associates.

Creating a name mapping

You can use the `vserver name-mapping create` command to create a name mapping. You use name mappings to enable Windows users to access UNIX security style volumes and the reverse.

About this task

For each Storage Virtual Machine (SVM), Data ONTAP supports up to 1,024 name mappings for each direction.

Step

1. To create a name mapping, enter the following command:

```
vserver name-mapping create -vserver vserver_name -direction {krb-unix|win-unix|unix-win} -position integer -pattern text -replacement text
```

`-vserver vserver_name` specifies the SVM name.

`-direction {krb-unix|win-unix|unix-win}` specifies the mapping direction.

`-position integer` specifies the desired position in the priority list of a new mapping.

`-pattern text` specifies the pattern to be matched, up to 256 characters in length.

`-replacement text` specifies the replacement pattern, up to 256 characters in length.

When Windows-to-UNIX mappings are created, any CIFS clients that have open connections to the Data ONTAP system at the time the new mappings are created must log out and log back in to see the new mappings.

Examples

The following command creates a name mapping on the SVM named vs1. The mapping is a mapping from UNIX to Windows at position 1 in the priority list. The mapping maps the UNIX user johnd to the Windows user ENG\John.

```
vs1::> vserver name-mapping create -vserver vs1 -direction unix-win
-position 1 -pattern johnd -replacement "ENG\John"
```

The following command creates another name mapping on the SVM named vs1. The mapping is a mapping from Windows to UNIX at position 1 in the priority list. The mapping maps every CIFS user in the domain ENG to users in the NIS domain associated with the SVM.

```
vs1::> vserver name-mapping create -vserver vs1 -direction win-unix
-position 1 -pattern "ENG\\(.+)"
-replacement "\1"
```

Configuring the default user

You can configure a default user to use if all other mapping attempts fail for a user, or if you do not want to map individual users between UNIX and Windows. Alternatively, if you want authentication of non-mapped users to fail, you should not configure a default user.

About this task

For CIFS authentication, if you do not want to map each Windows user to an individual UNIX user, you can instead specify a default UNIX user.

For NFS authentication, if you do not want to map each UNIX user to an individual Windows user, you can instead specify a default Windows user.

Step

- 1. Perform one of the following actions:

If you want to...	Enter the following command...
Configure the default UNIX user	vserver cifs options modify -default-unix-user <i>user_name</i>
Configure the default Windows user	vserver nfs modify -default-win-user <i>user_name</i>

Configuring local UNIX users and groups

You can use local UNIX users and groups for authentication and name mappings.

Creating a local UNIX user

You can use the `vserver services unix-user create` command to create local UNIX users. A local UNIX user is a UNIX user you create on the Storage Virtual Machine (SVM) as a UNIX name services option and to be used in the processing of name mappings.

Step

1. To create a local UNIX user, enter the following command:
`vserver services unix-user create -vserver vserver_name -user user_name -id integer -primary-gid integer -full-name full_name`
`-vserver vserver_name` specifies the SVM name.
`-user user_name` specifies the user name.
`-id integer` specifies the user ID.
`-primary-gid integer` specifies the primary group ID.
`-full-name full_name` specifies the full name of the user.

Example

The following command creates a local UNIX user named `johnm` on the SVM named `vs1`. The user has the ID 123 and the primary group ID 100. The user's full name is "John Miller".

```
node::> vserver services unix-user create -vserver vs1 -user johnm -
id 123
-primary-gid 100 -full-name "John Miller"
```

Loading local UNIX users from a URI

You can use the `vserver services unix-user load-from-uri` command to load one or more local UNIX users into Storage Virtual Machines (SVMs) from a uniform resource identifier (URI).

About this task

The URI must contain user information in the UNIX `/etc/passwd` format:

`user_name: password: user_ID: group_ID: full_name`

The command discards the value of the `password` field and of the fields after the `full_name` field (`home_directory` and `shell`).

Step

1. To load one or more local UNIX users into SVMs from a URI, enter the following command:

```
vserver services unix-user load-from-uri -vserver vserver_name -uri {ftp|http}://uri -overwrite {true|false}
```

`-vserver vserver_name` specifies the SVM name.

`-uri {ftp|http}://uri` specifies the URI to load from.

`-overwrite {true|false}` specifies whether to overwrite entries. The default is `false`.

Example

The following command loads user information from the URI `ftp://ftp.example.com/passwd` into the SVM named `vs1`. Existing users on the SVM are not overwritten by information from the URI.

```
node::> vserver services unix-user load-from-uri -vserver vs1
-uri ftp://ftp.example.com/passwd -overwrite false
```

Creating a local UNIX group

You can use the `vserver services unix-group create` command to create UNIX groups that are local to the Storage Virtual Machine (SVM). Local UNIX groups are used with local UNIX users.

Step

1. To create a local UNIX group, enter the following command:

```
vserver services unix-group create -vserver vserver_name -name group_name -id integer
```

`-vserver vserver_name` specifies the SVM name.

`-name group_name` specifies the group name.

`-id integer` specifies the group ID.

Example

The following command creates a local group named `eng` on the SVM named `vs1`. The group has the ID 101.

```
vs1::> vserver services unix-group create -vserver vs1 -name eng -
id 101
```

Loading local UNIX groups from a URI

You can use the `vserver services unix-group load-from-uri` command to load one or more local UNIX groups into Storage Virtual Machines (SVMs) from a uniform resource identifier (URI).

About this task

The URI must contain user information in the UNIX `/etc/group` format:

group_name: password: group_ID: comma_separated_list_of_users

The command discards the value of the *password* field.

Step

1. To load one or more local UNIX groups into SVMs from URI, enter the following command:

```
vserver services unix-group load-from-uri -vserver vserver_name -uri {ftp|http}://uri -overwrite {true|false}
```

`-vserver vserver_name` specifies the SVM name.

`-uri {ftp|http}://uri` specifies the URI to load from.

`-overwrite {true|false}` specifies whether to overwrite entries. The default is `false`.

Example

The following command loads group information from the URI `ftp://ftp.example.com/group` into the SVM named `vs1`. Existing groups on the SVM are not overwritten by information from the URI.

```
vs1::> vserver services unix-group load-from-uri -vserver vs1  
-uri ftp://ftp.example.com/group -overwrite false
```

Adding a user to a local UNIX group

You can use the `vserver services unix-group adduser` command to add a user to a UNIX group that is local to the Storage Virtual Machine (SVM).

Step

1. To add a user to a local UNIX group, enter the following command:

```
vserver services unix-group adduser -vserver vserver_name -name group_name -username user_name
```

`-vserver vserver_name` specifies the SVM name.

`-name group_name` specifies the name of the UNIX group to add the user to.

`-username user_name` specifies the user name of the user to add to the group.

Example

The following command adds a user named max to a local UNIX group named eng on the SVM named vs1.

```
vs1::> vserver services unix-group adduser -vserver vs1 -name eng
-username max
```

Loading netgroups into SVMs

You can use the `vserver services netgroup load` command to load netgroups into Storage Virtual Machines (SVMs) from a uniform resource identifier (URI).

Step

1. To load netgroups into SVMs from an FTP or HTTP URI, enter the following command:

```
vserver services netgroup load -vserver vserver_name -source {ftp|
http}://uri
```

`-vserver vserver_name` specifies the SVM name.

`-source {ftp|http}://uri` specifies the URI to load from.

The file must use the same proper netgroup text file format that is used to populate NIS.

Example

The following command loads netgroup definitions into the SVM named vs1 from the HTTP URL `http://intranet/downloads/corp-netgroup`:

```
vs1::> vserver services netgroup load -vserver vs1
-source http://intranet/downloads/corp-netgroup
```

Creating a NIS domain configuration

If you specified NIS as a name service option during Storage Virtual Machine (SVM) setup, you must create a NIS domain configuration for the SVM. You can use the `vserver services nis-domain create` command to create a NIS domain configuration.

About this task

You can create multiple NIS domains. However, you can only use one that is set to `active`.

Step

1. Use the `vserver services nis-domain create` command to create a NIS domain configuration.

Example

The following command creates a NIS domain configuration for a NIS domain called `nisdomain` on the SVM named `vs1` with a NIS server at IP address `192.0.2.180` and makes it `active`.

```
vs1::> vserver services nis-domain create -vserver vs1 -domain  
nisdomain -active true -servers 192.0.2.180
```

Support for NFS over IPv6

Starting with Data ONTAP 8.2, NFS clients can access files on your storage system over an IPv6 network.

Enabling IPv6 for NFS

If you want NFS clients to be able to access data on the storage system over an IPv6 network, Data ONTAP requires several configuration changes.

Steps

1. Enable IPv6 on the cluster.

This step must be performed by a cluster administrator. For more information about enabling IPv6 on the cluster, see the *Clustered Data ONTAP Network Management Guide*.

2. Configure data LIFs for IPv6.

This step must be performed by a cluster administrator. For more information about configuring LIFs, see the *Clustered Data ONTAP Network Management Guide*.

3. Configure NFS export policies and rules for NFS client access.

If you want to match clients by IPv6 address, be sure to configure the export rules accordingly.

Where to find information about setting up file access to Infinite Volumes

For more information about how to set up NFS file access to Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

Managing file access using NFS

After you have enabled NFS on the Storage Virtual Machine (SVM) and configured it, there are a number of tasks you might want to perform to manage file access using NFS.

Controlling NFS requests from nonreserved ports

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

About this task

By default, the option `-mount-rootonly` is enabled.

By default, the option `-nfs-rootonly` is disabled.

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>vserver nfs modify -vserver vserver_name -mount-rootonly disabled</code>
Reject NFS mount requests from nonreserved ports	<code>vserver nfs modify -vserver vserver_name -mount-rootonly enabled</code>
Allow all NFS requests from nonreserved ports	<code>vserver nfs modify -vserver vserver_name -nfs-rootonly disabled</code>
Reject all NFS requests from nonreserved ports	<code>vserver nfs modify -vserver vserver_name -nfs-rootonly enabled</code>

Considerations for clients that mount NFS exports using a nonreserved port

The `-mount-rootonly` option must be disabled on a storage system that must support clients that mount NFS exports using a nonreserved port even when the user is logged in as root. Such clients include Hummingbird clients and Solaris NFS/IPv6 clients.

If the `-mount-rootonly` option is enabled, Data ONTAP does not allow NFS clients that use nonreserved ports, meaning ports with numbers higher than 1,023, to mount NFS exports.

Commands for managing NFS servers

There are specific Data ONTAP commands for managing NFS servers.

If you want to...	Use this command...
Create an NFS server	<code>vserver nfs create</code>
Display NFS servers	<code>vserver nfs show</code>
Modify an NFS server	<code>vserver nfs modify</code>
Delete an NFS server	<code>vserver nfs delete</code>

See the man page for each command for more information.

Commands for managing name mappings

There are specific Data ONTAP commands for managing name mappings.

If you want to...	Use this command...
Create a name mapping	<code>vserver name-mapping create</code>
Insert a name mapping at a specific position	<code>vserver name-mapping insert</code>
Display name mappings	<code>vserver name-mapping show</code>
Exchange the position of two name mappings	<code>vserver name-mapping swap</code>
Modify a name mapping	<code>vserver name-mapping modify</code>
Delete a name mapping	<code>vserver name-mapping delete</code>

See the man page for each command for more information.

Commands for managing local UNIX users

There are specific Data ONTAP commands for managing local UNIX users.

If you want to...	Use this command...
Create a local UNIX user	<code>vserver services unix-user create</code>
Display local UNIX users	<code>vserver services unix-user show</code>
Modify a local UNIX user	<code>vserver services unix-user modify</code>
Delete a local UNIX user	<code>vserver services unix-user delete</code>

See the man page for each command for more information.

Commands for managing local UNIX groups

There are specific Data ONTAP commands for managing local UNIX groups.

If you want to...	Use this command...
Add a user to a local UNIX group	<code>vserver services unix-group adduser</code>
Create a local UNIX group	<code>vserver services unix-group create</code>
Display local UNIX groups	<code>vserver services unix-group show</code>
Modify a local UNIX group	<code>vserver services unix-group modify</code>
Delete a user from a local UNIX group	<code>vserver services unix-group deluser</code>
Delete a local UNIX group	<code>vserver services unix-group delete</code>

See the man page for each command for more information.

Verifying the status of netgroup definitions

After loading netgroups into the Storage Virtual Machine (SVM), you can use the `vserver services netgroup status` command to verify the status of netgroup definitions. This enables you to determine whether netgroup definitions are consistent on all of the nodes that back the SVM.

Steps

1. Set the privilege level to advanced:
`set -privilege advanced`

2. Verify the status of netgroup definitions:

```
vserver services netgroup status
```

The command displays the following information:

- SVM name
- Node name
- Load time for netgroup definitions
- Hash value of the netgroup definitions

You can display additional information in a more detailed view. See the reference page for the command for details.

3. Return to the admin privilege level:

```
set -privilege admin
```

Example

The following command displays netgroup status for all SVMs.

```
vsl:> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only when
        directed to do so by technical support.
Do you wish to continue? (y or n): y

vsl:> vserver services netgroup status
Virtual
Server      Node           Load Time           Hash Value
-----
vsl
            node1        9/20/2006 16:04:53   e6cb38ec1396a280c0d2b77e3a84eda2
            node2        9/20/2006 16:06:26   e6cb38ec1396a280c0d2b77e3a84eda2
            node3        9/20/2006 16:08:08   e6cb38ec1396a280c0d2b77e3a84eda2
            node4        9/20/2006 16:11:33   e6cb38ec1396a280c0d2b77e3a84eda2
```

Commands for managing NIS domain configurations

There are specific Data ONTAP commands for managing NIS domain configurations.

If you want to...	Use this command...
Create a NIS domain configuration	vserver services nis-domain create
Display NIS domain configurations	vserver services nis-domain show
Modify a NIS domain configuration	vserver services nis-domain modify
Delete a NIS domain configuration	vserver services nis-domain delete

See the man page for each command for more information.

Commands for managing LDAP client configurations

There are specific Data ONTAP commands for managing LDAP client configurations.

If you want to...	Use this command...
Create an LDAP client configuration	<code>vserver services ldap client create</code>
Display LDAP client configurations	<code>vserver services ldap client show</code>
Modify an LDAP client configuration	<code>vserver services ldap client modify</code>
Delete an LDAP client configuration	<code>vserver services ldap client delete</code>

Note: SVM administrators cannot modify or delete LDAP client configurations that were created by cluster administrators.

See the man page for each command for more information.

Commands for managing LDAP configurations

There are specific Data ONTAP commands for managing LDAP configurations.

If you want to...	Use this command...
Create an LDAP configuration	<code>vserver services ldap create</code>
Display LDAP configurations	<code>vserver services ldap show</code>
Modify an LDAP configuration	<code>vserver services ldap modify</code>
Delete an LDAP configuration	<code>vserver services ldap delete</code>

See the man page for each command for more information.

Commands for managing LDAP client schema templates

There are specific Data ONTAP commands for managing LDAP client schema templates.

Note: The `vserver services ldap client schema copy`, `modify`, and `delete` commands are only available at the advanced privilege level and higher.

If you want to...	Use this command...
Copy an existing LDAP schema template	<code>vserver services ldap client schema copy</code>
Display LDAP schema templates	<code>vserver services ldap client schema show</code>
Modify an LDAP schema template	<code>vserver services ldap client schema modify</code>
Delete an LDAP schema template	<code>vserver services ldap client schema delete</code>

Note: SVM administrators cannot modify or delete LDAP client schemas that were created by cluster administrators.

See the man page for each command for more information.

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 volume or qtree. This results in performance improvements due to less time used for DNS lookups.

Whenever an NFS client attempts to access a volume or qtree, Data ONTAP must determine whether to grant or deny access. Except in the most simple cases (for example, when a volume or qtree 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 volume or qtree
- 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 for this particular NFS client volume or qtree 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 volume or qtree export rules

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

Displaying information about NFS Kerberos configurations

You can use the `vserver nfs kerberos-config show` command to display information about NFS Kerberos configurations. This enables you to determine how NFS Kerberos is configured.

Step

1. To display information about NFS Kerberos configurations, enter the following command:

```
vserver nfs kerberos-config show
```

The command displays the following information:

- Virtual server name
- Logical interface name
- Logical interface IP address
- Whether Kerberos is enabled or disabled
- Kerberos SPN
- Numeric ID of the configuration

You can display additional information in a more detailed view. See the reference page for the command for details.

Example

The following command displays detailed information about an NFS Kerberos configuration:

```
vs1::> vserver nfs kerberos-config show -vserver vs1
-lif datalif1

    Virtual Server: vs1
    Logical Interface: datalif1
        Ip Address: 172.19.4.1
    Kerberos Enabled: Disabled
    Service Principal Name: nfs/security.example.com@ENG.EXAMPLE.COM
```

Modifying an NFS Kerberos configuration

You can use the `vserver nfs kerberos-config modify` command to modify a Kerberos configuration for NFS. This enables you to enable or disable the NFS-enabled Storage Virtual Machine (SVM) to use Kerberos authentication.

Step

1. Use the `vserver nfs kerberos-config modify` command to modify an NFS Kerberos configuration

Examples

The following command enables an NFS Kerberos configuration on the SVM named `vs1` and a logical interface named `data1if1`. The SPN is `nfs/security.example.com@eng.example.com` and the keytab file to be loaded is at the URL `ftp://ftp.example.com/keytab`.

```
vs1::> vserver nfs kerberos-config modify -vserver vs1 -lif
data1if1 -kerberos enable -spn nfs/
security.example.com@ENG.EXAMPLE.COM -admin-username admin -keytab-
uri ftp://ftp.example.com/keytab
```

Data ONTAP then prompts the user for the password for the admin-user. The admin-user should have permission on the KDC to add the principal to the principal's database.

The following command disables the NFS Kerberos configuration that was created in the previous example.

```
vs1::> vserver nfs kerberos-config modify -vserver vs1 -lif
data1if1 -kerberos disable
```

Commands for managing Kerberos realm configurations

There are specific Data ONTAP commands for managing Kerberos realm configurations.

If you want to...	Use this command...
Create a Kerberos realm configuration	<code>vserver services kerberos-realm create</code>
Display Kerberos realm configurations	<code>vserver services kerberos-realm show</code>

If you want to...	Use this command...
Modify Kerberos realm configurations	<code>vserver services kerberos-realm modify</code>
Delete a Kerberos realm configuration	<code>vserver services kerberos-realm delete</code>

See the man page for each command for more information.

Commands for managing export policies

There are specific Data ONTAP commands for managing export policies.

If you want to...	Use this command...
Display information about export policies	<code>vserver export-policy show</code>
Rename an export policy	<code>vserver export-policy rename</code>
Copy an export policy	<code>vserver export-policy copy</code>
Delete an export policy	<code>vserver export-policy delete</code>

See the man page for each command for more information.

Commands for managing export rules

There are specific Data ONTAP commands for managing export rules.

If you want to...	Use this command...
Create an export rule	<code>vserver export-policy rule create</code>
Display information about export rules	<code>vserver export-policy rule show</code>
Modify an export rule	<code>vserver export-policy rule modify</code>
Delete an export rule	<code>vserver export-policy rule delete</code>

See the man page for each command for more information.

Managing file locks

You can display information about the current locks for a Storage Virtual Machine (SVM) as a first step to determining why a client cannot access a volume or file. You can use this information if you need to break file locks.

For information about how file locks affect Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

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 an SMB client, locks are mandatory.

Because of differences between the NFS and SMB file locks, an NFS client might fail to access a file previously opened by an SMB application.

The following occurs when an NFS client attempts to access a file locked by an SMB application:

- In mixed or NTFS volumes, file manipulation operations such as `rm`, `rmdir`, and `mv` can cause the NFS application to fail.
- NFS read and write operations are denied by SMB deny-read and deny-write open modes, respectively.
- NFS write operations fail when the written range of the file is locked with an exclusive SMB `bytelock`.

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).

SMB 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 an SMB 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 SMB clients. Some software, when used by NFS clients and SMB 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 SMB 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 an SMB 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 SMB clients, but might not take effect immediately on NFS clients if the NFS client enables attribute caching.

Displaying information about locks

You can display information about the current file locks, including what types of locks are held and what the lock state is, details about byte-range locks, sharelock modes, delegation locks, and opportunistic locks, and whether locks are opened with durable or persistent handles.

About this task

The client IP address cannot be displayed for locks established through NFSv4 or NFSv4.1.

By default, the command displays information about all locks. You can use command parameters to display information about locks for a specific Storage Virtual Machine (SVM) or to filter the command's output by other criteria. If you do not specify any parameter, the command displays the following information:

- SVM name
- Volume name of the FlexVol volume or the name of the namespace constituent for the Infinite Volume
- Path of the locked object
- Logical interface name
- Protocol by which the lock was established
- Type of lock
- Client

The `vserver locks show` command displays information about four types of locks:

- Byte-range locks, which lock only a portion of a file.
- Share locks, which lock open files.
- Opportunistic locks, which control client-side caching over SMB.
- Delegations, which control client-side caching over NFSv4.x.

By specifying optional parameters, you can determine important information about each of these lock types. See the man page for the command for more information.

Step

1. Display information about locks by using the `vserver locks show` command.

Examples

The following example displays summary information for an NFSv4 lock on a file with the path `/vol1/file1`. The sharelock access mode is `write-deny_none`, and the lock was granted with write delegation:

```
cluster1::> vserver locks show

Vserver: vs0
Volume  Object Path          LIF          Protocol Lock Type  Client
-----
vol1    /vol1/file1               lif1         nfsv4    share-level -
                Sharelock Mode: write-deny_none
                Delegation Type: write
```

The following example displays detailed oplock and sharelock information about the SMB lock on a file with the path `/data2/data2_2/intro.pptx`. A durable handle is granted on the file with a share lock access mode of `write-deny_none` to a client with an IP address of 10.3.1.3. A lease oplock is granted with a batch oplock level:

```
cluster1::> vserver locks show -instance -path /data2/data2_2/intro.pptx

Vserver: vs1
Volume: data2_2
Logical Interface: lif2
Object Path: /data2/data2_2/intro.pptx
Lock UUID: 553cf484-7030-4998-88d3-1125adbba0b7
Lock Protocol: cifs
Lock Type: share-level
Node Holding Lock State: node3
Lock State: granted
Bytelock Starting Offset: -
Number of Bytes Locked: -
Bytelock is Mandatory: -
Bytelock is Exclusive: -
Bytelock is Superlock: -
Bytelock is Soft: -
Oplock Level: -
Shared Lock Access Mode: write-deny_none
Shared Lock is Soft: false
Delegation Type: -
Client Address: 10.3.1.3
SMB Open Type: durable
SMB Connect State: connected
SMB Expiration Time (Secs): -
SMB Open Group ID:
78a90c59d45ae211998100059a3c7a00a007f70da0f8ffffcd445b0300000000

Vserver: vs1
Volume: data2_2
Logical Interface: lif2
Object Path: /data2/data2_2/test.pptx
Lock UUID: 302fd7b1-f7bf-47ae-9981-f0dcb6a224f9
Lock Protocol: cifs
Lock Type: op-lock
Node Holding Lock State: node3
Lock State: granted
Bytelock Starting Offset: -
Number of Bytes Locked: -
Bytelock is Mandatory: -
```

```
Bytelock is Exclusive: -
Bytelock is Superlock: -
Bytelock is Soft: -
Oplock Level: batch
Shared Lock Access Mode: -
Shared Lock is Soft: -
Delegation Type: -
Client Address: 10.3.1.3
SMB Open Type: -
SMB Connect State: connected
SMB Expiration Time (Secs): -
SMB Open Group ID:
78a90c59d45ae211998100059a3c7a00a007f70da0f8ffffcd445b0300000000
```

Breaking locks

When file locks are preventing client access to files, you can display information about currently held locks, and then break specific locks. Examples of scenarios in which you might need to break locks include debugging applications.

About this task

The `vserver locks break` command is available only at the advanced privilege level and higher. The man page for the command contains detailed information.

Steps

1. To find the information you need to break a lock, use the `vserver locks show` command.
The man page for the command contains detailed information.
2. Set the privilege level to advanced:
`set -privilege advanced`
3. Perform one of the following actions:

If you want to break a lock by specifying...	Enter the command...
The SVM name, volume name, LIF name, and file path	<code>vserver locks break -vserver <i>vserver_name</i> -volume <i>volume_name</i> -path <i>path</i> -lif <i>lif</i></code>
The lock ID	<code>vserver locks break -lockid <i>UUID</i></code>

`-vserver vserver_name` specifies the SVM name.

`-volume volume_name` specifies the volume name of the FlexVol volume or the name of the namespace constituent for the Infinite Volume.

`-path path` specifies the path.

`-lif lif` specifies the logical interface.

`-lockid` specifies the universally unique identifier for the lock.

4. Return to the admin privilege level:

```
set -privilege admin
```

Modifying the NFSv4.1 server implementation ID

The NFSv4.1 protocol includes a server implementation ID that documents the server domain, name, and date. You can modify the server implementation ID default values. Changing the default values can be useful, for example, when gathering usage statistics or troubleshooting interoperability issues. For more information, see RFC 5661.

About this task

The default values for the three options are as follows:

Option	Option name	Default value
NFSv4.1 Implementation ID Domain	-v4.1-implementation-domain	ibm.com
NFSv4.1 Implementation ID Name	-v4.1-implementation-name	Cluster version name
NFSv4.1 Implementation ID Date	-v4.1-implementation-date	Cluster version date

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Perform one of the following actions:

If you want to modify the NFSv4.1 implementation ID...		Enter the command...
Domain		<code>vserver nfs modify -v4.1-implementation-domain <i>domain</i></code>
Name		<code>vserver nfs modify -v4.1-implementation-name <i>name</i></code>
Date		<code>vserver nfs modify -v4.1-implementation-date <i>date</i></code>

3. Return to the admin privilege level:

```
set -privilege admin
```

Managing NFSv4 ACLs

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

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

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 maximum number of ACEs for each ACL is 1,024, as defined by the `-v4-acl-max-aces` parameter. The default for this parameter is 400. When the `-v4-acl-preserve` parameter is enabled and an object has both UNIX mode bits and ACLs, the ACL requires three mandatory ACEs to grant the proper permissions to OWNER, GROUP, and EVERYONE. These mandatory ACEs are applied automatically by Data ONTAP when a user performs a `chmod` operation to set mode bits. If you set an ACL that is missing any of the mandatory ACEs, Data ONTAP limits the number of ACEs for that ACL to reserve space for any missing mandatory ACEs to ensure that they can be added later should you decide to set mode bits as well.

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 `-v4.0-acl` 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.

Enabling or disabling modification of NFSv4 ACLs

When Data ONTAP receives a `chmod` command for a file or directory with an ACL, by default the ACL is retained and modified to reflect the mode bit change. You can disable the `-v4-acl-preserve` parameter to change the behavior if you want the ACL to be dropped instead.

About this task

When using unified security style, this parameter also specifies whether NTFS file permissions are preserved or dropped when a client sends a `chmod`, `chgroup`, or `chown` command for a file or directory.

The default for this parameter is enabled.

Steps

1. Set the privilege level to advanced:
`set -privilege advanced`
2. Perform one of the following actions:

If you want to...	Enter the following command
Enable retention and modification of existing NFSv4 ACLs (default)	<code>vserver nfs modify -vserver vserver_name -v4-acl-preserve enabled</code>
Disable retention and drop NFSv4 ACLs when changing mode bits	<code>vserver nfs modify -vserver vserver_name -v4-acl-preserve disabled</code>

3. Return to the admin privilege level:
`set -privilege admin`

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 `-v4.0-acl` and `-v4.1-acl` options. These options are disabled by default.

About this task

The `-v4.0-acl` or `-v4.1-acl` option controls the setting and viewing of NFSv4 ACLs; it does not control enforcement of these ACLs for access checking.

NFSv4.0 ACLs are not supported for Storage Virtual Machines (SVMs) with Infinite Volume.

Step

1. Perform one of the following actions:

If you want to...	Then...
Enable NFSv4.0 ACLs	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.0-acl enabled</code>
Disable NFSv4.0 ACLs	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.0-acl disabled</code>
Enable NFSv4.1 ACLs	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.1-acl enabled</code>
Disable NFSv4.1 ACLs	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.1-acl disabled</code>

Modifying the maximum ACE limit for NFSv4 ACLs

You can modify the maximum number of allowed ACEs for each NFSv4 ACL by modifying the parameter `-v4-acl-max-aces`. By default, the limit is set to 400 ACEs for each ACL. Increasing

this limit can help ensure successful migration of data with ACLs containing over 400 ACEs to storage systems running Data ONTAP.

About this task

Increasing this limit might impact performance for clients accessing files with NFSv4 ACLs.

Steps

1. Set the privilege level to advanced:
`set -privilege advanced`
2. Modify the maximum ACE limit for NFSv4 ACLs by entering the following command:
`vserver nfs modify -v4-acl-max-aces max_ace_limit`
The valid range of `max_ace_limit` is 192 to 1024.
3. Return to the admin privilege level:
`set -privilege admin`

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 file 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.

Note: NFSv4 File delegations are not supported on Storage Virtual Machines (SVMs) with Infinite Volume.

Enabling or disabling NFSv4 read file delegations

To enable or disable NFSv4 read file delegations, you can modify the `-v4.0-read-delegation` or `-v4.1-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.

NFSv4 file delegations are not supported on Storage Virtual Machines (SVMs) with Infinite Volume.

Step

1. Perform one of the following actions:

If you want to...	Then...
Enable NFSv4 read file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.0-read-delegation enabled</code>
Enable NFSv4.1 read file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.1-read-delegation enabled</code>
Disable NFSv4 read file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.0-read-delegation disabled</code>
Disable NFSv4.1 read file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.1-read-delegation disabled</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 `-v4.0-write-delegation` or `-v4.1-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.

NFSv4 file delegations are not supported on Storage Virtual Machines (SVMs) with Infinite Volume.

Step

- 1. Perform one of the following actions:

If you want to...	Then...
Enable NFSv4 write file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.0-write-delegation enabled</code>
Enable NFSv4.1 write file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.1-write-delegation enabled</code>
Disable NFSv4 write file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.0-write-delegation disabled</code>
Disable NFSv4.1 write file delegations	Enter the following command: <code>vserver nfs modify -vserver vserver_name -v4.1-write-delegation disabled</code>

Result

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

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 `-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.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Enter the following command:

```
vserver nfs modify -vserver vserver_name -v4-lease-seconds  
number_of_seconds
```

3. Return to the admin privilege level:

```
set -privilege admin
```

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 `-v4-grace-seconds` option.

About this task

By default, this option is set to 45.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Enter the following command:

```
vserver nfs modify -vserver vserver_name -v4-grace-seconds  
number_of_seconds
```

3. Return to the admin privilege level:

```
set -privilege admin
```

How NFSv4 referrals work

When you enable NFSv4 referrals, Data ONTAP provides “intra-SVM” referrals to NFSv4 clients. Intra-SVM referral is when a cluster node receiving the NFSv4 request refers the NFSv4 client to another LIF on the Storage Virtual Machine (SVM).

The NFSv4 client should access the path that received the referral at the target LIF from that point onward. The original cluster node gives such a referral when it determines that there exists a LIF in the SVM that is resident on the cluster node on which the data volume resides, thereby allowing the clients faster access to the data and avoiding extra cluster communication.

Support for NFSv4 referrals is not uniformly available in all NFSv4 clients. In an environment where not all clients support this feature, you should not enable referrals. If the feature is enabled and a client that does not support it receives a referral from the server, the client cannot access the volume and experiences failures.

See RFC3530 for details about referrals.

Note: Referrals are not supported on SVMs with Infinite Volume.

Enabling or disabling NFSv4 referrals

You can enable NFSv4 referrals on Storage Virtual Machines (SVMs) with FlexVol volumes by enabling the options `-v4-fsid-change` and `-v4.0-referrals` or `-v4.1-referrals`. Enabling NFSV4 referrals can result in faster data access for NFSv4 clients that support this feature.

Before you begin

If you want to enable NFS referrals, you must first disable parallel NFS. You cannot enable both at the same time.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Perform one of the following actions:

If you want to...	Enter the command...
Enable NFSv4 referrals	<pre>vserver nfs modify -vserver vserver_name -v4-fsid-change enabled</pre> <pre>vserver nfs modify -vserver vserver_name -v4.0-referrals enabled</pre>
Disable NFSv4 referrals	<pre>vserver nfs modify -vserver vserver_name -v4.0-referrals disabled</pre>
Enable NFSv4.1 referrals	<pre>vserver nfs modify -vserver vserver_name -v4-fsid-change enabled</pre> <pre>vserver nfs modify -vserver vserver_name -v4.1-referrals enabled</pre>
Disable NFSv4.1 referrals	<pre>vserver nfs modify -vserver vserver_name -v4.1-referrals disabled</pre>

3. Return to the admin privilege level:

```
set -privilege admin
```

Displaying NFS statistics

You can display NFS statistics for Storage Virtual Machines (SVMs) on the storage system to monitor performance and diagnose issues.

Steps

1. Use the `statistics catalog object show` command to identify the NFS objects from which you can view data.

Example

```
statistics catalog object show -object nfs*
```

For more information about the `statistics` commands, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Use the `statistics start` and optional `statistics stop` commands to collect a data sample from one or more objects.
3. Use the `statistics show` command to view the sample data.

Example: Monitoring NFSv3 performance

The following example shows performance data for the NFSv3 protocol.

The following command starts data collection for a new sample:

```
vs1::> statistics start -object nfsv3 -sample-id nfs_sample
```

The following command shows data from the sample by specifying counters that show the number of successful read and write requests versus the total number of read and write requests:

```
vs1::> statistics show -sample-id nfs_sample -counter read_total|
write_total|read_success|write_success
```

```
Object: nfsv3
Instance: vs1
Start-time: 2/11/2013 15:38:29
End-time: 2/11/2013 15:38:41
Cluster: cluster1
```

Counter	Value
read_success	40042
read_total	40042

write_success	1492052
write_total	1492052

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.
- vStorage is not supported on Storage Virtual Machines (SVMs) with Infinite Volume.
- Copy offload operations can fail in the following scenarios:
 - While running wafiron on the source or destination volume because it temporarily takes the volume offline
 - While moving either the source or destination volume
 - While moving either the source or destination LIF
 - While performing takeover or giveback operations

Enabling or disabling VMware vStorage over NFS

You can enable or disable support for VMware vStorage over NFS on Storage Virtual Machines (SVMs) with FlexVol volumes by using the `vserver nfs modify` command.

About this task

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

Steps

1. Display the current vStorage support status for SVMs by entering the following command:

```
vserver nfs show -vserver vserver_name -instance
```

2. Perform one of the following actions:

If you want to...	Enter the following command...
Enable VMware vStorage support	vserver nfs modify -vserver vserver_name -vstorage enabled
Disable VMware vStorage support	vserver nfs modify -vserver vserver_name -vstorage disabled

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 IBM N Series NFS Plug-in for VMware VAAI*.

Enabling or disabling rquota support

Data ONTAP supports the remote quota protocol version 1 (rquota v1). The rquota protocol enables NFS clients to obtain quota information for users and groups from a remote machine. You can enable rquota on Storage Virtual Machines (SVMs) with FlexVol volumes by using the `vserver nfs modify` command.

About this task

By default, rquota is disabled.

Step

1. Perform one of the following actions:

If you want to...	Enter the following command...
Enable rquota support for SVMs	vserver nfs modify -vserver vserver_name -rquota enable
Disable rquota support for SVMs	vserver nfs modify -vserver vserver_name -rquota disable

For more information about quotas, see the *Clustered Data ONTAP Logical Storage Management Guide*.

NFSv3 performance improvement by modifying the TCP maximum read and write size

You might be able to improve the performance of NFSv3 clients connecting to storage systems over a high-latency network by modifying the TCP maximum read and write size.

When clients access storage systems over a high-latency network, such as a wide area network (WAN) or metro area network (MAN) with a latency over 10 milliseconds, you might be able to improve the connection performance by modifying the TCP maximum read and write size. Clients accessing storage systems in a low-latency network, such as a local area network (LAN), can expect little to no benefit from modifying these parameters. If the throughput improvement does not outweigh the latency impact, you should not use these parameters.

To determine whether your storage environment would benefit from modifying these parameters, you should first conduct a comprehensive performance evaluation of a poorly performing NFS client. Review whether the low performance is due to excessive round trip latency and small request on the client. Under these conditions, the client and server cannot fully use the available bandwidth because they spend the majority of their duty cycles waiting for small requests and responses to be transmitted over the connection.

By increasing the NFSv3 request size, the client and server can use the available bandwidth more effectively to move more data per unit time, therefore increasing the overall efficiency of the connection.

Keep in mind that the configuration between the storage system and the client might vary. If you configure the storage system to support 1 MB maximum read size but the client only supports 64 KB, then the mount read size is limited to 64 KB or less.

Before modifying these parameters, you must be aware that it results in additional memory consumption on the storage system for the period of time necessary to assemble and transmit a large response. The more high-latency connections to the storage system, the higher the additional memory consumption. Storage systems with high memory capacity might experience very little effect from this change. Storage systems with low memory capacity might experience noticeable performance degradation.

The successful use of these parameter relies on the ability to retrieve data from multiple nodes of a cluster. The inherent latency of the cluster network might increase the overall latency of the response. Overall latency tends to increase when using these parameters. As a result, latency sensitive workloads might show negative impact.

Modifying the NFSv3 TCP maximum read and write size

You can modify the `-v3-tcp-max-read-size` and `-v3-tcp-max-write-size` options to change the NFSv3 TCP maximum read and write size. Modifying these options can help improve NFSv3 performance over TCP in some storage environments.

Before you begin

All nodes in the cluster must be running Data ONTAP 8.1 or later.

About this task

You can modify these options individually for each Storage Virtual Machine (SVM).

Steps

1. Set the privilege level to advanced:
`set -privilege advanced`
2. Perform one of the following actions:

If you want to...	Enter the command...
Modify the NFSv3 TCP maximum read size	<code>vserver nfs modify -vserver vserver_name -v3-tcp-max-read-size integer_max_read_size</code>
Modify the NFSv3 TCP maximum write size	<code>vserver nfs modify -vserver vserver_name -v3-tcp-max-write-size integer_max_write_size</code>

Option	Range	Default
<code>-v3-tcp-max-read-size</code>	8192 to 1048576 bytes	65536 bytes
<code>-v3-tcp-max-write-size</code>	8192 to 65536 bytes	65536 bytes

Note: The maximum read or write size you enter must be a multiple of 4 KB (4096 bytes). Requests that are not properly aligned negatively affect performance.

3. Return to the admin privilege level:
`set -privilege admin`
4. Use the `vserver nfs show` command to verify the changes.
5. If any clients use static mounts, unmount and remount for the new parameter size to take effect.

Example

The following command sets the NFSv3 TCP maximum read size to 1048576 bytes on the SVM named vs1:

```
vs1::> vserver nfs modify -vserver vs1 -v3-tcp-max-read-size 1048576
```

Auditing NAS file access events on SVMs with FlexVol volumes

Auditing for NAS file access events is a security measure that enables you to track and log SMB and NFS file and folder access events on objects stored on Storage Virtual Machines (SVMs) with FlexVol volumes. This helps you track potential security problems and provides evidence of any file access security breaches.

How auditing works

Before you plan and configure your auditing configuration, you should understand how auditing works.

Basic auditing concepts

To understand auditing in Data ONTAP, you should be aware of some basic auditing concepts.

Staging files The intermediate binary files on individual nodes where audit records are stored prior to consolidation and conversion. Staging files are contained in staging volumes.

Staging volume A dedicated volume created by Data ONTAP to store staging files. There is one staging volume per aggregate. Staging volumes are shared by all audit-enabled Storage Virtual Machines (SVMs) with volumes in that particular aggregate. Staging volumes are a type of system volume.

System volumes are FlexVol volumes that contain special metadata, such as metadata for file services audit logs. System volumes are owned by the admin SVM and are visible across the cluster; therefore, there is no multi-tenancy for staging volumes. Only cluster administrators can view staging volumes. Additionally, cluster administrators can modify, or delete staging volumes, but they cannot create staging volumes.

Consolidation task A task that takes the audit records from staging files across the member nodes of the SVM on a per-SVM basis and merges them in sorted chronological order and then converts them to a user-readable event log format specified in the auditing configuration — either the EVTX or XML file format. The converted event logs are stored in the audit event log directory that is specified in the SVM auditing configuration.

How the Data ONTAP auditing process works

The Data ONTAP auditing process is different than the Microsoft auditing process. Before you configure auditing, you should understand how the Data ONTAP auditing process works.

Audit records are initially stored in binary staging files on individual nodes. If auditing is enabled on an SVM, every member node maintains staging files for that SVM. Periodically, they are consolidated and converted to user-readable event logs, which are stored in the audit event log directory for the SVM.

Process when auditing is enabled on an SVM

Auditing can only be enabled on SVMs with FlexVol volumes. When the storage administrator enables auditing on the SVM, the auditing subsystem checks to determine if staging volumes are present. A staging volume must exist for each aggregate containing data volumes owned by the SVM. The auditing subsystem creates any needed staging volumes if they do not exist.

The auditing subsystem also completes other prerequisite tasks before auditing is enabled:

- The auditing subsystem verifies that the log directory path is available and does not contain symlinks.

The log directory must already exist. The auditing subsystem does not assign a default log file location. If the log directory path specified in the auditing configuration is not a valid path, auditing configuration creation fails with the following error:

The specified path "<path>" does not exist in the namespace belonging to Vserver "<Vserver_name>"

Configuration creation fails if the directory exists but contains symlinks.

- Auditing schedules the consolidation task.

After these tasks are completed successfully, auditing is enabled. The SVM auditing configuration and the log files persist across a reboot or if the NFS or CIFS servers are stopped or restarted.

Event log consolidation

Log consolidation is a scheduled task that runs on a routine basis until auditing is disabled. When auditing is disabled, the final run of the consolidation task ensures that all the remaining logs are consolidated.

Guaranteed auditing

By default, auditing is guaranteed. Data ONTAP guarantees that all auditable file access events (as specified by configured audit policy ACLs) are recorded, even if a node is unavailable. A requested file operation cannot complete until the audit record for that operation is saved to the staging volume on persistent storage. If audit records cannot be committed to the disk in the staging files, either because of insufficient space or because of other issues, client operations are denied.

Consolidation process when a node is unavailable

If a node containing volumes belonging to an SVM with auditing enabled is unavailable, the behavior of the auditing consolidation task depends on whether the node's SFO partner (or the HA partner in the case of a two-node cluster) is available.

- If the staging volume is available through the SFO partner, the staging volumes last reported from the node are scanned, and consolidation proceeds normally.
- If the SFO partner is not available, the task creates a partial log file.
When a node is not reachable, the consolidation task consolidates the audit records from the other available nodes of that SVM. But to identify that it is not complete, the task adds the suffix `.partial` to the consolidated file name.
- After the unavailable node is available, the audit records in that node are consolidated with the audit records from the other nodes at that point of time.
- All audit records are preserved.

Event log rotation

Audit event log files are rotated when they reach a configured threshold log size or on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file and creates a new active converted event log file.

Process when auditing is disabled on the SVM

When auditing is disabled on the SVM, the consolidation task is triggered one final time. All outstanding, recorded audit records are logged in user-readable format. Existing event logs stored in the event log directory are not deleted when auditing is disabled on the SVM and are available for viewing.

After all existing staging files for that SVM are consolidated, the consolidation task is removed from the schedule. Disabling the auditing configuration for the SVM does not remove the auditing configuration. A storage administrator can reenabling auditing at any time.

Aggregate space considerations when enabling auditing

When an auditing configuration is created and auditing is enabled on at least one Storage Virtual Machine (SVM) in the cluster, the auditing subsystem creates staging volumes on all existing aggregates and on all new aggregates that are created. You need to be aware of certain aggregate space considerations when you enable auditing on the cluster.

Staging volume creation might fail due to non-availability of space in an aggregate. This might happen if you create an auditing configuration and existing aggregates do not have enough space to contain the staging volume.

You should ensure that there is enough space on existing aggregates for the staging volumes before enabling auditing on an SVM.

Auditing requirements and considerations

Before you configure and enable auditing on your Storage Virtual Machine (SVM) with FlexVol volumes, you need to be aware of certain requirements and considerations.

- Before you can enable auditing on your SVM, all nodes in the cluster must be running Data ONTAP 8.2 or later.
- The maximum number of audit-enabled SVMs supported in a cluster is 50.
- Auditing is not tied to CIFS or NFS licensing.
You can configure and enable auditing even if CIFS and NFS licenses are not installed on the cluster.
- NFS auditing supports security ACEs (type U).
- For NFS auditing, there is no mapping between mode bits and audit ACEs.
When converting ACLs to mode bits, audit ACEs are skipped. When converting mode bits to ACLs, audit ACEs are not generated.
- The directory specified in the auditing configuration must exist.
If it does not exist, the command to create the auditing configuration fails.
- The directory specified in the auditing configuration must meet the following requirements:
 - The directory must not contain symbolic links.
If the directory specified in the auditing configuration contains symbolic links, the command to create the auditing configuration fails.
 - You must specify the directory by using an absolute path.
You should not specify a relative path, for example, `/vs1/. . /.`
- Auditing is dependent on having available space in the staging volumes.
You must be aware of and have a plan for ensuring that there is sufficient space for the staging volumes in aggregates that contain audited volumes.
- Auditing is dependent on having available space in the volume containing the directory where converted audit event logs are stored.
You must be aware of and have a plan for ensuring that there is sufficient space in the volumes used to store event logs. You can specify the number of audit logs to retain in the auditing directory by using the `-rotate-limit` parameter when creating an auditing configuration, which can help to ensure that there is enough available space for the audit logs in the volume.

What the supported audit event log formats are

Supported file formats for the converted audit event logs are EVTX and XML file formats.

You can specify the type of file format when you create the auditing configuration. By default, Data ONTAP converts the binary logs to the EVTX file format.

Viewing audit event logs

You can use audit event logs to determine whether you have adequate file security and whether there have been improper file and folder access attempts. You can view and process audit event logs saved in the EVTX or XML file formats.

- **EVTX file format**

You can open the converted EVTX audit event logs as saved files using Microsoft Event Viewer. There are two options that you can use when viewing event logs using Event Viewer:

- **General view**

Information that is common to all events is displayed for the event record. In this version of Data ONTAP, the event-specific data for the event record is not displayed. You can use the detailed view to display event-specific data.

- **Detailed view**

A friendly view and a XML view are available. The friendly view and the XML view display both the information that is common to all events and the event-specific data for the event record.

- **XML file format**

You can view and process XML audit event logs on third-party applications that support the XML file format. XML viewing tools can be used to view the audit logs provided you have the XML schema and information about definitions for the XML fields. For more information about obtaining the XML schema and documents related to XML definitions, contact technical support or your account team.

SMB file and folder access events that can be audited

Data ONTAP can audit certain SMB file and folder access events. Knowing what access events can be audited is helpful when interpreting results from the converted audit event logs.

The following SMB file and folder access events can be audited:

Event ID (EVT/EVTX)	Event	Description	Category
560/4656	Open Object/ Create Object	OBJECT ACCESS: Object (file or directory) open.	File Access

Event ID (EVT/EVTX)	Event	Description	Category
567/4663	Read Object/ Write Object/Get Object Attributes/Set Object Attributes	OBJECT ACCESS: Object access attempt (read, write, get attribute, set attribute). Note: For this Event, Data ONTAP audits only the first SMB read and first SMB write operation (success or failure) on an object. This prevents Data ONTAP from creating excessive log entries when a single client opens an object and performs many successive read or write operations to the same object.	File Access
N/A/4664	Hard link	OBJECT ACCESS: An attempt was made to create a hard link.	File Access
N/A/N/A Data ONTAP Event ID 9999	Rename Object	OBJECT ACCESS: Object renamed. This is a Data ONTAP event. It is not currently supported by Windows as a single event.	File Access
N/A/N/A Data ONTAP Event ID 9998	Unlink Object	OBJECT ACCESS: Object unlinked. This is a Data ONTAP event. It is not currently supported by Windows as a single event.	File Access

Note: The object path printed in an audit record is the relative path from the root of the containing volume. For example, consider the following volume information:

Vserver	Volume	Junction Active	Junction Path	Junction Path Source
vs1	data	true	/data	RW_volume
vs1	data1	true	/data/data1	RW_volume

If a user accesses a file with the path `/data/data1/dir1/file.txt`, the path used in the `<ObjectName>` tag in the event contained in the audit logs is `/data1/dir1/file.txt`.

NFS file and directory access events that can be audited

Data ONTAP can audit certain NFS file and directory access events. Knowing what access events can be audited is helpful when interpreting results from the converted audit event logs.

You can audit the following NFS file and directory access events:

- READ
- OPEN

- CLOSE
- READDIR
- WRITE
- SETATTR
- CREATE
- LINK
- OPENATTR
- REMOVE
- GETATTR
- VERIFY
- NVERIFY
- RENAME

To reliably audit NFS RENAME events, you should set audit ACEs on directories instead of files because file permissions are not checked for a RENAME operation if the directory permissions are sufficient.

Planning the auditing configuration

Before you configure auditing on Storage Virtual Machines (SVMs) with FlexVol volumes, you must understand which configuration options are available and plan the values that you want to set for each option. This information can help you configure the auditing configuration that meets your business needs.

There are certain configuration parameters that are common to all auditing configurations.

Additionally, there are certain parameters that you can use to specify which of two methods are used when rotating the consolidated and converted audit logs. You can specify one of the two following methods when you configure auditing:

- Rotate logs based on log size
This is the default method used to rotate logs.
- Rotate logs based on a schedule

Parameters common to all auditing configurations

There are two required parameters that you must specify when you create the auditing configuration. There are also two optional parameters that you can specify. The first optional parameter determines how many audit logs are retained in the audit log directory. The second optional parameter specifies which log file format to use for the audit logs.

You can use the following list to determine what values to use for the parameters that are common to all auditing configurations:

Type of information	Option	Required	Include	Your values
<i>SVM name</i> Name of the SVM on which to create the auditing configuration. The SVM must already exist.	<code>-vserver vserver_name</code>	Yes	Yes	
<i>Log destination path</i> Specifies where the converted audit logs are stored. The path must already exist on the SVM. If the path is not valid, the audit configuration command fails.	<code>-destination text</code>	Yes	Yes	
<i>Log file output format</i> Determines the output format of the audit logs. The output format can be either Data ONTAP-specific XML or Microsoft Windows EVTX log format. By default, the output format is EVTX.	<code>-format {xml evtx}</code>	No		
<i>Log files rotation limit</i> Determines how many audit log files to retain before rotating the oldest log file out. A value of 0 indicates that all the log files are retained. The default value is 0. For example, if you enter a value of 5, the last five log files are retained.	<code>-rotate-limit integer</code>	No		

Parameters used for determining when to rotate audit event logs

Rotate logs based on log size

The default is to rotate audit logs based on size. The default log size is 100 MB. If you want to use the default log rotation method and the default log size, you do not need to configure any specific parameters for log rotation. If you do not want to use the default log size, you can configure the `-rotate-size` parameter to specify a custom log size:

Type of information	Option	Required	Include	Your values
<i>Log file size limit</i> Determines the audit log file size limit.	<code>-rotate-size {integer[KB MB GB TB PB]}</code>	No		

Rotate logs based on a schedule

If you choose to rotate the audit logs based on a schedule, you can schedule log rotation by using the time-based rotation parameters in any combination.

- If you configure time-based log rotation parameters, logs are rotated based on the configured schedule instead of log size.
- If you use time-based rotation, the `-rotate-schedule-minute` parameter is mandatory.
- All other time-based rotation parameters are optional.
- The rotation schedule is calculated by using all the time-related values.

For example, if you specify only the `-rotate-schedule-minute` parameter, the audit log files are rotated based on the minutes specified on all days of the week, during all hours on all months of the year.

- If you specify only one or two time-based rotation parameters (for example, `-rotate-schedule-month` and `-rotate-schedule-minutes`), the log files are rotated based on the minute values that you specified on all days of the week, during all hours, but only during the specified months.

For example, you can specify that the audit log is to be rotated during the months January, March, and August on all Mondays, Wednesdays, and Saturdays at 10:30 a.m.

- If you specify values for both `-rotate-schedule-dayofweek` and `-rotate-schedule-day`, they are considered independently.

For example, if you specify `-rotate-schedule-dayofweek` as Friday and `-rotate-schedule-day` as 13, then the audit logs would be rotated on every Friday and on the 13th day of the specified month, not just on every Friday the 13th.

You can use the following list of available auditing parameters to determine what values to use for configuring a schedule for audit event log rotations:

Type of information	Option	Required	Include	Your values
<p><i>Log rotation schedule: Month</i></p> <p>Determines the monthly schedule for rotating audit logs.</p> <p>Valid values are January through December, and all. For example, you can specify that the audit log is to be rotated during the months January, March, and August.</p>	<code>-rotate-schedule-month</code> <i>chron_month</i>	No		

Type of information	Option	Required	Include	Your values
<p><i>Log rotation schedule: Day of week</i></p> <p>Determines the daily (day of week) schedule for rotating audit logs.</p> <p>Valid values are January through December, and all. For example, you can specify that the audit log is to be rotated on Tuesdays and Fridays, or during all the days of a week.</p>	<p>-rotate-schedule-dayofweek</p> <p>chron_dayofweek</p>	No		
<p><i>Log rotation schedule: Day</i></p> <p>Determines the day of the month schedule for rotating the audit log.</p> <p>Valid values range from 1 through 31. For example, you can specify that the audit log is to be rotated on the 10th and 20th days of a month, or all days of a month.</p>	<p>-rotate-schedule-day</p> <p>chron_dayofmonth</p>	No		
<p><i>Log rotation schedule: Hour</i></p> <p>Determines the hourly schedule for rotating the audit log.</p> <p>Valid values range from 0 (midnight) to 23 (11:00 p.m.). Specifying all rotates the audit logs every hour. For example, you can specify that the audit log is to be rotated at 6 (6 a.m.) and 18 (6 p.m.).</p>	<p>-rotate-schedule-hour</p> <p>chron_hour</p>	No		
<p><i>Log rotation schedule: Minute</i></p> <p>Determines the minute schedule for rotating the audit log.</p> <p>Valid values range from 0 to 59. For example, you can specify that the audit log is to be rotated at the 30th minute.</p>	<p>-rotate-schedule-minute</p> <p>chron_minute</p>	Yes, if configuring schedule-based log rotation; otherwise, no.		

Creating a file and directory auditing configuration on SVMs

Creating a file and directory auditing configuration on your Storage Virtual Machine (SVM) with FlexVol volumes includes understanding the available configuration options, planning the configuration, and then configuring and enabling the configuration. You can then display information

about the auditing configuration to confirm that the resultant configuration is the desired configuration.

Steps

1. [Creating the auditing configuration](#) on page 120
Before you can begin auditing file and directory events, you must create an auditing configuration on the Storage Virtual Machine (SVM).
2. [Enabling auditing on the SVM](#) on page 121
After you finish setting up the auditing configuration, you must enable auditing on the Storage Virtual Machine (SVM).
3. [Verifying the auditing configuration](#) on page 122
After completing the auditing configuration, you should verify that auditing is configured properly and is enabled.

Creating the auditing configuration

Before you can begin auditing file and directory events, you must create an auditing configuration on the Storage Virtual Machine (SVM).

Step

1. Using the information in the planning worksheet, create the auditing configuration by using the appropriate command:

If you want to create an auditing configuration that rotates audit logs based on...	Enter the command...
Log size	<pre>vserver audit create -vserver vservers_name - destination path [-format {xml evtx}] [-rotate-limit integer] [-rotate-size {integer[KB MB GB TB PB]}]</pre>
A schedule	<pre>vserver audit create -vserver vservers_name - destination path [-format {xml evtx}] [-rotate-limit integer] [-rotate-schedule-month chron_month] [- rotate-schedule-dayofweek chron_dayofweek] [-rotate- schedule-day chron_dayofmonth] [-rotate-schedule- hour chron_hour] -rotate-schedule-minute chron_minute</pre> <p>Note: The <code>-rotate-schedule-minute</code> parameter is required if configuring time-based audit log rotation.</p>

Examples

The following example creates an audit configuration for SVM vs1. The log format is EVTX (the default). The logs are stored in the /audit_log directory. The log file size limit is 200 MB. The logs are rotated when they reach 200 MB in size:

```
cluster1::> vserver audit create -vserver vs1 -destination /audit_log -
rotate-size 200MB
```

The following example creates an audit configuration for SVM vs1 using size-based rotation. The log format is EVTX (the default). The log file size limit is 200 MB, and the log rotation limit is 5:

```
cluster1::> vserver audit create -vserver vs1 -destination /audit_log -
rotate-size 200MB -rotate-limit 5
```

The following example creates an audit configuration for SVM vs1 using time-based rotation. The log format is EVTX (the default). The audit logs are rotated monthly, at 12:30 p.m. on all days of the week:

```
cluster1::> vserver audit create -vserver vs1 -destination /audit_log -
rotate-size 200MB -rotate-schedule-month all -rotate-schedule-dayofweek all
-rotate-schedule-hour 12 -rotate-schedule-minute 30
```

Enabling auditing on the SVM

After you finish setting up the auditing configuration, you must enable auditing on the Storage Virtual Machine (SVM).

Before you begin

The SVM audit configuration must already exist.

Step

1. Enable auditing on the SVM:

```
vserver audit enable -vserver vserver_name
```

Example

```
vserver audit enable -vserver vs1
```

Verifying the auditing configuration

After completing the auditing configuration, you should verify that auditing is configured properly and is enabled.

Step

1. Verify the auditing configuration:

```
vserver audit show -instance -vserver vserver_name
```

Example

The following example displays in list form all audit configuration information for Storage Virtual Machine (SVM) vs1. The EVTX-formatted logs are stored in the /audit_log directory. The log file size limit is 200 MB, and the logs are rotated when they reach 200 MB in size. Auditing is enabled:

```
vserver audit show -instance -vserver vs1
      Vserver: vs1
      Auditing state: true
      Log Destination Path: /audit_log
      Log Format: evtx
      Log File Size Limit: 200MB
      Log Rotation Schedule: Month: -
      Log Rotation Schedule: Day of Week: -
      Log Rotation Schedule: Day: -
      Log Rotation Schedule: Hour: -
      Log Rotation Schedule: Minute: -
      Rotation Schedules: -
      Log Files Rotation Limit: 0
```

Configuring file and folder audit policies

Implementing auditing on file and folder access events is a two-step process. First you must create and enable an auditing configuration on Storage Virtual Machines (SVMs) with FlexVol volumes. Second, you must configure audit policies on the files and folders that you want to monitor. You can configure audit policies to monitor both successful and failed access attempts.

You can configure both SMB and NFS audit policies. SMB and NFS audit policies have different configuration requirements and audit capabilities.

If the appropriate audit policies are configured, Data ONTAP monitors SMB and NFS access events as specified in the audit policies only if the SMB or NFS servers are running.

Configuring audit policies on NTFS security-style files and directories

Before you can audit file and directory operations, you must configure audit policies on the files and directories for which you want to collect audit information. This is in addition to setting up and

enabling the audit configuration. You can configure NTFS audit policies by using the Windows Security tab or by using the Data ONTAP CLI.

Configuring NTFS audit policies using the Windows Security tab

You can configure audit policies on files and directories by using the **Windows Security** tab in the Windows Properties window. This is the same method used when configuring audit policies on data residing on a Windows client, which enables customers to use the same GUI interface that they are accustomed to using.

Before you begin

Auditing must be configured on the Storage Virtual Machine (SVM) that contains the data to which you are applying SACLs.

About this task

Configuring NTFS audit policies is done by adding entries to NTFS system access control lists (SACLs) that are associated with an NTFS security descriptor. The security descriptor is then applied to NTFS files and directories. These tasks are automatically handled by the Windows GUI. The security descriptor can contain discretionary access control lists (DACLS) for applying file and folder access permissions, system access control lists (SACLs) for file and folder auditing, or both SACLs and DACLS.

You can set NTFS audit policies for auditing access on individual files and folders using the Windows Security tab in the Windows Properties window by completing the following steps on a Windows host:

Steps

1. From the **Tools** menu in Windows Explorer, select **Map network drive**.
2. Complete the **Map Network Drive** box:
 - a) Select a **Drive** letter.
 - b) In the **Folder** box, type the CIFS server name that contains the share holding the data you would like to audit and the name of the share.

Example

If your CIFS server name is "CIFS_SERVER" and your share is named "share1", you should enter `\\CIFS_SERVER\share1`.

Note: You can specify the IP address of the data interface for the CIFS server instead of the CIFS server name.

- c) Click **Finish**.

The drive you selected is mounted and ready with the Windows Explorer window displaying files and folders contained within the share.

3. Select the file or directory for which you want to enable auditing access.
4. Right-click on the file or directory, and select **Properties**.
5. Select the **Security** tab.
6. Click **Advanced**.
7. Select the **Auditing** tab.
8. Perform the desired actions:

If you want to....	Do the following
Set up auditing for a new user or group	<ol style="list-style-type: none"> a. Click Add. b. In the Enter the object name to select box, type the name of the user or group that you want to add. c. Click OK.
Remove auditing from a user or group	<ol style="list-style-type: none"> a. In the Enter the object name to select box, select the user or group that you want to remove. b. Click Remove. c. Click OK. d. Skip the rest of this procedure.
Change auditing for a user or group	<ol style="list-style-type: none"> a. In the Enter the object name to select box, select the user or group that you want to change. b. Click Edit. c. Click OK.

If you are setting up auditing on a user or group or changing auditing on an existing user or group, the Auditing Entry for <object> box opens.

9. In the **Apply to** box, select how you want to apply this auditing entry.

You can select one of the following:

- **This folder, subfolders and files**
- **This folder and subfolders**
- **This folder only**
- **This folder and files**
- **Subfolders and files only**
- **Subfolders only**
- **Files only**

If you are setting up auditing on a single file, the **Apply to** box is not active. The **Apply to** defaults to **This object only**.

Note: Since auditing takes SVM resources, select only the minimal level that provides the auditing events that meet your security requirements.

10. In the **Access** box, select what you want audited and whether you want to audit successful events, failure events or both.

- To audit successful events, select the **Success** box.
- To audit failure events, select the **Failure** box.

You can audit the following events:

- **Full control**
- **Traverse folder / execute file**
- **List folder / read data**
- **Read attributes**
- **Read extended attributes**
- **Create files / write data**
- **Create folders / append data**
- **Write attributes**
- **Write extended attributes**
- **Delete subfolders and files**
- **Delete**
- **Read permissions**
- **Change permissions**
- **Take ownership**

Note: Select only the actions that you need to monitor to meet your security requirements. For more information on these auditable events, see your Windows documentation.

11. If you do not want the auditing setting to propagate to subsequent files and folders of the original container, select **Apply these auditing entries to objects and/or containers within this container only** box.

12. Click **Apply**.

13. After you finish adding, removing, or editing auditing entries, click **OK**.

The Auditing Entry for <object> box closes.

14. In the **Auditing** box, select the inheritance settings for this folder.

You can choose one of the following:

- Select the **Include inheritable auditing entries from this object's parent** box.
- Select the **Replace all existing inheritable auditing entries on all descendants with inheritable auditing entries from this object** box.
- Select both boxes.
- Select neither box.

If you are setting SACLs on a single file, the **Replace all existing inheritable auditing entries on all descendants with inheritable auditing entries from this object** box is not present in the Auditing dialog box.

Note: Select only the minimal level that provides the auditing events that meet your security requirements.

15. Click **OK**.

The Auditing box closes.

How to configure NTFS audit policies using the Data ONTAP CLI

You can configure audit policies on files and folders using the Data ONTAP CLI. This enables you to configure NTFS audit policies without needing to connect to the data using an SMB share on a Windows client.

You can configure NTFS audit policies by using the `vserver security file-directory` command family.

You can only configure NTFS SACLs using the CLI. Configuring NFSv4 SACLs is not supported with this Data ONTAP command family. See the man pages for more information about using these commands to configure and add NTFS SACLs to files and folders.

Configuring auditing for UNIX security style files and directories

You configure auditing for UNIX security style files and directories by adding audit ACEs to NFSv4.x ACLs. This allows you to monitor certain NFS file and directory access events for security purposes.

About this task

For NFSv4.x, both discretionary and system ACEs are stored in the same ACL. They are not stored in separate DACLs and SACLs. Therefore, you must exercise caution when adding audit ACEs to an existing ACL to avoid overwriting and losing an existing ACL. The order in which you add the audit ACEs to an existing ACL does not matter.

Steps

1. Retrieve the existing ACL for the file or directory by using the `nfs4_getfacl` or equivalent command.

For more information about manipulating ACLs, see the man pages of your NFS client.

2. Append the desired audit ACEs.
3. Apply the updated ACL to the file or directory by using the `nfs4_setfacl` or equivalent command.

Displaying information about audit policies applied to files and directories

Displaying information about audit policies applied to files and directories enables you to verify that you have the appropriate system access control lists (SACLs) set on specified files and folders.

Displaying information about audit policies using the Windows Security tab

You can display information about audit policies that have been applied to files and directories by using the Security tab in the Windows Properties window. This is the same method used for data residing on a Windows server, which enables customers to use the same GUI interface that they are accustomed to using.

About this task

To display information about SACLs that have been applied to NTFS files and folders, complete the following steps on a Windows host.

Steps

1. From the **Tools** menu in Windows Explorer, select **Map network drive**.
2. Complete the **Map Network Drive** dialog box:
 - a) Select a **Drive** letter.
 - b) In the **Folder** box, type the IP address or CIFS server name of the Storage Virtual Machine (SVM) containing the share that holds both the data you would like to audit and the name of the share.

Example

If your CIFS server name is "CIFS_SERVER" and your share is named "share1", you should enter \\CIFS_SERVER\share1.

Note: You can specify the IP address of the data interface for the CIFS server instead of the CIFS server name.

- c) Click **Finish**.

The drive you selected is mounted and ready with the Windows Explorer window displaying files and folders contained within the share.

3. Select the file or directory for which you display auditing information.
4. Right-click on the file or directory, and select **Properties**.
5. Select the **Security** tab.
6. Click **Advanced**.

7. Select the **Auditing** tab.
8. Click **Continue**.
The Auditing box opens. The **Auditing entries** box displays a summary of users and groups that have SACLs applied to them.
9. In the **Auditing entries** box select the user or group whose SACL entries you want displayed.
10. Click **Edit**.
The Auditing entry for <object> box opens.
11. In the **Access** box, view the current SACLs that are applied to the selected object.
12. Click **Cancel** to close the **Auditing entry for <object>** box.
13. Click **Cancel** to close the **Auditing** box.

Displaying information about NTFS audit policies on FlexVol volumes using the CLI

You can display information about NTFS audit policies on FlexVol volumes, including what the security styles and effective-security styles are, what permissions are applied, and information about system access control lists. You can use the results to validate your security configuration or to troubleshoot auditing issues.

About this task

You must supply the name of the Storage Virtual Machine (SVM) that contains the path to the files or directories whose audit information you want to display. If you want to customize the output, you can use the following optional parameters to display information only about file and directory security that matches the specified parameters:

Optional parameter	Description
-fields <i>fieldsname, ...</i>	You can use this parameter to display information on the fields you specify. You can use this parameter either alone or in combination with other optional parameters.
-instance	Displays detailed information about all entries.
-volume-name <i>volume_name</i>	Displays information where the specified path is relative to the specified volume. If this parameter is not specified, the SVM root volume is taken as default.
-share-name <i>share_name</i>	Displays information where the specified path is relative to the root of the specified share. If this parameter is not specified, the SVM root volume is taken as default.

Optional parameter	Description
<code>-lookup-names</code> { <code>true</code> <code>false</code> }	Displays information where the information about owner and group is set to one of the following: <ul style="list-style-type: none"> <code>true</code> displays information where the lookup name is stored as a name. <code>false</code> displays information where the lookup name is stored as a SID.
<code>-expand-mask</code> { <code>true</code> <code>false</code> }	Displays information where the hexadecimal bit mask entry is set to one of the following: <ul style="list-style-type: none"> <code>true</code> displays information where the bit mask entries are store in expanded form. <code>false</code> displays information where the bit mask entries are store in collapsed form.
<code>-security-style</code> { <code>unix</code> <code>ntfs</code> <code>mixed</code> <code>unified</code> }	Displays information for files and directories with paths in volumes of the specified security style. This command is not supported for SVMs with Infinite Volumes; therefore, the <code>unified</code> value is not valid for this release. This is the associated security type of the volume or qtree.
<code>-effective-style</code> { <code>unix</code> <code>ntfs</code> <code>mixed</code> <code>unified</code> }	Displays information for files and directories with the specified effective security style on the path. This command is not supported for SVMs with Infinite Volumes; therefore, the <code>unified</code> value is not valid for this release. This is the security scheme in effect for a given file or directory. A file or directory can have one of two security styles, either NTFS or UNIX. The effective security style is important with mixed security-style volumes and qtrees since a file or directory can have either NTFS-effective or UNIX-effective security (but not both).
<code>-dos-attributes</code> <i>hex_integer</i>	Displays information only for files and directories with the specified DOS attributes.
<code>-text-dos-attr</code> <i>text</i>	Displays information only for files and directories with the specified text DOS attributes.
<code>-expanded-dos-attr</code> <i>text</i>	Displays information only for files and directories with the specified extended DOS attributes.
<code>-user-id</code> <i>unix_user_ID</i>	Displays information only for files and directories with the specified UNIX user ID.
<code>-group-id</code> <i>unix_group_ID</i>	Displays information only for files and directories with the specified UNIX group ID.

Optional parameter	Description
<code>-mode-bits</code> <i>octal_permissions</i>	Displays information only for files and directories with the specified UNIX mode bits in Octal form.
<code>-text-mode-bits</code> <i>text</i>	Displays information only for files and directories with the specified UNIX mode bits in text form.
<code>-acls system_acls</code>	Displays information only for files and directories with the specified ACLs. You can enter the following information: <ul style="list-style-type: none"> • Type of ACL, which can be NTFS or NFSv4 • Control bits in the security descriptors • Owner, which applies only in the case of NTFS security descriptors. • Group, which applies only in the case of NTFS security descriptors. • Access Control Entries (ACEs) which includes both discretionary access control list (DACL) and system access control list (SACL) access control entries (ACEs) in the ACL.

Note: NTFS security-style volumes and qtrees use only NTFS system access control lists for audit policies. Mixed security-style volumes and qtrees can contain some files and directories that are of NTFS security style, which can have NTFS audit policies applied to them.

Step

1. Display audit policy settings:

```
vserver security file-directory show -vserver vserver_name -path path  
optional_parameters
```

Example

The following example displays the audit policy information about the path `/corp` in SVM vs1. This NTFS-security-style path has a NTFS-effective security style. The NTFS security descriptor contains both a SUCCESS and a SUCCESS/FAIL SACL entry:

```
vserver security file-directory show -vserver vs1 -path /corp
```

```
Vserver: vs1
  File Path: /corp
  Security Style: ntfs
  Effective Style: ntfs
  DOS Attributes: 10
  DOS Attributes in Text: ----D---
  Expanded Dos Attributes: -
    Unix User Id: 0
    Unix Group Id: 0
    Unix Mode Bits: 777
  Unix Mode Bits in Text: rwxrwxrwx
    ACLs: NTFS Security Descriptor
          Control:0x8014
          Owner:DOMAIN\Administrator
          Group:BUILTIN\Administrators
          SACL - ACEs
```

```

ALL-DOMAIN\Administrator-0x100081-OI|CI|SA|FA
SUCCESSFUL-DOMAIN\user1-0x100116-OI|CI|SA
DACL - ACEs
ALLOW-BUILTIN\Administrators-0x1f01ff-OI|CI
ALLOW-BUILTIN\Users-0x1f01ff-OI|CI
ALLOW-CREATOR OWNER-0x1f01ff-OI|CI
ALLOW-NT AUTHORITY\SYSTEM-0x1f01ff-OI|CI

```

Managing auditing configurations

You can manage Storage Virtual Machine (SVM) auditing configurations by manually rotating the audit logs, enabling or disabling auditing, displaying information about auditing configurations, modifying auditing configurations, and deleting auditing configurations. You also need to understand what happens when reverting to a release where auditing is not supported.

Manually rotating the audit event logs

Before you can view the audit event logs, the logs must be converted to user-readable formats. If you want to view the event logs for a specific Storage Virtual Machine (SVM) before Data ONTAP automatically rotates the log, you can manually rotate the audit event logs on an SVM.

Step

1. Rotate the audit event logs by using the `vserver audit rotate-log` command.

Example

```
vserver audit rotate-log -vserver vs1
```

The audit event log is saved in the SVM audit event log directory with the format specified by the auditing configuration (XML or EVTX), and can be viewed by using the appropriate application.

Enabling and disabling auditing on SVMs

You can enable or disable auditing on Storage Virtual Machines (SVMs) with FlexVol volumes. You might want to temporarily stop file and directory auditing by disabling auditing. You can enable auditing at any time (if an auditing configuration exists).

Before you begin

The Storage Virtual Machine (SVM) auditing configuration must already exist before you enable auditing. Disabling auditing does not delete the auditing configuration.

Steps

1. Perform the appropriate command:

If you want auditing to be...	Enter the command...
Enabled	<code>vserver audit enable -vserver vserver_name</code>
Disabled	<code>vserver audit disable -vserver vserver_name</code>

2. Verify that auditing is in the desired state:

```
vserver audit show -vserver vserver_name
```

Examples

The following example enables auditing for SVM vs1:

```
cluster1::> vserver audit enable -vserver vs1
cluster1::> vserver audit show -vserver vs1
```

Vserver	State	Log Format	Target Directory
vs1	true	evtx	/audit_log

The following example disables auditing for SVM vs1:

```
cluster1::> vserver audit disable -vserver vs1
```

Vserver	State	Log Format	Target Directory
vs1	false	evtx	/audit_log

Displaying information about auditing configurations

You can display information about auditing configurations for Storage Virtual Machines (SVMs) with FlexVol volumes. The information can help you determine whether the configuration is what you want in place for each SVM. The displayed information also enables you to verify whether an auditing configuration is enabled.

About this task

You can display detailed information about auditing configurations on all SVMs or you can customize what information is displayed in the output by specifying optional parameters. If you do not specify any of the optional parameters, the following is displayed:

- SVM name to which the auditing configuration applies
- The audit state, which can be `true` or `false`
If the audit state is `true`, auditing is enabled. If the audit state is `false`, auditing is disabled.
- The audit log format
- The target directory where the auditing subsystem stores consolidated and converted audit logs

Step

1. Display information about the auditing configuration by using the `vserver audit show` command.

For more information about using the command, see the man pages.

Examples

The following example displays the name, audit state, and target directory for all SVMs:

```
cluster1::> vserver audit show
```

Vserver	State	Log Format	Target Directory
vs1	false	evtx	/audit_log

The following example displays SVM names and details about the audit log for all SVMs:

```
cluster1::> vserver audit show -log-save-details
```

Vserver	Rotation File Size	Rotation Schedule	Rotation Limit
vs1	100MB	-	0

The following example displays, in list form, all audit configuration information about all SVMs:

```
cluster1::> vserver audit show -instance
```

```

Vserver: vs1
Auditing state: true
Log Destination Path: /audit_log
Log Format: evtx
Log File Size Limit: 100MB
Log Rotation Schedule: Month: -
Log Rotation Schedule: Day of Week: -
Log Rotation Schedule: Day: -
Log Rotation Schedule: Hour: -
Log Rotation Schedule: Minute: -
Rotation Schedules: -
Log Files Rotation Limit: 0

```

Commands for modifying auditing configurations

If you want to change an auditing setting for your Storage Virtual Machine (SVM), you can modify the current configuration at any time.

If you want to...	Use this command...
Modify the log destination path	<code>vserver audit modify</code> with the <code>-destination</code> parameter
Enabling automatic saves based on internal log file size	<code>vserver audit modify</code> with the <code>-rotate-size</code> parameter
Enabling automatic saves based on a time interval	<code>vserver audit modify</code> with the <code>-rotate-schedule-month</code> , <code>-rotate-schedule-dayofweek</code> , <code>-rotate-schedule-day</code> , <code>-rotate-schedule-hour</code> , and <code>-rotate-schedule-minute</code> parameters
Specifying the maximum number of saved log files	<code>vserver audit modify</code> with the <code>-rotate-limit</code> parameter

See the man page for the `vserver audit modify` command for more information.

Deleting an auditing configuration

If you no longer want to audit file and directory events on the Storage Virtual Machine (SVM) and do not want to maintain an auditing configuration on the SVM, you can delete the auditing configuration.

Steps

1. Disable the auditing configuration:

```
vserver audit disable -vserver vserver_name
```

Example

```
vserver audit disable -vserver vs1
```

2. Delete the auditing configuration:

```
vserver audit delete -vserver vserver_name
```

Example

```
vserver audit delete -vserver vs1
```

What the process is when reverting

If you plan to revert the cluster you should be aware of the process Data ONTAP follows when reverting and there are auditing-enabled Storage Virtual Machines (SVMs) in the cluster. You must take certain actions before reverting.

Reverting to a version of Data ONTAP that supports auditing, but does not support the EVTX log format

Support for the EVTX log format starts with Data ONTAP 8.2.1 in the 8.2 release family. If you are reverting to Data ONTAP 8.2, a version that supports auditing, but does not support the EVTX log format, you do not need to disable auditing on auditing-enabled SVMs before you revert. However, for each auditing configuration on the cluster (enabled or disabled), you must change the log format to the XML log format prior to reverting.

Reverting to a version of Data ONTAP that does not supports auditing

Support for auditing starts with Data ONTAP 8.2. If you plan to revert the cluster to a Data ONTAP release that does not support auditing and you have audit-enabled Storage Virtual Machines (SVMs), you should be aware of the process Data ONTAP follows when reverting.

- Prior to revert, you must manually disable and delete all auditing configurations on all SVMs in the cluster.

When you disable auditing on all SVMs in the cluster, Data ONTAP consolidates and converts all auditing logs in the staging files for all SVMs. All converted audit logs are stored in the event log directory location specified in the auditing configuration for each audit-enabled SVM. The converted event logs are available post-revert.

- When you delete all auditing configurations across the cluster, Data ONTAP deletes all staging volumes.

There is no need to manually delete staging volumes.

- During the revert, each file that has an NFSv4.x ACL is checked to determine whether the ACL contains an audit ACE.

If it does, the complete ACL is dropped.

Troubleshooting auditing and staging volume space issues

Issues can arise when there is insufficient space on either the staging volumes or on the volume containing the audit event logs. If there is insufficient space, new audit records cannot be created, which prevents clients from accessing data, and access requests fail. You should know how to troubleshoot and resolve these volume space issues.

How to troubleshoot space issues related to the event log volumes

If volumes containing event log files run out of space, auditing cannot convert log records into log files. This results in client access failures. You need to know how to troubleshoot space issues related to event log volumes.

- Storage Virtual Machine (SVM) and cluster administrators can determine whether there is insufficient volume space by displaying information about volume and aggregate usage and configuration.
- If there is insufficient space in the volumes containing event logs, SVM and cluster administrators can resolve the space issues by either removing some of the event log files or by increasing the size of the volume.

Note: If the aggregate that contains the event log volume is full, then the size of the aggregate must be increased before you can increase the size of the volume. Only a cluster administrator can increase the size of an aggregate.

- The destination path for the event log files can be changed to a directory on another volume by modifying the auditing configuration.

For more information about viewing information about volumes and increasing volume size, see the *Clustered Data ONTAP Logical Storage Management Guide*.

For more information about viewing information about aggregates and managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

How to troubleshoot space issues related to the staging volumes (cluster administrators only)

If any of the volumes containing staging files for your Storage Virtual Machine (SVM) runs out of space, auditing cannot write log records into staging files. This results in client access failures. To troubleshoot this issue, a cluster administrator needs to determine whether any of the staging volumes used in the SVM are full by displaying information about volume usage.

If the volume containing the consolidated event log files has sufficient space but there are still client access failures due to insufficient space, then the staging volumes might be out of space. The SVM administrator must contact the cluster administrator to determine whether the staging volumes that contain staging files for the SVM have insufficient space. The auditing subsystem generates an EMS event if auditing events cannot be generated due to insufficient space in a staging volume. The following message is displayed: `No space left on device`. Only the cluster administrator can view information about staging volumes.

If there is insufficient space in the staging volumes, the cluster administrators can resolve the space issues by increasing the size of the volume.

Note: If the aggregate that contains the staging volume is full, then the size of the aggregate must be increased before the cluster administrator can increase the size of the volume. Only a cluster administrator can increase the size of an aggregate.

For more information about viewing information about volumes and increasing volume size, see the *Clustered Data ONTAP Logical Storage Management Guide*.

For more information about viewing information about aggregates and managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Using FPolicy for file monitoring and management on SVMs with FlexVol volumes

FPolicy is a file access notification framework that is used to monitor and manage file access events on Storage Virtual Machines (SVMs) with FlexVol volumes.

The framework generates notifications that are sent to either external FPolicy servers or to Data ONTAP. FPolicy supports event notifications for files and directories that are accessed using NFS and SMB.

Note: FPolicy is not supported on SVM with Infinite Volume.

How FPolicy works

Before you plan and create your FPolicy configuration, you should understand the basics of how FPolicy works.

What the two parts of the FPolicy solution are

There are two parts to an FPolicy solution. There is the Data ONTAP FPolicy framework that manages activities on the cluster and sends notifications to external FPolicy servers and there are external FPolicy servers that process notifications sent by Data ONTAP FPolicy.

The Data ONTAP framework creates and maintains the FPolicy configuration, monitors file events, and sends notifications to external FPolicy servers. Data ONTAP FPolicy provides the infrastructure that allows communication between external FPolicy servers and Storage Virtual Machine (SVM) nodes.

The FPolicy framework connects to external FPolicy servers and sends notifications for certain file system events to the FPolicy servers when these events occur as a result of client access. The external FPolicy servers process the notifications and send responses back to the node. What happens as a result of the notification processing depends on the application and whether the communication between the node and the external servers is asynchronous or synchronous.

What synchronous and asynchronous communications are

FPolicy sends notifications to external FPolicy servers via the FPolicy interface. The notifications are sent either in synchronous or asynchronous mode. The notification mode determines what Data ONTAP does after sending notifications to FPolicy servers.

Asynchronous notifications	With asynchronous notifications, the node does not wait for a response from the FPolicy server, which enhances overall throughput of the system. This type of notification is suitable for applications where the FPolicy server does not require
-----------------------------------	---

that any action be taken as a result of notification evaluation. For example, asynchronous notifications are used when the Storage Virtual Machine (SVM) administrator wants to monitor and audit file access activity.

Synchronous notifications

When configured to run in synchronous mode, the FPolicy server must acknowledge every notification before the client operation is allowed to continue. This type of notification is used when an action is required based on the results of notification evaluation. For example, synchronous notifications are used when the SVM administrator wants to either allow or deny requests based on criteria specified on the external FPolicy server.

Synchronous and asynchronous applications

There are many possible uses for FPolicy applications, both asynchronous and synchronous.

Asynchronous applications are ones where the external FPolicy server does not alter access to files or directories or modify data on the Storage Virtual Machine (SVM) and include the following:

- File access and audit logging
- Storage resource management

Synchronous uses cases are ones where data access is altered or data is modified by the external FPolicy server and include the following:

- Quota management
- File access blocking
- File archiving and hierarchical storage management
- Encryption and decryption services
- Compression and decompression services

These applications are in no way all-encompassing, and by using the SDK for FPolicy, implementation of other applications are possible.

Roles that cluster components play with FPolicy

The cluster, the contained Storage Virtual Machines (SVMs), and data LIFs all play a role in an FPolicy implementation.

cluster The cluster contains the FPolicy management framework and maintains and manages information about all FPolicy configurations in the cluster.

SVM An FPolicy configuration is defined at the SVM level. The scope of the configuration is the SVM, and it only operates on SVM resources. One SVM configuration cannot monitor and send notifications for file access requests that are made for data residing on another SVM.

FPolicy configurations can be defined on the admin SVM. Once configurations are defined on the admin SVM, they can be seen and used in all SVMs.

data LIFs Connections to the FPolicy servers are made through data LIFs belonging to the SVM with the FPolicy configuration. The data LIFs used for these connections can fail over in the same manner as data LIFs used for normal client access.

How FPolicy works with external FPolicy servers

After FPolicy is configured and enabled on the Storage Virtual Machine (SVM), FPolicy runs on every node on which the SVM participates. FPolicy is responsible for establishing and maintaining connections with external FPolicy servers (FPolicy servers), for notification processing, and for managing notification messages to and from FPolicy servers.

Additionally, as part of connection management, FPolicy has the following responsibilities:

- Ensures that file notification flows through the correct LIF to the FPolicy server.
- Ensures that when multiple FPolicy servers are associated with a policy, load balancing is done when sending notifications to the FPolicy servers.
- Attempts to reestablish the connection when a connection to an FPolicy server is broken.
- Sends the notifications to FPolicy servers over an authenticated session.

How control channels are used for FPolicy communication

FPolicy initiates a control channel connection to an external FPolicy server from the data LIFs of each node participating on a Storage Virtual Machine (SVM). FPolicy uses control channels for transmitting file notifications; therefore, an FPolicy server might see multiple control channel connections based on SVM topology.

How privileged data access channels are used for synchronous communication

With synchronous use cases, the FPolicy server accesses data residing on the Storage Virtual Machine (SVM) through a privileged data access path. Access through the privileged path exposes the complete file system to the FPolicy server. It can access data files to collect information, to scan files, read files, or write into files.

Because the external FPolicy server can access the entire file system from the root of the SVM through the privileged data channel, the privileged data channel connection must be secure.

How FPolicy connection credentials are used with privileged data access channels

The FPolicy server makes privileged data access connections to cluster nodes by using a specific Windows user credential that is saved with the FPolicy configuration. SMB is the only supported protocol for making a privileged data access channel connection.

If the FPolicy server requires privileged data access, the following conditions must be met:

- A CIFS license must be enabled on the cluster.
- The FPolicy server must run under the credentials configured in the FPolicy configuration.

When making a data channel connection, FPolicy uses the credential for the specified Windows user name. Data access is made over the admin share `ONTAP_ADMIN$`.

What granting super user credentials for privileged data access means

Data ONTAP uses the combination of the IP address and the user credential configured in the FPolicy configuration to grant super user credentials to the FPolicy server.

Super user status grants these privileges when the FPolicy server accesses data:

- Avoid permission checks
The user avoids checks on files and directory access.
- Special locking privileges
Data ONTAP allows read, write, or modify access to any file regardless of existing locks. If the FPolicy server takes byte range locks on the file, it results in immediate removal of existing locks on the file.
- By-pass any FPolicy checks
Access does not generate any FPolicy notifications.

How FPolicy manages policy processing

There might be multiple FPolicy policies assigned to your Storage Virtual Machine (SVM); each with a different priority. To create an appropriate FPolicy configuration on the SVM, it is important to understand how FPolicy manages policy processing.

Each file access request is initially evaluated to determine which policies are monitoring this event. If it is a monitored event, information about the monitored event along with interested policies is passed to FPolicy where it is evaluated. Each policy is evaluated in order of the assigned priority.

You should consider the following recommendations when configuring policies:

- When you want a policy to always be evaluated before other policies, configure that policy with a higher priority.
- If the success of requested file access operation on a monitored event is a prerequisite for a file request that is evaluated against another policy, give the policy that controls the success or failure of the first file operation a higher priority.
For example, if one policy manages FPolicy file archiving and restore functionality and a second policy manages file access operations on the online file, the policy that manages file restoration must have a higher priority so that the file is restored before the operation managed by the second policy can be allowed.
- If you want all policies that might apply to a file access operation to be evaluated, give synchronous policies a lower priority.

You can reorder policy priorities for existing policies by modifying the policy sequence number. However, to have FPolicy evaluate policies based on the modified priority order, you must disable and reenable the policy with the modified sequence number.

What the node-to-external FPolicy server communication process is

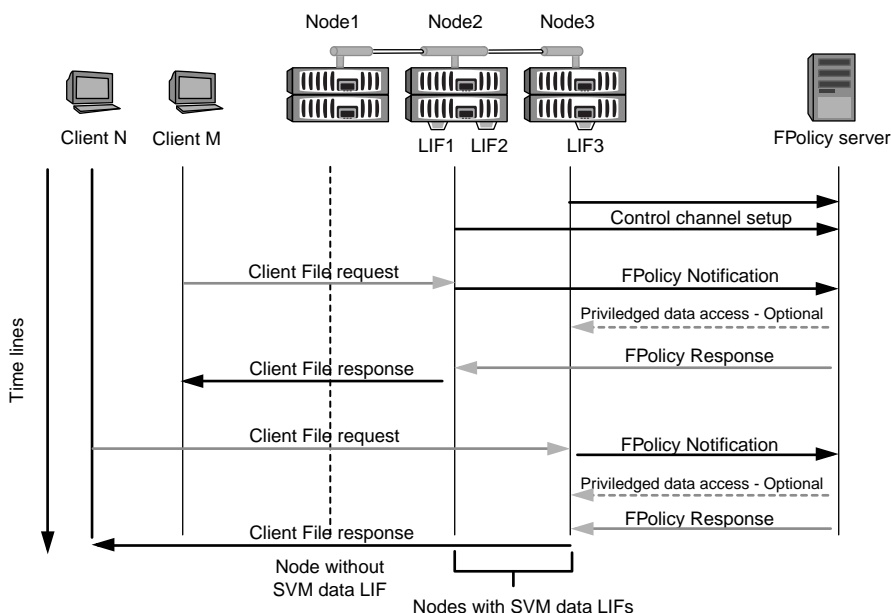
To properly plan your FPolicy configuration, you should understand what the node-to-external FPolicy server communication process is.

Every node that participates on each Storage Virtual Machine (SVM) initiates a connection to an external FPolicy server (FPolicy server) using TCP/IP. Connections to the FPolicy servers are set up using node data LIFs; therefore, a participating node can set up a connection only if the node has an operational data LIF for the SVM.

Each FPolicy process on participating nodes attempts to establish a connection with the FPolicy server when the policy is enabled. It uses the IP address and port of the FPolicy external engine specified in the policy configuration.

The connection establishes a control channel from each of the nodes participating on each SVM to the FPolicy server through the data LIF. In addition, if IPv4 and IPv6 data LIF addresses are present on the same participating node, FPolicy attempts to establish connections for both IPv4 and IPv6. Therefore, in a scenario where the SVM extends over multiple nodes or if both IPv4 and IPv6 addresses are present, the FPolicy server must be ready for multiple control channel setup requests from the cluster after the FPolicy policy is enabled on the SVM.

For example, if a cluster has three nodes—Node1, Node2, and Node3—and SVM data LIFs are spread across only Node2 and Node3, control channels are initiated only from Node2 and Node3, irrespective of the distribution of data volumes. Say that Node2 has two data LIFs—LIF1 and LIF2—that belong to the SVM and that the initial connection is from LIF1. If LIF1 fails, FPolicy attempts to establish a control channel from LIF2.



How FPolicy manages external communication during LIF migration or failover

Data LIFs can be migrated to data ports in the same node or to data ports on a remote node.

When a data LIF fails over or is migrated, a new control channel connection is made to the FPolicy server. FPolicy can then retry SMB and NFS client requests that timed out, with the result that new notifications are sent to the external FPolicy servers. The node rejects FPolicy server responses to original, timed-out SMB and NFS requests.

How FPolicy manages external communication during node failover

If the cluster node that hosts the data ports used for FPolicy communication fails, Data ONTAP breaks the connection between the FPolicy server and the node.

The impact of cluster failover to the FPolicy server can be mitigated by configuring the LIF manager to migrate the data port used in FPolicy communication to another active node. After the migration is complete, a new connection is established using the new data port.

If the LIF manager is not configured to migrate the data port, the FPolicy server must wait for the failed node to come up. After the node is up, a new connection is initiated from that node with a new Session ID.

Note: The FPolicy server detects broken connections with the keep-alive protocol message. The timeout for purging the session ID is determined when configuring FPolicy. The default keep-alive timeout is two minutes.

How FPolicy services work across SVM namespaces

Data ONTAP provides a unified Storage Virtual Machine (SVM) namespace. Volumes across the cluster are joined together by junctions to provide a single, logical file system. The FPolicy server is aware of the namespace topology and provides FPolicy services across the namespace.

The namespace is specific to and contained within the SVM; therefore, you can see the namespace only from the SVM context. Namespaces have the following characteristics:

- A single namespace exists in each SVM, with the root of the namespace being the root volume, represented in the namespace as slash (/).
- All other volumes have junction points below the root (/).
- Volume junctions are transparent to clients.
- A single NFS export can provide access to the complete namespace; otherwise, export policies can export specific volumes.
- SMB shares can be created on the volume or on qtrees within the volume or on any directory within the namespace.
- The namespace architecture is flexible.

Examples of typical namespace architectures are as follows:

- A namespace with a single branch off of the root
- A namespace with multiple branches off of the root

- A namespace with multiple unbranched volumes off of the root

FPolicy configuration types

There are two basic FPolicy configuration types. One configuration uses external FPolicy servers to process and act upon notifications. The other configuration does not use external FPolicy servers; instead, it uses the Data ONTAP internal, native FPolicy server for simple file blocking based on extensions.

External FPolicy server configuration	The notification is sent to the FPolicy server, which screens the request and applies rules to determine whether the node should allow the requested file operation. For synchronous policies, the FPolicy server then sends a response to the node to either allow or block the requested file operation.
Native FPolicy server configuration	The notification is screened internally. The request is allowed or denied based on file extensions settings configured in the FPolicy scope.

When to create a native FPolicy configuration

Native FPolicy configurations use the Data ONTAP internal FPolicy engine to monitor and block file operations based on the file's extension. This solution does not require external FPolicy servers (FPolicy servers). Using a native file blocking configuration is appropriate when this simple solution is all that is needed.

Native file blocking enables you to monitor any file operations that match configured operation and filtering events and then deny access to files with particular extensions. This is the default configuration.

This configuration provides a means to block file access based only on the file's extension. 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 for operations that generate notifications.

The following applies to native FPolicy configurations:

- The same set of filters and protocols that are supported by FPolicy server-based file screening are also supported for native file blocking.
- Native file blocking and FPolicy server-based file screening applications can be configured at the same time.
To do so, you can configure two separate FPolicy policies for the Storage Virtual Machine (SVM), with one configured for native file blocking and one configured for FPolicy server-based file screening.
- The native file blocking feature only screens files based on the extensions and not on the content of the file.
- In the case of symbolic links, native file blocking uses the file extension of the root file.

When to create a configuration that uses external FPolicy servers

FPolicy configurations that use external FPolicy servers to process and manage notifications provide robust solutions for use cases where more than simple file blocking based on file extension is needed.

You should create a configuration that uses external FPolicy servers when you want to do such things as monitor and record file access events, provide quota services, perform file blocking based on criteria other than simple file extensions, provide data migration services using hierarchical storage management applications, or provide a fine-grained set of policies that monitor only a subset of data in the Storage Virtual Machine (SVM).

Requirements, considerations, and best practices for configuring FPolicy

Before you create and configure FPolicy configurations on your Storage Virtual Machines (SVMs) with FlexVol volumes, you need to be aware of certain requirements, considerations, and best practices for configuring FPolicy.

Ways to configure FPolicy

FPolicy features are configured either through the command line interface (CLI) or through APIs. This guide uses the CLI to create, manage, and monitor an FPolicy configuration on the cluster.

Requirements for setting up FPolicy

Before you configure and enable FPolicy on your Storage Virtual Machine (SVM), you need to be aware of certain requirements.

- All nodes in the cluster must be running a version of Data ONTAP that supports FPolicy.
- If you are not using the Data ONTAP native FPolicy engine, you must have external FPolicy servers (FPolicy servers) installed.
- The FPolicy servers must be installed on a server accessible from the data LIFs of the SVM where FPolicy policies are enabled.
- The IP address of the FPolicy server must be configured as a primary or secondary server in the FPolicy policy external engine configuration.
- If the FPolicy servers access data over a privileged data channel, the following additional requirements must be met:
 - CIFS must be licensed on the cluster.
Privileged data access is accomplished using SMB connections.
 - A user credential must be configured for accessing files over the privileged data channel.
 - The FPolicy server must run under the credentials configured in the FPolicy configuration.

Best practices and recommendations when setting up FPolicy

When setting up FPolicy on Storage Virtual Machines (SVMs), you need to be familiar with configuration best practices and recommendations to ensure that your FPolicy configuration provides robust monitoring performance and results that meet your requirements.

- External FPolicy servers (FPolicy servers) should be placed in close proximity to the cluster with high-bandwidth connectivity to provide minimal latency and high-bandwidth connectivity.
- The FPolicy external engine should be configured with more than one FPolicy server to provide resiliency and high availability of FPolicy server notification processing, especially if policies are configured for synchronous screening.
- It is recommended to disable the FPolicy policy before making any configuration changes. For example, if you want to add or modify an IP address in the FPolicy external engine configured for the enabled policy, you should first disable the policy.

- If you configure FPolicy to monitor FlexCache volumes, it is recommended that you do not configure FPolicy to monitor `read` and `get attr` file operations on the FlexCache volumes. This is because Data ONTAP needs to retrieve inode-to-path (I2P) data with these operations, and this data cannot be retrieved from the FlexCache volume. Instead, the I2P request is forwarded to the origin volume, with the result that the performance benefits from FlexCache are not realized when FPolicy is used to monitor `read` and `get attr` operations on FlexCache volumes.
- The cluster node-to-FPolicy server ratio should be optimized to ensure that FPolicy servers are not overloaded, which can introduce latencies when the SVM responds to client requests. The optimal ratio depends on the application for which the FPolicy server is being used.

Important revert considerations

You must understand and act on some important revert considerations before reverting to a Data ONTAP release that does not support FPolicy.

Before reverting to a version of Data ONTAP that does not support FPolicy, the following conditions must be met:

- Every file on which FPolicy servers set the offline bit must be either deleted or replaced with the original files before disabling FPolicy and reverting to a version of Data ONTAP that does not support FPolicy.

Note: If you do not replace the files with the offline bit set with the original files prior to reverting, clients access the stub files instead of the files to which the stub refers.
- FPolicy functionality must be disabled on the cluster by disabling every FPolicy policy on the cluster.

What the steps for setting up an FPolicy configuration are

Before FPolicy can monitor file access, an FPolicy configuration must be created and enabled on the Storage Virtual Machine (SVM) for which FPolicy services are required.

The steps for setting up and enabling an FPolicy configuration on the SVM are as follows:

1. Create an FPolicy external engine.

The FPolicy external engine identifies the external FPolicy servers (FPolicy servers) that are associated with a specific FPolicy configuration. If the internal “native” FPolicy engine is used to create a native file-blocking configuration, you do not need to create an FPolicy external engine.

2. Create an FPolicy event.

An FPolicy event describes what the FPolicy policy should monitor. Events consist of the protocols and file operations to monitor, and can contain a list of filters. Events use filters to narrow the list of monitored events for which the FPolicy external engine must send notifications. Events also specify whether the policy monitors volume operations.

3. Create an FPolicy policy.

The FPolicy policy is responsible for associating, with the appropriate scope, the set of events that need to be monitored and for which of the monitored events notifications must be sent to the designated FPolicy server (or to the native engine if no FPolicy servers are configured). The policy also defines whether the FPolicy server is allowed privileged access to the data for which it receives notifications. An FPolicy server needs privileged access if the server needs to access the data. Typical use cases where privileged access is needed include file blocking, quota management, and hierarchical storage management. The policy is where you specify whether the configuration for this policy uses an FPolicy server or the internal “native” FPolicy server.

A policy specifies whether screening is mandatory. If screening is mandatory and all FPolicy servers are down or no response is received from the FPolicy servers within a defined timeout period, then file access is denied.

A policy's boundaries are the SVM. A policy cannot apply to more than one SVM. However, a specific SVM can have multiple FPolicy policies, each with the same or different combination of scope, event, and external server configurations.

4. Configure the policy scope.

The FPolicy scope determines which volumes, shares, or export-policies the policy acts on or excludes from monitoring. A scope also determines which file extensions should be included or excluded from FPolicy monitoring.

Note: Exclude lists take precedence over include lists.

5. Enable the FPolicy policy.

When the policy is enabled, the control channels and, optionally, the privileged data channels are connected. The FPolicy process on the nodes on which the SVM participates begin monitoring file and folder access and, for events that match configured criteria, sends notifications to the FPolicy servers (or to the native engine if no FPolicy servers are configured).

Note: If the policy uses native file blocking, an external engine is not configured or associated with the policy.

Planning the FPolicy configuration

Before you create an FPolicy configuration, you must understand what is involved in each step of the configuration. You need to decide what settings you need to use when performing the configuration and record them in the planning worksheets.

You need to plan for the following configuration tasks:

- Creating the FPolicy external engine
- Creating the FPolicy policy event
- Creating the FPolicy policy
- Creating the FPolicy policy scope

FPolicy is supported on Storage Virtual Machines (SVMs) with FlexVol volumes. FPolicy is not supported on SVMs with Infinite Volume.

Planning the FPolicy external engine configuration

Before you configure the FPolicy external engine (external engine), you must understand what it means to create an external engine and which configuration parameters are available. This information helps you to determine which values to set for each parameter.

What it means to create an external engine

Creating the external engine means defining the information that FPolicy needs to make and manage connections to the external FPolicy servers (FPolicy servers). The external engine configuration defines the following configuration information:

- Storage Virtual Machine (SVM) name
- Engine name
- The IP addresses of the primary and secondary FPolicy servers and the TCP port number to use when making the connection to the FPolicy servers
- Whether the engine type is asynchronous or synchronous
- How to authenticate the connection between the node and the FPolicy server
If you choose to configure mutual SSL authentication, then you must also configure parameters that provide SSL certificate information.
- How to manage the connection (advanced privilege settings)
This includes parameters that define such things as timeout values, retry values, keep-alive values, and maximum request values.

What the basic external engine parameters are

You can use the following table of basic FPolicy configuration parameters to help you plan your configuration:

Type of information	Option
<p><i>SVM</i></p> <p>Specifies the SVM name that you want to associate with this external engine.</p> <p>Each FPolicy configuration is defined within a single SVM. The external engine, policy event, policy scope, and policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.</p>	<p>-vserver vserver_name</p>
<p><i>Engine name</i></p> <p>Specifies the name to assign to the external engine configuration. You must specify the engine name later when you create the FPolicy policy. This associates the external engine with the policy.</p>	<p>-engine-name engine_name</p>
<p><i>Primary FPolicy servers</i></p> <p>Specifies the primary FPolicy servers to which the node sends notifications for a given FPolicy policy. The value is specified as a comma-delimited list of IP addresses.</p> <p>If more than one primary server IP address is specified, every node on which the SVM participates creates a control connection to every specified primary FPolicy server at the time the policy is enabled. If you configure multiple primary FPolicy servers, notifications are sent to the FPolicy servers in a round-robin fashion.</p>	<p>-primary-servers IP_address,...</p>
<p><i>Port number</i></p> <p>Specifies the port number of the FPolicy service.</p>	<p>-port integer</p>
<p><i>Secondary FPolicy servers</i></p> <p>Specifies the secondary FPolicy servers to which to send file access events for a given FPolicy policy. The value is specified as a comma-delimited list of IP addresses.</p> <p>Secondary servers are used only when none of the primary servers are reachable. Connections to secondary servers are established when the policy is enabled, but notifications are sent to secondary servers only if none of the primary servers are reachable. If you configure multiple secondary servers, notifications are sent to the FPolicy servers in a round-robin fashion.</p>	<p>-secondary-servers IP_address,...</p>

Type of information	Option
<p><i>External engine type</i></p> <p>Specifies whether the external engine operates in synchronous or asynchronous mode. By default, FPolicy operates in synchronous mode.</p> <p>When set to <i>synchronous</i>, file request processing sends a notification to the FPolicy server, but then does not continue until after receiving a response from the FPolicy server. At that point, request flow either continues or processing results in denial, depending on whether the response from the FPolicy server permits the requested action.</p> <p>When set to <i>asynchronous</i>, file request processing sends a notification to the FPolicy server, and then continues.</p>	<p><code>-extern-engine-type</code> <code>external_engine_type</code></p> <p>The value for this parameter can be one of the following:</p> <ul style="list-style-type: none"> • <code>synchronous</code> • <code>asynchronous</code>
<p><i>SSL option for communication with FPolicy server</i></p> <p>Specifies the SSL option for communication with the FPolicy server. This is a required parameter. You can choose one of the options based on the following information:</p> <ul style="list-style-type: none"> • When set to <code>no-auth</code>, no authentication takes place. The communication link is established over TCP. • When set to <code>server-auth</code>, the SVM authenticates the FPolicy server. If you choose this value, before creating the external engine, you must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate. • When set to <code>mutual-auth</code>, mutual authentication takes place between the SVM and the FPolicy server; the SVM authenticates the FPolicy server, and the FPolicy server authenticates the SVM. If you choose this value, before creating the external engine, the administrator must install the public certificate of the CA that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. <p>The public certificate of CA that is used to sign the FPolicy server certificate is installed by using the <code>security certificate install</code> command with the <code>-type</code> parameter set to <code>client_ca</code>. The private key and public certificate required for authentication of the SVM is installed by using the <code>security certificate install</code> command with the <code>-type</code> parameter set to <code>server</code>.</p> <p>If you choose to configure mutual SSL authentication, then you must also configure the <code>-certificate-common-name</code>, <code>-certificate-serial</code>, and <code>-certificate-ca</code> parameters.</p>	<p><code>-ssl-option {no-auth server-auth mutual-auth}</code></p>

Type of information	Option
<p><i>Certificate FQDN or custom common name</i></p> <p>Specifies the certificate name used if SSL authentication between the SVM and the FPolicy server is configured. You can specify the certificate name as an FQDN or as a custom common name.</p> <p>If you specify <code>mutual-auth</code> for the <code>-ssl-option</code> parameter, you must specify a value for the <code>-certificate-common-name</code> parameter.</p>	<code>-certificate-common-name text</code>
<p><i>Certificate serial number</i></p> <p>Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured.</p> <p>If you specify <code>mutual-auth</code> for the <code>-ssl-option</code> parameter, you must specify a value for the <code>-certificate-serial</code> parameter.</p>	<code>-certificate-serial text</code>
<p><i>Certificate authority</i></p> <p>Specifies the CA name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured.</p> <p>If you specify <code>mutual-auth</code> for the <code>-ssl-option</code> parameter, you must specify a value for the <code>-certificate-ca</code> parameter.</p>	<code>-certificate-ca text</code>

What the advanced external engine options are

You can use the following table of advanced FPolicy configuration parameters as you plan whether to customize your configuration with advanced parameters. You use these parameters to modify communication behavior between the cluster nodes and the FPolicy servers:

Type of information	Option
<p><i>Timeout for canceling a request</i></p> <p>Specifies the time interval in hours (h), minutes (m), or seconds (s) that the node waits for a response from the FPolicy server.</p> <p>If the timeout interval passes, the node sends a cancel request to the FPolicy server. The node then sends the notification to an alternate FPolicy server. This timeout helps in handling an FPolicy server that is not responding, which can improve SMB/NFS client response. Also, canceling requests after a timeout period can help in releasing system resources because the notification request is moved from a down/bad FPolicy server to an alternate FPolicy server.</p> <p>The range for this value is 0 through 100. If the value is set to 0, the option is disabled and cancel request messages are not sent to the FPolicy server. The default is 20s.</p>	<code>-reqs-cancel-timeout integer[h m s]</code>

Type of information	Option
<p><i>Timeout for aborting a request</i></p> <p>Specifies the timeout in hours (h), minutes (m), or seconds (s) for aborting a request.</p> <p>The range for this value is 0 through 200.</p>	<p>-reqs-abort-timeout <i>integer</i>[h m s]</p>
<p><i>Interval for sending status requests</i></p> <p>Specifies the interval in hours (h), minutes (m), or seconds (s) after which a status request is sent to the FPolicy server.</p> <p>The range for this value is 0 through 50. If the value is set to 0, the option is disabled and status request messages are not sent to the FPolicy server. The default is 10s.</p>	<p>-status-req-interval <i>integer</i>[h m s]</p>
<p><i>Maximum outstanding requests on the FPolicy server</i></p> <p>Specifies the maximum number of outstanding requests that can be queued on the FPolicy server.</p> <p>The range for this value is 1 through 10000. The default is 50.</p>	<p>-max-server-reqs <i>integer</i></p>
<p><i>Timeout for disconnecting a nonresponsive FPolicy server</i></p> <p>Specifies the time interval in hours (h), minutes (m), or seconds (s) after which the connection to the FPolicy server is terminated. The connection is terminated after the timeout period only if the FPolicy server's queue contains the maximum allowed requests and no response is received within the this timeout period. The maximum allowed number of requests is either 50 (the default) or the number specified by the <code>max-server-reqs</code> parameter.</p> <p>The range for this value is 1 through 100. The default is 60s.</p>	<p>-server-progress-timeout <i>integer</i>[h m s]</p>
<p><i>Interval for sending keep-alive messages to the FPolicy server</i></p> <p>Specifies the time interval in hours (h), minutes (m), or seconds (s) at which keep-alive messages are sent to the FPolicy server. Keep-alive messages detect half-open connections.</p> <p>The range for this value is 10 through 600. If the value is set to 0, the option is disabled and keep-alive messages are prevented from being sent to the FPolicy servers. The default is 120s.</p>	<p>-keep-alive-interval <i>integer</i>[h m s]</p>
<p><i>Maximum reconnect attempts</i></p> <p>Specifies the maximum number of times the SVM attempts to reconnect to the FPolicy server after the connection has been broken.</p> <p>The range for this value is 0 through 20. The default is 5.</p>	<p>-max-connection-retries <i>integer</i></p>

Completing the FPolicy external engine configuration worksheet

You can use this worksheet to record the values that you need during the FPolicy external engine configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the external engine.

Information for a basic external engine configuration

You should record whether you want to include each parameter setting in the external engine configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage Virtual Machine (SVM) name	Yes	Yes	
Engine name	Yes	Yes	
Primary FPolicy servers	Yes	Yes	
Port number	Yes	Yes	
Secondary FPolicy servers	No		
External engine type	No		
SSL option for communication with external FPolicy server	Yes	Yes	
Certificate FQDN or custom common name	No		
Certificate serial number	No		
Certificate authority	No		

Information for advanced external engine parameters

To configure an external engine with advanced parameters, you must enter the configuration command while in advanced privilege mode.

Type of information	Required	Include	Your values
Timeout for canceling a request	No		
Timeout for aborting a request	No		
Interval for sending status requests	No		

Type of information	Required	Include	Your values
Maximum outstanding requests on the FPolicy server	No		
Timeout for disconnecting a nonresponsive FPolicy server	No		
Interval for sending keep-alive messages to the FPolicy server	No		
Maximum reconnect attempts	No		

Planning the FPolicy event configuration

Before you configure FPolicy events, you must understand what it means to create an FPolicy event. You must determine which protocols you want the event to monitor, which events to monitor, and which event filters to use. This information helps you plan the values that you want to set.

What it means to create an FPolicy event

Creating the FPolicy event means defining information that the FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server. The FPolicy event configuration defines the following configuration information:

- Storage Virtual Machine (SVM) name
- Event name
- Which protocols to monitor
FPolicy can monitor SMB, NFSv3, and NFSv4 file access operations.
- Which file operations to monitor
Not all file operations are valid for each protocol.
- Which file filters to configure
Only certain combinations of file operations and filters are valid. Each protocol has its own set of supported combinations.
- Whether to monitor volume operations

Note: There is a dependency with three of the parameters (`-protocol`, `-file-operations`, `-filters`). The following are the valid combinations for the three parameters:

- You can specify the `-protocol` and `-file-operations` parameters.
- You can specify all three of the parameters.
- You can specify none of the parameters.

What the FPolicy event configuration contains

You can use the following list of available FPolicy event configuration parameters to help you plan your configuration:

Type of information	Option
<p><i>SVM</i></p> <p>Specifies the SVM name that you want to associate with this FPolicy event.</p> <p>Each FPolicy configuration is defined within a single SVM. The external engine, policy event, policy scope, and policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.</p>	<p><code>-vserver</code> <code>vserver_name</code></p>
<p><i>Event name</i></p> <p>Specifies the name to the FPolicy event configuration. When you create the FPolicy policy you associate the FPolicy event with the policy using the event name.</p>	<p><code>-event-name</code> <code>event_name</code></p>
<p><i>Protocol</i></p> <p>Specifies which protocol to configure for the FPolicy event. The list for <code>-protocol</code> can include one of the following values:</p> <ul style="list-style-type: none"> • <code>cifs</code> • <code>nfsv3</code> • <code>nfsv4</code> <p>Note: If you specify <code>-protocol</code>, then you must specify a valid value in the <code>-file-operations</code> parameter. As the protocol version changes, the valid values might change.</p>	<p><code>-protocol protocol</code></p>

Type of information	Option
<p><i>File operations</i></p> <p>Specifies the list of file operations for the FPolicy event.</p> <p>The event checks the operations specified in this list from all client requests using the protocol specified in the <code>-protocol</code> parameter. You can list one or more file operations by using a comma-delimited list. The list for <code>-file-operations</code> can include one or more of the following values:</p> <ul style="list-style-type: none"> • <code>close</code> for file close operations • <code>create</code> for file create operations • <code>create-dir</code> for directory create operations • <code>delete</code> for file delete operations • <code>delete_dir</code> for directory delete operations • <code>getattr</code> for get attribute operations • <code>link</code> for link operations • <code>lookup</code> for lookup operations • <code>open</code> for file open operations • <code>read</code> for file read operations • <code>write</code> for file write operations • <code>rename</code> for file rename operations • <code>rename_dir</code> for directory rename operations • <code>setattr</code> for set attribute operations • <code>symlink</code> for symbolic link operations <p>Note: If you specify <code>-file-operations</code>, then you must specify a valid protocol in the <code>-protocol</code> parameter.</p>	<p><code>-file-operations</code> <code>file_operations,...</code></p>

Type of information	Option
<p><i>Filters</i></p> <p>Specifies the list of filters for a given file operation for the specified protocol. The values in the <code>-filters</code> parameter are used to filter client requests. The list can include one or more of the following:</p> <ul style="list-style-type: none"> • <code>monitor-ads</code> to filter the client request for alternate data stream • <code>close-with-modification</code> to filter the client request for close with modification • <code>close-without-modification</code> to filter the client request for close without modification • <code>first-read</code> to filter the client request for first read • <code>first-write</code> to filter the client request for first write • <code>offline-bit</code> to filter the client request for offline bit set Setting this filter results in the FPolicy server receiving notification only when offline files are accessed. • <code>open-with-delete-intent</code> to filter the client request for open with delete intent Setting this filter results in the FPolicy server receiving notification only when an attempt is made to open a file with the intent to delete it. This is used by file systems when the <code>FILE_DELETE_ON_CLOSE</code> flag is specified. • <code>open-with-write-intent</code> to filter client request for open with write intent Setting this filter results in the FPolicy server receiving notification only when an attempt is made to open a file with the intent to write something in it. • <code>write-with-size-change</code> to filter the client request for write with size change <p>Note: If you specify the <code>-filters</code> parameter, then you must also specify valid values for the <code>-file-operations</code> and <code>-protocol</code> parameters.</p>	<p><code>-filters <i>filter</i>, ...</code></p>
<p><i>Is volume operation required</i></p> <p>Specifies whether volume operation monitoring is required. The default is <code>false</code>.</p>	<p><code>-volume-operation {true false}</code></p>

List of supported file operation and filter combinations that FPolicy can monitor for SMB

When you configure your FPolicy event, you need to be aware that only certain combinations of file operations and filters are supported for monitoring SMB file access operations.

The list of supported file operation and filter combinations for FPolicy monitoring of SMB file access events is provided in the following table:

Supported file operations	Supported filters
close	monitor-ads, offline-bit, close-with-modification, close-without-modification
create	monitor-ads, offline-bit
create_dir	Currently no filter is supported for this file operation.
delete	monitor-ads, offline-bit
delete_dir	Currently no filter is supported for this file operation.
getattr	offline-bit
open	monitor-ads, offline-bit, open-with-delete-intent, open-with-write-intent
read	monitor-ads, offline-bit, first-read
write	monitor-ads, offline-bit, first-write, write-with-size-change
rename	monitor-ads, offline-bit
rename_dir	Currently no filter is supported for this file operation.
setattr	monitor-ads, offline-bit

List of supported file operation and filter combinations that FPolicy can monitor for NFSv3

When you configure your FPolicy event, you need to be aware that only certain combinations of file operations and filters are supported for monitoring NFSv3 file access operations.

The list of supported file operation and filter combinations for FPolicy monitoring of NFSv3 file access events is provided in the following table:

Supported file operations	Supported filters
<i>create</i>	offline-bit
<i>create_dir</i>	Currently no filter is supported for this file operation.
<i>delete</i>	offline-bit
<i>delete_dir</i>	Currently no filter is supported for this file operation.

Supported file operations	Supported filters
<i>link</i>	offline-bit
<i>lookup</i>	offline-bit
<i>read</i>	offline-bit
<i>write</i>	offline-bit, write-with-size-change
<i>rename</i>	offline-bit
<i>rename_dir</i>	Currently no filter is supported for this file operation.
<i>setattr</i>	offline-bit
<i>symlink</i>	offline-bit

List of supported file operation and filter combinations that FPolicy can monitor for NFSv4

When you configure your FPolicy event, you need to be aware that only certain combinations of file operations and filters are supported for monitoring NFSv4 file access operations.

The list of supported file operation and filter combinations for FPolicy monitoring of NFSv4 file access events is provided in the following table:

Supported file operations	Supported filters
<i>close</i>	offline-bit
<i>create</i>	offline-bit
<i>create_dir</i>	Currently no filter is supported for this file operation.
<i>delete</i>	offline-bit
<i>delete_dir</i>	Currently no filter is supported for this file operation.
<i>getattr</i>	offline-bit
<i>link</i>	offline-bit
<i>lookup</i>	offline-bit
<i>open</i>	offline-bit
<i>read</i>	offline-bit
<i>write</i>	offline-bit, write-with-size-change
<i>rename</i>	offline-bit
<i>rename_dir</i>	Currently no filter is supported for this file operation.
<i>setattr</i>	offline-bit

Supported file operations	Supported filters
<i>symlink</i>	offline-bit

Completing the FPolicy event configuration worksheet

You can use this worksheet to record the values that you need during the FPolicy event configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the FPolicy event.

You should record whether you want to include each parameter setting in the FPolicy event configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage Virtual Machine (SVM) name	Yes	Yes	
Event name	Yes	Yes	
Protocol	No		
File operations	No		
Filters	No		
Is volume operation required	No		

Planning the FPolicy policy configuration

Before you configure the FPolicy policy, you must understand what it means to create an FPolicy policy. You must understand what configuration options are available. You also need to understand why you might want to attach more than one event to an FPolicy policy. This information helps you as you determine what values that you want to set.

What it means to create an FPolicy policy

Creating the FPolicy policy means associating a specific Storage Virtual Machine (SVM), an FPolicy event, and an FPolicy external engine (external engine) to an FPolicy policy. You also specify the following:

- Whether mandatory screening is required for this policy.
 - Whether to use the Data ONTAP native external engine for simple file blocking or whether to specify an external engine that is configured to use external FPolicy servers (FPolicy servers) for more sophisticated file blocking and file management.
 - Whether you want to associate more than one FPolicy event to the policy.
- An event is specific to a protocol. You can use a single FPolicy policy to monitor file access events for more than one protocol by creating an event for each protocol that you want the policy to monitor, and then associating the events to the policy.

- Whether you want the FPolicy server to have privileged access to the monitored files and folders by using a privileged data connection.
If you want to configure the policy to allow privileged access, you must also specify the user name for the account that you want the FPolicy server to use for privileged access.

What the FPolicy policy configuration contains

You can use the following list of available FPolicy policy configuration parameters to help you plan your configuration:

Type of information	Option
<p><i>SVM</i></p> <p>Specifies the SVM name on which you want to create an FPolicy policy. Each FPolicy configuration is defined within a single SVM. The external engine, FPolicy event, FPolicy scope, and FPolicy policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.</p>	<p>-vserver <i>vserver_name</i></p>
<p><i>Policy name</i></p> <p>Specifies the name of the FPolicy policy.</p> <p>The name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters (a through z, A through Z, and 0 through 9), “_”, and “.”.</p>	<p>-policy-name <i>policy_name</i></p>
<p><i>Event names</i></p> <p>Specifies a comma-delimited list of events to associate with the FPolicy policy. The events must already exist.</p>	<p>-events <i>event_name, ...</i></p>
<p><i>External engine name</i></p> <p>Specifies the name of the external engine to associate with the FPolicy policy. The external engine must already exist.</p> <p>An external engine contains information required by the node to send notifications to an FPolicy server. The default value for this parameter is <i>native</i>. This means that, if you do not specify a value for the external engine, the default native external engine is used. The native external engine is internal to Data ONTAP and is used if you want to configure native file blocking and you do not want to use FPolicy servers. If you want to use the native external engine, you can either not specify a value for this parameter or you can specify <i>native</i> as the value.</p>	<p>-engine <i>engine_name</i></p>

Type of information	Option
<p><i>Is mandatory screening required</i></p> <p>Specifies whether mandatory file access screening is required.</p> <p>This parameter specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When set to <code>true</code>, file access events are denied. When set to <code>false</code>, file access events are allowed. The default is <code>true</code>.</p>	<p><code>-is-mandatory</code> <code>{true false}</code></p>
<p><i>Allow privileged access</i></p> <p>Specifies whether the FPolicy servers can have privileged access to monitored data.</p> <p>With this option set to <code>yes</code>, FPolicy servers can access files from the root of the SVM containing the monitored data using the privileged data channel. The default is <code>no</code>.</p>	<p><code>-allow-privileged-access</code> <code>{yes no }</code></p>
<p><i>Privileged user name</i></p> <p>Specifies the user name of the account the FPolicy servers use for privileged data access.</p> <p>The value for this parameter should use the “domain\user name” format. If <code>-allow-privileged-access</code> is set to <code>no</code>, any value set for this parameter is ignored.</p>	<p><code>-privileged-user-name</code> <i>user_name</i></p>

Completing the FPolicy policy worksheet

You can use this worksheet to record the values that you need during the FPolicy policy configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the FPolicy policy.

You should record whether you want to include each parameter setting in the FPolicy policy configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage Virtual Machine (SVM) name	Yes	Yes	
Policy name	Yes	Yes	
Event names	Yes	Yes	
External engine name	Yes	Yes	
Is mandatory screening required	No		
Allow privileged access	No		

Type of information	Required	Include	Your values
Privileged user name	No		

Planning the FPolicy scope configuration

Before you configure the FPolicy scope, you must understand what it means to create a scope. You must understand what the scope configuration contains. You also need to understand what the scope rules of precedence are. This information can help you plan the values that you want to set.

What it means to create an FPolicy scope

Creating the FPolicy scope means defining the boundaries on which the FPolicy policy applies. The Storage Virtual Machine (SVM) is the basic boundary. When you create a scope for an FPolicy policy, you must define the FPolicy policy to which it will apply, and you must designate to which SVM you want to apply the scope.

There are a number of parameters that further restrict the scope within the specified SVM. You can restrict the scope by specifying what to include in the scope or by specifying what to exclude from the scope. After you apply a scope to an enabled policy, policy event checks get applied to the scope defined by this command.

Notifications are generated for file access events where matches are found in the “include” options. Notifications are not generated for file access events where matches are found in the “exclude” options.

The FPolicy scope configuration defines the following configuration information:

- SVM name
- Policy name
- The shares to include or exclude from what gets monitored
- The export policies to include or exclude from what gets monitored
- The volumes to include or exclude from what gets monitored
- The file extensions to include or exclude from what gets monitored
- Whether to do file extension checks on directory objects

Note: There are special considerations for the scope for a cluster FPolicy policy. The cluster FPolicy policy is a policy that the cluster administrator creates for the admin SVM. If the cluster administrator also creates the scope for that cluster FPolicy policy, the SVM administrator cannot create a scope for that same policy. However, if the cluster administrator does not create a scope for the cluster FPolicy policy, then any SVM administrator can create the scope for that cluster policy. In the event that the SVM administrator creates a scope for that cluster FPolicy policy, the cluster administrator cannot subsequently create a cluster scope for that same cluster policy. This is because the cluster administrator cannot override the scope for the same cluster policy.

What the scope rules of precedence are

The following rules of precedence apply to scope configurations:

- When a share is included in the `-shares-to-include` parameter and the parent volume of the share is included in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has precedence over `-shares-to-include`.
- When an export policy is included in the `-export-policies-to-include` parameter and the parent volume of the export policy is included in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has precedence over `-export-policies-to-include`.
- An administrator can specify both `-file-extensions-to-include` and `-file-extensions-to-exclude` lists. The `-file-extensions-to-exclude` parameter is checked first before the `-file-extensions-to-include` parameter is checked.

What the FPolicy scope configuration contains

You can use the following list of available FPolicy scope configuration parameters to help you plan your configuration:

Note: When configuring what shares, export policies, volumes, and file extensions to include or exclude from the scope, the include and exclude parameters can contain regular expressions and can include metacharacters such as “?” and “*”.

Type of information	Option
<i>SVM</i> Specifies the SVM name on which you want to create an FPolicy scope. Each FPolicy configuration is defined within a single SVM. The external engine, policy event, policy scope, and policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.	<code>-vserver</code> <i>vserver_name</i>
<i>Policy name</i> Specifies the name of the FPolicy policy to which you want to attach the scope. The FPolicy policy must already exist.	<code>-policy-name</code> <i>policy_name</i>
<i>Shares to include</i> Specifies a comma-delimited list of shares to monitor for the FPolicy policy to which the scope is applied.	<code>-shares-to-include</code> <i>share_name, ...</i>
<i>Shares to exclude</i> Specifies a comma-delimited list of shares to exclude from monitoring for the FPolicy policy to which the scope is applied.	<code>-shares-to-exclude</code> <i>share_name, ...</i>
<i>Volumes to include</i> Specifies a comma-delimited list of volumes to monitor for the FPolicy policy to which the scope is applied.	<code>-volumes-to-include</code> <i>volume_name, ...</i>

Type of information	Option
<i>Volumes to exclude</i> Specifies a comma-delimited list of volumes to exclude from monitoring for the FPolicy policy to which the scope is applied.	-volumes-to-exclude volume_name, ...
<i>Export policies to include</i> Specifies a comma-delimited list of export policies to monitor for the FPolicy policy to which the scope is applied.	-export-policies-to-include export_policy_name , ...
<i>Export policies to exclude</i> Specifies a comma-delimited list of export policies to exclude from monitoring for the FPolicy policy to which the scope is applied.	-export-policies-to-exclude export_policy_name , ...
<i>File extensions to include</i> Specifies a comma-delimited list of file extensions to monitor for the FPolicy policy to which the scope is applied.	-file-extensions-to-include file_extensions, ...
<i>File extension to exclude</i> Specifies a comma-delimited list of file extensions to exclude from monitoring for the FPolicy policy to which the scope is applied.	-file-extensions-to-exclude file_extensions, ...
<i>Is file extension check on directory enabled</i> Specifies whether the file name extension checks apply to directory objects as well. If this parameter is set to <code>true</code> , the directory objects are subjected to the same extension checks as regular files. If this parameter is set to <code>false</code> , the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match.	-is-file-extension-check-on-directories-enabled {true false }

Completing the FPolicy scope worksheet

You can use this worksheet to record the values that you need during the FPolicy scope configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the FPolicy scope.

You should record whether you want to include each parameter setting in the FPolicy scope configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage Virtual Machine (SVM) name	Yes	Yes	
Policy name	Yes	Yes	

Type of information	Required	Include	Your values
Shares to include	No		
Shares to exclude	No		
Volumes to include	No		
Volumes to exclude	No		
Export policies to include	No		
Export policies to exclude	No		
File extensions to include	No		
File extension to exclude	No		
Is file extension check on directory enabled	No		

Creating the FPolicy configuration

There are several steps you must perform to creating an FPolicy configuration. First, you must plan your configuration. Then, you create an FPolicy external engine, an FPolicy event, and an FPolicy policy. You then create an FPolicy scope and attach it to the FPolicy policy, and then enable the FPolicy policy.

FPolicy is supported on Storage Virtual Machines (SVMs) with FlexVol volumes. FPolicy is not supported on SVMs with Infinite Volume.

Steps

1. [Creating the FPolicy external engine](#) on page 167

The first step to creating an FPolicy configuration is to create an external engine. The external engine defines how FPolicy makes and manages connections to external FPolicy servers. If your configuration uses the native external engine for simple file blocking, you do not need to configure an external engine.

2. [Creating the FPolicy policy event](#) on page 167

As part of creating an FPolicy policy configuration, you need to create an FPolicy event. You associate the event with the FPolicy policy when it is created. An event defines which protocol to monitor and which file access events to monitor and filter.

3. [Creating the FPolicy policy](#) on page 168

After creating an FPolicy external engine and FPolicy events, you create the FPolicy policy. The policy associates an external engine and one or more events to the policy. The FPolicy policy also specifies whether mandatory screening is required and whether the external FPolicy servers (FPolicy servers) have privileged access to data on the Storage Virtual Machine (SVM).

4. *Creating the FPolicy policy scope* on page 168

After creating the FPolicy policy, you need to create an FPolicy scope. When creating the scope, you associate the scope with an FPolicy policy. A scope defines the boundaries on which the FPolicy policy applies. Scopes can include or exclude files based on shares, export policies, volumes, and file extensions.

5. *Enabling the FPolicy policy* on page 169

After you are through configuring an FPolicy policy configuration, you enable the FPolicy policy. Enabling the policy sets its priority and starts file access monitoring for the policy.

Creating the FPolicy external engine

The first step to creating an FPolicy configuration is to create an external engine. The external engine defines how FPolicy makes and manages connections to external FPolicy servers. If your configuration uses the native external engine for simple file blocking, you do not need to configure an external engine.

Before you begin

The external engine worksheet should be completed.

Steps

1. Create the FPolicy external engine:

```
vserver fpolicy policy external-engine create -vserver-name vserver_name
-engine-name engine_name -primary-servers IP_address,... -port integer -
ssl-option {no-auth|server-auth|mutual-auth} optional_parameters
```

2. Verify the FPolicy external engine configuration:

```
vserver fpolicy policy external-engine show -vserver vserver_name
```

Creating the FPolicy policy event

As part of creating an FPolicy policy configuration, you need to create an FPolicy event. You associate the event with the FPolicy policy when it is created. An event defines which protocol to monitor and which file access events to monitor and filter.

Before you begin

The FPolicy event worksheet should be completed.

Steps

1. Create the FPolicy event by using the following command:

```
vserver fpolicy policy event create -vserver-name vserver_name -event-
name event_name optional_parameters
```

2. Verify the FPolicy event configuration:

```
vserver fpolicy policy event show -vserver vserver_name
```

Creating the FPolicy policy

After creating an FPolicy external engine and FPolicy events, you create the FPolicy policy. The policy associates an external engine and one or more events to the policy. The FPolicy policy also specifies whether mandatory screening is required and whether the external FPolicy servers (FPolicy servers) have privileged access to data on the Storage Virtual Machine (SVM).

Before you begin

- The FPolicy policy worksheet should be completed.
- If you plan on configuring the policy to use FPolicy servers, the external engine must exist.
- At least one FPolicy event that you plan on associating with the FPolicy policy must exist.

Steps

1. Create the FPolicy policy by entering the following command:

```
vserver fpolicy policy create -vserver-name vserver_name -policy-name  
policy_name -events event_name,... -engine engine_name  
optional_parameters
```

2. Verify the FPolicy policy configuration:

```
vserver fpolicy policy show -vserver vserver_name
```

Creating the FPolicy policy scope

After creating the FPolicy policy, you need to create an FPolicy scope. When creating the scope, you associate the scope with an FPolicy policy. A scope defines the boundaries on which the FPolicy policy applies. Scopes can include or exclude files based on shares, export policies, volumes, and file extensions.

Before you begin

The FPolicy scope worksheet must be completed. The FPolicy policy must exist with an associated external engine (if the policy is configured to use external FPolicy servers) and must have at least one associated FPolicy event.

Steps

1. Create the FPolicy scope by entering the following command:

```
vserver fpolicy policy scope create -vserver-name vserver_name -policy-  
name policy_name optional_parameters
```

2. Verify the FPolicy scope configuration:

```
vserver fpolicy policy scope show -vserver vserver_name
```


Enabling the FPolicy policy

After you are through configuring an FPolicy policy configuration, you enable the FPolicy policy. Enabling the policy sets its priority and starts file access monitoring for the policy.

Before you begin

The FPolicy policy must exist with an associated external engine (if the policy is configured to use external FPolicy servers) and must have at least one associated FPolicy event. The FPolicy policy scope must exist and must be assigned to the FPolicy policy.

About this task

The priority is used when multiple policies are enabled on the Storage Virtual Machine (SVM) and more than one policy has subscribed to the same file access event. Policies that use the native engine configuration have a higher priority than policies for any other engine, regardless of the sequence number assigned to them when enabling the policy.

Note: A policy cannot be enabled on the admin SVM.

Steps

1. Enable the FPolicy policy by entering the following command:

```
vserver fpolicy enable -vserver-name vserver_name -policy-name  
policy_name -sequence-number integer
```

2. Verify that the FPolicy policy is enabled:

```
vserver fpolicy show -vserver vserver_name
```

Modifying FPolicy configurations

You can modify FPolicy configurations by modifying the elements that make up the configuration. You can modify external engines, FPolicy events, FPolicy scopes, and FPolicy policies. You can also enable or disable FPolicy policies. When you disable the FPolicy policy, file monitoring is discontinued for that policy.

It is recommended to disable the FPolicy policy before modifying the configuration.

Commands for modifying FPolicy configurations

You can modify FPolicy external engines, events, scopes, and policies.

If you want to modify...	Use this command...
External engines	<code>vserver fpolicy policy external-engine modify</code>

If you want to modify...	Use this command...
Events	<code>vserver fpolicy policy event modify</code>
Scopes	<code>vserver fpolicy policy scope modify</code>
Policies	<code>vserver fpolicy policy modify</code>

See the man pages for the commands for more information.

Enabling or disabling FPolicy policies

You can enable FPolicy policies after the configuration is complete. Enabling the policy sets its priority and starts file access monitoring for the policy. You can disable FPolicy policies if you want to stop file access monitoring for the policy.

Before you begin

Before enabling FPolicy policies, the FPolicy configuration must be completed.

About this task

- The priority is used when multiple policies are enabled on the Storage Virtual Machine (SVM) and more than one policy has subscribed to the same file access event.
- Policies that use the native engine configuration have a higher priority than policies for any other engine, regardless of the sequence number assigned to them when enabling the policy.
- If you want to change the priority of an FPolicy policy, you must disable the policy and then reenable it using the new sequence number.

Step

1. Perform the appropriate action:

If you want to...	Enter the following command...
Enable an FPolicy policy	<code>vserver fpolicy enable -vserver-name <i>vserver_name</i> - policy-name <i>policy_name</i> -sequence-number <i>integer</i></code>
Disable an FPolicy policy	<code>vserver fpolicy disable -vserver-name <i>vserver_name</i> - policy-name <i>policy_name</i></code>

Displaying information about FPolicy configurations

You might want to display information about FPolicy configurations to determine whether the configuration for each Storage Virtual Machine (SVM) is correct or to verify that an FPolicy policy configuration is enabled. You can display information about FPolicy external engines, FPolicy events, FPolicy scopes, and FPolicy policies.

How the show commands work

It is helpful when displaying information about the FPolicy configuration to understand how the show commands work.

A `show` command without additional parameters displays information in a summary form. Additionally, every `show` command has the same two mutually exclusive optional parameters, `-instance` and `-fields`.

When you use the `-instance` parameter with a `show` command, the command output displays detailed information in a list format. In some cases, the detailed output can be lengthy and include more information than you need. You can use the `-fields fieldname[,fieldname...]` parameter to customize the output so that it displays information only for the fields you specify. You can identify which fields that you can specify by entering `?` after the `-fields` parameter.

Note: The output of a `show` command with the `-fields` parameter might display other relevant and necessary fields related to the requested fields.

Every `show` command has one or more optional parameters that filter that output and enable you to narrow the scope of information displayed in command output. You can identify which optional parameters are available for a command by entering `?` after the `show` command.

The `show` command supports UNIX-style patterns and wildcards to enable you to match multiple values in command-parameters arguments. For example, you can use the wildcard operator (`*`), the NOT operator (`!`), the OR operator (`|`), the range operator (integer...integer), the less-than operator (`<`), the greater-than operator (`>`), the less-than or equal to operator (`<=`), and the greater-than or equal to operator (`>=`) when specifying values.

For more information about using UNIX-style patterns and wildcards, see the “Using the Data ONTAP command-line interface” section of the *Clustered Data ONTAP System Administration Guide for SVM Administrators*.

Commands for displaying information about FPolicy configurations

You use the `fpolicy show` commands to display information about the FPolicy configuration, including information about FPolicy external engines, events, scopes, and policies.

If you want to display information about FPolicy...	Use this command...
External engines	<code>vserver fpolicy policy external-engine show</code>
Events	<code>vserver fpolicy policy event show</code>
Scopes	<code>vserver fpolicy policy scope show</code>
Policies	<code>vserver fpolicy policy show</code>

See the man pages for the commands for more information.

Displaying information about FPolicy policy status

You can display information about the status for FPolicy policies to determine whether a policy is enabled, what external engine it is configured to use, what the sequence number is for the policy, and to which Storage Virtual Machine (SVM) the FPolicy policy is associated.

About this task

If you do not specify any parameter, the command displays the following information:

- SVM name
- Policy name
- Policy sequence number
- Policy status

In addition to displaying information about policy status for FPolicy policies configured on the cluster or a specific SVM, you can use command parameters to filter the command's output by other criteria.

You can specify the `-instance` parameter to display detailed information about listed policies. Alternatively, you can use the `-fields` parameter to display only the indicated fields in the command output, or `-fields ?` to determine what fields you can use.

Step

1. Display filtered information about FPolicy policy status by using the appropriate command:

If you want to display status information about policies...	Enter the command...
On the cluster	<code>vserver fpolicy show</code>
That have the specified status	<code>vserver fpolicy show -status {on off}</code>
On a specified SVM	<code>vserver fpolicy show -vserver <i>vserver_name</i></code>
With the specified policy name	<code>vserver fpolicy show -policy-name <i>policy_name</i></code>
With the specified sequence number	<code>vserver fpolicy show -sequence-number <i>integer</i></code>
That use the specified external engine	<code>vserver fpolicy show -engine <i>engine_name</i></code>

The following example displays the information about FPolicy policies on the cluster:

```
cluster1::> vserver fpolicy show
```

Vserver	Policy	Sequence Number	Status	Engine
-----	-----	-----	-----	-----
FPolicy	cserver_policy	-	off	eng1
vs1	vlp1	-	off	eng2
vs1	vlp2	-	off	native
vs1	vlp3	-	off	native
vs1	cserver_policy	-	off	eng1
vs2	vlp1	3	on	native
vs2	vlp2	1	on	eng3
vs2	cserver_policy	2	on	eng1

Displaying information about enabled FPolicy policies

You can display information about enabled FPolicy policies to determine what FPolicy external engine it is configured to use, what the priority is for the policy, and to which Storage Virtual Machine (SVM) the FPolicy policy is associated.

About this task

If you do not specify any parameters, the command displays the following information:

- SVM name
- Policy name
- Policy priority

You can use command parameters to filter the command's output by specified criteria.

Step

1. Display information about enabled FPolicy policies by using the appropriate command:

If you want to display information about enabled policies...	Enter the command...
On the cluster	vserver fpolicy show-enabled
On a specified SVM	vserver fpolicy show-enabled -vserver <i>vserver_name</i>
With the specified policy name	vserver fpolicy show-enabled -policy-name <i>policy_name</i>
With the specified sequence number	vserver fpolicy show-enabled -priority <i>integer</i>

The following example displays the information about enabled FPolicy policies on the cluster:

```
cluster1::> vserver fpolicy show-enabled
```

Vserver	Policy Name	Priority
-----	-----	-----
vs1	pol_native	native
vs1	pol_native2	native
vs1	pol1	2
vs1	pol2	4

Managing FPolicy server connections

You can manage your FPolicy server connections by connecting to external FPolicy servers, disconnecting from external FPolicy servers, or displaying information about connections and connection status.

Connecting to external FPolicy servers

To enable file processing, you might need to manually connect to an external FPolicy server if the connection has previously been terminated. A connection is terminated after the server timeout is reached or due to some error. Alternatively, the administrator might manually terminate a connection.

About this task

If a fatal error occurs, the connection to the FPolicy server can be terminated. After resolving the issue that caused the fatal error, you must manually reconnect to the FPolicy server.

Steps

1. Connect to the external FPolicy server by using the `vserver fpolicy engine-connect` command.

For more information about the command, see the man pages.

2. Verify that the external FPolicy server is connected by using the `vserver fpolicy show-engine` command.

For more information about the command, see the man pages.

Disconnecting from external FPolicy servers

You might need to manually disconnect from an external FPolicy server. This might be desirable if the FPolicy server has issues with notification request processing or if you need to perform maintenance on the FPolicy server.

Steps

1. Disconnect from the external FPolicy server by using the `vserver fpolicy engine-disconnect` command.

For more information about the command, see the man pages.

2. Verify that the external FPolicy server is disconnected by using the `vserver fpolicy show-engine` command.

For more information about the command, see the man pages.

Displaying information about connections to external FPolicy servers

You can display status information about connections to external FPolicy servers (FPolicy servers) for the cluster or for a specified Storage Virtual Machine (SVM). This information can help you determine which FPolicy servers are connected.

About this task

If you do not specify any parameter, the command displays the following information:

- SVM name
- Node name
- FPolicy policy name
- FPolicy server IP address
- FPolicy server status
- FPolicy server type

In addition to displaying information about FPolicy connections on the cluster or a specific SVM, you can use command parameters to filter the command's output by other criteria.

You can specify the `-instance` parameter to display detailed information about listed policies. Alternatively, you can use the `-fields` parameter to display only the indicated fields in the command output. You can enter `?` after the `-fields` parameter to find out which fields you can use.

Step

1. Display filtered information about connection status between the node and the FPolicy server by using the appropriate command:

If you want to display connection status information about...	Enter the command...
FPolicy servers that you specify	<code>vserver fpolicy show-engine -server <i>IP_address</i></code>
FPolicy servers for a specified SVM	<code>vserver fpolicy show-engine -vserver <i>vserver_name</i></code>
FPolicy servers that are attached with a specified policy	<code>vserver fpolicy show-engine -policy-name <i>policy_name</i></code>
FPolicy servers with the server status that you specify	<code>vserver fpolicy show-engine -server-status <i>status</i></code> The server status can be one of the following: <ul style="list-style-type: none"> • <code>connected</code> • <code>disconnected</code> • <code>connecting</code> • <code>disconnecting</code>
FPolicy servers with the specified type	<code>vserver fpolicy show-engine -server-type <i>type</i></code> The FPolicy server type can be one of the following: <ul style="list-style-type: none"> • <code>primary</code> • <code>secondary</code>
FPolicy servers that were disconnected with the specified reason	<code>vserver fpolicy show-engine -disconnect-reason <i>text</i></code> Disconnect can be due to multiple reasons. The following are common reasons for disconnect: <ul style="list-style-type: none"> • Disconnect command received from CLI. • Error encountered while parsing notification response from FPolicy server. • FPolicy Handshake failed. • SSL handshake failed. • TCP Connection to FPolicy server failed. • The screen response message received from the FPolicy server is not valid.

This example displays information about external engine connections to FPolicy servers on SVM vs1:

```
cluster1::> vserver fpolicy show-engine -vserver vs1
FPolicy          Server-      Server-
Vserver Policy    Node          Server      status      type
```


vs1	policy1	node1	1.1.1.1	connected	primary
-----	---------	-------	---------	-----------	---------

This example displays information only about connected FPolicy servers:

```
cluster1::> vserver fpolicy show-engine -fields server -server-status
connected
node          vserver policy-name server
-----
node1         vs1      policy1      1.1.1.1
```

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.
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.
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.
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 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	
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 <i>adapter</i> .
expansion slot	The slots on the storage system board into which you insert expansion cards.
F	
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.
G	
GID	See <i>Group ID (GID)</i> .
Group ID (GID)	The number used by UNIX systems to identify groups.
H	
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.

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.

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 *NMS*.

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.

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 gateway 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.
PCI	Peripheral Component Interconnect. The bus architecture used in newer storage system models.
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	
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.
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), Fibre Channel (FC), and Internet SCSI (iSCSI). The gateway treats an IBM volume as a disk.

- For IBM, the area on the storage array that is available for a gateway or non gateway host to read data from or write data to. The gateway documentation uses the term *array LUN* 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.

Copyright and trademark information

Copyright ©1994 - 2014 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2014 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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.A. 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

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 <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

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.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

8.3-format file names

creating [23](#)

A

access

how security types determine levels of client [45](#)

access cache

explained [86](#)

access control lists

See ACLs

access control lists (ACLs)

NFSv4, benefits of enabling [95](#)

NFSv4, managing [95](#)

access control lists)

See ACLs

access events

SMB file and folder, that can be audited [114](#)

access levels

how security types determine client [45](#)

access requests

mapping to anonymous [42](#)

ACEs

limit for NFSv4 ACLs [97](#)

ACLs

enabling or disabling modification of NFSv4 [96](#)

enabling or disabling NFSv4 [97](#)

limit of ACEs for NFSv4 [97](#)

NFSv4, how they work [95](#)

ACLs (access control lists)

NFSv4, benefits of enabling [95](#)

NFSv4, managing [95](#)

adding

rules to export policies [49](#)

users to local UNIX groups [77](#)

aggregates

space considerations when staging volumes are created by enabled auditing subsystem [112](#)

anonymous

mapping clients to [42](#)

anonymous access

how to configure with export rules [46](#)

APIs

supported VMware vStorage, for NFS [105](#)

architectures

typical NAS namespace [16](#)

assigning

export policies to qtrees [54](#)

asynchronous

FPolicy applications [139](#)

FPolicy communication notifications, defined [138](#)

audit event logs

manually rotating [131](#)

audit policies

configuring using the Windows Security tab [123](#)

displaying using the Windows Security tab [127](#)

NTFS, how to configure using the Data ONTAP CLI [126](#)

using the Data ONTAP CLI to display information about NTFS [128](#)

audit-enabled SVMs

actions you must take before revert [135](#)

auditing

actions you must take on audit-enabled SVMs before revert [135](#)

actions you must take prior to revert [135](#)

aggregate space considerations when enabling [112](#)

commands for modifying configuration [134](#)

configuring for NFS [126](#)

consolidation and conversion tasks [110](#)

creating configuration [120](#)

creating file and directory, configuration [119](#)

deleting configuration [134](#)

displaying information about configuration [132](#)

displaying information about NTFS audit policies

using the Data ONTAP CLI [128](#)

enabling and disabling on SVMs [131](#)

enabling on the SVM [121](#)

event log consolidation [111](#)

event log consolidation when a node is unavailable [111](#)

event log rotation [111](#)

how implementing is a two-step process [122](#)

how staging volumes are created on aggregates [112](#)

how the Data ONTAP process works [111](#)

how to troubleshoot event log volume space issues [136](#)

how to troubleshoot staging volume space issues [136](#)

list of NFS events [115](#)

manually converting the audit event logs [131](#)

NFS and SMB file and folder access [110](#)

- partial event log consolidation [111](#)
- planning the configuration [116](#)
- process when enabling or disabling [111](#)
- requirements and considerations for configuring [113](#)
- revert process when there are audit-enabled SVMs [135](#)
- SMB file and folder access events that can be audited [114](#)
- staging files and volumes [110](#)
- supported audit event log formats [113](#)
- verifying configuration [122](#)
- verifying that it is enabled [132](#)
- viewing audit event logs [114](#)
- auditing for NFS [126](#)
- authentication
 - how Data ONTAP handles NFS client [25](#)
 - Kerberos [56](#)
- authentication-based
 - restrictions [13](#)

B

- basic concepts
 - introduction to how Data ONTAP secures LDAP communication using LDAP over SSL/TLS [63](#)
- best practices
 - FPolicy setup [146](#)
- bits
 - how Data ONTAP treats read-only [90](#)
- breaking
 - locks [93](#)

C

- CA certificates
 - installing self-signed root, on the SVM [65](#)
- case-sensitivity
 - of file names [23](#)
- certificates
 - installing self-signed root, on the SVM [65](#)
- CIFS
 - file naming dependencies [23](#)
 - how Data ONTAP grants file access from NFS clients [25](#)
- client access
 - how security types determine levels of [45](#)
- client authentication
 - how Data ONTAP handles [25](#)
- client configurations
 - creating LDAP [66](#)

- client schema templates
 - commands for managing LDAP [85](#)
- clients
 - validating qtree IDs for file operations [56](#)
- clusters
 - role with FPolicy implementations [139](#)
- commands
 - for managing LDAP client schema templates [85](#)
 - for modifying SVM auditing configurations [134](#)
 - name mapping management [82](#)
- concepts
 - introduction to how Data ONTAP secures LDAP communication using LDAP over SSL/TLS [63](#)
- configuration requirements
 - LIF file access management [14](#)
- configuration types
 - FPolicy, defined [144](#)
- configurations
 - creating Kerberos realm [61](#)
 - creating LDAP client [66](#)
- configuring
 - audit policies using the Windows Security tab [123](#)
 - auditing [120](#)
 - default users [74](#)
 - FPolicy [166](#)
 - local UNIX users and groups [75](#)
 - NIS domains [84](#)
 - security style on FlexVol volumes [35](#)
 - security style on qtrees [36](#)
 - security style on SVM root volumes [35](#)
 - SVMs to use LDAP [62](#)
- connecting
 - external FPolicy servers [174](#)
- connection credentials
 - FPolicy, how used with privileged data access channels [140](#)
- considerations
 - aggregate space, for staging volumes when enabling auditing [112](#)
 - auditing configuration [113](#)
 - for FPolicy before reverting [146](#)
- consolidation task
 - auditing [110](#)
- control channels
 - how FPolicy uses [140](#)
- conversion task
 - auditing [110](#)
- copying
 - export policies [89](#)
 - LDAP client schema templates [85](#)

creating

- auditing configuration [120](#)
- export policies [48](#)
- export rules [89](#)
- file and directory auditing configuration [119](#)
- file names [23](#)
- FPolicy configurations [166](#)
- FPolicy events [167](#)
- FPolicy external engines [167](#)
- FPolicy policies [168](#)
- FPolicy scopes [168](#)
- Kerberos realm configurations [61, 88](#)
- LDAP client configurations [66](#)
- LDAP client configurations, command for [85](#)
- LDAP configurations [85](#)
- local UNIX groups [76, 83](#)
- local UNIX users [75, 83](#)
- name mappings [73](#)
- new LDAP client schema [65](#)
- NFS servers [39](#)
- NIS domain configuration [79](#)
- NIS domains [84](#)

D

data access

- introduction to how security styles affect [19](#)

data access channels

- how FPolicy connection credentials are used with privileged [140](#)
- how FPolicy uses privileged [140](#)

data LIFs

- how control channels are used with FPolicy communication [140](#)
- how FPolicy handle migrations and failovers for [143](#)
- role with FPolicy implementations [139](#)

Data ONTAP

- how the auditing process works [111](#)

Data ONTAP CLI

- how to configure NTFS audit policies using [126](#)

default users

- configuring [74](#)

definitions

- FPolicy [138](#)

deleting

- audit configuration [134](#)
- export policies [89](#)
- export rules [89](#)
- Kerberos realm configurations [88](#)
- LDAP client configurations, command for [85](#)

LDAP client schema templates [85](#)

LDAP configurations [85](#)

local UNIX groups [83](#)

local UNIX users [83](#)

name mappings, command for [82](#)

NFS servers [82](#)

NIS domains [84](#)

users from local UNIX groups [83](#)

disabling

auditing on SVMs [131](#)

FPolicy policies [170](#)

NFSv3 [38](#)

NFSv4 [38](#)

NFSv4 referrals [103](#)

NFSv4.1 [38](#)

parallel NFS [39](#)

pNFS [39](#)

rquota support [106](#)

vStorage over NFS [105](#)

disconnecting

from external FPolicy servers [175](#)

displaying

audit policy information using the Windows Security tab [127](#)

export policies [89](#)

export rules [89](#)

FPolicy configuration, commands for [171](#)

FPolicy configuration, how show commands work when [171](#)

information about auditing configurations [132](#)

information about connections to FPolicy servers [175](#)

information about enabled FPolicy policies [173](#)

information about FPolicy configurations [170](#)

information about FPolicy policy status [172](#)

information about locks [91](#)

Kerberos realm configurations [88](#)

LDAP client configurations, command for [85](#)

LDAP client schema templates [85](#)

LDAP configurations [85](#)

local UNIX groups [83](#)

local UNIX users [83](#)

name mappings, command for [82](#)

NFS Kerberos configurations, information about [87](#)

NFS servers [82](#)

NFS statistics [104](#)

NIS domains [84](#)

NTFS auditing information on FlexVol volumes using the Data ONTAP CLI [128](#)

volume mount and junction point information [33](#)

E

effects on file permissions [20](#)

enabling

- auditing on SVMs [131](#)
- auditing on the SVM [121](#)
- FPolicy policies [169](#), [170](#)
- IPv6 for NFS [79](#)
- LDAP on SVMs [68](#)
- NFSv3 [38](#)
- NFSv4 [38](#)
- NFSv4 referrals [103](#)
- NFSv4.1 [38](#)
- parallel NFS [39](#)
- pNFS [39](#)
- rquota support [106](#)
- vStorage over NFS [105](#)

event log formats

- support for EVT X file format [113](#)
- support for XML file format [113](#)

event logs

- manually rotating audit [131](#)
- supported file formats for audit [113](#)
- viewing audit [114](#)

events

- creating FPolicy [167](#)
- information to gather for configuring FPolicy [160](#)
- planning the configuration for FPolicy [154](#)
- SMB file and folder access, that can be audited [114](#)
- supported combinations of file operations and filters that FPolicy can monitor for NFSv3 [158](#)
- supported combinations of file operations and filters that FPolicy can monitor for NFSv4 [159](#)
- supported combinations of file operations and filters that FPolicy can monitor for SMB [158](#)

EVT X

- file format, viewing audit event logs with [114](#)
- supported audit event log file format [113](#)

exchanging

- name mappings, command for [82](#)

export policies

- adding rules to [49](#)
- assigning to qtrees [54](#)
- associating with a FlexVol volume [53](#)
- creating [48](#)
- default, for SVMs [40](#)
- how they control client access to qtrees [40](#)
- how they control client access to volumes [40](#)
- managing [89](#)
- removing from qtrees [55](#)

restrictions and nested junctions [56](#)

setting index numbers for rules [52](#)

export rules

- how they work [41](#)
- how to configure anonymous access [46](#)
- how to configure superuser access [46](#)
- managing [89](#)

exporting

- qtrees [54](#)

external engines

- creating FPolicy [167](#)
- information to gather for configuring FPolicy [153](#)
- planning the configuration for FPolicy [148](#)

external FPolicy servers

- configuration type defined [144](#)
- connecting to [174](#)
- disconnecting from [175](#)
- displaying information about connections to [175](#)
- how FPolicy works with external FPolicy servers [140](#)
- when to create FPolicy configurations that use [145](#)

F

file access

- how Data ONTAP controls [13](#)
- LIF configuration requirements for managing [14](#)
- NFS, managing [81](#)
- setting up for NFS [30](#)
- to Infinite Volumes, where to find information about setting up for NFS [80](#)

file access events

- SMB, that can be audited [114](#)

file and directory auditing

- creating configuration on SVMs [119](#)

file and folder access

- auditing NFS and SMB [110](#)

file and record locking

- NFSv4, described [101](#)

file delegations

- enabling or disabling NFSv4 read [99](#)
- enabling or disabling NFSv4 write [100](#)
- how they work for NFSv4 [98](#)
- NFSv4, managing [98](#)

file formats

- viewing audit event logs with XML or EVT X [114](#)

file locking

- between protocols, explained [90](#)

file locks

- breaking [93](#)

- displaying information about [91](#)
- introduction to managing [90](#)
- file names
 - case-sensitivity [23](#)
 - creating [23](#)
 - dependencies for NFS and CIFS [23](#)
 - valid characters for [23](#)
- file operations
 - supported combinations of file operations and filters for NFSv4 FPolicy events [159](#)
 - supported combinations of file operations and filters for SMB FPolicy events [158](#)
 - supported combinations with filters for NFSv3 FPolicy events [158](#)
 - validating qtree IDs for [56](#)
- file permissions
 - effect of security styles on [20](#)
- file-based
 - restrictions [14](#)
- filters
 - supported combinations of file operations and filters for NFSv4 FPolicy events [159](#)
 - supported combinations of file operations and filters for SMB FPolicy events [158](#)
 - supported combinations with file operations for NFSv3 FPolicy events [158](#)
- FlexVol volumes
 - associating export policy to [53](#)
 - configuring security style on [35](#)
- folder access events
 - SMB, that can be audited [114](#)
- FPolicy
 - commands for displaying configuration information [171](#)
 - commands for modifying configurations [169](#)
 - configuration types defined [144](#)
 - connecting to external FPolicy servers [174](#)
 - creating events [167](#)
 - creating external engines [167](#)
 - creating scopes [168](#)
 - definition [138](#)
 - disconnecting from external FPolicy servers [175](#)
 - events, supported combinations of file operations and filters for NFSv3 [158](#)
 - events, supported combinations of file operations and filters for NFSv4 [159](#)
 - how communications are managed with node failovers [143](#)
 - how services work across SVM namespaces [143](#)
 - important revert considerations [146](#)
 - parts explained [138](#)
 - protocols that can be monitored [138](#)
 - roles that cluster components play with [139](#)
 - super user credentials for privileged data access [141](#)
 - synchronous and asynchronous applications [139](#)
 - synchronous and asynchronous communication defined [138](#)
- FPolicy best practices
 - for setup [146](#)
- FPolicy configuration types
 - when to create a native FPolicy configuration [144](#)
 - when to create configurations that use external FPolicy servers [145](#)
- FPolicy configurations
 - creating [166](#)
 - displaying information about [170](#)
 - how show commands work when displaying information about [171](#)
 - information about requirements, considerations, and best practices [145](#)
 - overview of configuration planning [148](#)
 - steps to setup [147](#)
- FPolicy connections
 - displaying information about server connections [175](#)
 - FPolicy connection management responsibilities when connecting to external FPolicy servers [140](#)
 - how connection credentials are used with privileged data access channels [140](#)
 - how control channels are used with [140](#)
 - how data LIF migrations and failovers are handled [143](#)
 - how privileged data access channels are used [140](#)
 - what the node-to-external FPolicy server communication process is [142](#)
- FPolicy events
 - information to gather for configuring [160](#)
 - planning the configuration for [154](#)
 - supported combinations of file operations and filters that FPolicy can monitor for SMB [158](#)
- FPolicy external engines
 - information to gather for configuring [153](#)
 - planning the configuration for [148](#)
- FPolicy external servers
 - what they do [138](#)
- FPolicy framework
 - what it does [138](#)
- FPolicy policies
 - creating [168](#)
 - displaying information about enabled [173](#)
 - displaying information about status [172](#)

- enabling [169](#)
- enabling or disabling [170](#)
- how FPolicy manages processing multiple [141](#)
- information to gather for configuration [162](#)
- planning the configuration for [160](#)

FPolicy recommendations

- for setup [146](#)

FPolicy requirements

- for setup [145](#)

FPolicy scopes

- configuration information to gather [165](#)
- planning the configuration for [163](#)

FPolicy servers

- connecting to [174](#)
- disconnecting from [175](#)
- displaying information about connections to [175](#)
- how FPolicy works with external FPolicy servers [140](#)
- what the communication process to nodes is [142](#)
- when to create FPolicy configurations that use external [145](#)

G

glossary [178](#)

group IDs

- limitation for NFS RPCSEC_GSS [57](#)

groups

- local, UNIX, configuring [75](#)
- UNIX, adding users to local [77](#)
- UNIX, creating local [76](#)
- UNIX, loading local from URIs [77](#)

guaranteed auditing

- how Data ONTAP ensures [111](#)

H

hard mounts [24](#)

I

implement

- auditing, two-steps to implement [122](#)

Infinite Volumes

- where to find information about NFS support on [27](#)
- where to find information about setting up NFS file access to [80](#)
- where to get information about security styles of [35](#)

inserting

- name mappings, command for [82](#)

installing

- self-signed root CA certificate on the SVM [65](#)

IPv6

- enabling for NFS [79](#)
- NFS support for [79](#)

J

junction points

- creating volumes with specified [30](#)
- creating volumes without specified [31](#)
- displaying information about volume [33](#)
- volume, how used to create namespaces [15](#)

junctions

- defined [15](#)
- volume, how used in SMB and NFS namespaces [16](#)

K

Kerberos

- authentication [56](#)
- creating configuration for SVMs [62](#)
- creating realm configurations [61](#)
- displaying configuration information for NFS [87](#)
- modifying configuration for NFS servers [88](#)
- realms, managing [88](#)
- requirements for configuring with NFS [57](#)

L

LDAP

- commands for managing [85](#)
- commands for managing client configurations [85](#)
- commands for managing client schema templates [85](#)
- configuring SVMs to use [62](#)
- creating client configurations [66](#)
- creating new client schema [65](#)
- enabling on SVMs [68](#)

LDAP over SSL/TLS

- installing self-signed root CA certificate on the SVM [65](#)
- introduction to configuring and using to secure communication [63](#)
- introduction to how Data ONTAP uses to secure LDAP communication [63](#)

LIFs

- configuration requirements for file access management [14](#)
- data, role with FPolicy implementations [139](#)

- how FPolicy handle migrations and failovers for data [143](#)
- limitations
 - of Data ONTAP support for NFSv4 [26](#)
 - of group IDs for NFS RPCSEC_GSS [57](#)
- limits
 - of ACEs for NFSv4 ACLs [97](#)
- loading
 - local UNIX groups from URIs [77](#)
 - local UNIX users from URIs [75](#)
- locking grace period
 - NFSv4, specifying [102](#)
- locking lease period
 - specifying NFSv4 [101](#)
- locks
 - breaking [93](#)
 - displaying information about [91](#)
- logs
 - manually rotating audit logs [131](#)

M

- managing
 - export policies [89](#)
 - export rules [89](#)
 - Kerberos realm configurations [88](#)
 - LDAP client configurations, commands for [85](#)
 - LDAP client schema templates [85](#)
 - LDAP configurations [85](#)
 - local UNIX groups [83](#)
 - local UNIX users [83](#)
 - name mappings, commands for [82](#)
 - NFS servers [82](#)
- manually rotating
 - audit event logs [131](#)
- modifying
 - auditing configurations, commands for [134](#)
 - export rule index numbers [52](#)
 - export rules [89](#)
 - FPolicy configurations, commands for [169](#)
 - Kerberos realm configurations [88](#)
 - LDAP client configurations, command for [85](#)
 - LDAP client schema templates [85](#)
 - LDAP configurations [85](#)
 - local UNIX groups [83](#)
 - local UNIX users [83](#)
 - name mapping patterns, command for [82](#)
 - NFS Kerberos configuration [88](#)
 - NFS servers [82](#)
 - NFSv3 TCP maximum read and write size [108](#)

- NFSv4 ACLs, enabling or disabling capability of [96](#)
- NIS domains [84](#)
- protocols for SVMs [37](#)
- server implementation ID [94](#)
- mount requests
 - controlling NFS, from nonreserved ports [81](#)
- mounting
 - NFS exports using nonreserved ports [82](#)
 - volumes in NAS namespaces [32](#)
- mounts [24](#)

N

- name mapping
 - configuring SVMs to use LDAP [62](#)
 - explained [69](#)
- name mappings
 - commands for managing [82](#)
 - conversion rules [72](#)
 - creating [73](#)
 - how used [68](#)
 - multidomain searches for UNIX user to Windows user [70](#)
- name services
 - configuring SVMs to use LDAP [62](#)
- names
 - valid characters for file [23](#)
- namespaces
 - defined [15](#)
 - how FPolicy works across SVM [143](#)
 - how volume junctions are used for NAS access [16](#)
 - introduction to creating and managing data volumes in NAS [30](#)
 - mounting or unmounting volumes within NAS [32](#)
 - typical architectures for NAS [16](#)
- NAS
 - creating volumes with specified junction points [30](#)
 - creating volumes without specified junction points [31](#)
 - displaying volume mount and junction point information [33](#)
 - mounting or unmounting volumes in the namespace [32](#)
 - typical namespace architectures [16](#)
- NAS namespaces
 - introduction to creating and managing data volumes in [30](#)
- native FPolicy configurations
 - when to create [144](#)
- native FPolicy servers

- configuration type defined [144](#)
 - netgroups
 - loading into SVMs from URIs [78](#)
 - verifying status of definitions [83](#)
 - NFS
 - benefits of enabling v4 ACLs [95](#)
 - clients, how Data ONTAP grants CIFS file access from [25](#)
 - configuring auditing [126](#)
 - controlling requests from nonreserved ports [81](#)
 - displaying statistics [104](#)
 - enabling IPv6 for [79](#)
 - enabling or disabling v3 [38](#)
 - enabling or disabling v4 [38](#)
 - enabling or disabling v4 ACLs [97](#)
 - enabling or disabling v4 read file delegations [99](#)
 - enabling or disabling v4 write file delegations [100](#)
 - enabling or disabling v4.1 [38](#)
 - events that can be audited [115](#)
 - file access setup for Infinite Volumes, where to find information about [80](#)
 - file locking between protocols explained [90](#)
 - file naming dependencies [23](#)
 - group ID limitation for RPCSEC_GSS [57](#)
 - how Data ONTAP handles client authentication [25](#)
 - how Data ONTAP treats read-only bits [90](#)
 - how v4 ACLs work [95](#)
 - Infinite Volume support, where to find information about [27](#)
 - Kerberos configuration, displaying information about [87](#)
 - Kerberos configuration, modifying [88](#)
 - managing file access [81](#)
 - managing v4 ACLs [95](#)
 - modifying protocols for SVMs [37](#)
 - mounting exports using nonreserved ports [82](#)
 - process to access NTFS security style data [28](#)
 - process to access UNIX security style data [28](#)
 - requirements for configuring with Kerberos [57](#)
 - setting up file access for [30](#)
 - specifying user ID domain for v4 [60](#)
 - support for parallel [27](#)
 - support for v4.1 [27](#)
 - support for, over IPv6 [79](#)
 - supported versions and clients [26](#)
 - using with Kerberos [56](#)
 - v4 file and record locking described [101](#)
 - v4 support [26](#)
 - v4, determining file deletion [97](#)
 - v4, how file delegations work [98](#)
 - v4, limitations of Data ONTAP support [26](#)
 - v4, managing file delegations [98](#)
 - v4, specifying locking grace period [102](#)
 - v4, specifying locking lease period [101](#)
 - NFS exports
 - how volume junctions are used with [16](#)
 - NFS servers
 - creating [39](#)
 - managing [82](#)
 - NFSv3
 - improving performance for TCP [107](#)
 - modifying TCP maximum read and write size [108](#)
 - NFSv4
 - ACLs, enabling or disabling modification of [96](#)
 - enabling or disabling referrals [103](#)
 - how referrals work [102](#)
 - NIS domain
 - configuring [84](#)
 - creating [79](#), [84](#)
 - deleting [84](#)
 - displaying [84](#)
 - modifying [84](#)
 - nodes
 - how FPolicy handles communications with failovers [143](#)
 - what the communication process is for FPolicy-enabled [142](#)
 - NTFS
 - how to use the Data ONTAP CLI to configure audit policies for [126](#)
- ## O
- operations
 - validating qtree IDs for file [56](#)
- ## P
- parallel NFS
 - enabling or disabling [39](#)
 - performance
 - improving for NFSv3 TCP [107](#)
 - permissions
 - effect of security styles on file [20](#)
 - how Data ONTAP preserves UNIX [22](#)
 - UNIX, how to manage using Windows Security tab [22](#)
 - planning
 - auditing configuration [116](#)
 - FPolicy configuration overview [148](#)

- FPolicy event configuration [154](#)
- FPolicy external engine configurations [148](#)
- FPolicy policy configurations [160](#)
- FPolicy scope configurations [163](#)

pNFS

- enabling or disabling [39](#)
- support for [27](#)

policies

- adding rules to export [49](#)
- assigning export, to qtrees [54](#)
- creating export [48](#)
- creating FPolicy [168](#)
- displaying information about enabled FPolicy [173](#)
- enabling FPolicy [169](#)
- enabling or disabling FPolicy [170](#)
- FPolicy, information to gather for configuration [162](#)
- how FPolicy manages processing multiple FPolicy [141](#)
- planning the configuration for FPolicy [160](#)
- removing export, from qtrees [55](#)
- using the Data ONTAP CLI to display information about NTFS audit [128](#)

ports

- controlling NFS mount requests from nonreserved [81](#)

preferred trusted domains

- how used with multidomain searches for user name mapping [70](#)

priorities

- how FPolicy manages processing FPolicy policy [141](#)

privileged data access

- super user credentials for FPolicy [141](#)

protocols

- file locking between, explained [90](#)
- modifying for SVMs [37](#)
- supported [13](#)
- that FPolicy can monitor [138](#)

Q

qtrees

- assigning an export policy to [54](#)
- configuring security style on [36](#)
- exporting [54](#)
- how export policies control client access to [40](#)
- removing an export policy from [55](#)
- validating IDs for file operations [56](#)

R

read file delegations

- enabling or disabling NFSv4 [99](#)

read-only bits

- how Data ONTAP treats [90](#)

realm configurations

- creating Kerberos [61](#)

recommendations

- FPolicy setup [146](#)

referrals

- enabling or disabling NFSv4 [103](#)
- how they work for NFSv4 [102](#)

removing

- export policies from qtrees [55](#)

renaming

- export policies [89](#)

requirements

- auditing configuration [113](#)
- FPolicy setup [145](#)
- Kerberos with NFS configuration [57](#)

restrictions

- authentication-based [13](#)
- file-based [14](#)

reverting

- important FPolicy considerations before [146](#)
- process when there are audit-enabled SVMs [135](#)

root CA certificates

- installing on the SVM [65](#)

root volumes

- configuring security style on SVM [35](#)

rotating

- audit event logs, manually [131](#)

rquotas

- enabling or disabling [106](#)

S

schema templates

- commands for managing LDAP client [85](#)

schemas

- creating new LDAP client [65](#)

scopes

- configuration information to gather for FPolicy [165](#)
- creating FPolicy [168](#)
- planning the configuration for FPolicy [163](#)

security

- Kerberos with NFS [56](#)

security styles

- configuring on FlexVol volumes [35](#)

- configuring on qtrees [36](#)
 - configuring on SVM root volumes [35](#)
 - how inheritance works [21](#)
 - how to choose [21](#)
 - introduction to how they affect data access [19](#)
 - when and where to set [21](#)
 - security types
 - how client access levels are determined by [45](#)
 - how to handle clients with unlisted [42](#)
 - self-signed root CA certificates
 - installing on the SVM [65](#)
 - server implementation ID
 - modifying [94](#)
 - servers
 - creating NFS [39](#)
 - show commands
 - how they work when displaying FPolicy configuration [171](#)
 - SMB
 - file and folder access events that can be audited [114](#)
 - file locking between protocols explained [90](#)
 - how Data ONTAP treats read-only bits [90](#)
 - modifying protocols for SVMs [37](#)
 - SMB shares
 - how volume junctions are used with [16](#)
 - soft mounts [24](#)
 - staging files
 - auditing [110](#)
 - staging volumes
 - aggregate space considerations when enabling auditing for [112](#)
 - auditing [110](#)
 - statistics
 - displaying NFS [104](#)
 - super user credentials
 - for FPolicy privileged data access [141](#)
 - superuser access
 - how to configure with export rules [46](#)
 - supported protocols [13](#)
 - SVMs
 - actions you must take before revert when there are audit-enabled [135](#)
 - auditing NAS file access events [110](#)
 - commands for modifying auditing configurations [134](#)
 - configuring security style on root volume [35](#)
 - creating a file and directory auditing configuration on [119](#)
 - creating Kerberos configuration for [62](#)
 - creating the FPolicy policy [168](#)
 - default export policy for [40](#)
 - deleting an auditing configuration [134](#)
 - enabling and disabling auditing on [131](#)
 - enabling auditing on [121](#)
 - enabling LDAP on [68](#)
 - how FPolicy manages processing policies [141](#)
 - how FPolicy works across namespaces [143](#)
 - installing root CA self-signed certificate for LDAP over SSL/TLS on [65](#)
 - loading netgroups into [78](#)
 - modifying protocols [37](#)
 - revert process when there are audit-enabled [135](#)
 - role with FPolicy implementations [139](#)
 - using FPolicy for file monitoring and management [138](#)
 - SVMs with FlexVol volumes
 - default export policy for [40](#)
 - synchronous
 - communication, how privileged data access channels are used with [140](#)
 - FPolicy applications [139](#)
 - FPolicy notifications, defined [138](#)
- ## T
- TCP
 - modifying maximum read and write size for NFSv3 [108](#)
 - modifying NFSv3 maximum read and write size [107](#)
 - templates
 - commands for managing LDAP schema [85](#)
 - terminology
 - LDAP over SSL/TLS [63](#)
 - troubleshooting
 - auditing event log volume space issues [136](#)
 - staging volume space issues [136](#)
 - trusted domains
 - discovered, how used with multidomain searches for user name mapping [70](#)
- ## U
- UNIX
 - adding users to local groups [77](#)
 - configuring local users and groups [75](#)
 - creating local groups [76](#)
 - creating local users [75](#)
 - loading local groups from URIs [77](#)
 - loading local users from URIs [75](#)
 - managing local groups [83](#)

- managing local users [83](#)
- UNIX permissions
 - how Data ONTAP preserves [22](#)
 - how to manage using Windows security tab [22](#)
- unmounting
 - volumes in NAS namespaces [32](#)
- URIs
 - loading local UNIX groups from [77](#)
 - loading local UNIX users from [75](#)
 - loading netgroups into SVMs from [78](#)
- user ID domains
 - specifying for NFSv4 [60](#)
- user name mappings
 - preferred trusted domains used with multidomain search for [70](#)
- users
 - adding to local UNIX groups [77](#)
 - local, UNIX, configuring [75](#)
 - UNIX, creating local [75](#)
 - UNIX, loading local from URIs [75](#)

V

- verifying
 - auditing configuration [122](#), [132](#)
 - status of netgroup definitions [83](#)
- viewing
 - audit event logs [114](#)
- VMware
 - enabling or disabling vStorage over NFS [105](#)
 - supported vStorage APIs for NFS [105](#)
- volume junctions
 - defined [15](#)
 - how used in SMB and NFS namespaces [16](#)
- volumes
 - aggregate space considerations when enabling auditing for staging [112](#)

- associating export policy to FlexVol [53](#)
 - configuring security style on FlexVol [35](#)
 - creating with specified junction points [30](#)
 - creating without specified junction points [31](#)
 - displaying mount and junction point information [33](#)
 - how export policies control client access to [40](#)
 - how junction points are used to create namespaces with [15](#)
 - introduction to creating and managing in NAS namespaces [30](#)
 - mounting and unmounting in NAS namespaces [32](#)
- vStorage
 - enabling or disabling over NFS [105](#)
 - supported APIs for NFS [105](#)

W

- worksheets
 - for recording information needed to configure FPolicy events [160](#)
 - for recording information needed to configure FPolicy external engines [153](#)
 - for recording information needed to configure FPolicy policies [162](#)
 - for recording information needed to configure FPolicy scopes [165](#)
- write file delegations
 - enabling or disabling NFSv4 [100](#)

X

- XML
 - file format, viewing audit event logs with [114](#)
 - supported audit event log file format [113](#)



NA 210-06380_A0, Printed in USA

SC27-6604-00

